occupies 7.4 mL of volume at a mass of 23 g, while holding about 0.003 kWh of energy. How much volume and how heavy would a collection of AA batteries be in order to match the energy in a liter of gasoline, and by what factors (in volume and mass) is gasoline superior?

17. A gallon of gasoline contains about 37 kWh of energy and costs about $4, while a typical AA battery holds about 0.003 kWh and costs about $0.50 each, in bulk. By what factor are batteries more expensive, for the same amount of energy?

18. Putting the cheapness of fossil fuels into perspective, a gallon of gasoline purchased for $4 might deliver 6 kWh of mechanical energy after accounting for efficiency of the associated engine. A laborer might be expected to export 100 W of mechanical power, on average, and be limited to 8 hours per day. How many hours would it take for the laborer to accomplish the equivalent output of a gallon of gasoline? At a rate of $15/hr, how much will this cost you?

19. A number of attempts to estimate the energy investment in our food arrive at the conclusion that every kcal of food we eat took 10 kcal[85] of fossil fuel input energy, so that we are effectively eating our fossil fuels! As a sanity check, what fraction of our fossil fuel energy would have to go into food production in the U.S. if diets are typically 2,000 kcal/day and we use fossil fuels at a rate of 8,000 W?[86] Does the answer seem plausible?

20. List at least five ways in which your life benefits from fossil fuels.

21. List at least three negative impacts of fossil fuels that most concern you (or explain why not, if they don't concern you).

22. Which of the possibilities from the list on page 125 (or combination thereof) seem most likely to play out, to your mind? Explain what makes you think so.

23. Let's say that Earth was originally endowed with one million flerbits,[87] and that we have already used up 400,000 of them. We currently extract 15,000 per year. How long does the R/P ratio suggest the resource will last?

24. Proven remaining reserves of oil, gas, and coal are 10, 8, and 20 ZJ,[88] while we have used 8, 4, and 8 ZJ of each. What fraction of the *original* total fossil fuel resource have we already used, then?

25. Explain, both in practical and mathematical terms, why the R/P ratio overestimates the time remaining for a resource if the rate of production (use) of that resource is continually increasing.

26. It is hard for many people to appreciate that fossil fuels will not just

85: Before fossil fuels—when food-driven muscle power was used instead—if we put in more energy than we got out, we would have starved and died out.

86: 80% of the 10,000 W American energy rate per person

87: . . . a made-up thing that we presume is irreplaceable

88: ZJ is zetta-Joules, or 10^{21} J.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

"run out one day," because they don't appreciate the substantial amount of work that must go into extracting the resource from a reluctant ground. What common, day-to-day personal experiences do you imagine contributes to this disconnect?[89]

89: In other words, what is your personal experience obtaining resources that can suddenly just run out?

27. Simultaneously considering Table 8.5 and the lessons from Figure 8.6, what sort of timescale might you guess for when the world might see a more-or-less permanent downturn in oil production? Say how many years you think we have until a downturn and explain your reasoning.

28. What analogy from everyday life can you think of that would help someone understand the idea that extraction of oil from the ground must be preceded by exploration and discovery of the resource, and that we can't produce more than we discover?

29. Explain why the area under the red curve in Figure 8.7 or Figure 8.9 cannot be larger than the total area under the blue curve.[90]

90: Note that the shaded blue area is not the entire area under the blue curve, but has been set to equal the red area.

30. If the inevitable decline in fossil fuel availability is a potentially important disrupter of the status quo in the decades to come, what are some reasons it gets little attention compared to, say, climate change? No right answer here, but what do you think contributes?

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.



9 Climate Change

Climate change stands tall among the global scale problems created by our energy appetite—caused by the accumulation of carbon dioxide (CO_2) in our atmosphere from the burning of fossil fuels. This chapter aims to provide a no-nonsense account of the basis for climate change that leaves little room for the kind of uncertainty often injected by (alarmingly successful) disinformation campaigns. While the *response* of our complex climate system is more difficult to predict in detail, the core physics is unassailable. We will see that the rise in CO_2 is not at all mysterious, stemming from fossil fuels. We will also explore a few scenarios and connect the CO_2 rise to temperature consequences.

9.1 The Source of CO_2 138
CO_2 Measurements 139
CO_2 Expectations 140
Chief Contributors 143
9.2 Warming Mechanism 143
9.3 Possible Trajectories 148
9.4 Consequences 150
Heating Up 151
Heating Earth's Skin 152
Sea Level Rise 155
9.5 What Can Be Done? 157
9.6 Upshot: A Serious Problem 158
9.7 Problems 158

9.1 The Source of CO_2

The climate change forces at play today are primarily due to increased concentration of CO_2 in our atmosphere as a result of burning fossil fuels.

The chemistry is unambiguous (Eq. 8.1; p. 121): energy is released when fossil fuels are combusted with oxygen (O_2) to get CO_2 and H_2O. Table 9.1 extends properties of fossil fuels first presented in Chapter 8, adding CO_2 attributes.

Table 9.1: Combustion properties of fossil fuels, including CO_2 emission per gram of input and per MJ of energy out.

Fuel	Representative	molar mass	kJ/mol	kJ/g	kcal/g	CO_2 g/g	CO_2 g/MJ
coal	C	12	393.5	32.8	7.8	3.67	112
natural gas	CH_4	16	890.3	55.6	13.3	2.75	49
petroleum	C_8H_{18}	114	5,471	48.0	11.5	3.09	64

Of chief interest in Table 9.1 for this chapter are the last two columns:

Elephants doing their best. Photo courtesy of O'Connell & Rodwell.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

1. **mass ratio**: how many grams of CO_2 are produced per gram[1] of input fuel;
2. **carbon intensity**: how many grams of CO_2 are produced per unit of energy delivered.

1: . . . or any mass/weight measure you prefer

While all forms produce a mass ratio of ***approximately* 3 units of CO_2 for every unit of fossil fuels**, the lower energy density of coal together with its slightly higher mass ratio[2] make it *more than twice* as carbon-intense as natural gas.

2: Coal produces more CO_2 per gram of fuel because the other fossil fuels also contain mass in the form of hydrogen, which not only adds to energy production but also does not end up in CO_2.

Example 9.1.1 Roughly how much CO_2 is produced from each full tank of gasoline in a car?

A typical tank might hold about 50 L of gasoline (13 gallons). The density of gasoline is 0.75 kg/L, so that one tank has a gasoline mass of about 38 kg. Applying the simple and convenient factor-of-three ratio of CO_2 mass to input fuel mass, we see that one tank of gasoline will produce something like 110 kg of CO_2—not a small amount!

9.1.1 CO_2 Measurements

Beginning in 1958, Dave Keeling of the Scripps Institution of Oceanography began recording CO_2 concentration in the atmosphere from the top of Mauna Loa in the middle of the Pacific ocean.[3] In addition to seeing annual variation due to the seasonal cycle of photosynthesis,[4] he began to see a steady year-by-year increase in the level. The measurements have continued to the present, now known as the "Keeling Curve," shown in Figure 9.1.

3: . . . far from continental influences

4: Plants seasonally absorb and then release CO_2 as leaves grow and then die.

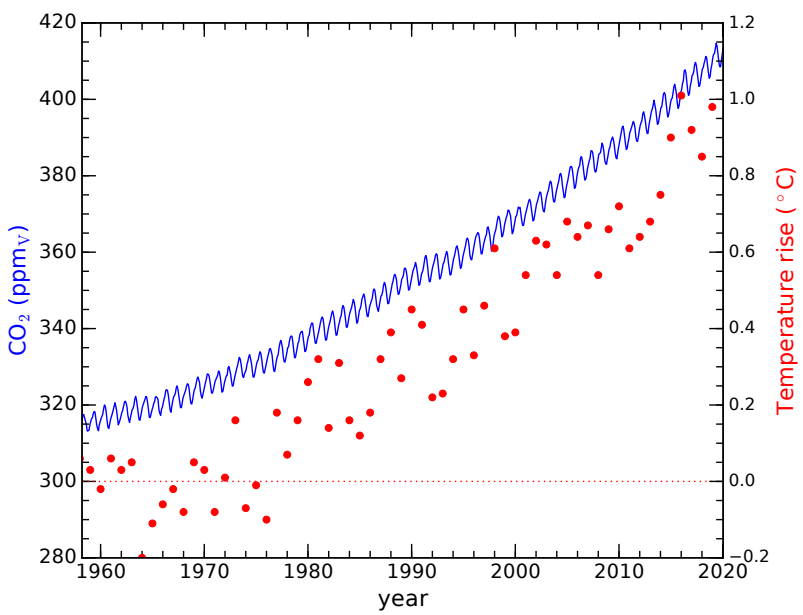


Figure 9.1: In blue (left axis), CO_2 measurements from Mauna Loa (Hawaii) for the last 60 years, showing a relentless and accelerating upward trend now at abut 2.6 ppm_v per year [50]. Seasonal variations due to photosynthesis are seen on top of this trend. Pre-industrial levels were around 280 ppm_v, so that we have added about 130 parts per million (ppm). Red dots (right axis) show global average mean temperature records over the same period [51]. Thus far, global average temperature has risen about 1°C. Note that the Kyoto Protocol in 1997 and Paris Agreement in 2015 (Box 19.4; p. 320) do not visibly curb the upward trajectory of CO_2 emissions.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

When the measurements started, the atmosphere consisted of less than 320 parts per million by volume (ppm_v), or < 0.032%. By now, we are beyond 410 ppm_v.

Measurements of trapped air bubbles in the Greenland ice sheet going back about 100,000 years and the Antarctic ice sheet going back 800,000 years indicate that CO_2 has fluctuated between 180–280 ppm_v, reaching the higher end of the range during the warmer periods between ice ages (interglacial periods). For at least the last thousand years before the Fossil Fuel Age, CO_2 held steady around 280 ppm_v.

9.1.2 CO_2 Expectations

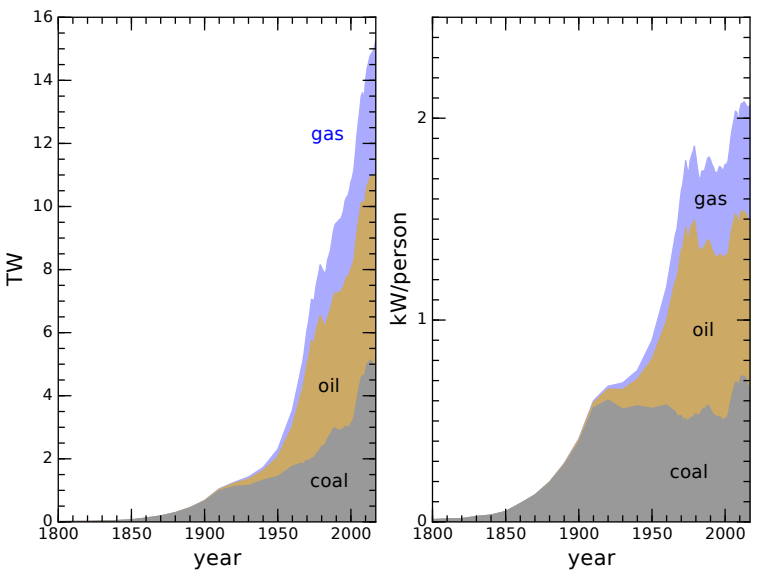


Figure 9.2: Historical use of fossil fuels worldwide, repeated from Figure 8.2 [16]. On the left is the raw usage rate expressed in terawatts, while the right is a per-capita measure. The three fossil fuel types are stacked on top of one another, so that gas makes the smallest contribution, not the largest. For instance, on the right-hand-edge of the left panel, coal goes from 0 to 5 TW, oil from 5 to 11 (thus 6 TW from oil), and gas from 11 to 15, indicating 4 TW from gas. The left-hand panel makes clear that fossil fuel use is still rising dramatically, and thus CO_2 emissions.

We saw in Figure 8.2—repeated as Figure 9.2—a history of fossil fuel usage for the world, in coal, oil, and natural gas. Meanwhile, Table 9.1 indicates how much CO_2 each fuel contributes per kilogram or Joule used. These two pieces can be combined to make an estimate of how much CO_2 is emitted globally each year, and to track total CO_2 emission over time. Table 9.2 and Box 9.1 elucidate how to go from the fossil fuel power figures (TW) in Figure 9.2 to CO_2 atmospheric concentrations.

Box 9.1: Computing CO_2 ppm_v from TW

We will use oil as an example. In Figure 9.2, we appear to get about 6 TW from oil (5 TW from coal, 4 TW from natural gas). Multiplying by 10^{12} puts this in Watts (J/s) and by 3.156×10^7 seconds per year results in the annual global energy from oil in Joules: 1.9×10^{20} J/yr. Table 9.1 indicates that oil contains about 11.5 kcal/g, so the number of grams of oil used per can be determined by first converting J to kcal

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

Operation	Resulting Units	Coal	Oil	Gas
Starting value	TW			
$\times 10^{12}$ W/TW	W (J/s)			
$\times 3.16 \times 10^{7}$ s/yr	J/yr			
$\div$ 4,184 J/kcal	kcal/yr			
$\div\, x$ kcal/g	FF g/yr	$x = 6.5$	$x = 11.5$	$x = 13.3$
$\div$ 1,000 g/kg	FF kg/yr			
$\times\, y$ CO_2 kg/kg	CO_2 kg/yr	$y = 3.67$	$y = 3.09$	$y = 2.75$
$\div 5 \times 10^{18}$ kg	CO_2 frac/yr			
$\times 10^{6}$	CO_2 ppm_m/yr			
$\div$ 44/29	CO_2 ppm_v/yr	total emissions		
$\div\ \sim 2$	CO_2 ppm_v/yr	stays in atmosphere		

Table 9.2: Stepwise procedure to convert TW to ppm_v of CO_2. FF means fossil fuel, which can be coal, oil, or gas—each computed separately using the various values provided on the right.

(divide by 4,184), yielding 4.5×10^{16} kcal/yr, and thus 3.9×10^{15} g/yr of oil, or 3.9×10^{12} kg/yr.[5] Table 9.1 shows that each kilogram of oil combusted produces 3.09 kg of CO_2, translating to 1.2×10^{13} kg/yr of CO_2 from oil. Incidentally, we get the same answer using 64 g/MJ from Table 9.1 and 1.9×10^{14} MJ/yr from a few lines up. The remainder of the work, converting to parts per million, is covered in the text and the lower part of Table 9.2.

The mass of the atmosphere is about 5×10^{18} kg, obtained by multiplying 10,000 kg of air[6] sitting over every square meter by the $4\pi R_\oplus^2$ surface area of the earth. Dividing the 1.2×10^{13} kg[7] of oil-generated CO_2 by the mass of the atmosphere yields 2.4×10^{-6}, or 2.4 parts per million.[8] We are almost there. The quantity we have calculated is parts per million by mass (ppm_m), not the conventional parts per million by volume (ppm_v). Since air[9] averages 29 g/mol, and CO_2 is 44 g/mol, the mass concentration of CO_2 is higher than the volume occupation in air by a factor of 44/29, or 1.52. Thus we divide our 2.4 ppm_m result by 1.52 to get 1.6 ppm_v. A final correction is that only half of this stays in the atmosphere, so that today we are putting 0.8 ppm_v into the atmosphere each year from oil.

Can you validate this number in the left panel of Figure 9.3 for oil?

Figure 9.3 shows the result of this computation as a function of time in terms of annual rate and cumulative emission. We find that even though oil and coal deliver about the same global annual energy now, the carbon-intensity of coal is much higher so that its CO_2 emission dominates[10] over the other sources. In fact, coal *always* has been the dominant CO_2 producer, seen by the fact that the black curve in the left-hand panel of Figure 9.3 has always been higher than the other two CO_2 sources. In total, this estimate suggests that we should see CO_2 levels rising by 2.6 ppm_v per year, a little over half of which is due to coal (1.4 ppm_v/year).

Summing up the rising contributions over time, this method of estimation suggests that we have increased the atmospheric CO_2 by 123 ppm_v,[11] and that 75 ppm_v of this (61%) is attributed to coal (right-hand panel of Figure 9.3).

5: At 120 kg per barrel, this turns into the expected 30 billion barrels per year as a check to see that we're on the right track.

6: This is a close approximation to the actual value, obtained by dividing standard atmospheric pressure of 101,325 Pa by $g \approx 9.8$ m/s^2.

7: ... the result of Box 9.1

8: Just multiply by one million, or 10^6 to get ppm.

9: Air is about 75% N_2 at 28 g/mol plus 25% O_2 at 32 g/mol.

10: Why do we keep using coal if it's the worst? Because replacement infrastructure is very expensive, and fossil fuel extraction does not work like a bank account allowing withdrawals at an arbitrary rate. We could not suddenly switch over and continue to satisfy demand, even if everyone wanted to—which they don't.

11: ... very close to the ~130 ppm_v we observe!

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

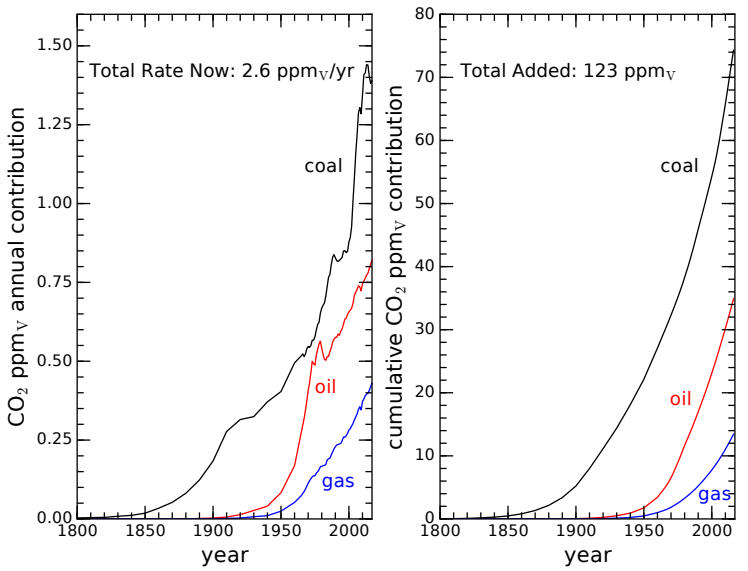


Figure 9.3: Estimated CO_2 contributions from known fossil fuel expenditures based on chemistry and the assumption that half of CO_2 stays in the atmosphere, while the rest is absorbed by the ocean and land. Units are parts per million by volume. The left-hand panel shows the annual addition, adding to 2.6 ppm$_V$ per year and accounting for the slope in Figure 9.1. The right-hand panel is the cumulative emission to date as a function of time—essentially adding up all the annual emissions from the left-hand panel. These curves are not stacked as are the ones in Figure 9.2, so each can be read directly from the vertical axis. Note that oil and gas are still on the rise in the left-hand panel: we emit more CO_2 each year than we did the year before.

Adding all three contributions from the right-hand panel of Figure 9.3 and plotting the result on top of the Keeling Curve,[12] we find astounding overlap in the shape—as shown in Figure 9.4.

12: . . . actual CO_2 measurements

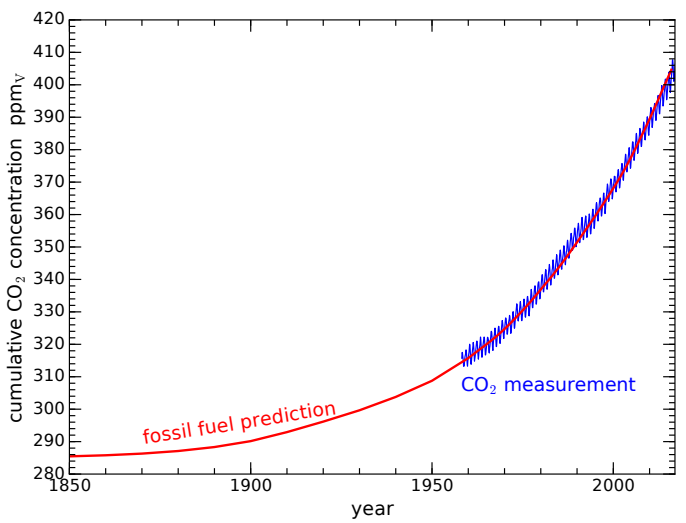


Figure 9.4: Fossil fuel contribution to CO_2 (red) on top of CO_2 measurements (blue). The red curve uses a starting point of 285 ppm$_V$ and has 49% of CO_2 emissions staying in the atmosphere. The overlap is remarkably good and convincing.

The curve computed from fossil fuel use overlaps the Keeling Curve so faithfully that little mystery is left as to where the excess CO_2 in our atmosphere originates. The chemistry and historical use of fossil fuels are not in dispute. The only "fudge" is in what fraction of the CO_2 emitted from fossil fuel combustion remains in the atmosphere vs. being absorbed by the ocean and other "sinks." Empirically, about half stays in the atmosphere, while the rest disappears into the ocean,[13] and into plant matter that gets buried in the ground. If unaware of the oceanic and land absorption mechanisms, we would have *overestimated* the amount of CO_2 due to fossil fuels by a factor of two (see Box 9.2).

13: . . . acidifying the ocean's water

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

Box 9.2: If We Ignored Oceans

Reflect on what our reaction might have been had we not corrected for oceanic and land absorption: we would conclude that we have *no trouble* quantitatively accounting for the CO_2 rise based on fossil fuels, and would be left asking why we don't see an even *larger* rise. In other words, it is no stretch to account for the cause of excess CO_2—making it far from mysterious. If our CO_2 increase were *not* due to humans,[14] we would be left with a *real* mystery: then where is all of the *known* CO_2 emission from burning fossil fuels disappearing to?

14: . . . via fossil fuels

9.1.3 Chief Contributors

Climate change is a global phenomenon. Even if all emissions came from one country or region, atmospheric circulation would spread the result around the globe—albeit more slowly across the equator. It is, therefore, a global problem. All the same, it is interesting to look at chief emitters.

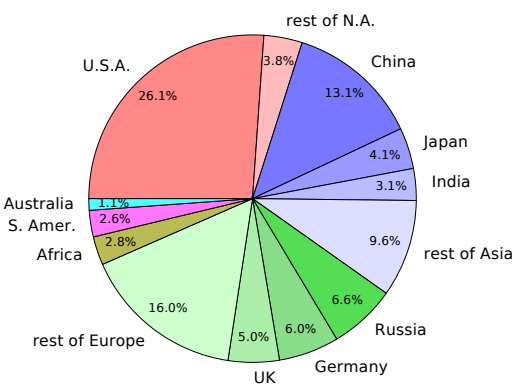


Figure 9.5: Cumulative contributions to CO_2 emissions to date, grouped by continent. Major contributing countries are called out explicitly, and the remaining contributions from each continent are grouped as a "rest of" category [52].

Table 9.3: Leading CO_2 emitters, 2018. Gt means gigatons, or 10^{12} kg, as 1 ton is 1,000 kg.

Country	Gt/yr	% share
China	9.43	27.8
U.S.	5.15	15.2
India	2.48	7.3
Russia	1.55	4.6
Japan	1.15	3.4
Germany	0.73	2.1
Top 6 Total	20.49	60.4
World	34	100

Figure 9.5 indicates that the U.S. is the single country bearing the largest responsibility for cumulative CO_2 emissions: roughly twice that of the second-largest (China) [52]. Today, China is the largest emitter, at 9.4 Gt per year, while the U.S. is in second place at 5.15 Gt/yr. Table 9.3 lists the top six emitters lately, accounting for about 60% of the 34 Gt per year [53].

[52]: Ritchie (2019), *Who has contributed most to global* CO_2 *emissions*

[53]: Rapier (2019), *The World's Top 10 Carbon Dioxide Emitters*

9.2 Warming Mechanism

The presence of excess CO_2 in our atmosphere is undeniably from fossil fuel combustion. But how does this alter our climate? How can such a minor constituent of the atmosphere (now 0.04%) cause so much trouble? The answer lies in infrared radiation [54]. Recall from Sec. 1.3 (p. 10) that

[54]: Pierrehumbert (2011), "Infrared radiation and planetary temperature"

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*; *Freely available at*: *https://escholarship.org/uc/energy_ambitions*.

this is the mechanism by which energy leaves the earth, the power of which is governed by the Stefan–Boltzmann law $P = A_{\text{surf}}\sigma T^4$, where the Stefan–Boltzmann constant,[15] $\sigma = 5.67 \times 10^{-8}$ W/m^2/K^4, and T is the temperature of the radiating surface, in Kelvin.

15: . . . easy as 5-6-7-8

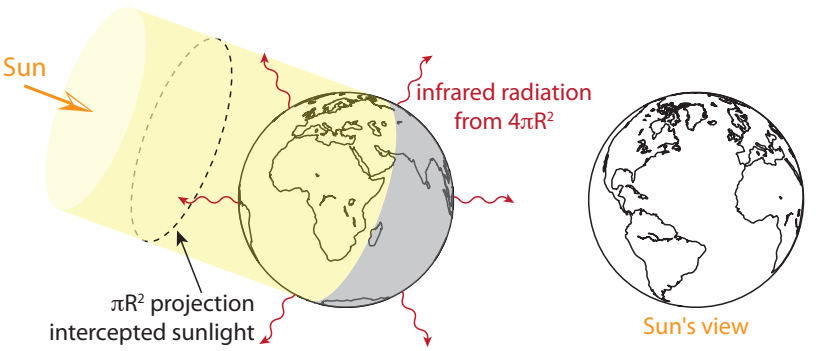


Figure 9.6: Earth intercepts sunlight across the projected area of the Earth's disk (πR^2), while radiating from the entire surface area, which is four times larger ($4\pi R^2$).

The sun delivers energy to the *top* of the earth's atmosphere at a rate[16] of 1,360 W/m^2. About 30% of this light—29.3%, to be precise—is immediately reflected by clouds, snow, and to a lesser extent water and terrain. The remaining 70.7% of the light intercepts the earth in a projected disk of area $A_{\text{proj}} = \pi R_\oplus^2$ (Figure 9.6). But the total surface area of the earth is four times this, all of it contributing to infrared radiation to space. In perfect balance,[17] energy absorbed equals energy radiated:

16: This is called the solar constant [4], and will appear again in Chapter 10 and Chapter 13.

17: An imbalance would mean energy is accumulating or being lost, leading to warming or cooling. Even under present conditions, the balance is good to within 1 W/m^2.

$$0.707 \times 1360\,\text{W/m}^2 \times \pi R_\oplus^2 = 4\pi R_\oplus^2 \sigma T^4. \tag{9.1}$$

The $\pi R_\oplus^2$ factors cancel, and we can rearrange to isolate temperature:

$$T^4 = \frac{0.707 \times 1360\,\text{W/m}^2}{4\sigma}, \tag{9.2}$$

solving to $T \approx 255$ K, or −18°C (about 0°F). This is about 33°C colder than the 288 K (15°C; 59°F) we actually observe as the average temperature of Earth. The 33°C difference[18] is due to greenhouse gases—mostly H_2O—impacting the thermal balance by preventing most radiation from escaping directly to space.

Verify for yourself. Two square roots in a row accomplish a fourth root.

18: Life on Earth is adapted to and reliant upon this 33°C greenhouse effect. Abruptly *changing* it is what causes problems.

We understand this mechanism perfectly. Being at a temperature of 288 K, the surface emission peaks at a wavelength around 10 μm.[19] The atmosphere is not transparent at all wavelengths, its various absorption features depicted in Figure 9.7. The blue curve at upper right in this figure is the emission spectrum associated with infrared radiation.

19: We will see how/why in in Section 13.2 (Eq. 13.5; p. 199).

Of the greenhouse gases contributing to absorption as pictured in Figure 9.7, water vapor is the dominant player, followed by CO_2. Notice that the blue solid portion in the figure[20] is mirrored in white in the total absorption panel just below,[21] and that the window is mostly defined by water vapor. But the right-hand—longer wavelength—side of the water window is stepped on by the CO_2 absorption feature, seen more clearly in Figure 9.8. This CO_2 feature is responsible for the sharp cutoff on the

20: . . . the infrared radiation that directly escapes to space

21: Thus the white portions indicate the open "windows."

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*; *Freely available at*: *https://escholarship.org/uc/energy_ambitions*.

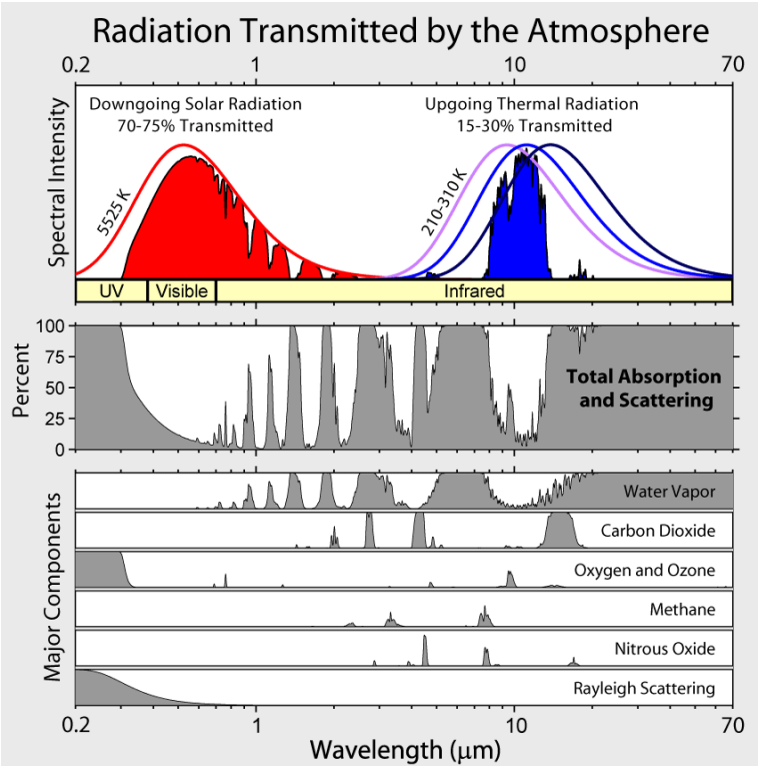


Figure 9.7: Atmospheric transmission/absorption spectra [55]. The top panel shows solar input in red and infrared (thermal) radiation output in blue. The smooth curves are the theoretical blackbody Planck spectra at solar and terrestrial temperatures. Thus the smooth red curve represents the distribution of solar energy arriving at the top of Earth's atmosphere, while the solid red filled feature is what survives the path to the ground. The smooth blue curve (middle of the set of three) represents radiation from the ground, but only a small fraction (blue filled region) passes directly through the atmosphere—the rest absorbed by greenhouse gases. The lower panels detail where light gets absorbed or scattered. Gray regions indicate absorption and scattering, so that white portions can be thought of as the transmitted part—often called "windows." Key contributors (greenhouse gases) are broken out in panels below the total absorption panel. Notice that ozone blocks ultraviolet (UV), and Rayleigh scattering is what makes the sky blue—by being effective at scattering blue light from the sun (blue is on the left side of the band labeled "Visible," while red is on the right). Created by Robert Rohde.

right side of the solid blue[22] shape in Figure 9.7. As CO_2 concentration in the atmosphere increases, this absorption feature gets wider, cutting deeper into the right edge of the escaping radiation (solid blue feature), allowing less radiation to escape.

22: . . . escaped radiation

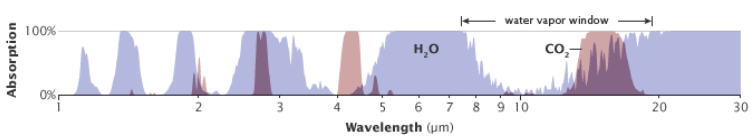


Figure 9.8: Another view of just the water and carbon dioxide absorption spectra, better showing the overlapping role of each in the 10 μm window. From Robert Rohde (NASA).

If some portion of the infrared radiation does not escape to space but is absorbed by the atmosphere, the planet does not cool as effectively, adding some offset to Eq. 9.2—in Earth's case 33°C. It is like the earth is wearing a blanket that raises its temperature by 33°C. Figure 9.9 illustrates the mechanism. Water vapor is responsible for ~20°C of this 33°C, and CO_2 is responsible for another ~8°C, leaving ~5°C for ozone, methane, and other minor contributors (Table 9.4). Incidentally, methane (CH_4) is about 80 times more potent than the same amount of CO_2 as a greenhouse gas, but is at a far lower concentration than CO_2, and also shorter-lived in the atmosphere before being chemically destroyed.[23] We focus on CO_2 because this is what human activity is changing rapidly by burning fossil fuels. The vast ocean–air interface means water concentration is impossible to control, and simply responds to temperature due to the fact that warmer air holds more moisture—becoming an important feedback agent. Water is not the *driver* of climate

Table 9.4: Greenhouse contributions [56].

Molecule	ΔT (°C)
H_2O	20
CO_2	8.6
O_3 (ozone)	2.6
CH_4 (methane)	1.5
N_2O (nitrous oxide)	0.5
Total	33

23: Methane emission becomes important via leaks from drill sites and also permafrost melt.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*; *Freely available at*: *https://escholarship.org/uc/energy_ambitions*.

change, but a hefty *passenger*.

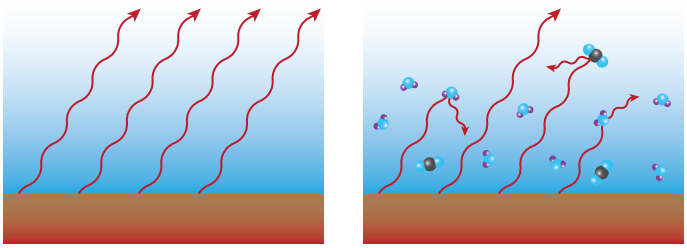


Figure 9.9: In the absence of greenhouse gases, infrared radiation has no difficulty escaping to space (left). When greenhouse gases (GHGs) are present (right), most of the infrared radiation is absorbed by GHG molecules. These molecules later release the absorbed energy as new infrared radiation, but in a randomized direction, so that some energy is returned to the ground, thereby keeping the surface warmer than it would be without greenhouse gases.

Naïvely speaking, going from 280 ppm_v to 420 ppm_v (an increase by a factor of 1.5), might be expected to turn the 8.6°C greenhouse effect from CO_2 into 12.9°C (1.5×) for a 4.3 degree human-caused warming.[24] But the CO_2 absorption feature at ~ 15 μm is *saturated*, so as CO_2 is added, it becomes *wider*, but logarithmically rather than linearly as a function of CO_2 concentration. Climate scientists often express the impact of various influences as a radiative forcing, measured in W/m^2.

24: ⚠ This is not how it works.

Definition 9.2.1 ***Radiative forcing*** *is used to describe the areal power (in W/m^2) of absorbed solar energy and infrared radiation to space. Various influences or constituents each contribute their own radiative forcing. In equilibrium, a balance exists so that the net*[25] *radiative forcing is zero.*

25: . . . sum of all: positive and negative

The average solar forcing is

$$RF_\odot = 1360\,W/m^2 \times 0.707/4 \approx 240\,W/m^2, \tag{9.3}$$

sharing substantial overlap with Eq. 9.2—the factors having been explained in that context. Without adding to the pre-industrial set of greenhouse gases (GHGs), we would solve for temperature[26] as

$$T = \left(\frac{RF_\odot}{\sigma}\right)^{0.25} + 33, \tag{9.4}$$

26: . . . adding the baseline 33°C GHG contribution

evaluating to 288 K, or about 15°C. If we add (or deduct) radiative forcing from another source, it would add (subtract) in the numerator of Eq. 9.4. The addition of CO_2 over the original amount ($CO_{2,orig}$; 280 ppm_v) generates a radiative forcing of

$$RF_{CO_2} = 5.35 \ln\left(\frac{CO_2}{CO_{2,orig}}\right) W/m^2, \tag{9.5}$$

where the ln() function is the natural logarithm.[27] At our present ~420 ppm_v, the ratio of current CO_2 to $CO_{2,orig}$ is 1.5 and $RF_{CO_2} \approx 2.2\,W/m^2$, so that the new temperature (adding the new forcing to Eq. 9.4) is

$$T = \left(\frac{RF_\odot + RF_{CO_2}}{\sigma}\right)^{0.25} + 33 \approx 288.6, \tag{9.6}$$

27: . . . capturing the logarithmic broadening of the width of the CO_2 absorption feature at 15 μm in Figure 9.7

© 2022 T. W. Murphy, Jr.; Creative Commons Attribution-NonCommercial 4.0 International Lic.; Freely available at: https://escholarship.org/uc/energy_ambitions.

which is about 0.6°C larger than it was before increasing the amount of CO_2. We can express this as a *sensitivity*: how much ΔT do we get for a given imbalance in radiative forcing? In this case, 0.6°C from a 2.2 W/m^2 increase is 0.27°C per W/m^2. But when the known feedback mechanisms are included—most of them in the positive direction—the temperature sensitivity becomes 0.8°C for every W/m^2 of additional forcing: called the climate sensitivity parameter.[28]

Definition 9.2.2 *The **climate sensitivity parameter** connects the amount of warming expected for a given amount of radiative forcing. Current understanding puts this at 0.8°C per W/m^2 of radiative forcing.*

Therefore, our current 2.2 W/m^2 of additional (fossil-fuel added) radiative forcing translates[29] to a 1.7°C temperature increase (Figure 9.10), which is about three times what the non-feedback calculation would provide.

Example 9.2.1 If we double our CO_2 concentration from pre-industrial levels, what would we expect the temperature increase to be?

Pre-industrial CO_2 was 280 ppm$_v$, so doubling it adds 280 ppm$_v$ for a total of 560 ppm$_v$. The radiative forcing is then $5.35 \ln(2) \approx 3.7$ W/m^2. Multiply this by the climate sensitivity of 0.8°C per W/m^2 to get a temperature increase of about 3.0°C (would be 1°C without feedback).

The positive feedbacks are important, and include factors like:

1. A warmer planet means less ice (glaciers, Arctic cap), resulting in less reflected sunlight, increasing the 0.707 absorption factor in Eq. 9.3 to increase solar forcing.
2. Warmer air can hold more water vapor—the principal greenhouse gas, thus driving up the nominal 33°C greenhouse gas contribution.
3. A warmer environment leads to additional CO_2 loss from drying forests, desertification, and accelerated decomposition of plant matter and peat.

A few negative feedback mechanisms[30] exist as well, but are outweighed by the positive feedback terms.

Global temperature increase is already about 1.0°C [58]. Note that even if we never added another CO_2 molecule to the atmosphere, the temperature would continue to rise as the ocean[31] slowly catches up to the new equilibrium. We would expect the temperature to stabilize around 1.7°C higher for today's CO_2 excess, according to the calculation above. Thus the climb is about 55% done. Of course, more CO_2 will be added, so the eventual temperature rise is destined to be higher still.

28: See [57] for a good synopsis and references to primary material within.

29: Just multiply 2.2 W/m^2 by 0.8 °C per W/m^2.

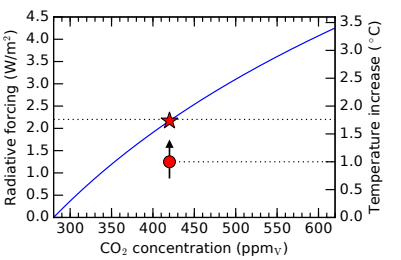


Figure 9.10: As CO_2 concentration increases, the radiative forcing (left axis) increases, driving the temperature (right axis) up. We are now at 420 ppm$_v$, corresponding to a radiative forcing of 2.2 W/m^2 and 1.7°C of ultimate temperature increase (red star). Presently, the temperature has only increased by 1.0°C (red circle), but will catch up to a new equilibrium once oceans warm and ice melts.

30: ⚠ By far the most important negative feedback mechanism is the infrared radiation itself, increasing dramatically as temperature increases (as T^4), thus opposing the temperature rise by a cooling influence. Here, we mean negative feedback influences *in addition to* this main one.

[58]: National Oceanic and Atmospheric Administration (NOAA) (2019), *Global Climate Report*

31: . . . lots of thermal mass, or heat capacity

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

9.3 Possible Trajectories

Launching from the data used to generate Figure 9.3, we can now play a few games to understand what our future might hold in terms of total CO_2 rise and corresponding ΔT increases by the year 2100 under various contrived scenarios.[32]

First, let's imagine that we suddenly arrest the upward climb characteristic of fossil-fuel usage to date[33] and maintain present-day levels of fossil fuel use from now until 2100. Figure 9.11 shows what happens. The total added CO_2 rises to 2.75 times the current excess, to 339 ppm_v[34] above pre-industrial levels. The associated radiative forcing would be 4.25 W/m^2 and result in a 3.4°C temperature increase. Table 9.5 summarizes this scenario and the three to follow.

Verify yourself following a procedure like Example 9.2.1 as good practice.

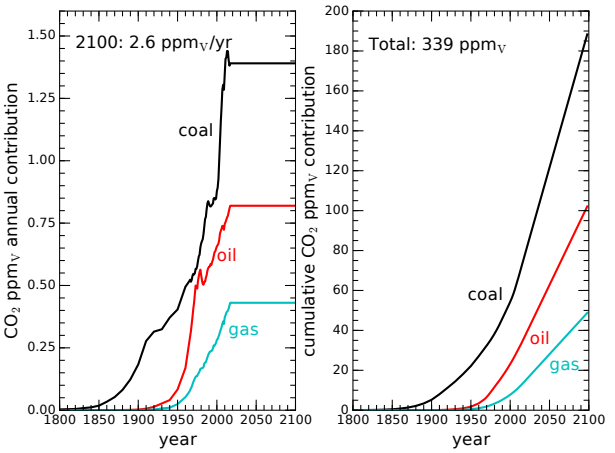


Figure 9.11: CO_2 rise if fixing fossil fuel use at today's levels for the rest of the century, following the conventions of Figure 9.3. We would still add 2.6 ppm_v per year from now until the end of the run in 2100, and would have accumulated a total rise of 339 ppm_v, or 2.75 times the problematic amount already accumulated to date. The associated temperature rise would be 3.4°C. For this and all subsequent scenarios, the plots show only the half of emitted CO_2 that remains in the atmosphere.

Table 9.5: Summary of scenario outcomes.

Scenario	ΔCO_2 (ppm_v)	vs. today	CO_2 (ppm_v)	RF_{CO_2} (W/m^2)	ΔT (°C)
Arrest FF climb	339	2.75 ×	620	4.25	3.4
Arrest; no coal	268	2.18 ×	548	3.6	2.9
Curtail by 2100	235	1.91 ×	515	3.3	2.6
Curtail by 2050	169	1.37 ×	450	2.5	2.0

We're already seeing serious problems emerging today, at about 1°C increase, so this 3.4°C scenario is not desirable.[35] And reflect for a moment how unlikely it is that we can even arrest the climb of fossil fuel use so suddenly. It would seem that the *rate* of CO_2 emission is destined to climb higher than it is today: we have not yet found the peak!

The second scenario focuses on eliminating coal, since it is the highest intensity CO_2 emitter,[36] as Figure 9.3 makes clear. What if natural gas—the fossil fuel having the lowest carbon intensity—could replace all coal applications? This is already happening—gradually—in the electricity generation sector. Countless advocates encourage such a transition as rapidly as can be accomplished. The pretend world of

32: None of the scenarios we will fabricate are *realistic*, exactly, but help us establish boundaries of possible outcomes. Mathematical models need not capture all the nuances to still be useful guides to understanding.

33: . . . reflected in the left-hand panel of Figure 9.3

34: . . . resulting in about 620 ppm_v; up 339 ppm_v from the pre-industrial 280 ppm_v

35: On the other hand, we may not have enough fossil fuel resource to realize this scenario, in which case it can be treated as an upper limit.

36: . . . based on its lower energy density (5–8 kcal/g vs. 13 for natural gas) and its higher CO_2-to-fuel mass ratio (3.67 vs. 2.75 for natural gas)

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

simulation allows us to evaluate the best-case (and unrealistic) scenario of *instant*, complete replacement, to put a limit on how much benefit such a move brings. Figure 9.12 shows what happens. The rate of CO_2 emission would immediately drop to 1.8 ppm_v/year.[37] That definitely helps, but the total emission by 2100—if carrying on at today's energy demand via fossil fuels—would climb to 268 ppm_v. The effect would more than double the 123 ppm_v that we've *already* contributed to the atmosphere, and would approximately double the pre-industrial CO_2 level in the atmosphere, leading to a forcing of 3.6 W/m^2 and $\Delta T \approx 2.9°C$ (summarized in Table 9.5). So as beneficial as the termination of coal would be, any path that involves carrying our fossil fuel use forward at today's levels—even substituting the best form for the worst form—does not look promising.

37: This would be about 70% the present rate of 2.6 ppm_v per year.

See if you can replicate these numbers!

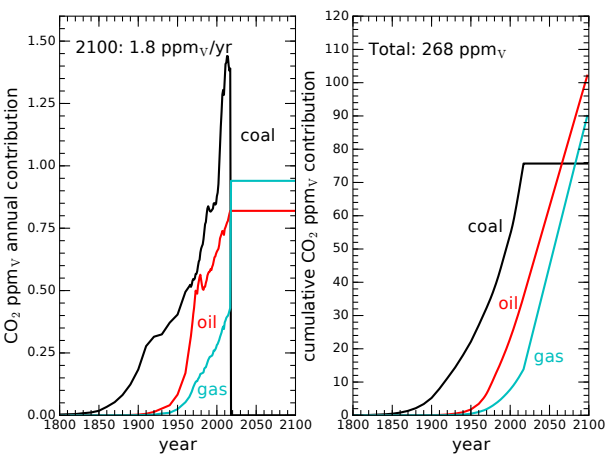


Figure 9.12: CO_2 rise if instantly substituting coal (worst CO_2 intensity) with natural gas (best CO_2 intensity) and then maintaining current levels for the rest of this century. Our annual contribution would drop from 2.6 ppm_v/yr to 1.8 ppm_v/yr based on this substitution, and the total accumulation would be 268 ppm_v by century's end (2.2 times the accumulation to date). The associated temperature rise would be 2.9°C.

The emphasis, then, should be to taper off fossil fuel use so that we wean ourselves of dependency. The transition could be fast or slow. A slower version might target the year 2100 for a full termination of fossil fuels. Figure 9.13 shows an idealization of what this might look like. Notice that the resulting curves are *roughly* symmetric, in that the downslope is not terribly different from the upslope. Let's pause to reflect on how incredible and fast the rise of fossil fuels has been. A descent as steep as the rise represents change at an astounding pace—which would be pretty disruptive in the best circumstances. In the absence of suitable substitutions, this would be a tremendously difficult journey, but one we may be forced to make by any number of paths.[38] In any case, the eventual added CO_2 would end up at 235 ppm_v—almost doubling what we have already emitted, and nearly doubling the pre-industrial CO_2 level in the atmosphere. The forcing in this case would be 3.3 W/m^2 and $\Delta T \approx 2.6°C$.

38: . . . driven by policies, markets or—more certainly—resource limits

Again, follow along for good practice.

Reducing fossil fuel use even more quickly, tapering to zero by 2050, results in Figure 9.14. The descent is alarmingly steep, and difficult to imagine happening in practice unless major disruptions[39] force this upon us. In any case, should we manage such a feat, our total CO_2

39: . . . resource wars, devastating climate change repercussions

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

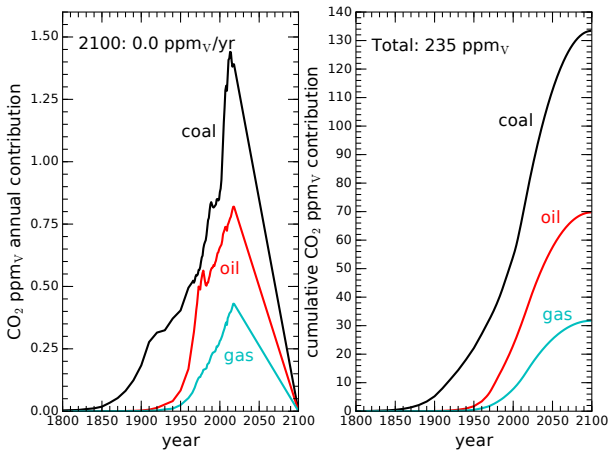


Figure 9.13: CO_2 rise if immediately reversing fossil fuel use in an ambitious decline reaching zero by the year 2100. The accumulation would come to 235 ppm_V at the end, which is almost twice the current level (1.9 times). The associated temperature rise would be 2.6°C.

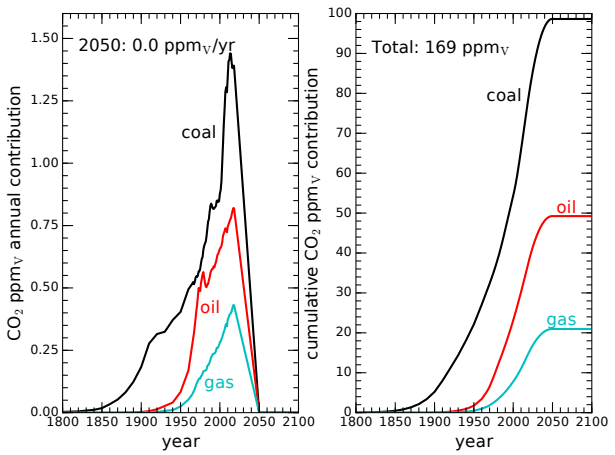


Figure 9.14: CO_2 rise if we mounted a super-aggressive weaning of fossil fuels, completed by the year 2050. In this extreme case, the total CO_2 emission would be 169 ppm_V, or 37% more than we have produced to date. The associated temperature rise would be 2°C.

contribution to the atmosphere would be an increase of 169 ppm_V, which is 37% more than we have emitted to date. Adding another ~ 40%, then, seems like the best we could hope to see, but possibly accompanied by serious hardships in adapting. The radiative forcing associated with this scenario is 2.5 W/m^2, corresponding to $\Delta T \approx 2$°C.

Try replicating!

9.4 Climate Change Consequences

Turning up the thermostat on the planet results in too many effects to chronicle here. Obviously, the climate is impacted—in terms of storm frequency and intensity, rainfall, snowfall and water supply, season durations, and the ability of plant and animal species to adapt to the changes. The timescale over which we are changing the climate is far faster than evolution can track, except for microbial life and maybe insects, whose shorter generational turnover permits a more dynamic response. Humans are late arrivals to a long evolutionary sequence, which laid a foundation to support our lives in complex interconnected

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

ways that we are not close to fully understanding. Climate change messes with the system in such a way to prevent accurate prediction of the long-term consequences of one species or another disappearing from the web of life.

The consequences of climate change are elaborated in many sources that are not difficult to find. Rather than try to add to the general awareness, this section—in the spirit of the book—aims to provide students with some tools[40] to be able to quantitatively understand how the physical world reacts to changes in radiative forcing. Specifically, we concentrate on the process of heating up[41] elements of the planet, and on sea level rise.

40: . . . and make connections to earlier content

41: For amusement, try substituting the phrase "climate change" with "hotting up."

9.4.1 Heating Up

Recall that the radiative forcing of 2.2 W/m^2 arising from a 50% enhancement to the pre-industrial CO_2 concentration[42] is expected to result in 1.7°C of eventual warming. But measurements indicate only 1.0°C of warming so far. Is our understanding wrong?

42: . . . from 280 ppm_v to 420 ppm_v

As we saw in Sec. 6.2 (p. 85), it takes energy to change something's temperature. When the *rate* of energy input[43] is limited, it takes *time* to accomplish a temperature rise.[44]

43: . . . which we know as power

44: A burrito in the microwave does not heat up instantly, for instance.

Earth is by-and-large in thermodynamic equilibrium. The sun deposits energy onto Earth at a rate of 240 W/m^2, when averaged over the surface (Eq. 9.3). Prior to the modern increase in CO_2 concentration, we had no additional radiative forcing from CO_2 and had an average surface temperature of 288 K (15°C), as in Eq. 9.4. Because it was in equilibrium, we know that the infrared radiation from Earth must have *also* totaled 240 W/m^2 to match the solar input.

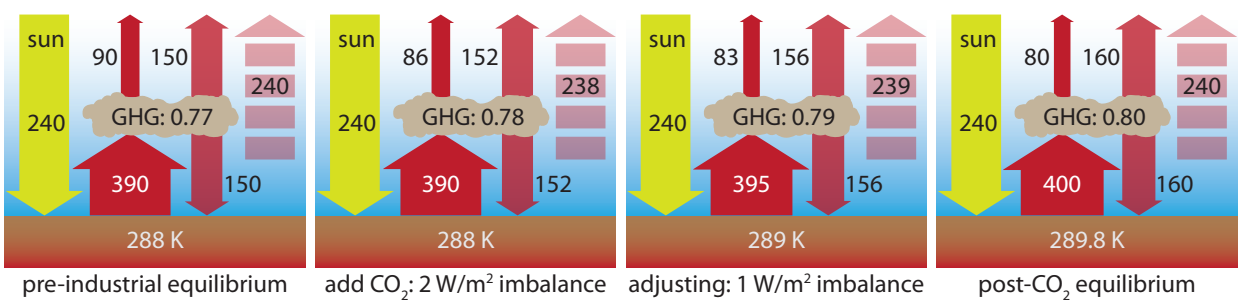


Figure 9.15: Four steps to illustrate (in a grossly simplified way) the process of Earth adapting to an increase in greenhouse gas (GHG). Starting from the left within each panel, solar input is held constant at 240 W/m^2. Most of the radiation leaving the ground—quantitatively adhering to σT^4—is absorbed by GHGs (fraction absorbed indicated in GHG "cloud"), the rest escaping directly to space. Half the absorbed energy is radiated up (escaping) and half back down. The dashed arrow at right is the net radiation escaping. Integer numbers are in W/m^2, and arrow widths are scaled accordingly. Ground temperature is indicated at bottom. See text for narrative sequence.

Figure 9.15 sums up the story.[45] The **first panel** shows the pre-industrial equilibrium condition, in which 77% of the infrared radiation from the ground was intercepted by the greenhouse gases, while 23% (90 W/m^2)

45: Note that in each panel, adding the two top arrows or subtracting the two bottom ones both yield the same number—matching the dashed arrow on the right.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

directly escaped. The 77% absorbed (300 W/m^2) is re-radiated equally up and down (150 W/m^2 each). The net outgoing radiation matches the 240 W/m^2 solar input[46] exactly, signaling equilibrium.

Pretend that we suddenly increased the CO_2 concentration to 420 ppm$_v$ in one instant,[47] which takes us to the **second panel** of Figure 9.15. The ground has not had time to change temperature yet, but the added GHG absorbs more of the outgoing radiation (78%). Now the numbers are not balanced. Only 238 W/m^2 radiates away, leaving a 2 W/m^2 net influx of energy. This is what we have been calling the radiative forcing.

After some time (**third panel**; representing our current status, roughly), the extra input energy starts to heat up the earth environment to 289 K (16°C),[48] which cranks up the radiation leaving the ground according to the σT^4 radiation law. At the same time, the higher temperature drives some positive feedback effects, putting more GHG into the atmosphere[49] and raising the absorption fraction. Meanwhile, the imbalance has moderated to 1 W/m^2 as the system edges toward a new equilibrium.

Try re-creating the upward radiation numbers using σT^4

Finally, equilibrium is re-established in the **last panel** of Figure 9.15, at which point outgoing energy and incoming energy are again matched at 240 W/m^2. The feedback mechanisms essentially tripled the change in GHG absorption: initially a 1% bump ultimately becoming 3%. Without this effect, the system would equilibrate at 288.6 K (15.6°C),[50] in which case 393 W/m^2 leaves the ground, and 153 W/m^2 is re-radiated from the atmosphere.

9.4.2 Heating Earth's Skin

Let's now understand a bit more about the process of heating up the Earth's air, water, land, and ice during the time when outgoing energy does not match incoming energy.[51] We know from Sec. 6.2 (p. 85) that it takes energy to change something's temperature. For instance, it takes 4,184 J to raise the temperature of a kilogram (liter) of water by 1°C, or about 1,000 J for many other substances like air and rock. Also relevant, and not introduced before, is that it takes a substantial amount of energy to melt ice.

Definition 9.4.1 *The **heat of fusion** of ice is 334 J per gram, meaning that every gram of ice that goes from just below freezing to just above[52] requires an input of 334 J.*

To put this in context, that same 334 J would heat a gram of liquid water by 80°C, but gets eaten up in the phase change without changing the temperature really at all.[53]

Using the properties of various constituents in Table 9.6, we can construct Table 9.7 to describe how much energy is needed to "charge up" different planetary components to a different temperature.[54] Note that the ocean

46: In reality, the solar absorption can also change as surface reflectivity changes—like when Arctic ice caps melt and expose dark water.

47: Obviously, it took time, but this approach is illustrative.

48: . . . as a result of no longer being in temperature equilibrium

49: H_2O, CH_4, for example.

50: . . . just as in Eq. 9.6

51: . . . middle two panels of Figure 9.15

52: . . . e.g., from −0.001°C to +0.001°C

53: This is why ice in a glass of water does not all melt suddenly—keeping the water around it basically right at 0°C as the ice slowly melts, limited by the rate of heat transfer.

54: Two entries in the table are blank, because it does not make much sense to talk about heating just a layer of the atmosphere, or heating the "whole" ground: how deep makes sense?

© 2022 T. W. Murphy, Jr.; Creative Commons Attribution-NonCommercial 4.0 International Lic.;
Freely available at: https://escholarship.org/uc/energy_ambitions.

Component	area (10^{12} m^2)	volume (10^{15} m^3)	mass (10^{18} kg)	density (kg/m^3)	thermal (J/kg)
Atmosphere	510	4,000	5.3	1.3	1,000/°C
Land/Rock	150	—	—	2,000	1,000/°C
Ocean	360	1,400	1,400	1,000	4,184/°C
Ice	16	29	2.6	917	334,000

Table 9.6: Properties of surface mass components of Earth. About 90% of the ice volume is in the Antarctic ice sheet, 10% in the Greenland ice sheet, and less than half-percent in glaciers [59, 60]. The volume of air corresponds to what it would be if compressed to uniform (sea-level) density. The last column captures specific heat capacity, or heat of fusion for ice.

takes about 1,000 times more energy than the atmosphere to heat it up by the same amount.

Component	charge per meter (10^{21} J/m)	total charge (10^{24} J)
Atmosphere	—	0.0053/°C
Land/Rock	0.3/°C	—
Ocean	1.5/°C	5.9/°C
Ice	4.9	8.8

Table 9.7: Energy requirements to heat up (thermally "charge") Earth components, derived from Table 9.6. The energy investment for the first three components depends on the temperature change sought, while melting ice is independent of temperature change.

Example 9.4.1 We will use the ocean as an example of how to interpret and use Table 9.7. The two numbers tell us what it takes to heat up the ocean per meter of depth and to heat the entire volume.

If we ask how much energy it would take to raise the temperature of the upper 10 m of the ocean by 2.5°C, we multiply 1.5×10^{21} J/m/°C by 10 m and 2.5°C to get 3.75×10^{22} J. Heating the entire volume of the ocean by 0.5°C would require 5.9×10^{24} J/°C times 0.5°C for 3×10^{24} J of energy.

Similarly—but without any temperature element—it would take $\sim 49\times10^{21}$ J to melt 10 m of Earth's ice, and would take 8.8×10^{24} J to melt it all.

We are now in a position to appreciate how long it can take to change temperatures on a planetary scale for a certain imbalance in radiative forcing. If, for instance, the imbalance is 1 W/m^2,[55] then Earth receives an extra 5.1×10^{14} J each second, or 1.6×10^{22} J in a year.[56] We can compare this annual surplus energy to the values in Table 9.7 to understand how deeply the components would be heated or melted per year for a 1 W/m^2 radiative forcing imbalance.

55: . . . as it roughly is now; third panel in Figure 9.15

56: 1 W/m^2 times the 5.1×10^{14} m^2 area of Earth, then 3.16×10^{7} seconds in a year.

Example 9.4.2 If we could direct *all* of the annual surplus 1.6×10^{22} J arising from a 1 W/m^2 imbalance into one component only,[57] we could ask: to what depth will each component be heated by 1°C or melt the ice?

57: ⚠ not how it really works

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

component	math	depth (m)
land	$1.6\times10^{22}\,\mathrm{J}/(0.3\times10^{21}\,\mathrm{J/m/°C}\cdot 1°\mathrm{C})$	54
ocean	$1.6\times10^{22}\,\mathrm{J}/(1.5\times10^{21}\,\mathrm{J/m/°C}\cdot 1°\mathrm{C})$	11
ice	$1.6\times10^{22}\,\mathrm{J}/(4.9\times10^{21}\,\mathrm{J/m})$	3.3

From this, we see that land is more easily heated, and ice is most resistant—only shaving about 3 m per year if somehow the entire wrath of radiative imbalance were directed upon it.

If the radiative forcing imbalance is larger or smaller than 1 $\mathrm{W/m^2}$, the effect will be proportionally deeper or shallower. Likewise, if we ask about a temperature increase larger than 1°C, the depth would be smaller in proportion.

Likewise, we can explore how long it would take to raise the temperature of *entire* components—or melt *all* of the ice—for an imbalance of 1 $\mathrm{W/m^2}$.

Example 9.4.3 If we could direct *all* of the annual surplus 1.6×10^{22} J arising from a 1 $\mathrm{W/m^2}$ imbalance into one component only,[58] we could ask: how much will the temperature rise of the entire body be per year, or how much ice would melt?

component	math	ΔT (°C)	years to 1°C
atmosphere	$1.6\times10^{22}\,\mathrm{J}/(5.3\times10^{21}\,\mathrm{J/°C})$	3.0	0.33
ocean	$1.6\times10^{22}\,\mathrm{J}/(5.9\times10^{24}\,\mathrm{J/°C})$	0.0027	367
ice	$1.6\times10^{22}\,\mathrm{J}/(8.8\times10^{24}\,\mathrm{J})$	0.185%	545 to melt

The atmosphere is a wimp: it takes very little to change its temperature. The ocean, however, is very sluggish to change temperature. For ice, we look at fractional loss per year instead of temperature increase.

58: ⚠ not how it really works

As before, this pretends that the entire radiative imbalance is directed upon a single component, in turn. If the radiative imbalance is larger or smaller than 1 $\mathrm{W/m^2}$, the effect will be proportionally more or less severe.

What we learn from the examples is that the ocean and ice are both substantial thermal brakes on fast heating. Even though ice has a much larger energy cost per kg,[59] its total mass is substantially smaller than that of the entire ocean—the two effects roughly balancing. In reality, we might expect oceanic circulation patterns to concentrate heating in the upper layers rather then distributing uniformly to the full depth. So the upper layer of the ocean—which then controls air temperature—can reach a 1°C increase well before 367 years elapse. Indeed, we are already seeing warming at this scale in less than a century.

59: . . . 334 kJ vs. ~4 kJ to heat water by 1°C

Tip: don't treat sections like this as the textbook "telling" you stuff. It is far more productive to follow the logic, reasoning, and numbers yourself as an active participant in the story. Otherwise, OMG is it boring! The value is empowerment: bestowing tools for transparent understanding of fundamental processes.

The way things *really* work is that the excess 1.6×10^{22} J annually associated with a radiative excess of 1 $\mathrm{W/m^2}$ gets distributed into lots of channels at once. If ice only gets 3% of the attention in proportion to its area, only 5×10^{20} J goes into the ice in a year. Dividing by 4.9×10^{21} J/m (from Table 9.7), we find that we might expect 0.1 m of ice to disappear each year. Since the ocean is roughly 25 times larger than the ice areas,[60] the associated sea level rise from redistributing the ice melt across the entire ocean surface would be about 25 times less, or about 4 mm/yr.[61]

60: Ice covers 3% of the globe, while the ocean covers 71%.

61: This is not far from reality, as will be seen in Section 9.4.3.

Meanwhile, the constant, swirling contact between water and air, and between air and land, keeps all three in sync with each other: one will not race off to get hot without the others. And in this case, the ocean—with its giant thermal mass and extensive air-water interface—is the limiting

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*; *Freely available at*: *https://escholarship.org/uc/energy_ambitions*.

factor on how fast things can heat up. If we confine ocean heating to the top 300 m[62] of water, the excess 1.6×10^{22} J per year leads to an annual temperature rise of about 0.035°C per year, or about 30 years to climb 1°C (at a 1 W/m^2 imbalance).

See if you can replicate this result.

Although this section may seem to be long, convoluted, and perhaps even boring, it accomplished a number of things for us:

1. it showed how a change in GHG absorption fraction leads to a radiative forcing imbalance;
2. it indicated how a radiative imbalance changes the surface temperature until the earth re-establishes a new equilibrium (balance) at a higher temperature, including feedback effects;
3. it assessed energy requirements for heating up relevant masses of material and melting ice via straightforward physics;
4. it showed that the two most important thermal masses on the planet are the ocean (first) and the ice sheets (second);
5. it established approximate timescales over which we might expect temperature to climb, and why the ocean in particular is important in slowing down the consequences.

An additional insight is that even if we stopped CO_2 emissions today, Earth's temperature will continue to climb as the oceans slowly adjust to the new radiative reality imposed by 420 ppm_v of CO_2 in the atmosphere.

62: The ocean is such a dominant thermal player that the rate at which temperature rises depends critically on how well and deeply mixed the thermal influence is.

9.4.3 Sea Level Rise

The previous section covered the energetics of melting ice sheets. The resulting melt-water flows to the ocean[63] and contributes to sea level rise. Besides melting ice, sea level also rises as a consequence of thermal expansion of water as it warms up. Figure 9.16 shows the recent history.

63: Ice floating on the ocean is already displacing water, so its melting does not impact sea level.

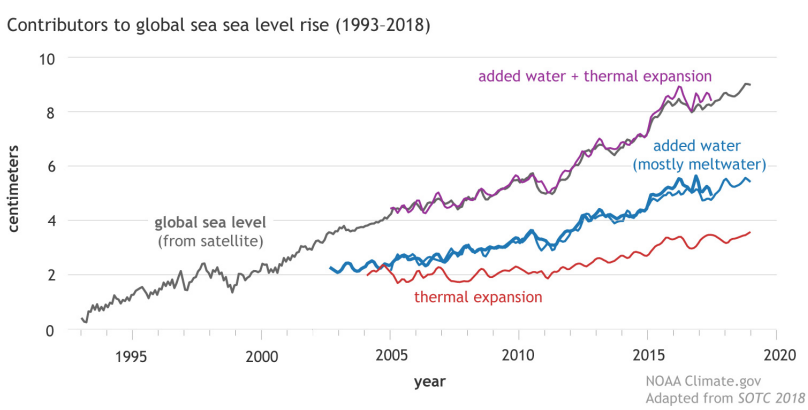


Figure 9.16: Satellite measurements of sea level since 1993, showing a rise of 3.6 mm per year. Melting ice is the largest contribution, although thermal expansion plays a role as well [61]. From NOAA.

Melting ice contributes about 2.4 mm/yr of rise, while thermal expansion accounts for about 1.2 mm/yr for a total rate of 3.6 mm/yr [61]. Since

[61]: Lindsey (2020), *Climate Change: Global Sea Level*

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

1880, sea level has risen about 230 mm At the *current* rate, we would expect a comparable addition by 2100 for a total of 0.5 m.

But the current rate is not likely to be the right measure, since warming air temperatures result in a faster rate of ice melting. Positive feedbacks also accelerate ice melt. For instance, melted pools of water on top of the ice are darker than ice, increasing the rate at which solar energy is absorbed.

We can get a quick handle on how much sea level rise might possibly be in store, based on the fact that the vast majority of "permanent" ice on the planet is in Antarctica and Greenland. These two ice sheets constitute 2.7% and 0.3% of the globe's surface area, respectively. From there, it is easy to estimate sea level rise, because the ocean (71% of the globe) has an area 26 times bigger than the Antarctic ice sheet and 210 times larger than the Greenland ice. What this means is that it takes 26 meters of ice melting from Antarctica[64] to raise sea level by 1 meter, and 210 meters off Greenland to do the same.[65] See Figure 9.17 to understand the logic here.

64: . . . if uniformly distributed across the continent

65: In the spirit of an approximate estimate, we ignore the 10% difference in the density of ice vs. water and assume that one cubic meter of ice displaces one cubic meter of water.

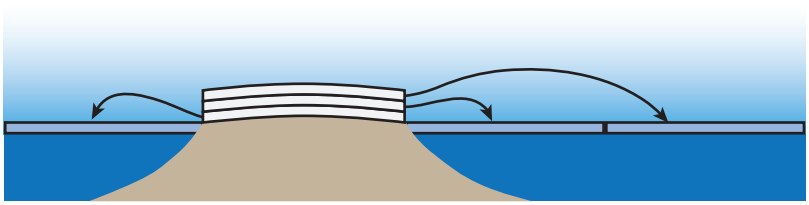


Figure 9.17: If sea area is x times that of an ice-covered island, water level will rise by $1/x$ times the ice thickness if it all melts. The diagram shows a way to think about this, for $x = 3$: by slicing the ice into $x = 3$ layers to redistribute the volume on top of the water. In this case, a 30 m ice sheet would raise sea level by 10 m if melted.

Now if we just knew the average thickness of each ice sheet, we could figure out how much sea level would rise if all the ice melted. The Greenland ice sheet is estimated to be 2.85 million cubic kilometers, translating to an average depth of 1.7 km.[66] 210 m goes into this 8 times, so we might expect something like 8 m of sea level rise if all the ice in Greenland melts. For Antarctica, the 26.5 million cubic kilometers corresponds to an average ice thickness of 1.9 km, which is about 70 of our 26 meter units, so we would expect about 70 meters of sea level rise in the extreme case of Antarctica losing all of its ice.

The preceding paragraph and Table 9.6 have enough information (absolute and relative areas and volumes) to figure out.

66: Divide volume by area.

Box 9.3: A more rigorous estimate. . .

It is relatively easy to find references estimating sea level rise potential from complete melting of Greenland and Antarctica. One such [60] arrives at a 7.4 m rise from Greenland melting, 58 m from Antarctica, and 0.3 m from glaciers.

So why did we go through the crude estimation process? The goal was to remove the mystery.[67] Once we have an estimate of the depth and a ratio of surface areas, it is within our grasp to estimate the rise ourselves.

[60]: Davies (2020), *Calculating glacier ice volumes and sea level equivalents*

67: One goal of this book is to empower students to independently check more authoritative sources—much like 2+2 = 4 can be personally verified and is not a matter of faith.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*; *Freely available at*: *https://escholarship.org/uc/energy_ambitions*.

Coastal cities are already struggling to deal with the present ~0.25 m rise. This may not sound like much, and on normal days does not cause problems. But the low pressure associated with storms results in a local swell in water level,[68] so that a storm plus high tide plus heavy waves adds to the climate-change rise to destroy human structures that otherwise may have escaped harm. Even if CO_2 emissions ceased today, the warming that has already happened will continue melting ice. Temperatures will also continue to climb as increased forcing from the already-present CO_2 continues to heat the ocean as we creep toward equilibrium—a slow process. So sea-level rise may turn out to be the gift that keeps on giving for centuries more.

68: High pressure elsewhere squeezes water into the low pressure areas.

Sea level in the distant past—tens of millions of years ago—has been as high as 200 m over where it is now. And 20,000 years ago, during the last glacial maximum, sea level was 120 m lower than it is today [62]. So Earth is no stranger to large fluctuations in sea level. By contrast, many coastal cities do not know how to handle even a one meter rise—which is within the projections for 2100 in a number of models.

[62]: Lambeck et al. (2014), "Sea level and global ice volumes from the Last Glacial Maximum to the Holocene"

9.5 What Can Be Done?

So far, we have simply described the climate change phenomenon as an unambiguous consequence of our fossil fuel habit, but have not addressed what we might do to combat it. In this section, the author will not attempt to conceal his personal take on the matter, and will keep it short.

First, clever geo-engineering ideas ring of hubris, and seem like solving a problem of hurtling toward the ground by digging a hole ahead of the fall fast enough to keep up. Our whole problem is that we have convinced ourselves that we can outsmart nature.[69] You would stop going to a doctor who tried to treat your illness by applying superficial remedies to a resulting skin condition without addressing the underlying cause.

69: See Sec. D.6 (p. 408) in Appendix D for an extended discussion of this notion.

In this case, the underlying cause is very simple: unchecked[70] human ambitions. The combination of fossil fuels, an incessant drive for growth, destruction of forests and habitat,[71] population pressures, and an industrial approach to agriculture all play a role. This is why it has been hard to combat: it's us. It pervades all the things we do. We are our own adversaries. How do we fight ourselves? To do so requires honesty and even a collective willingness to sacrifice and prioritize planetary health over narrow short-term human interests. What is more important: that individuals attempt to fulfill all their dreams now, or that civilization endures for the long haul? Are we willing to dial back our own desires so that billions of future humans who we will never meet and countless other species on the planet can also enjoy life? We have never had to make such a difficult choice as a global whole, so it is hard to say if *Homo sapiens* can pull it off.

70: Nature just called and left this cryptic message: "Check."

71: . . . eliminating carbon sinks and compromising nature's ability to adapt

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*; *Freely available at*: *https://escholarship.org/uc/energy_ambitions*.

The last chapters in this book, beginning with Chapter 18, deal with human factors that contribute to our prospects, ending at Chapter 20 discussing mitigation strategies at the individual level. It's not geo-engineering, but behavioral adaptation leading to reduction can have a huge impact, if broadly adopted.

9.6 Upshot: Climate Change is Serious

Despite the ease with which we can demonstrate that the measured CO_2 is from fossil fuels, and the straightforward physics governing the equilibrium temperature in the face of greenhouse gases, it is astounding that we are not fully embracing reality. But perhaps it should not be surprising. Climate change is a clear-and-present affront to some economic and political ideologies that wish we would let markets set our path, with no consequence. But operating in denial and insisting that we keep to the way things worked in the past is not our smartest play. Climate change is a stark indication that we can't just do whatever makes us the most money. Alternatives to fossil fuels are expensive and less convenient. Climate change is bad for business, and also threatens capitalism by imposing limits on our ambitions. It is little wonder that the U.S. has been among the most resistant to accepting climate change realities, being the most proud capitalist country, boasting an enormous per-capita energy demand, and having contributed the most to global CO_2 emissions throughout history (Figure 9.5).

As real as climate change is, it is less clear whether it represents an existential threat to civilization.[72] As costly as it may be to adapt,[73] the changes are gradual enough on a human timescale to possibly be manageable—though decidedly less than fun.

Resource[74] disruptions, however, can send markets into free-fall and stimulate global military actions that could be far more devastating on a shorter timescale. We are in a sort of a race to see which causes the biggest problem the soonest. With luck, a third option presents itself that does not involve such hardship. But just keep in mind that human society is a highly nonlinear construct that could become unrecognizably disrupted on a much faster scale than will be the case for the physics-governed unfolding of climate change.

72: Is it apocalyptic, or just very painful?

73: . . . abandoning some coastal areas, migrating away from unsurvivable climate regions, reducing population to deal with decline of agricultural resources, adapting to new ecosystems robbed of some species

74: . . . e.g., energy, water, food, mined substances

9.7 Problems

1. If you were asked to characterize a typical carbon cost of fossil fuels, in grams of CO_2 emitted per gram of input fuel, what single simple number would be a good approximation to all three types, referring to Table 9.1?

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

2. How many kilograms of CO_2 are produced from consuming a 40 L tank of gasoline? Compare the result to typical human mass. Gasoline is about 0.75 kg/L.

3. Typical household electricity use in the U.S. is about 30 kilowatt-hours (kWh) per day. If the electricity is produced by a natural gas plant operating at 40% efficiency, each household requires 75 kWh of energy from natural gas each day. Converting this to Joules, and using the intensity of 49 g/MJ from Table 9.1, how much CO_2 is produced daily to supply a house with electricity from a natural gas source?

4. Americans use energy at an average rate of 10,000 W. Convert this to Joules in a day, then MJ, then use a representative number[75] in the right-hand column of Table 9.1 to estimate an approximate CO_2 emission per American *per day*, in kg. Compare this to the mass of a typical person.

5. Getting more specific, Americans, on average, get about 320 MJ per person from oil each day, 265 MJ from gas, and 110 MJ from coal. Based on the CO_2 mass per energy column in Table 9.1, how much mass, in kg, of CO_2 does an American produce per day from each of these sources, and in total? Compare this daily emission to the mass of a person.

6. If carbon dioxide emission is a chief concern, switching from a coal-fired electricity plant to natural gas input is still a use of fossil fuels, but natural gas produces less CO_2 per energy unit than coal (Table 9.1). By what factor would CO_2 emissions be reduced if replacing all energy[76] we now get from coal with energy from natural gas?

7. Divide the total CO_2 increase by the annual CO_2 increase in Figure 9.3 to get a timescale in years. Compare this timescale to the period over which we have been burning fossil fuels (e.g., Figure 9.2). Do they agree? If not, what reason would you offer?

8. The right-hand panel of Figure 9.3 is an accumulation of the values on the left-hand side—representing the area under the left-hand curve, for instance. Since the gas trend on the left looks like a triangle, it is easy to approximate its area.[77] Does the result match the right-most value for gas in the right-hand plot?

9. According to the fossil fuel history shown in Figure 9.2, the world currently gets energy from coal at a rate of 5 TW. From this, figure out how many ppm_v/year we add in CO_2 from coal, following the outline in Table 9.2 and Box 9.1. Use a typical coal energy density of 6.5 kcal/g and a mass ratio of 3.67 CO_2 g/g.[78] The result should match information gleaned from Figure 9.3.

10. What is the most direct and convincing response to a "climate

75: . . . considering that we use a mix

76: E.g., Even swap, MJ for MJ

77: Multiplying ppm_v/yr (height) times years (base) gives units of ppm_v, matching the plot on the right.

78: . . . or equivalently, 135 g/MJ

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

skeptic" who names many natural sources of CO_2 and says we should not be too quick to label the CO_2 rise as *anthropogenic*?[79]

79: . . . of human origin

11. What letter grade would you give to the performance of the red curve in Figure 9.4 for its faithful tracking of the measured CO_2 blue curve? Describe reasons why you might deduct any points.

12. The allocation of CO_2 emissions among countries in Figure 9.5 differs pretty markedly from the distribution in Table 9.3. What does this tell us about past and present activity among countries?

13. Use a modified version of Eq. 9.2 to compute what the earth's temperature would be if it had no greenhouse gases and absorbed 100% of incident solar energy.[80] Is this warmer or cooler than the actual average surface temperature (given GHG and ~30% reflection)?[81]

80: . . . i.e., dark and no bright clouds/snow

81: Hint: the fourth root can be obtained by taking two square roots one after the other, or raising to the 0.25 power using the y^x button.

14. Based on Figures 9.7 and 9.8, is the atmosphere transparent or opaque at a wavelength of 1.0 μm? What about at 6 μm? And how would you characterize the situation at 2 μm?

15. If enough ice on Earth melted and resulted in 25% reflectance instead of 29.3%, what would the equilibrium temperature become, still applying the nominal 33 K GHG contribution? How much temperature rise is this compared to the equilibrium temperature at 29.3% reflection?

16. Table 9.4 has H_2O as the leading greenhouse gas, and all of them adding to 33 °C of warming effect on Earth. Should it be our goal to reduce all of these effects to the lowest numbers possible?[82] Why or why not? How would you characterize the greenhouse effect and why it is of concern to us?

82: For instance, if we could cut them all in half, would that be good?

17. Treating Figure 9.9 somewhat literally, in which one out of every four infrared photons escapes without being absorbed by a greenhouse gas molecule, what would the effective upward radiation be if 400 W/m^2 left the ground, and re-radiation of the absorbed fraction was split equally between upward radiation and radiation returning to the ground?[83]

83: Hint: Figure 9.15 offers clues.

18. If the "bad" news that we are about halfway through the fossil fuels is wrong, and we are only one-quarter of the way through, and we end up using all of it, what would the ultimate CO_2 concentration be in ppm$_v$, extrapolating the increase so far? How many degrees would this turn into based on our understanding of radiative forcing and the climate sensitivity parameter?

19. Using Eq. 9.5 and $CO_{2,orig}$ = 280 ppm$_v$, together with a climate sensitivity parameter of 0.8°C per W/m^2, how much would you predict Earth's temperature to rise for CO_2 levels of 330, 380, 430, 480, 530, and 580 ppm$_v$?[84] The inputs increase by steps of 50

84: A table would make sense, including the radiative forcing as one column.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

(linearly). Is the corresponding temperature trend linear?[85]

85: . . . ; same-size steps? A plot might help elucidate.

20. Which of the following are positive vs. negative feedback effects for climate change as a result of warming, and why?
    a) warming tundra releases methane trapped in the permafrost
    b) more water in the atmosphere creates more clouds and increases the amount of sunlight reflected from Earth
    c) an increase in forest fire activity burns lots of carbon-based material and leaves a blackened earth
    d) snow covers less of the land each winter as temperatures rise

21. Coal is, and always has been, the worst offender of the fossil fuels in terms of CO_2 emission. Based on the scenarios summarized in Table 9.5, what is the percentage reduction in final temperature rise (ΔT) based on emissions out to 2100 (in the case that we hold steady at today's fossil fuel levels) if we substitute natural gas for all uses of coal vs. if we keep things the same?

Comparing first two scenarios

22. Which of the scenarios in Section 9.3 do you deem to be most realistic, and why? Given this, how much more eventual warming is likely in store compared to the 1°C to date?

23. Verify for each panel in Figure 9.15 that the sum of the upper two arrows in the middle of each panel *and* the difference of the lower two arrows in the middle match the effective (dashed) arrow at the right (a total of 8 comparisons).

24. Construct your own panel in the manner of Figure 9.15 corresponding to the third case in Section 9.3 in which we curtail fossil fuel use by 2100, resulting in a final equilibrium ΔT of 2.6°C (thus 290.6 K ground temperature). Characterize the final equilibrium state, when net upward radiation equals solar input. After balancing the books, figure out what the GHG fractional absorption must be (the number in the "cloud" in Figure 9.15).

Tips: Upgoing radiation from the ground is computed from σT^4. Verify that you can match the first or last (equilibrium) panel of Figure 9.15 by the same technique to know if you are doing it right.

25. On a sunny day after a big snowfall, the (low winter) sun might put 500 W/m^2 onto the ground. Snow reflects most of it, but let's say that it absorbs 5% of the incoming energy. How much snow (ice) will melt in an hour from the absorbed energy if a cubic meter of snow has a mass of 100 kg?

Hint: It is easiest just to consider a single square meter of surface and figure out what fraction of a cubic meter (thus what fraction of a meter-depth) melts. Ignore any effect of air temperature.

26. If the ocean absorbs an additional 3 W/m^2 of forcing,[86] how long would it take to heat the ~4 km deep column of water directly under a particular square meter (thus 4×10^6 kg) by 1°C?[87]

86: . . . 3 J deposited every second in each square meter

87: Hint: use kcal (4,184 J) and solve for the temperature increase in one year, then figure out how many years for that to accumulate to 1°C.

27. Based on Table 9.7, if we could magically turn off infrared radiation to space, how long would it take for the sun's 240 W/m^2 average to heat the entire ocean by 1°C? It is fine to assume that the ocean covers the whole globe, for this limiting-case calculation.

Hint: multiply by area of Earth and seconds in a year to get Joules of input.

28. Let's say annual excess energy input from imbalanced radiative

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

forcing is 3×10^{22} J per year. How much would the temperature of the top 10 m of ocean rise in one year at this level?

29. If we somehow maintained a steady 2×10^{22} J/yr of unbalanced input indefinitely, how long would it take to:
    a) raise the entire ocean temperature by 1°C?
    b) melt all the ice?
    c) raise the entire atmosphere's temperature by 1°C?
    d) heat up the top 100 m of earth's land area?

    The sum of all these is how long it would take to get all of these tasks done, even though they would be happening in parallel.

ⓘ In practice, this is not how things will go. The ocean will heat up more near the surface, allowing air to get hotter and ice to melt faster while the deep ocean stubbornly resists warming for centuries.

30. On a planet 60% covered in ocean, and 5% covered in a 3 km thick ice sheet on top of land, how much would the sea level rise if all the ice melted?

31. On a planet two-thirds covered in ocean, and 1% covered by an ice sheet melting at a rate of 1 meter per year, how fast would sea level be rising?

32. How might you express the tradeoffs between the available fossil fuel resource and climate change? In other words, if it turns out that we don't have nearly as much fossil fuel as we think, what does this mean for climate change vs. our economic/geopolitical stability and viability? On the other hand, if we have centuries of fossil fuels left, what does that mean? Which scenario is least disruptive?

33. If it were certain that the only way to provide a comfortable existence to future generations involved substantial cutbacks toward a lifestyle having far fewer energy and material comforts today, do you think humanity would voluntarily do so? Can we leave treats on the shelf, within easy reach? If so, is it only uncertainty that prevents us? If not, what do you see as the barriers?

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

Part III

Alternative Energy

The options have not changed in 50 years.
There is nothing new under the sun.



10 Renewable Overview

10.1 The Players 164
10.2 Alternatives vs. Renewables 165
10.3 Renewable Energy Budget . 167
10.4 Renewable Snapshot 169
10.5 Upshot: Our Path Forward . 170
10.6 Problems 171

We now understand that the path human civilization has traveled to this point—a path paved by fossil fuels—cannot lead very much farther before a new path must take over. One distinct possibility is a trek back toward a low-tech existence, but most people would consider such a development to be a failure. What would success look like, then? Sec. D.5 (p. 404) in the Appendices takes a look at the big picture of long-term human success, but for present purposes we will focus on the critical issue of energy: where could we get enough energy to replace the prodigious one-time gift of fossil fuels?

This chapter is very brief, only setting the stage for upcoming chapters that go into alternative energy sources in greater detail. At the end of this effort, Chapter 17 summarizes the potential of all the energy resources. For now, we describe the origins and scales of Earth's energy inputs, and finish with a reminder of how much energy we get from these sources today.

10.1 The Players

Before launching into detailed attributes of energy sources beyond fossil fuels, it is helpful to list the cast of characters.

- **Hydroelectric** energy (Chap. 11) traps water from a river behind a dam, forced to flow through a turbine[1] that spins a generator to make electricity.
- **Wind** energy (Chap. 12) spins a turbine[2] that makes electricity via a generator.

1: A turbine is basically a set of fan blades.

2: Wind turbines are sometimes called rotors or windmills.

The remnant of a fallen tree—called a nurselog—provides nutrition for a row of new trees, arranged in a colonnade. Photo credit: Tom Murphy

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

- **Solar** energy (Chap. 13) can provide direct heat or make electricity via photovoltaic (PV) panels or in a utility-scale solar thermal installation.
- **Biological** energy (Chap. 14) can range from food to firewood, or biofuels. These are generally a source of thermal energy, via burning, able to drive heat engines.
- **Nuclear Fission** (Chap. 15) relies on mining a finite supply of radioactive elements in the Earth's crust whose fission (splitting) generates heat that can make steam to run a heat engine and generator to make electricity.
- **Nuclear Fusion** (Chap. 15), if successful, would use abundant hydrogen resources from water to build helium nuclei—a process that would release heat to make steam for running a heat engine and generator to make electricity.
- **Geothermal** energy (Chap. 16) originates in Earth's hot interior, and can be used for heating or to make steam to run a heat engine and generator to make electricity.
- **Tidal Capture** (Chap. 16) is very similar to hydroelectricity, but based on trapped tidal basins instead of dammed rivers.
- **Ocean Currents** (Chap. 16) behave much like wind, and can be used similarly to produce electricity, but underwater.
- **Waves** (Chap. 16) bring energy to shorelines that can drive specialized generators to make electricity.

One prevalent theme connects most of these entries: lots of electricity, often by way of heat engines and/or generators. Indeed, alternatives to fossil fuels tend to excel at electricity production. But as we saw in Fig. 7.2 (p. 105), electricity makes up only 38% of energy demand in the U.S., and only 17% of the energy *delivered* to the four end-use sectors.[3] One lesson is that current energy uses that are *not* electricity-based—like transportation and industrial processing[4]—will be more difficult to replace by the alternatives above.

3: While 38.3 qBtu of the 101.3 qBtu total energy input in the U.S. is dedicated to electricity production, only 13.0 qBtu emerges from the process, which is 17% of the 76 qBtu total flowing into these end-use sectors.

4: . . . which requires a lot of heat

10.2 Alternatives vs. Renewables

Before going further, we should clarify the difference between "alternative" and "renewable" resources.

Definition 10.2.1 ***Alternative Energy*** *is a* non-fossil *source of energy. Solar, wind, and nuclear would be examples.*

Definition 10.2.2 ***Renewable Energy*** *is a source of energy that is replenished by nature, so that its use may be sustained "indefinitely," without depletion. Solar energy, for instance is not "used up" by placing a panel in the sunlight.*

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

The sun will continue to shine no matter how many solar panels we set out. Wind is replenished daily by the sun heating the land and driving air currents. Solar energy drives the hydrological cycle, refilling the reservoirs behind hydroelectric dams. Plants grow back to replace harvested ones—again thanks to the sun. Ocean currents and waves are also driven by the sun, via wind.

Nothing, of course, lasts forever, but the sun will continue to operate in its current mode for billions of years more, and this is long enough to count as indefinite, for our purposes. Table 10.1 classifies various sources as to whether they are alternatives or renewables, along with justifications. The items with asterisks are *technically* not renewable, but last long enough that we can treat them as such in a practical sense (see Box 10.1).

Table 10.1: Energy classification. Asterisks indicate non-replenished, but long-lasting sources.

Resource	Chapter	Alternative?	Renewable?	Reason
Petroleum	8	No	No	finite supply in ground
Natural Gas	8	No	No	finite supply in ground
Coal	8	No	No	finite supply in ground
Hydroelectric	11	Yes	Yes	Sun generates rain and refills reservoirs
Wind	12	Yes	Yes	Sun generates daily by heating Earth surface
Solar	13	Yes	Yes	Sun will last billions of years
Biomass (wood)	14	Yes	Yes	Sun grows more
Nuclear Fission	15	Yes	No	finite supply of fissile material in ground
Nuclear Fusion	15	Yes	Yes*	billions of years of deuterium; not tritium/lithium
Geothermal	16	Yes	Yes*	finite, but large; rate-limited
Tidal Capture	16	Yes	Yes*	can drive Moon away eventually
Ocean Currents	16	Yes	Yes	Sun/wind-driven
Waves	16	Yes	Yes	Sun/wind-driven

Just because a resource is renewable does not mean it is *limitless*.[5] We only have so much land, nutrients, and fresh water to grow biomass, for example. Cutting trees down faster than they grow back would result in depleting the resource—possibly permanently if the land is altered severely enough that trees do not grow back. Installing turbines throughout the ocean to capture ocean currents[6] would eventually create enough impediment to the flow that it could stop altogether.

5: Ultimately, only so much sunlight strikes the earth.

6: This would be a hugely expensive and impractical undertaking, but helps illustrate the point that renewable does not mean unlimited.

Box 10.1: About Those Asterisks. . .

The items sporting asterisks in Table 10.1 deserve additional explanation as to why they are not technically renewable, even if the depletion timescales are extremely long.

Capturing all available tidal energy would end up accelerating the moon's egress from Earth,[7] eventually causing loss of the resource.[8]

About half of the geothermal energy store represents a one-time

7: . . . now at 3.8 cm/year

8: It would take hundreds of millions of years to "accomplish" this (see Sec. D.4; p. 402).

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

deposit of heat left over from the collapse/formation of the earth,[9] the other half coming from radioactive decays of elements ultimately tracing to ancient astrophysical cataclysms.[10] For both the formation and radioactive contributions, the supply is not replenished after its use, although the timescale for the radioactive decays to fade away is billions of years.

Nuclear fusion needs deuterium and tritium.[11] Roughly one out of every 10,000 hydrogen atoms is deuterium, so ocean water (H_2O) will have enough deuterium to last billions of years. Tritium, however, is not found naturally and must be synthesized from lithium, in finite supply. Details will follow in Chapter 15.

9: . . . a conversion of gravitational potential energy

10: . . . primarily supernova explosions and neutron star mergers

11: Eventually it is hoped that only deuterium could be used.

10.3 Renewable Energy Budget

Notice that all of the unqualified[12] "Yes" entries in Table 10.1 originate from the sun. For that matter, fossil fuels represent captured ancient solar energy, stored for all these years. The sun sends energy toward the earth at a rate of 1,360 W/m^2. Multiplying this by the projected area[13] of the earth ($\pi R_\oplus^2 \approx 1.28 \times 10^{14}$ m^2) results in 174,000 TW of solar power intercepting the earth. This number *absolutely dwarfs* the 18 TW societal energy budget of all humans on Earth. Figure 10.1 shows graphically what happens to this energy input.

12: . . . i.e., no asterisk

13: See Example 10.3.1.

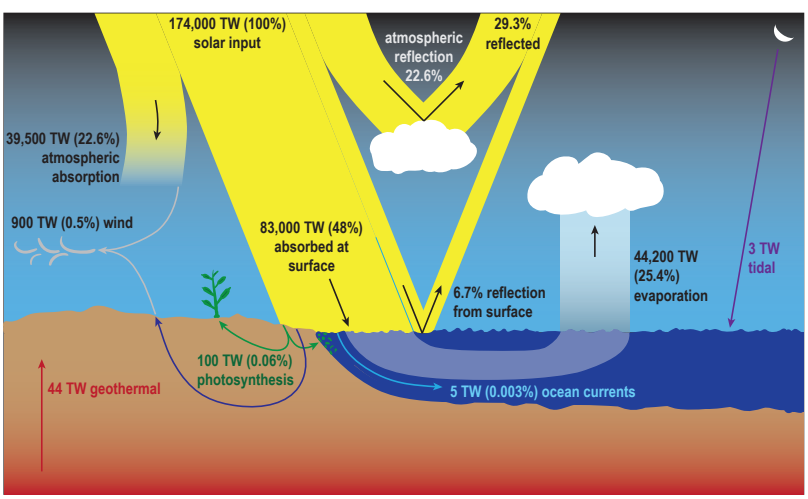


Figure 10.1: Energy inputs to the earth, ignoring the radiation piece (since that is an output channel). About 70% of incoming solar energy is absorbed by the atmosphere and land, while about 30% is immediately reflected back to space (mostly by clouds). About half of the energy absorbed at the surface goes into evaporating water, while smaller portions drive winds, photosynthesis (land and sea), and ocean currents. Additional non-solar inputs are geothermal and tidal in origin [63–65].

Example 10.3.1 Solar Input: Because we will encounter solar power flux many times in this textbook, this is a good opportunity to spell out some key numbers and concepts.

First, sunlight arriving at the top of Earth's atmosphere delivers energy at a rate of 1,360 Joules per second per square meter (1,360 W/m^2), which is known as the solar constant [4].

[4]: Kopp et al. (2011), "A new, lower value of total solar irradiance: Evidence and climate significance"

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*; *Freely available at*: *https://escholarship.org/uc/energy_ambitions*.

Treating Earth as a sphere of radius R, it has surface area $4\pi R^2$, but the sun doesn't "see" the whole surface at once. In fact, from the vantage point of the rays of sunlight intercepting Earth, what matters is the *projection* of Earth,[14] which just looks like a disk of area πR^2. Averaging the solar input across the entire planet therefore reduces the 1,360 W/m^2 by a factor of 4 to 340 W/m^2.

Not all the sunlight arriving at the top of the atmosphere makes it to the surface, so in practice, a typical location will receive an average[15] of about 200 W/m^2. This is a number that comes up often as a typical insolation, so is worth remembering.

14: Imagine taking a picture of a sphere situated across the room. The area the sphere takes up on the photo is πR^2, not the total curved surface area of $4\pi R^2$. See also Fig. 9.6 (p. 144).

15: . . . averaging over day, night, sun angles, and weather conditions.

Clouds and ice (mostly) reflect almost 30% of incoming sunlight, leaving 123,000 TW to be absorbed by land, water, and atmosphere in various forms (see Table 10.2). Virtually all of the energy hitting the surface goes to direct thermal absorption,[16] much of which then flows into evaporation of water—the starting point of the hydrological cycle. A tiny portion of the absorbed energy gives rise to wind, some of which will drive waves. An even smaller portion contributes to photosynthesis and supports essentially all life (biology) on the planet. And finally, a tiny fragment of the absorbed energy drives ocean currents. Table 10.2 tracks where the incoming solar energy goes, in several stages, also listing non-solar geothermal and tidal contributions. For comparison, the current energy scale of human activity is approximately 18 TW, while human metabolism,[17] is about 0.8 TW.

16: That is, heating (see Example 10.3.2); note that solar panels could intercept part of this energy flow.

17: Recall Ex. 5.5.2 (p. 74). 100 W per person times 8 billion people is 800 GW, or 0.8 TW.

Table 10.2: Earth's energy input budget. Symbols ⊙, ⊕, and ☾ represent Sun, Earth, and Moon, respectively. The second group breaks out the solar input into three pieces that add to the total in the row above. The third group all comes from absorbed energy—mostly at Earth's surface. The last group is not from radiant solar energy, so that percentages are in parentheses as they do not belong to the solar budget [63–65].

Category	Power (TW)	% solar	source	Comments
total solar input	174,000	100	⊙	the next 3 inputs come from here
surface absorption	83,000	47.9	⊙	heats surface; evaporates H_2O, powers life, wind, etc.
reflection to space	51,000	29.3	⊙	from clouds, ice; uncaptured energy
atmos. absorption	40,000	22.6	⊙	heats atmosphere, some to wind
evaporation	44,000	25.4	⊙ → ⊕ surf.	from surface absorption; hydrological cycle
wind	900	0.5	⊙ absorp.	from absorptions above, also makes waves
photosynthesis	100	0.06	⊙ → ⊕ surf.	fuels biology (life) on the planet
ocean currents	5	0.003	⊙ → ⊕ surf.	moves water around
geothermal	44	(0.025)	⊕	half original heat, half radioactive decays
tides	3	(0.0018)	☾ , ⊙	gravitational; mostly from Moon, some from Sun

Example 10.3.2 Solar Heating: How much would a black table[18] warm up sitting in full sun for ten minutes?

A nice round-number approximation of full overhead sunlight is that it delivers 1,000 W/m^2 to the ground. If we situate a table whose top surface area is 1 m^2, has a mass of 20 kg, and a specific heat capacity

18: . . . parameters defined below

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

of 1,000 J/kg/°C under full sun, we can proceed as follows.

The table will absorb 1,000 J per second,[19] and therefore receives 600,000 J over the course of ten minutes. Multiplying the specific heat capacity by the table mass means the table absorbs 20,000 J for every 1°C of temperature rise, and therefore would climb 30°C in 10 minutes, in this case. That's a little unrealistically high, because a real table would also have cooling influences from the air and infrared radiation. But the main point is to show how absorbed sunlight heats things up—like the Earth.

19: Because a Watt is a Joule per second and the table area is 1 m^2; the black property essentially means that it absorbs all light that hits it.

Box 10.2: Making New Fossil Fuels

We know from Chapter 8 that fossil fuels get their energy from ancient photosynthesis trapped in buried plant matter.[20] We also now have a figure for how much solar power goes into photosynthesis: 100 TW.

We can compare this to the power that goes into making *new* fossil fuels right now by noting that the entire fossil fuel resource contains roughly[21] 10^{23}J (page 127), and formed over something like 100 million years, or about 3×10^{15} seconds. Dividing the two gives a power of about 3×10^7 W, or 30 MW.[22]

Three neat insights come out of this. First, we currently burn fossil fuels at a rate of about 15 TW, which is 500,000 times faster than they are being replaced! It's like short-circuiting a battery in a dramatic explosion of power. Imagine charging a phone for 2 hours and discharging it 500,000 times faster: in 0.014 seconds! Now look at the extravagant lights of Las Vegas: should we be proud of the blaze of glory or appalled?[23]

Secondly, out of the total 100 TW photosynthetic budget on Earth, only 30 MW gets captured as fossil fuels, which is one part in three-million. Therefore, the chances that any given living matter on the planet today eventually ends up converted to fossil fuels is exceedingly slim.

Finally, if we only used fossil fuels at a rate of 30 MW,[24], then we *could* consider fossil fuels to be a renewable resource, as the sun/geology will slowly make more! So whether or not something is renewable also relies on the *rate* of use not exceeding the rate at which it is replenished.

20: In some cases, animals ate the plants first, but the energy starts in plants.

21: It is not important to nail down precise numbers for this exercise.

22: For reference, a single large university consumes energy at about this rate.

23: Also coming to mind is the Big Bay Boom in San Diego, July 4, 2012, when the entire fireworks display that was meant to last 15–20 minutes all went off in a few dazzling seconds. LMAO. Best ever!

24: This amount of power could supply only a single campus-sized consumer on Earth.

10.4 Renewable Snapshot

Table 7.1 (p. 106) already gave an account of the mix of energy use in the U.S., including many of the renewables. This section revisits those numbers, in slightly more detail.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

Source	qBtu	TW thermal equiv.	% of renewables	% of total
hydroelectricity	2.77	0.093*	24	2.7
wind	2.49	0.083*	22	2.5
wood	2.36	0.079	21	2.3
biofuels	2.28	0.076	20	2.3
solar	0.92	0.031*	8	0.9
waste (incineration)	0.49	0.016	4	0.5
geothermal	0.21	0.007	2	0.2
total	11.52	0.382	100	11.4

Table 10.3: U.S. Renewable energy consumption in 2018. The last column compares to the 101.3 qBtu total U.S. consumption in 2018 [34]. Asterisks denote thermal equivalent conversions. While the second column of numbers would be more naturally expressed in GW, the choice of TW is deliberate to emphasize how small renewable usage is compared to the available resource.

In 2018, roughly 11% of energy in the U.S. came from renewable resources. Table 10.3 lists the contributions from each, the data coming from the Annual Energy Review published by the U.S. Energy Information Administration for 2018. As introduced in Box 7.2 (p. 106), the EIA adopts the practice of assigning a thermal equivalent for each source, in qBtu, even if the source had nothing to do with a thermal process. The rationale is to put everything on the same footing as fossil fuels, for more direct quantitative comparisons. In doing so, they implicitly use the average thermal-to-delivered energy efficiency for fossil fuels of 37.5%.[25] In other words, it takes 100 units of thermal fossil energy to produce 37.5 units of useful work. If, then, a solar panel delivers 37.5 units of energy over some amount of time, it will be *called* 100 units of "input" (thermal equivalent), even though it only delivered 37.5 units.[26]

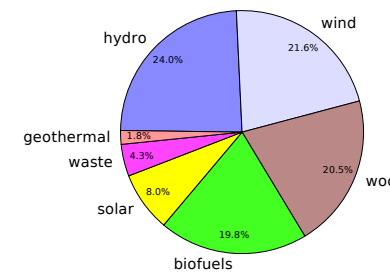


Figure 10.2: Same info as Table 10.3 in pie chart form.

Four forms dominate Table 10.3 (and Figure 10.2), each roughly equal in contribution. Lumping wood and biofuels into a generic "biomass" puts this aggregate category in the clear lead, accounting for nearly half of our renewable energy.[27] As noted, the entries in Table 10.3 are in qBtu of thermal equivalent, and these have been converted to TW, for easy comparison to Table 10.2.[28] A key take-away is how *tiny* the renewable energy numbers are compared to the natural flows in Earth's energy budget.

10.5 Upshot: Our Path Forward

We are now ready to plunge into learning about renewable energy resources. The topics are arranged according to ease of understanding the associated physics, which will be new to many students. So while solar is the most potent of the resources, its chapter follows those of hydroelectricity and wind, since its scheme for generating electricity[29] is likely the least intuitive of the three. Biologically-derived energy comes next—sharing its direct sunlight origin with solar. Following a foray into nuclear energy, a number of minor contributors that are unlikely to be important are relegated to a single chapter of misfits, for completeness.

After this, we will be in a position to assess the entire landscape of alternative energy options (Chapter 17). The book will then take a turn

25: 37.5% is the 2018 number from Appendix A6 of the AER.

26: Alternatively, it would have taken 100 units of fossil fuel energy to match the 37.5 units of energy delivered by the solar panel.

27: . . . the other half, roughly, coming from hydroelectricity and wind

28: ⚠ The items with asterisks in Table 10.3 will not match production numbers in later chapters for non-thermal resources that are not expressed in thermal-equivalent terms.

29: Hydroelectricity and wind share much in common with conventional power generation schemes, in that they use turbines and generators.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

away from physics and address how all this new information might fit into future plans at societal and personal levels.

10.6 Problems

1. Based on what you already know or suspect about the alternative energy sources listed in Section 10.1, which ones do you suspect[30] are pollution free (no emissions, no waste)?

2. Based on what you already know or suspect about the alternative energy sources listed in Section 10.1, which ones do you suspect[31] have *no* negative environmental impacts? Explain your reasoning.

3. Which of the truly renewable ("Yes" without asterisk) entries in Table 10.1 are limited (finite) in terms of the scale[32] we could achieve on Earth (explain the nature of the limit), and which are not limited (and why)?

4. Which four entries in Table 10.1 do not ultimately trace entirely to solar energy input to Earth?

5. If, for some terrible reason, the sun ceased shining and humans managed to survive for 1,000 years longer,[33] what options would be left for obtaining energy?

6. About 30% of the 1,360 W/m^2 solar power arriving at Earth is immediately bounced back without a trace. Of the part that remains, when distributed/averaged around the whole sphere,[34] what is the average energy deposition rate per square meter into the earth system?

7. Based on the numbers in Table 10.2, Figure 10.1, and in the text, what fraction (in percent) of all biological activity on Earth do (all) humans represent, from a metabolic energy standpoint?

8. Based on the numbers in Table 10.2, Figure 10.1, and in the text, what fraction (in percent) is human societal energy production (all activities; fossil fuels, etc.) compared to all biological activity on the planet?

9. What fraction (in percent) of the solar energy *that is absorbed by Earth's surface* goes into evaporation (the hydrological cycle; refer to Table 10.2, and Figure 10.1)?

10. If we add up the five smallest pieces in Table 10.2, what fraction (in percent) does this set of energy flows represent in comparison to the total solar energy *absorbed by the earth system*?

11. The percentages of absorption and reflection in Figure 10.1 represent averages over weather conditions. On a clear day with the

30: Don't worry about correctness, as this is an opportunity to think and gauge current understanding.

31: Don't worry about correctness, as this is an opportunity to think and gauge current understanding.

32: . . . not how long it can last; how much power is available

33: It's nearly impossible to imagine that survival is at all realistic in this dire situation.

34: ⚠ account for projected vs. total surface area

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

sun overhead, a solar panel has access to more light than the 48% depicted reaching the ground. If we can also capture the portion that on average is reflected by the ground and clouds,[35] what might we expect for the rate of energy reaching the ground (in W/m^2) if the input at the top of the atmosphere is 1,360 W/m^2?

12. Similar to Example 10.3.2, how much warmer would you expect a 10 cm deep puddle of water[36] to get after an hour in full sun, if the water absorbs all the energy and does not lose it to the environment?

13. Comparing Table 10.2 and Table 10.3, what is the most striking mismatch,[37] in terms of large potential vs. small contribution to U.S. energy consumption?

14. In one day, a typical residential solar installation might deliver about 10 kilowatt-hours of energy. Meanwhile, a gallon of gasoline contains about 37 kWh of thermal energy. But the two ought not be directly compared, as burning the gasoline inevitably loses a lot of energy as heat. Correcting the solar output to a thermal equivalent,[38] how many gallons per day of gasoline could it displace?

15. The U.S. occupies about 2% of Earth's surface area, thus collecting only part of the 83,000 TW absorbed at the surface. Compare the amount of solar power received by the U.S. to the *total* renewable power in Table 10.3.[39] By what factor are we "underutilizing" the solar input, under the simplifying assumption that we could use all of it?

35: Recall, it's a clear day; no clouds.

36: ⓘ Water has a density of 1,000 kg per cubic meter and a specific heat capacity of 4,184 J/kg/°C. Make up whatever area you want; the answer should be the same for any pick.

37: In other words, for those sources that can be matched within the tables, which stands out as the most underutilized relative to its input numbers?

38: ⓘ Use the 37.5% factor discussed in the text.

39: . . . since it's almost all solar in ultimate origin

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.



11 Hydroelectric Energy

11.1 Grav. Potential Energy . . . 173
Comparison to Other Forms 174
11.2 Hydroelectric Energy 175
Theoretical Potential 177
11.3 Hydropower in the U.S. . . 178
11.4 Global Hydropower 180
11.5 Upshot: Pros and Cons . . . 180
11.6 Problems 181

Energy has been harnessed from flowing water for ages. Milling operations were often located on streams so that water could turn a wheel attached to grinding machinery. Today, captured water flow is a significant contributor to electricity worldwide in the form of hydroelectricity. The U.S. gets about 2.8% of its current energy (and 7% of electricity) from hydropower. Globally, hydropower accounts for about 9% of energy, or 16% of electrical production (Table 7.2; p. 107).

Hydroelectricity taps into the solar-driven evaporation cycle, relying on the gravitational energy embodied in water lifted onto the land from lower bodies of water. In other words, solar energy lifts water, giving it gravitational potential energy, which is captured and converted to electrical energy.

While hydroelectric power is a simple and low-tech form of renewable energy that has been heavily exploited for over a hundred years, it is not one that is easy to expand beyond its current level of usage. This chapter will provide a better understanding of this mainstay of the renewable portfolio and its likely role in our future.

Tidal power works in exactly the same way as hydroelectricity, but is a niche player and will be deferred to Section 16.2.

11.1 Gravitational Potential Energy

Gravitational force is incredibly weak. It may not seem so from daily experience, but consider the fact that a magnet held in your hand can lift a paperclip—overwhelming the gravitational pull of the *entire* earth! By comparison, electromagnetic forces are forty orders-of-magnitude stronger than gravitational force. We don't tend to notice because electric charges tend to balance out so that gravity is the most obvious force in our daily lives.[1]

1: Somewhat ironically, we only *feel* gravity because of a much stronger electromagnetic force that prevents us from falling through the floor. Electrons in the atoms in the floor and our feet repel each other to prevent free-fall—a weightless state in which gravity can't be felt!

Boulder Dam. Photo Credit: Tom Murphy

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

We know intuitively that lifting a massive object requires work, and thus energy. In fact, since work is defined as a force times distance,[2] and the force of gravity on an object follows Newton's second law, $F = ma$,[3] the force we need to exert on an object to lift it against gravity is called its weight, and is $W = mg$, where $g = 9.8\,\text{m/s}^2 \approx 10\,\text{m/s}^2$ is the acceleration due to gravity on the surface of the earth. Thus, to lift an object through height, h, requires an energy input of this force, W, times the height, h. We call this gravitational potential energy since the energy put in to lift something can be released later if the mass is allowed to fall or be lowered. An early use of gravitational potential energy was in the form of weights on chains in old clocks.

2: Recall Def. 5.1.1 (p. 68).

3: Force is mass times acceleration.

Definition 11.1.1 *gravitational potential energy is computed as*

$$\text{G.P.E.} = mgh, \tag{11.1}$$

where m is the mass in kg, $g \approx 10\ m/s^2$ is the acceleration due to gravity,[4] and h is the height to which the mass is raised, in meters. The result is in Joules.

4: Some may remember it more pedantically as $9.8\ \text{m/s}^2$, but for the purposes of this book, $10\ \text{m/s}^2$ will do nicely. Note that choosing this number implies that we are concerned only with gravitational energy on the earth's surface.

Most commonly, gravitational potential energy is converted to kinetic energy as an object falls: slowly at first but accelerating as more potential energy is converted to kinetic energy while the object gains speed (see Fig. 5.1; p. 70). Only the vertical distance matters in computing gravitational potential energy: sideways motion does not operate against the gravitational force. Sliding a crate across a flat, level floor does take work to overcome friction, but that energy is converted to heat and cannot be later returned in useful form.[5] In this case, the crate has gained no gravitational potential energy, since its height never changed.

5: We often say in this case that the energy is "lost." But energy is strictly conserved—not created or destroyed—so it is never *really* lost, it just escapes into a non-useful form.

Example 11.1.1 Lifting a 20 kg box of books, whose weight is therefore $W = mg \approx 200$ N,[6] from the floor to a high shelf through a vertical distance of 2 m involves an energy expenditure of $mgh \approx 400$ J (Figure 11.1). We would say the box gained 400 J of potential energy.

If the person doing the work is exerting energy at a rate of 200 W (200 J/s), it will take two seconds to complete the action.

If the box later falls off the shelf and hits a 1.5 m tall person on the head, the box has *lost* 100 J (20 kg × 10 m/s^2× 0.5 m) of potential energy (now kinetic) by the time it hits the person's head.

6: ... again using $g \approx 10\,\text{m/s}^2$

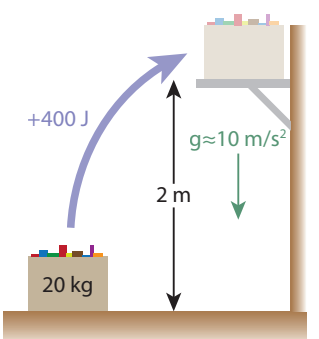


Figure 11.1: Diagram for Example 11.1.1.

11.1.1 Comparison to Other Forms

To give a flavor of how weak gravitational potential energy is compared to other familiar forms of energy storage, we will consider the energy content in a standard AA alkaline battery and in a similar volume of gasoline. So we're talking about something approximately the size of

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*; *Freely available at*: *https://escholarship.org/uc/energy_ambitions*.

a small finger. We want to know how much mass must be lifted to yield the same amount of gravitational potential energy as is contained in a battery or equivalent volume of gasoline. In the comparison, we will imagine having a hoist that can lift a large mass[7] 4 m high—about house-height.

A standard AA battery cell has a charge rating of 2.5 Ah[8] and operates at about 1.5 V. Following the development in Sec. 5.8 (p. 76), we multiply these two numbers to get 3.75 Wh, translating to 13.5 kJ. Equating this to mgh, where we know $g \approx 10\,\mathrm{m/s^2}$ and $h = 4\,\mathrm{m}$, we find that $m \approx 340\,\mathrm{kg}$. That's really heavy—about the mass of 4–5 people.[9] Meanwhile, the AA battery is a puny 0.023 kg. Reflect for a moment on this comparison, visualizing 340 kg lifted 4 m above the ground providing the same amount of energy as a AA battery held in your hand.

Gasoline is even more extreme. At an energy density around 34 kJ per mL of volume, filling a AA-sized cup[10] with gasoline yields about 250 kJ of energy.[11] Performing the same computation, we would need to lift over 6,000 kg (6 metric tons) to a height of 4 m to get the same energy content. Typical cars have masses in the 1,000–2,000 kg range, so we're talking about something like 4 cars! One caveat is that we are not typically able to convert the thermal energy in gasoline[12] into useful work at much better than 25%, while gravitational potential energy can be converted at nearly 100%. Still, being able to lift 1,500 kg[13] to a height of 4 m using the energy in 7 mL of gasoline is rather impressive, again emphasizing that gravitational potential energy is pretty weak. It only amounts to significance when the masses (volumes) of water are rather large.

7: . . . a rock, for instance

8: The number is usually given as, e.g., 2,500 mAh (milli-amp-hours).

9: Amuse yourself by picturing 4–5 people slung haphazardly into a net and hoisted to roof height—a very odd (and grumpy?) replacement for a AA battery.

10: . . . just over 7 mL

11: Thus, gasoline is nearly 20 times as energy-dense as a AA battery *by volume*. Usually, we will discuss energy density *by mass*, in which case the ~5× denser battery provides nearly 100× less energy per gram than does gasoline.

12: . . . via combustion; see Sec. 6.4 (p. 88)

13: . . . now just one car, rather than four; it means this small volume of gasoline can propel a car up a 4 m hill

11.2 Hydroelectric Energy

The basic idea behind hydroelectricity is that water in a reservoir behind a dam (Figure 11.2) creates pressure at the base of the dam that can force water to flow through a turbine that drives a generator to make electricity—sharing elements of Fig. 6.2 (p. 90) but spinning the turbine by water flow instead. The amount of energy available works out to be the gravitational potential energy corresponding to the height of water at the lake's surface relative to the water level on the other side. It's as if dropping the water from the surface to the turbine and asking how much potential energy it gave up in the process. In reality, water is not dropping from the lake surface, but the force on the water at the turbine is determined by the height of water above it: the "pressure head," as it is called. The process is highly efficient, approaching 90% capture of the potential energy in the water delivered as electrical power from the generator.

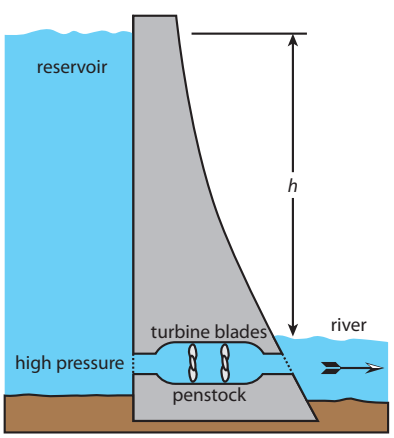


Figure 11.2: Cross section of a dam, holding back a reservoir of water at height, h, over the downstream river.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

Box 11.1: Why So Efficient?

Achieving 90% efficiency is superb! Electric motors and generators[14] can be > 90% efficient in converting between mechanical energy (rotation) and electrical energy. When coupled with low-friction turbines, dams just have very little loss—unlike thermal sources where most of the energy is unavoidably lost (for reasons covered in Sec. 6.4; p. 88).

14: Fundamentally, motors and generators are nearly identical in concept and construction.

Example 11.2.1 To compute the power available from a hydroelectric plant, we need to know the *height* of the reservoir and the *flow rate* of water—usually measured in cubic meters per second. The density of water is, conveniently, 1,000 kg/m^3 (Figure 11.3), so that if we consider a dam having a flow rate of 2,000 m^3/s and a reservoir height of 50 m, we can see that every second of time will pass 2×10^6 kg of water,[15] and the associated potential energy is $mgh \approx 10^9$ J. If each second delivers 1 GJ of energy, the power available is 1 GJ/s, or 1 GW. At an efficiency of 90%, we get to keep 900 MW of electrical power.

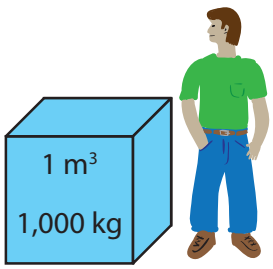


Figure 11.3: One cubic meter of water has a mass of 1,000 kg.

15: Flow rate times density gives mass per second: 2,000 m^3/s times 1,000 kg/m^3 = 2×10^6 kg/s

The largest hydroelectric facility in the world is the Three Gorges Dam in China, rated at an astounding 22.5 GW. The largest in the U.S. is the Grand Coulee on the Columbia River, producing a maximum of 6.8 GW. The iconic Boulder Dam (a.k.a. Hoover Dam) is just over 2 GW.

Look at the Wikipedia page on largest hydroelectric power stations [66] for a complete list.

Note that flow rates vary seasonally with rainfall, so that dams cannot always operate at full capacity. In fact, the U.S. has about 80 GW of *capacity* installed, but operates at an annual average of about 33 GW. This implies a typical "capacity factor" around 40%.

Definition 11.2.1 *A **capacity factor** is the ratio of actual performance over time to the peak possible performance—or average output divided by maximum output, expressed as a percentage.*

Example 11.2.2 Boulder (Hoover) Dam on the Colorado River is listed in [66] as having a capacity of 2,080 MW and an annual production of 4.2 TWh. What is its capacity factor?

We just need to turn the 4.2 TWh in a year into an average delivered power. Following the definition of a watt-hour, we note that all we really have to do is divide 4.2×10^{12} Wh[16] by the number of hours in a year: 24 times 365, or 8760.

4.2×10^{12} Wh/8760 h $\approx$ 480 MW average power. Dividing this by 2,080 MW (max capacity) gives a 23% capacity factor.

16: 1 TWh is 10^{12} Wh.

As we saw in Fig. 7.5 (p. 108) and Table 10.3 (p. 170), hydroelectricity in the U.S. accounts for 2.7% of the nation's total energy consumption, corresponding to about 33 GW of production. Globally, hydroelectric production averaged 477 GW in 2017. By comparison, Table 10.2 (p. 168)

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

indicates that 44,000 TW of solar input goes into evaporation and the hydrological cycle. Why, then, are we only able to use 0.477 TW (0.001%) of this bounty? Is this a great, untapped renewable resource?

11.2.1 Theoretical Potential

To understand the giant mismatch between solar input and hydroelectric development, we first need to study evaporation.

Definition 11.2.2 *The **heat of vaporization** of water is about 2,250 J per gram, meaning that every gram of water that goes from liquid to gas (vapor) requires an energy input of ~2,250 J.*

Box 11.2: Vaporization is Serious Energy

To put this in perspective, it takes 100 calories (418 J) to bring one gram of water from freezing to boiling temperature. Then it takes another 2,250 J to evaporate the water, which is a far larger quantity. This is why water in a pot does not all flash into steam once the water reaches 100°C, as it would if the evaporation energy was very small. Instead, a boiling pot will retain water for a good while as energy continues to be applied before all boiling away.

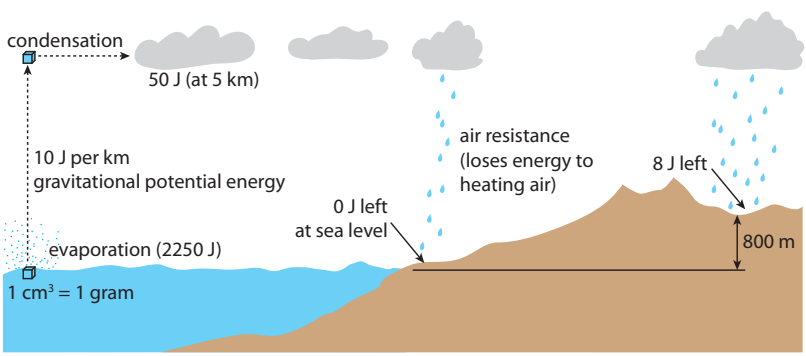


Figure 11.4: The hydrological cycle. Sunlight evaporates water from the surface, at a cost of 2,250 J per gram. Each kilometer of height the gram of water gains in forming clouds costs an additional 10 J. When rain falls on terrain, most of the gravitational potential energy is spent, but on average retains 8 J—based on an average land elevation of 800 m. The 2,250 J of evaporation energy is released as heat when the water condenses into clouds.

So let's follow the energetics of a gram of water[17] on its journey to a hydroelectric dam—most of which is represented in Figure 11.4. First, the sun injects 2,250 J to evaporate that gram. Then let's say it gets lofted to 5 km.[18] The gravitational potential energy, mgh, comes to $0.001 \times 10 \times 5000 = 50$ J. That's only 2% of the amount that went into evaporation.[19]

When the water condenses in the cloud, it *releases* 2,250 J of thermal energy into the cloud/air, then falls back to the ground as rain, offering 50 J of still-available energy. If it falls on the ocean, where it presumably started, it gives up all 50 J of gravitational potential energy into useless forms.[20] But if it falls on land—higher than sea level—it retains some gravitational potential, based on how high that land is above sea level.

17: . . . one cubic centimeter

18: . . . typical cloud height

19: The sun must, in total, supply 2,300 J to evaporate and lift the gram of water, and only 50 J of the 2,300 J is kept as potential energy.

20: . . . heat through air resistance and collision with the ocean surface

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*; *Freely available at*: *https://escholarship.org/uc/energy_ambitions*.

On average, terrain is about 800 m above sea level, so each gram that falls on land has an average of 8 J left as available energy. But only 29% of the earth's surface is land, so that the gram of water we're tracking preserves about 2 J of energy, on average.[21]

21: . . . reduced from 8 J since most rain falls back onto ocean

We're down to only 0.1% of the input solar energy—2 J out of 2,300 J input—so that the theoretical hydroelectric potential might be about 44 TW: reduced from the 44,000 TW input. But only a small fraction of rain flows into rivers suitable for damming. And once dammed, a typical dam height is in the neighborhood of 50 m, knocking us down even further. Much of the journey from terrain to reservoir involves losing elevation in streams too small to practically dam, or just seeping through the ground. In the end, perhaps it is not surprising that we end up in the sub-TW regime globally.

The 90% efficiency of a hydroelectric dam is now contextualized a bit better. That last step is pretty efficient, but the overall process is extremely inefficient. Still, it takes relatively little effort to exploit, and provides real power. Efficiency is not everything.

Detailed assessments [67] of hydroelectric potential globally estimate a *technically* feasible potential[22] around 2 TW, but only half of this is deemed to be economically viable. Recall that 477 GW, or about 0.5 TW, is delivered globally, which is therefore about half of what we believe to be the practical limit of ~1 TW. Thus we might not expect more than a factor-of-two expansion of current hydroelectricity as possible/practical. The low-hanging fruit has been plucked already, capturing about half of the total practical resource.

[67]: (1997), *Study on the Importance of Harnessing the Hydropower Resources of the World*

22: . . . if cost is no barrier

Compared to the 18 TW global scale of energy use, hydroelectricity is not poised to assume a large share at this level, unless the overall scale of energy use is reduced substantially. Let's say this more visibly: **hydroelectric power cannot possibly come close to satisfying present global power demand.**

11.3 Hydropower in the U.S.

Hydroelectric power is not available to the same degree everywhere. Geography and rainfall are key factors. This brief section serves to present a snapshot of the distribution and qualities of hydroelectric power generation in the United States. We start with Figure 11.5, showing the average hydroelectric power generated in each state, the top four states being listed in Table 11.1. These four states account for 56% of hydroelectricity in the U.S., and the next states on the ranked list drop to 1 GW or lower. Most of the California generation is in the northern part of the state, effectively as part of the Pacific Northwest region.

Table 11.1: Top hydroelectric states.

State	Production (GW)
Washington	8.9
Oregon	3.8
California	3.0
New York	2.9
U.S. Total	33

To get a sense for how concentrated different sources are, we will make a habit of examining power density for renewable resource implementations. Figure 11.6 indicates the state-by-state density of hydroelectric power generation,[23] just dividing generation by state area. No state exceeds 0.05 W/m^2, which can be contrasted to insolation values (see Ex. 10.3.1; p. 167) of ~200 W/m^2. Globally, total land area is about

23: . . . based on actual generation, not installed capacity

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*; *Freely available at*: *https://escholarship.org/uc/energy_ambitions*.

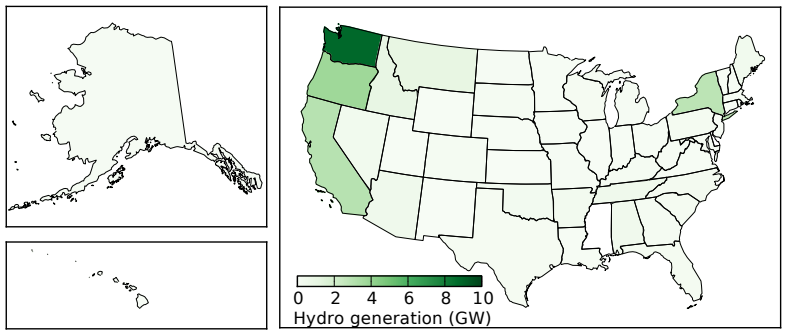


Figure 11.5: Average U.S. hydroelectric power delivered for each state, showing strongly along the west-coast, plus New York. And yes, Alaska really is that big.

1.25×10^{14} m^2, so that a total hydroelectric potential of 2.5 TW[24] would yield 0.02 W/m^2. Therefore, the state of Washington stands out as unusual, having *already developed* a generation capacity 2.5 times larger than the upper-end global average expectation. In other words, most of the world cannot emulate what nature has provided in Washington.[25] Not all places have the same available resources.

24: This is higher than estimated potential developed resources, but mathematically convenient here.

25: Washington's hydroelectric dominance owes largely to the presence of the mighty Columbia River, rather than human factors.

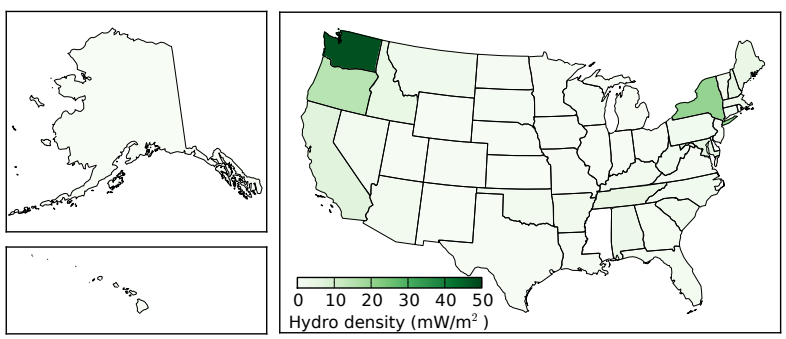


Figure 11.6: U.S. hydroelectric power per area delivered for each state, giving a sense of how concentrated the resource is. The units are milliwatts per square meter, peaking at 48 mW/m^2 for Washington.

Next, we look at hydroelectric generation per capita. Figure 11.7 shows the result. In this view, the states of the Pacific Northwest really pop up, and New York dims relative to its by-area showing. The contrast between Figure 11.6 and Figure 11.7 is effectively reflecting population density: large, sparsely-populated states[26] show up more prominently on the per-capita map than the per-area map.

26: Montana, Idaho, even Alaska

Finally, for completeness, we look at the capacity factors of hydroelectric installations, by state. The total installed capacity in the database used for these plots is 77.6 GW spread among 1,317 dams, while producing an annual average of 28.1 GW—corresponding to an overall capacity factor of 0.36. Figure 11.8 shows how this distributes around the country. Since the Pacific Northwest dominates in installed hydroelectric power, it largely determines the overall capacity factor. Iowa stands out as having a high capacity factor, but only has 0.153 GW of installed capacity.[27] Contrast this to Washington, having an installation capacity of 20.7 GW.[28]

27: . . . delivering an average of 0.114 GW in 8 dams, dominated by the 0.125 GW Keokuk dam

28: . . . delivering an average of 8.9 GW spread across 65 dams

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

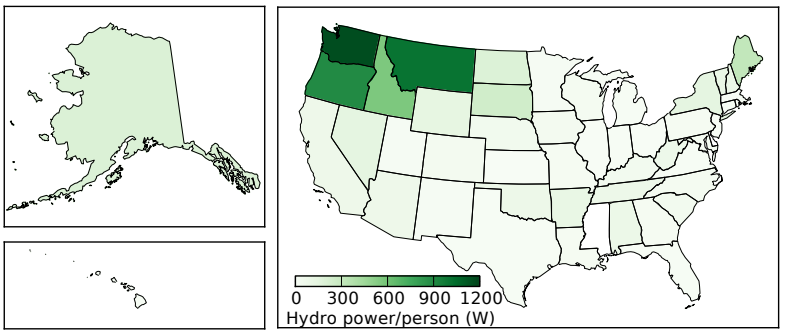


Figure 11.7: U.S. hydroelectric power per population within each state, indicating which residents receive the most hydroelectric power. The Pacific Northwest wins.

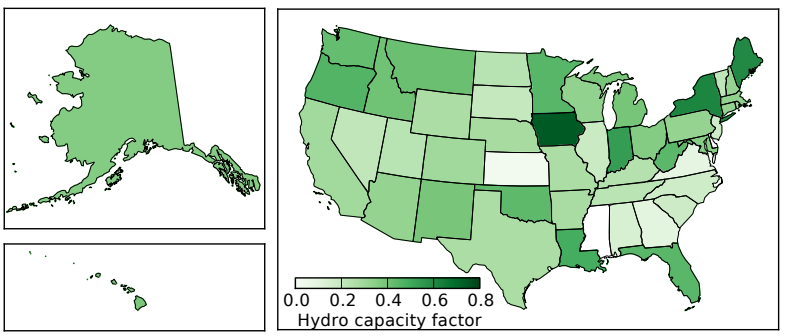


Figure 11.8: U.S. hydroelectric capacity factors by state. The national average (weighted by generation) is about 0.4, meaning that over the course of a year the dams deliver 40% of their rated capacity—subject to seasonal water flow.

11.4 Global Hydropower

This section provides a brief snapshot of hydroelectric production globally, which we saw at the end of Section 11.2.1 amounts to 477 GW. Figure 11.9 shows which countries have the most hydroelectricity, the corresponding numbers appearing in Table 11.2—including the percentage of electricity derived from hydroelectric sources within the country. Notice that Norway, Venezuela, Brazil, and Canada derive more than half their electricity demand from hydroelectricity. Keep in mind that electricity is not the whole energy story for a country, as Fig. 7.2 (p. 105) made clear.

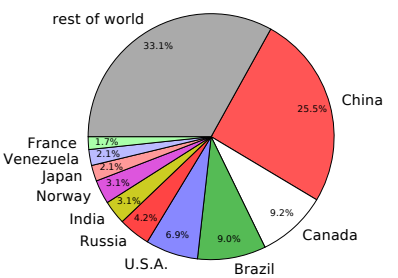


Figure 11.9: Global allocation of hydroelectric power.

Rank	Country	GW	% elec.	Rank	Country	GW	% elec.
1.	China	122	19	6.	India	15	10
2.	Canada	44	58	7.	Norway	15	96
3.	Brazil	43	63	8.	Japan	10	8
4.	U.S.	33	6.5	9.	Venezuela	10	68
5.	Russia	20	17	10	France	8	12

Table 11.2: Top ten hydroelectric producers in the world, accounting for two-thirds of the hydroelectricity production worldwide.

11.5 Upshot: Pros and Cons of Hydro

The two lists below provide some pros and cons to hydroelectric power that are relevant to our assessment of its value to our mix of renewable energy options. First, the positive attributes:

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

- Natural source, solar-driven, without waste products or pollution;[29]
- Technologically simple, therefore straightforward to implement and maintain;
- High efficiency, turning 90% of available energy into useful electricity;
- Good baseline of steady power over daily timescales;[30]
- Life-cycle CO_2 emissions only 4% that of traditional fossil fuel electricity [68];
- Facilities[31] can last a century or longer;
- While not directly related to energy, dams can provide flood control and reliable water supplies.

29: . . . aside from construction and decommissioning aspects

30: . . . no imposed short-term fluctuations in available power, as happens for solar or wind

31: . . . the dam itself, at least; turbines and generators will need periodic replacement

And some of the downsides that may discourage further development:

- Silt can build up behind dams displacing the reservoir, eventually rendering them useless and dangerous;
- Requires the semi-permanent flooding of an ecological habitat, of varying quality and value;
- Seasonal variability of available power, often by a factor of ten;
- Defunct or poorly maintained facilities represent a dam-burst danger to downstream residents;
- Blocks salmon runs and impacts the health of both oceanic and forest ecosystems;
- As the distribution maps show, hydroelectric is not a viable option everywhere: the combination of terrain[32] and rainfall is needed.

32: . . . mountains or canyons to hold the reservoir

On balance, our society has embraced hydroelectricity as a clean and reliable resource. One can view it as nature's low-hanging fruit, partly evidenced by how early it was adopted on a large scale. It is likely to remain an attractive form of energy as we face increasing pressures to migrate away from carbon-based fuels.[33] It will not, however, be able to provide an avenue for wholesale replacement of fossil fuels given its limited scale and electric-only nature. Only if our overall energy demand is reduced substantially will it form a large fraction of our portfolio.

33: . . . whether due to resource limits or climate change action

11.6 Problems

1. If a 70 kg person climbs 10 flights of stairs, each flight 3 m high, how much potential energy have they gained?

2. If an 80 kg person is capable of delivering external mechanical energy at a rate of 200 W sustained over several minutes,[34] how high would they be able to climb in two minutes?

34: It is hard to keep up 200 W for too long.

3. A 10 kg box is lifted 2 m off the floor and placed on a frictionless horizontal conveyor to take it 30 m across a warehouse. At the end of the conveyor, it is lowered 1 m where it ends up on a shelf.[35]

35: The shelf is therefore 1 m off the same (original) floor.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

How much *net* gravitational potential energy was given to the box from the start to the end of the process?

4. A standard AA battery cell stores about 13.5 kJ of energy. At a mass of 23 g each, how high would you have to lift the battery to get the same amount of gravitational potential energy?

ⓘ The result emphasizes how weak gravitational potential energy is.

5. A gallon of gasoline contains about 130 MJ of chemical energy at a mass of around 3 kg. How high would you have to lift the gallon of gasoline to get the same amount of gravitational potential energy? Compare the result to the radius of the earth.

ⓘ The result emphasizes how weak gravitational potential energy is.

6. Problem 5 used one gallon of gasoline to compute the equivalent height for gravitational potential energy. Does the result depend on what volume of gasoline we selected? Make an airtight argument as to why or why not. Solving symbolically[36] may be a helpful route—but not the only one.

Hint: the energy density in MJ per kg is a property of the gasoline.

36: . . . using variables/symbols

7. A typical American household uses approximately 30 kWh per day of electricity. Convert this to Joules and then imagine building a water tank 10.8 m above the house[37] to supply one day's worth of electricity.[38] How much mass of water is this, in kg? At a density of 1,000 kg/m^3, what is the volume in cubic meters, and what is the side length of a cube[39] having this volume? Take a moment to visualize (or sketch) this arrangement.

37: Pretend all the water is at this height.

38: Assume 100% conversion efficiency for mathematical convenience.

39: . . . cube root of volume

8. The biggest hydroelectric installation in the U.S. is the Grand Coulee dam on the Columbia River. The enormous flow rate reaches its maximum at 4,300 m^3/s, and the dam (reservoir) height is 168 m. At an efficiency of 90%, at what rate is this dam capable of producing hydroelectric power (in GW[40])? Don't forget the density of water and that $g \approx 10\,\mathrm{m/s^2}$.

40: ⓘ For comparison, a large nuclear reactor typically produces about 1 GW of electrical power.

9. The Hoover Dam (also called Boulder Dam) on the Colorado River is rated at 2.08 GW when the flow is at its maximum rate of 1,280 m^3/s. How high is the reservoir if the efficiency of the installation is 90%?

10. A dam 50 meters high is constructed on a river and is delivering 180 MW at some moment in time. What is the flow rate of water, in cubic meters per second, if the facility converts gravitational potential energy into electricity at 90% efficiency?

11. A hydroelectric facility is built to deliver a peak power of 1 GW, which it manages to do for three months of the year during the spring snow-melt. But for three months in summer, it drops to 700 MW, then 500 MW for three months in fall. In winter, it drops way down to 200 MW for three months. Using the concept of the capacity factor (Definition 11.2.1), what is the annual average capacity factor for this facility?

This stepwise behavior is not exactly realistic, but need not be to apply the concept correctly.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

12. While the Chief Joseph Dam on the Columbia River can generate as much as 2.62 GW (2.62×10^9 W) of power at full flow, the river does not always run at full flow. The average annual production is 10.7 TWh per year (10.7×10^{12} Wh/yr). What is the capacity factor of the dam: the amount delivered vs. the amount if running at 100% capacity the whole year?

    Hint: Multiplying peak power by hours in a year will result in units similar to Wh/yr for direct comparison.

13. The Robert Moses Niagara dam in New York is rated at 2,429 MW[41] and has a high capacity factor of 0.633. How many kWh does it produce in an average day, and how many homes would this serve at the national average of 30 kWh/day?

    41: . . . peak power capacity

    Hint: Convert average power to kW then multiply by hours.

14. The Robert Moses Niagara dam from Problem 13 is 30 m high. What is the peak flow rate, in m^3/s, if it can produce full capacity power (2.43 GW electrical output) while converting gravitational potential energy to electricity at 90% efficiency?

15. It takes 2,250 J to evaporate each gram of water, while only taking about 330 J to raise the temperature of water from room temperature to the boiling point. If it takes 10 minutes to bring a pot of water from room temperature to boiling, how much additional time will it take to boil off (evaporate) all the water if injecting energy at the same rate the whole time?

16. Starting at 44,000 TW of solar input to the hydrologic cycle, parallel the development in Section 11.2.1 by computing the power remaining at each stage[42] if, for each gram of water:

    a) water is evaporated and lifted to 5 km height;[43]
    b) 30% of the water falls on land where collection is possible;
    c) the typical land height is 800 m;
    d) only 20% of the water makes it to dammable locations;
    e) only 50 m of height (of the original 800 m average) is left for the dam.

    By this analysis, how much hydroelectric power is theoretically possible, globally?

    42: Each stage will knock down the number further; report each in TW.

    43: This is the largest jump, keeping only 50 J out of every 2,300 J.

17. Fig. 10.1 (p. 167) indicated that about 44,000 TW globally goes into evaporating water. We can turn this into an estimate of how much rain we expect per year, on average. The simplest way to do this is to think of a single square meter of ocean surface, receiving an average evaporation input power of 120 W.[44] Each millimeter of of water depth across our square meter has a volume of 1 L, or a mass of 1 kg. At a steady input of 120 W,[45] how many millimeters of water are drawn off in a year? That same amount will come back down *somewhere* as precipitation.

    44: $44{,}000 \times 10^{12}$ W divided by 3.7×10^{14} m^2 of ocean surface is 120 W/m^2.

    45: ⓘ The steady 120 W is already accounting for day/night: this is a time average.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.



12 Wind Energy

12.1 Kinetic Energy 184
12.2 Wind Energy 185
Wind Turbines 187
12.3 Wind Installations 191
12.4 Upshot: Not Overblown . . 193
12.5 Problems 194

Wind energy has made significant inroads for electricity production in the U.S. and globally. Today, the U.S. gets about 2.5% of its energy (and 6.5% of electricity) from wind. Globally, wind accounts for about 2.6% of energy, or 4.8% of electrical production (Table 7.2; p. 107).

The basic technology of harnessing wind power is rather old, powering ships, milling operations, and water pumps for centuries. Today, the predominant use of wind drives generators to make electricity.

Just as hydroelectric power is related to the very basic form of gravitational potential energy, wind is connected to another simple and easy-to-understand form: kinetic energy. This chapter first develops familiarity with kinetic energy, then explores how we get energy from wind, how much we get, and future prospects.

12.1 Kinetic Energy

An object in motion carries kinetic energy equal to one-half its mass times the velocity squared.

Definition 12.1.1 *The **kinetic energy** of a mass, m, moving at velocity, v, is*

$$\text{K.E.} = \frac{1}{2}mv^2. \tag{12.1}$$

Example 12.1.1 A 70 kg person walking at a brisk 2 m/s would have a kinetic energy of 140 J.

Pushing a 50 kg ice skater from rest at a power of 75 W for 3 seconds would impart 225 J of kinetic energy, resulting in a speed of 3 m/s.

For examples like these, framed as statements and not questions, you can practice solving several types of problems by covering up one number and then solving for it using still-available information. So they provide several examples in one!

Photo credit: Tom Murphy

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

Often, we use energy sources to *deliver* kinetic energy, as in moving planes, trains and automobiles.

Example 12.1.2 A 1,500 kg car moving at freeway speeds (30 m/s) has a kinetic energy around 675 kJ.

Getting up to this speed from rest in 5 seconds would require a power of 135 kW,[1] equating to 180 horsepower.[2]

1: Do the calculation yourself to follow along.

2: Recall that 1 hp is 746 W. Indeed, it takes a powerful engine to provide this level of acceleration.

But we can also go the other direction and convert kinetic energy into different forms of energy[3] for versatile use. Most commonly, we turn kinetic energy into electrical potential energy (a voltage) that can drive a circuit. At this point, the energy can be used to toast a bagel, charge a phone, or wash clothes.

3: See Table 5.2 (p. 70), for examples.

The method for converting kinetic energy into electricity is usually accomplished by transferring kinetic energy in a moving fluid[4] into rotation of a shaft by way of a turbine—essentially fan blades. The spinning shaft then turns an electric generator, which consists of the relative motion between magnets and coils of wire, and is essentially the same construction/concept as an electric motor run in reverse.

4: In this sense, "fluid" is a general term that can mean a liquid or even air.

Hydroelectric installations do the same thing—turning a shaft via blades of a turbine—even though we framed the energy source as one of gravitational potential energy. Within the dam's turbine, the water acquires kinetic energy as it flows from the reservoir to the outlet. Wind energy acts in much the same way, converting kinetic energy in the moving air into rotational motion of a fan/turbine whose shaft is connected to a generator located behind the blades.

12.2 Wind Energy

It is tempting to think of air as "empty" space, but at sea level air has a density of 1.25 kg per cubic meter ($\rho_{\text{air}} \approx 1.25\,\text{kg/m}^3$). Let this sink in visually: imagine a cubic meter sitting next to you (as in Figure 12.1). The air within has a mass of 1.25 kg (about 2.75 lb). Now draw a square meter on the ground—either literally or in your imagination. The many kilometers of air extending vertically over the top of that square meter has a mass of ~10,000 kg! For context, figure out how many cars that would be (typ. 1,500 kg ea.), or what kind of animal would be this massive.

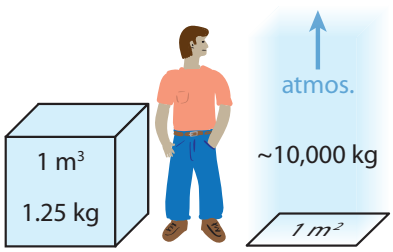


Figure 12.1: The mass of a cubic meter of air is 1.25 kg, and the mass of atmosphere over one square meter is an astounding 10 metric tons.

What this means is that air in motion can carry a significant amount of kinetic energy, since neither its mass nor velocity are zero. If the entire earth's atmosphere moved at 5 m/s—a noticeable breeze—at a total mass[5] of 5×10^{18} kg, we'd have 6×10^{19} J of kinetic energy in air currents. If we somehow pulled all this energy out of the air—stopping its motion entirely—we might expect the atmosphere to revive its normal wind

5: 10^4 kg per square meter times the surface area of Earth ($4\pi R_\oplus^2$)

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*; *Freely available at*: *https://escholarship.org/uc/energy_ambitions*.

patterns over the course of 24 hours: a full day of the driving solar input around the globe. The associated power works out to 700 TW. Notice that the value for wind in Table 10.2 (p. 168) is pretty-darned close to this, at 900 TW.[6] As the margin note indicates, we should be pleased to get within a factor of two for so little work and very off-the-cuff assumptions about global average air speed (see Box 12.1 for related thoughts). Figure 12.2 shows the annual average wind velocity at a height of 80 m (typical wind turbine height) for the U.S. Note that the 5 m/s we used above falls comfortably within the 4–8 m/s range seen in the map.

See if you can confirm; own it yourself!

6: In the first draft of this textbook, a different data source was used for Table 10.2 that had wind at 370 TW. Even so, the 700 TW estimate corroborated the order-of-magnitude scale and was deemed a satisfactory check: within a factor of two.

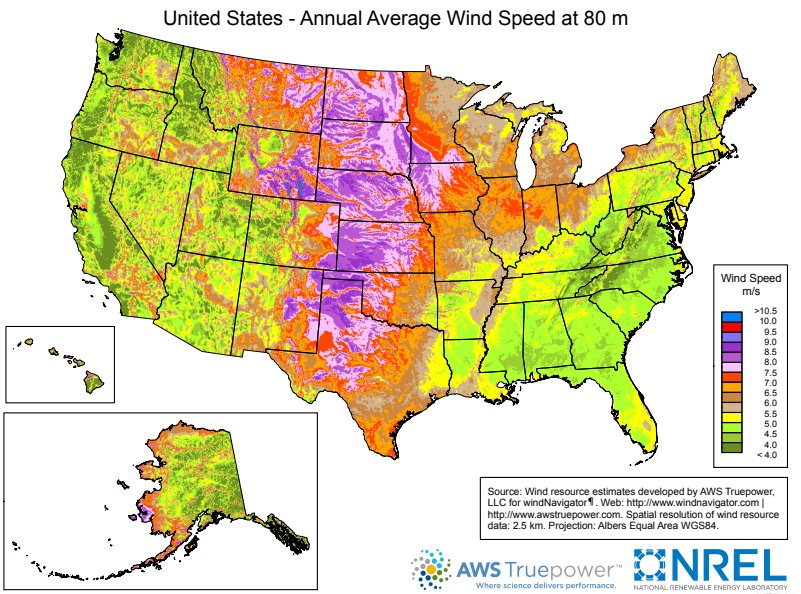


Figure 12.2: Average wind velocity at a height of 80 m across the U.S. [69]. Boundaries between colored boxes are every 0.5 m/s from 4.0 m/s to 10.0 m/s. Nothing on this map exceeds 9 m/s, and the deepest green is below 4 m/s. The plains states are the hot ticket. Note that Alaska is not to scale. From NREL.

Box 12.1: The Value of Estimation/Checking

Calculations like the one above offer a way to see if something at least checks out and seems plausible. If we had found that the whole atmosphere would have to be moving at 50 m/s to get the 900 TW figure in Table 10.2 (p. 168), we would suspect a problem, and either distrust the 900 TW number—seeking another source to confirm—or re-evaluate our own understanding. If we could get to 900 TW by only having wind speeds of 0.1 m/s,[7] we would also have cause for skepticism. When crude estimates of this type land in the general vicinity[8] of a number we see in a table, we can at least be assured that the number is not outlandish, and that our basic understanding of the phenomenon is okay. Checking understanding against presented data is excellent practice!

7: . . . or requiring weeks rather than a day to re-establish, once sapped

8: . . . say, within a factor of ten

But we can't capture the *entire* atmospheric wind, because doing so would require wind turbines throughout the volume, up to 10 km high! In fact, some estimates [70] of practical global wind installations come in as low as 1 TW—well below our 18 TW demand. Wind alone is unlikely

[70]: Castro et al. (2011), "Global Wind Power Potential: Physical and Technological Limits"

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*; *Freely available at*: *https://escholarship.org/uc/energy_ambitions*.

able to replace the energy currently derived from fossil fuels.

12.2.1 Wind Turbines

To understand practically-available energy, we back up and consider how much air hits a wind turbine whose rotor diameter is R. Figure 12.3 illustrates the concept. If the wind speed is v, the air travels a distance $v\Delta t$ in time interval Δt.[9]

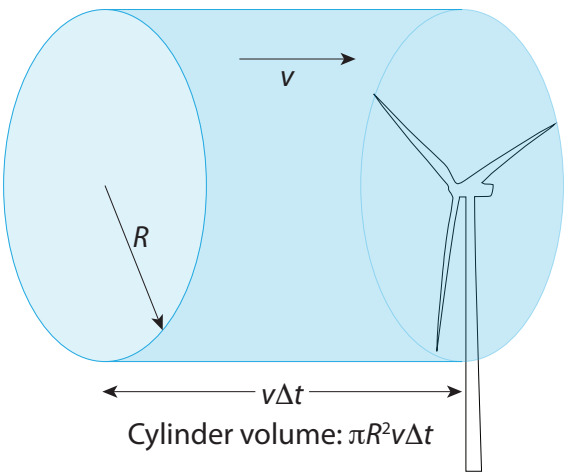


Figure 12.3: Wind power concept. In time interval Δt at wind speed v, a volume of air encounters the rotor having the shape of a cylinder of radius R and length $v\Delta t$. Note that most wind turbines are designed to pivot about a vertical axis to face into the wind, whatever the direction.

9: We can pick any value for Δt: a long interval makes a very long cylinder, while a small Δt results in a short, stubby cylinder. In the end, the value we chose for Δt will cancel out, so as not to matter.

The cross-sectional area of the wind turbine (rotor) is defined as the area swept out by the blades, so πR^2. Thus the volume of the cylinder of air interacting with the turbine over time interval Δt is the "base" (circular) area of the cylinder times its "height" (straight length, $v\Delta t$), or $V = \pi R^2 v\Delta t$. We know the density of the air,[10] so the mass of the cylinder is $m = \rho_{\text{air}} V = \rho_{\text{air}} \pi R^2 v\Delta t$. The kinetic energy *contained* in this cylinder of air is therefore K.E. $= \frac{1}{2}mv^2 = \frac{1}{2}\rho_{\text{air}}\pi R^2 v^3 \Delta t$. Now let's get rid of that pesky Δt. Think for a moment what happens if we divide both sides by Δt: it will definitely get rid of the Δt on the right-hand-side, but what does the left-hand side mean: energy over time? Hopefully, this is familiar by now as the concept of power.

10: $\rho_{\text{air}} \approx 1.25\,\text{kg/m}^3$

Definition 12.2.1 *The power delivered by a wind turbine of radius R in wind speed v and operating at efficiency ε is*

$$P_{\text{wind}} = \frac{1}{2}\varepsilon\rho_{\text{air}}\pi R^2 v^3. \tag{12.2}$$

where $\rho_{\text{air}} \approx 1.25\,\text{kg/m}^3$ at sea level. Efficiency has been inserted as ε and tends to be 40–50% for modern turbines.

Notice that the delivered power scales, sensibly, with the area of the wind turbine's blade path, but more importantly and perhaps surprisingly as the velocity *cubed* (Table 12.1). The cubed part should make you sit

Table 12.1: Wind power scales as the cube of wind speed.

Speed	Power
0	0
1	1
2	8
3	27
4	64
5	125
⋮	⋮
10	1,000

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

up straight: that's a very strong function of velocity! It means that if the wind changes from a gentle 5 m/s to a brisk 10 m/s, the power available goes up by a factor of 8. A strong wind at 20 m/s has 64 times as much power as the 5 m/s breeze.[11] We can understand the three powers of velocity thusly: two powers come from kinetic energy, and one from the length of the cylinder. As wind speed increases, not only does the oncoming air have more kinetic energy per fixed volume, but also a larger volume encounters the turbine in a given time.

Setting $\varepsilon = 1$ in Eq. 12.2 corresponds to the *total* power present in the wind. But we can't be greedy and grab *all* of it. In fact, if we did, it would mean *stopping* the air at the wind turbine: pulling out all the kinetic energy means no velocity is left. If we did this, newly arriving air would divert around the stopped mass of air, and the turbine would no longer have access to oncoming energy. The theory has all been worked out:[12] a turbine is limited to $\varepsilon \leq 16/27$ (59%) of the available energy, known as the Betz limit [72]. This is not a technological limitation, but comes from the physics of fluid flow. A second consideration enters for low-speed rotor motion, known as the Glauert limit [73], resulting in diminishing efficiency as wind speed drops.

11: Now it may be easier to understand why hurricanes can be so destructive, if their power scales as the cube of wind velocity, and velocities exceed 50 m/s.

12: A recent derivation is in [71].

[72]: Betz (1926), "Wind-Energie und ihre Ausnützung durch Windmühlen"

[73]: Glauert (1993), *The elements of airfoil and airscrew theory*

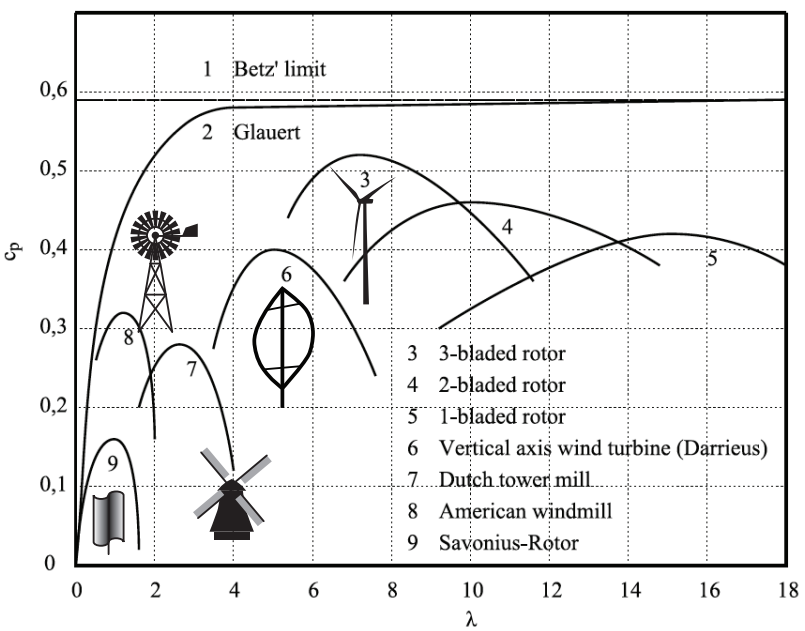


Figure 12.4: Theoretical and practical wind turbine efficiencies (ε, or c_p in the plot), for various designs. The parameter λ is the ratio of tip speed to wind speed: higher λ means a faster tip speed [74]. All designs must be below the Betz limit (horizontal line near top). At slower speeds, the Glauert limit confines performance to occupy the region to the right of the curve marked 2. Each of the 7 designs shown have arched curves, achieving maximum efficiency at a particular tip speed. Too slow, and the turbine is not transmitting much energy; too fast and drag/friction begins to dominate. Adapted from ©2010 WIT Press.

Figure 12.4 shows these theoretical limits, along with design limits from various rotor configurations. Curves reflect an optimum rotor speed for each design: speeding up produces more generator output until it gets fast enough that air drag on the blades starts to dominate. The most common modern turbine is the 3-blade design,[13] able to get roughly 50% of the energy out of the wind. Notice that the tip speed can be quite high: 6–8 times the wind speed. This can be quite alarming to birds in the area, whose cruising speed is nearer wind speed, and they

13: Not only is the three-blade design the most efficient, its lower tip speed is less dangerous than for 2 or 1-blade designs, according to Figure 12.4.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

have never met something so fast before. For a modern derivation of the Betz limit and how efficiency depends on tip speed, see [71]. The largest turbines—having 150 m diameter rotors—are rated for up to 10 MW of electrical power production.

[71]: Ragheb et al. (2011), "Wind Turbines Theory - The Betz Equation and Optimal Rotor Tip Speed Ratio"

Example 12.2.1 How much power could you expect a small (4 m *diameter*) 3-blade wind turbine situated atop your house to deliver in a respectable 5 m/s breeze?

The radius is 2 m and we'll pick a middle-of-the-road efficiency of 45%: $P = \frac{1}{2} \cdot 0.45 \cdot (1.25\,\mathrm{kg/m^3}) \cdot \pi \cdot (2\,\mathrm{m})^2 \cdot (5\,\mathrm{m/s})^3$ comes to about 450 W.[14]

14: Not too impressive: hard to get much wind power on a household scale, although 10 m/s would give 3.6 kW.

Besides the limit on how much power can be pulled out of the air by a single turbine, we also find limits on how *densely* they may be populated in a given area: how much space is required between turbines so that one does not disrupt the other. Obviously, it would not serve to put one turbine directly behind another, as they would at best split the available power arriving as wind. Even side by side, it is best to leave room between windmills so that additional rows are not deprived of wind power. A rule of thumb is to separate turbines by at least 5–8 diameters side-to-side, and 7–15 diameters[15] along the (prevailing) wind direction. For the sake of illustration, Figure 12.5 shows a spacing on the denser side of the range, but otherwise we adopt the more recent recommendations and use 8 diameters side-to-side and 15 diameters deep [75]. This works out to a 0.65% "fill factor," meaning that 0.65% of the land area contains an associated rotor cross section.[16]

15: An older "rule of thumb" was 5 side-to-side and 7–8 deep, but newer work suggests as much as 8 diameters side-to-side and 15 diameters deep.

16: ... one πR^2 rotor area for every $8D \times 15D = 120D^2 = 120 \times (2R)^2 = 480R^2$ of land area

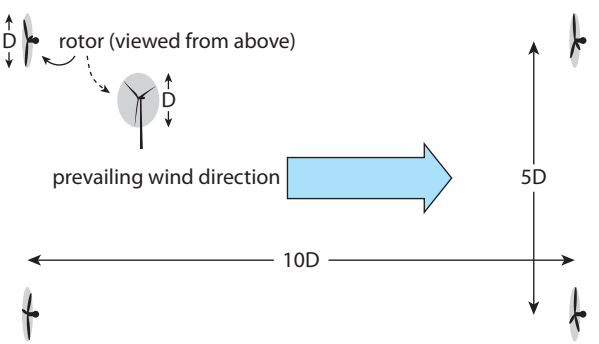


Figure 12.5: Overhead view of wind farm turbine locations, for the case where separations are 10 rotor-diameters along the wind direction, and 5 rotor diameters in the cross-wind direction—a geometry that yields 1.6% area "fill factor." Current recommendations are for 15 and 8 rotor diameters, which is significantly more sparse than even this depiction, leading to 0.65% area fill. Note that most wind turbines can turn to face the wind direction, for times when its direction is not the prevailing one.

In order to compare to other forms of renewable energy, we can evaluate a power per unit land area (in $\mathrm{W/m^2}$) by the following approach:

$$\frac{\text{power}}{\text{area}} = \frac{\frac{1}{2}\varepsilon\rho_{\text{air}}\pi R^2 v^3}{480R^2} = \frac{\pi}{960}\varepsilon\rho_{\text{air}}v^3, \tag{12.3}$$

employing the rule-of-thumb 8×15 turbine placement scheme. Using an efficiency of 40% and $v = 5\,\mathrm{m/s}$,[17] we get 0.2 $\mathrm{W/m^2}$—which is 1,000 times smaller than solar's ~200 $\mathrm{W/m^2}$ insolation (Ex. 10.3.1; p. 167).

17: Recall that this choice gave sensible global wind power estimates lining up with Table 10.2 (p. 168).

A final general note about wind generation is somewhat obvious: the

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*; *Freely available at*: *https://escholarship.org/uc/energy_ambitions*.

wind is not always blowing, and its speed varies over wide ranges. In this sense, wind is an *intermittent* power source. Just as for hydroelectric installations, wind resources are characterized by a regionally-dependent capacity factor, which is the ratio of energy delivered to what would have been delivered if the generation facility operated at full capacity at all times. Typical capacity factors for wind in the U.S.[18] are around 33%, and Figure 12.6 provides a visual sense for how this manifests in the real world: pretty erratic.

18: Capacity factors for wind are smaller than for hydroelectricity due to wind being more variable than river flow.

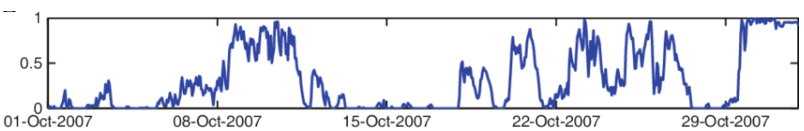


Figure 12.6: One month of wind generation from a 20 MW wind farm, illustrating the intermittent nature and why capacity factors are low [76]. The facility saturates at maximum power late in the month, self-limiting to avoid damage to the turbines. ©2010 Springer.

For very low wind speeds,[19] wind turbines do not have enough wind to turn at all and sit still at zero output. Furthermore, a turbine is rated at some maximum power output, which occurs at some moderately high wind speed,[20] beyond which the generator risks damage—like "redlining" a car's engine. When the wind climbs above this maximum-rated speed, the turbine is pegged at its maximum power—no longer following a v^3 relation—and deliberately twists its blades[21] to be *less* efficient as the wind speed grows so that it maintains constant (maximum) power output. When the wind speed becomes large enough to endanger the turbine, it will twist its blades parallel to the wind to allow the air to pass without turning the rotor at all, so that it no longer spins while it "rides out" the high winds.[22]

19: . . . less than about 3 m/s; called the "cut-in" velocity

20: . . . usually around 12–15 m/s

21: The blades are like long airplane wings and are mounted so that they can be rotated on an axis running the length of the blade, allowing them to engage the wind at any angle, thus varying efficiency.

22: A typical shut-off wind speed for turbines is 20–30 m/s.

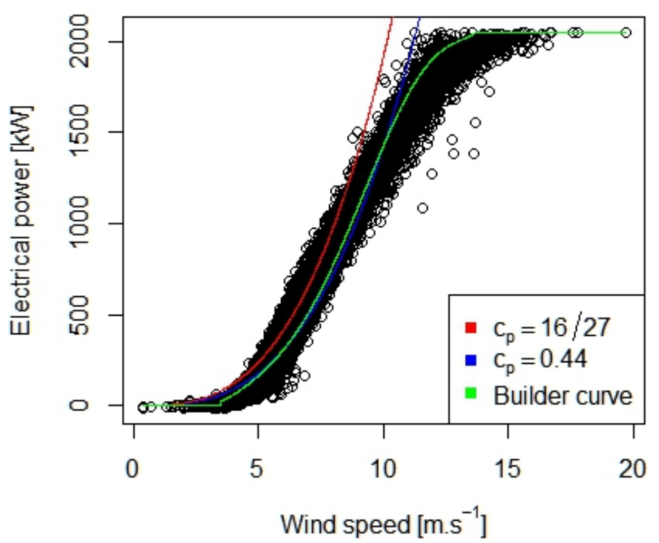


Figure 12.7: Actual data (thickly-clustered black circles) of power delivered by a turbine rated at 2 MW, as a function of wind velocity. The red curve represents the theoretical Betz limit of 59%, appearing as a cubic function of velocity—as Eq. 12.2 dictates. The better-matching blue curve corresponds to an overall efficiency $\varepsilon = c_p = 0.44$ (44%), and the green curve—which rolls over from the cubic function and saturates at higher velocities—is the manufacturer's expectation for the unit [77]. The "cut-in" velocity for this turbine is around 3.5 m/s: note the small step up from zero output in the green curve. This turbine saturates around 12 m/s: the green curve flattens out and no black circles appear above the cutoff. From ©2017 Wiley.

Figure 12.7 shows a typical power curve for a 2 MW turbine, on top of which are drawn a cubic function of velocity at the theoretical Betz limit (red curve), a cubic (blue) at 44% efficiency ($\varepsilon = 0.44$), and the green manufacturer's curve [77]. Notice that the turbine performance

[77]: Fischer et al. (2017), "Statistical learning for windpower: a modeling and stability study towards forecasting"

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*; *Freely available at*: *https://escholarship.org/uc/energy_ambitions*.

demonstrates the aspects covered in the previous paragraph: "cutting in" just above 3 m/s and maxing out (saturating) beyond about 12 m/s. In between, it closely follows a cubic function at an overall efficiency of 44% (blue curve).

12.3 Wind Installations

Global wind installations are rising rapidly, currently (as of 2020) above 600 GW of installed capacity.[23] Table 12.2 lists the major players, in terms of installed capacity, average generation, fraction of total energy,[24] capacity factor, and share of global wind generation. The amount of wind energy in each country depends on a combination of how much wind is available in the country, how fast electricity demand is growing, electrical infrastructure, and political interest in renewable energies.

23: A small fraction of this is realized, due to the capacity factor.

24: To compare wind to total energy, we follow the thermal equivalent convention discussed for Table 10.3 (p. 170)

Table 12.2: Global wind installations in 2018 [78–84]. The top six countries capture 85% of the global total.

Country	GW installed	GW average	cap. fac. (%)	energy fraction (%)	global share (%)
China	184	41.8	22.7	3.0	33
U.S.	97	31.4	32.4	2.7	25
Germany	59	12.7	21.4	8.3	10
India	35	6.5	18.5	2.3	5.2
Spain	23	5.4	23.5	8.3	4.3
UK	21.7	6.5	30.0	6.9	5.2
World Total	592	~ 125	21.1	2.0	100

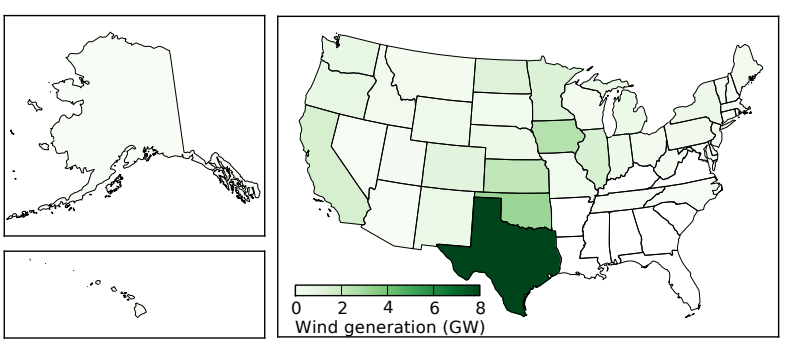


Figure 12.8: Wind power by state, in terms of average generation, in GW, in 2018. The color scale may seem unhelpful, but the unavoidable truth is that many states don't have a lot going on, and Texas is so dominant as to render other states almost insignificant. A logarithmic color scale could help, but then the important lesson on the gross disparity might go unappreciated.

In 2018, the U.S. had about 94 GW of *installed* wind capability.[25] This number has recently surpassed hydroelectric installed capacity (about 80 GW). Both are impacted by capacity factors, which for wind averages 33% in the U.S., while hydropower is just over 40%. The net effect is that the generation for the two is pretty comparable.[26] Where is the wind power in the U.S. installed? Figure 12.8 shows that Texas wins, at 8.7 GW. Oklahoma is a distant second at 3.2 GW, Iowa at 2.5 GW. California is in fifth place at 1.6 GW.

25: From tables 1.14.B and 6.2.B in [85]

26: . . . as we also saw in Table 10.3 (p. 170)

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

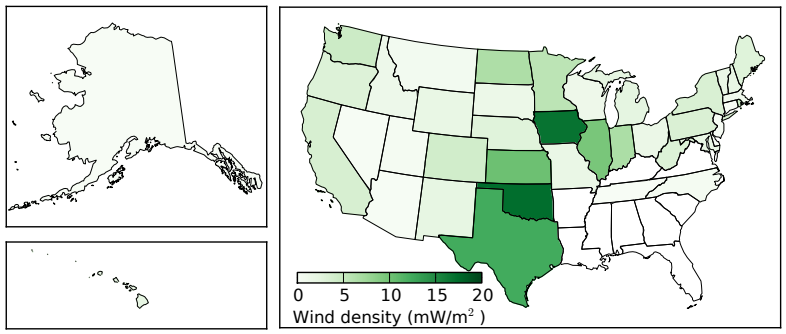


Figure 12.9: Average wind power by state, divided by state area to indicate a density of the developed resource, in milliwatts per square meter (based on 2018 data). We might expect some resemblance to Figure 12.2, based on where the resource is most favorable.

Following the flow we used in Sec. 11.3 (p. 178), we show wind generation as a function of area (Figure 12.9), to get a sense for how concentrated the installations are. Oklahoma and Iowa jump in front of Texas by this measure. Texas has more total generation than all others, but is a very large state in terms of area. Iowa, for instance produces about 30% as much wind power as Texas, but at only 20% the size. The numbers, reaching about 0.017 W/m^2 are a bit smaller than those for hydroelectricity, where two states exceeded this value.[27] We can compare these numbers to the 0.2 W/m^2 *fully-developed* potential we estimated in the paragraph following Eq. 12.3 to conclude that in principle we could expand wind by a large factor.[28]

27: . . . 0.05 W/m^2 in Washington; followed by New York at 0.02 W/m^2

28: . . . about 10 times, in the case of Iowa, but bear in mind the 0.2 W/m^2 estimate was based on 5 m/s, while Iowa scores a bit higher, according to Figure 12.2

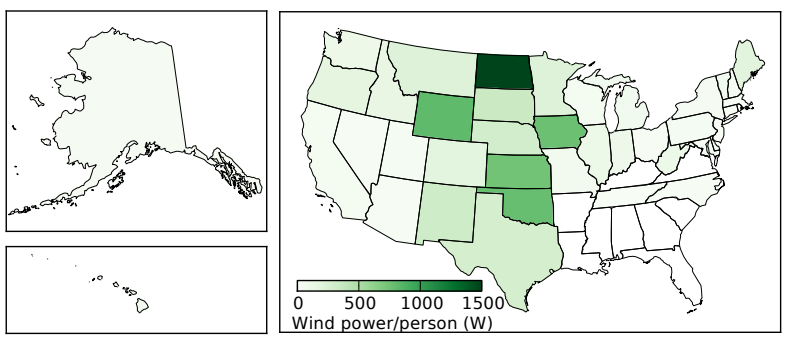


Figure 12.10: Average wind power generation by state divided by state population for an average power per person (based on 2018 data).

Next, we look at wind generation per capita in states, in Figure 12.10. Now North Dakota blows away the rest, at 1.6 kW per person, followed by four states at about half of this value. We can put this in context by noting that the average power consumption in the U.S. per capita is around 10 kW.

Finally, Figure 12.11 shows wind capacity factors, indicating where the wind is most reliable. It peaks around 41% in Kansas, but all of the plains states in general do well. The southeastern U.S. has almost no wind development,[29] as is evident in any one of these figures.

29: Lack of wind makes it a poor fit: see also Figure 12.2.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

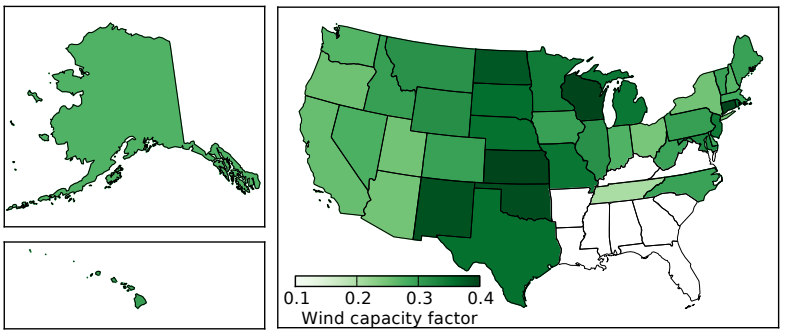


Figure 12.11: Capacity factor for wind installations by state (based on 2018 data).

12.4 Upshot: Wind is not Overblown

Wind has surged tremendously in the last decade (Fig. 7.5; p. 108), proving to be an economically viable and competitive resource. But how much could we expect to get from wind?

Putting a few of the previous results together, If the entire contiguous U.S. (area $\sim 10^{13}$ m^2) were developed for wind at an estimated power density of 0.2 W/m^2—which was based on a 5 m/s average wind speed—and a capacity factor of 33%, the U.S. could theoretically produce 0.7 TW[30] from wind—roughly 20 times what is produced today. We should take this crude estimate as an extreme upper end, since it is inconceivable that we would develop wind so fully as to never be more than a few hundred meters—a few rotor diameters—away from a wind turbine, no matter where we go. Also, many areas are sub-threshold and would not support investment in wind development.

Even so, the inflated 0.7 TW estimate falls short of the current 3.3 TW energy demand in the U.S., has major intermittency problems, and is not in a form that can be well-used in all sectors, like transportation and industrial processing. While wind *alone* cannot replace fossil fuels at the current level of demand, it *can* doubtless be a significant contributor.

Globally, estimates for wind potential tend to be in the few-terawatt range, though can be as low as 1 TW for a number of practical reasons [70]. As was the case for hydroelectricity, wind is a viable player in the renewables mix, but is unable to shoulder the entire load.

Wind energy is not free of environmental concerns, disturbing landscapes and habitats. Its impact on birds[31] and bats is most worrisome, as the rotors move far faster than anything to which the wildlife is habituated. Still, compared to the environmental toll from fossil fuels, it is fairly clean—similar to the impact of hydroelectric power.

A pros and cons list will help summarize. First, the positive attributes:

- Wind is replenished on the planet daily by solar illumination;

30: ⚠ We're fudging things a bit here for the sake of simplicity. If turbines are built for 12 m/s, the capacity factor already has some averaging built in, so using 5 m/s *and* a capacity factor of 0.33 is unfair. On the other hand, much of the country spends substantial time below the cut-in speed for turbines, and the cruel cubic function of velocity greatly suppresses much of the land area as impractical for wind development. So the approach is a compromise that might balance out reasonably.

[70]: Castro et al. (2011), "Global Wind Power Potential: Physical and Technological Limits"

31: Domestic cats turn out to kill far more birds than wind turbines do, currently.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

- Harnessing wind is relatively low-tech and straightforward;
- Wind has decent efficiency—typically 40–50%—in extracting energy from the oncoming wind;
- Life-cycle CO_2 emissions for wind is only 2% that of traditional fossil fuel electricity [68];
- Growth in the wind sector points to economic viability;
- Wind is able to scale up to cover a meaningful fraction of energy demand.

[68]: (2020), *Life Cycle GHG Emissions*

And the downsides:

- Wind is intermittent: power when nature allows, not when people demand;
- Wind is regionally variable: many places do not produce enough wind to support development;
- Wind can cause environmental disruption to habitats—especially dangerous to birds and bats;
- Esthetic objections to noise and degradation of scenery hamper expansion.

12.5 Problems

1. A modest slap[32] might consist of about 1 kg of mass moving at 2 m/s. How much kinetic energy is this?

2. A *hard* slap might consist of about 1 kg of mass moving at 10 m/s. How much kinetic energy is this, and how much warmer would 10 g of skin[33] get if the skin has the heat capacity properties of water, as in the definition of a calorie (Sec. 5.5; p. 73 and Sec. 6.2; p. 85 are relevant)?

3. A 10 kg bowling ball falls from a height of 5 m. Using the convenient $g \approx 10\,\mathrm{m/s^2}$, how much gravitational potential energy does it have? Just before it hits the ground, all of this potential energy has gone into kinetic energy.[34] What is the speed of the bowling ball when it reaches the ground, based on kinetic energy?

4. Did the final answer for the speed of the bowling ball at the end of its drop depend on the mass?[35] Write out the math symbolically[36] and solve for velocity, v. Does the result depend on mass?

5. Thermal energy is just randomized kinetic energy on a microscopic scale. To gain some insight into this, consider one liter (1 kg) of water, and figure out how much energy it would take to heat it from absolute zero temperature[37] to 300 K assuming that the definition of the calorie (Sec. 5.5; p. 73) applies across this entire range. If this same amount of energy went into kinetic energy—hurling the water across the room—what would the corresponding velocity be?

32: . . . how painful can a few Joules be?

33: . . . corresponding to a volume of 10 mL appropriate to a slap area of 10 cm by 10 cm and to a depth of 1 mm

34: . . . neglecting any energy flow to air resistance

35: Try it using a different mass.

36: . . . using variables/symbols

37: . . . 0 K, when the kinetic energy is effectively frozen out, or stopped

ⓘ As large as the number is, it is representative of the speeds of individual molecules within the water, and is, not coincidentally, similar to the speed of sound in water.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

6. A typical house may have a floor area around 150 m^2 (1,600 square feet). If the floor–to-ceiling distance is typically 2.5 m, how much mass is in the air within the house? Could you lift this much mass if handed to you as bags of rocks?

7. Atmospheric pressure is about 10^5 N/m^2, meaning that a 100,000 N weight of air—corresponding to a mass of 10,000 kg—sits atop very square meter of the ground (at or near sea level). If the air density were constant at 1.25 kg/m^3—rather than decreasing with height as it actually does—how high would the atmosphere extend to result in this weight (mass)?

ⓘ This is called the *scale height* of the atmosphere, h_0, which you may wish to compare to the tallest mountains on Earth or the height at which airplanes fly. The actual density of the atmosphere decreases exponentially as a function of height, with h_0 being the characteristic scale.

8. Comparing the pale green region in the southeastern U.S. to the purple region of the plains states in Figure 12.2, how much more power would we expect out of the same rotor placed in the plains than in the southeast (by what factor is it bigger in the plains)?

9. How much more powerful is a hurricane-strength wind of 50 m/s hitting your house than is a light breeze of 5 m/s?

10. How much *power* would a moderate-sized 50%–efficient wind turbine produce whose radius is 10 m at wind speeds of 5 m/s, 10 m/s, 15 m/s, and 20 m/s? Express the answers in kW or MW, depending on what is most natural.

11. The Betz limit says that we get to keep no more than 59% of the available wind power. If 59% of the kinetic energy in a lump of air moving at speed [38] v is removed, how fast is it going afterwards,[39] as a fraction of the original speed?

38: Make up your own velocity or solve in symbols/variables: same either way.

39: This relates to wind speed just behind a wind turbine

12. The largest wind turbines have rotor *diameters*[40] around 150 m. Using a sensible efficiency of 50%, what power does such a jumbo turbine deliver at a maximum design wind speed of 13 m/s?

40: ⚠ not radius

13. A recent news article announces the largest wind turbine yet, measuring 220 m in *diameter* and having a maximum power output of 13 MW. Using a reasonable efficiency, calculate the velocity of the wind at which maximum power is reached.

14. Compare the tip speed of a three-blade turbine operating at its optimal efficiency (as per Figure 12.4) in a moderate wind of 7 m/s to typical freeway driving speeds in the same units.

15. Traveling down the road, you carefully watch a three-bladed wind turbine, determining that it takes two seconds to make a full revolution. Assuming it's operating near the peak of its efficiency curve[41] according to Figure 12.4, how fast do you infer the wind speed to be if the blade length[42] appears to be 15 m long?

41: Hint: focus on tip speed.

42: . . . corresponding to radius of the rotor

16. Building from the result in Problem 15, how much power is this windmill delivering if its efficiency is about 50%?

Reflect on the fact that just estimating the rotor blade length and timing its revolution is enough for you to produce an estimate of power being generated!

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

17. In a way similar to Figure 12.5, replicate the statement in the text that the fraction of land covered per rotor area is 0.65% if turbines are separated by 15 rotor diameters along one direction and 8 rotor diameters along the cross direction.

18. Check that the units of Eq. 12.3[43] indeed are equivalent to Watts per square meter (W/m^2).

19. Provide a clear explanation of why the area under the blue curve in Figure 12.6 compared to the area of the whole rectangular box is an appropriate way to assess the capacity factor of the depicted wind farm?

20. What capacity factor would you estimate for the wind farm performance depicted in Figure 12.6? In other words, what is the approximate area under the curve compared to the entire box area, as explored in Problem 19? An approximate answer is fine.

21. Referring to Figure 12.7, examine performance at 5 m/s and at 10 m/s, picking a representative power for each in the middle of the cluster of black points, and assigning a power value from the left-hand axis. What is the ratio of power values you read off the plot, and how does this compare to theoretical expectations for the ratio going like the cube of velocity?

22. Figure 12.7 surprisingly has all the information required to deduce the rotor diameter! The turbine appears to produce 1,400 kW when the wind velocity is 10 m/s, and we also know it appears to operate at $\varepsilon = 0.44$. What is the rotor diameter?

23. Considering that wind turbines are rated for the maximum-tolerable wind speed around 12 m/s, and tend to operate at about 30% capacity factor, how much *average* power[44] would a 100 m diameter turbine operating at 45% efficiency be expected to produce?

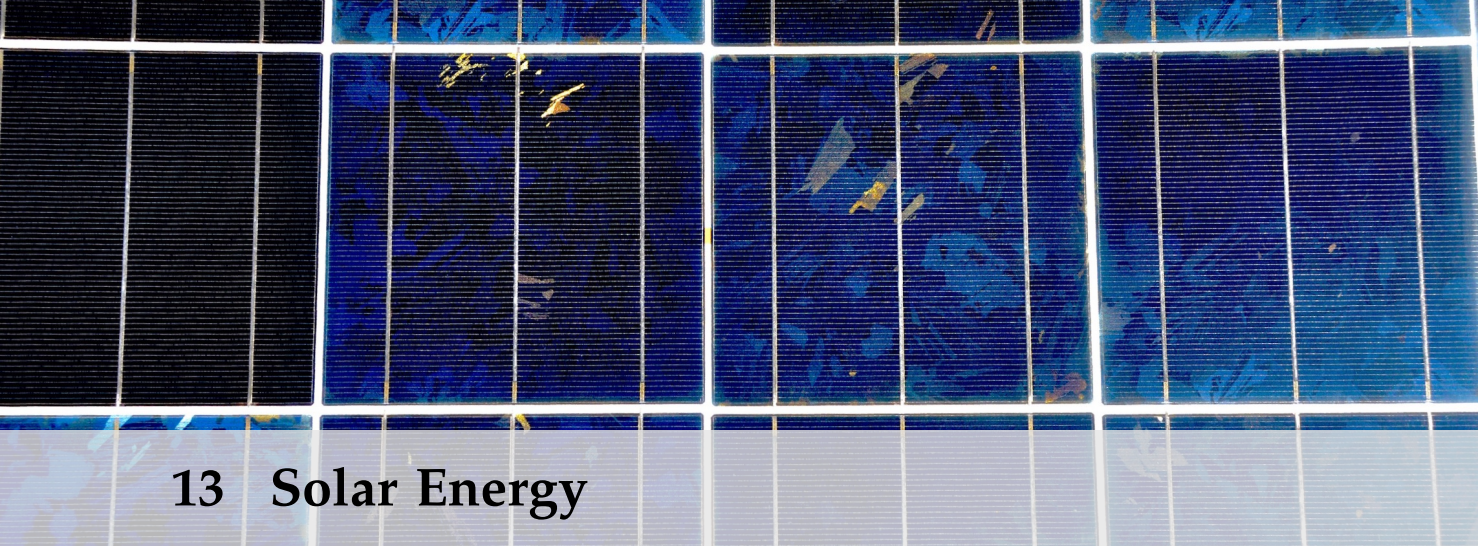
24. Table 12.2 shows Germany having more than twice the wind capability as Spain, yet each gets 8.3% of its power from wind. What do you infer the difference to be between the countries?

43: ... essentially ρv^3

44: Hint: compute power at 12 m/s then apply capacity factor

Hint: no external research necessary: what do the numbers mean?

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.



13 Solar Energy

13.1 The Energy of Light 197
13.2 The Planck Spectrum 199
13.3 Photovoltaics 201
Theoretical Efficiency 202
13.4 Insolation 206
13.5 Incredible Solar Potential . . 211
13.6 Residential Considerations 213
Configurations 213
Sizing and Cost 213
13.7 Photovoltaic Installations . 215
Pros and Cons of PV 217
13.8 Solar Thermal 218
Passive Solar Heat 218
Solar Thermal Electricity . 219
Pros and Cons of ST 221
13.9 Upshot for Solar 221
13.10 Problems 222

Now we come to the main attraction. As we saw in Chapter 10, all of the strictly renewable energy options are ultimately derived from the sun. The two resources of the last two chapters—hydroelectricity and wind—represent tiny crumbs of the overall solar input to the planet.

Practically speaking, it is difficult to concoct ways to harness more than a few terawatts of power from hydroelectric or wind-based resources. Similar limitations apply to biologically-derived energy, geothermal, tidal, ocean currents, wave energy, etc. This may be worrisome considering that human society currently demands 18 TW. Meanwhile, the sun delivers 83,000 TW to the earth's surface (Fig. 10.1; p. 167, Table 10.2; p. 168). That's almost 5,000 times more than the demand. By the numbers, then, the sun seems to offer all we might ever need. In fact, the quantitative imbalance is so extreme as to make one wonder why we would ever mess around doing anything else.

Yet at present, the U.S. gets only about 0.9% of its energy from solar power, or 2.3% of its electricity. Similarly for the world, 1.2% of global energy is solar (2.1% of the electricity). It would seem to be vastly underutilized.

This chapter explains the nature of solar energy, its potential for use, practical considerations, and looks at the state of installations. While most of the focus is on photovoltaic (PV) panels that directly generate electricity from light, solar thermal power generation is also covered.

13.1 The Energy of Light

Sec. 5.10 (p. 79) introduced the basics of the energy of light. This section acts as a refresher and lays the foundation for the rest of the chapter.

Photovoltaic cells, showing grid contacts and crystal domains. Photo Credit: Tom Murphy

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

Light—a form of electromagnetic radiation—is composed of photons—individual particles of energy each having a characteristic wavelength—what we might call color. Photons are such tiny quanta of energy that familiar environments are awash in unfathomably large numbers of them. A typical light bulb, for instance, emits quintillions[1] of photons every second.

1: A quintillion is 10^{18}.

Definition 13.1.1 *The **energy of a single photon**, in various forms, is*

$$E_{\text{photon}} = h\nu = \frac{hc}{\lambda} \approx \frac{2\times10^{-19}\,\text{J}}{\lambda(\text{in }\mu\text{m})} = \frac{1.24\,\text{eV}}{\lambda(\text{in }\mu\text{m})}, \tag{13.1}$$

where $h = 6.626\times10^{-34}$ J · s (Planck's constant), and ν is the frequency of the light in Hertz (Hz, or inverse seconds).

The second form (hc/λ) is useful, as we more commonly characterize the "color" of light by its wavelength, λ. The speed of light, $c \approx 3\times10^8$ m/s, connects frequency to wavelength via

$$\lambda\nu = c. \tag{13.2}$$

The third form in Definition 13.1.1 makes it easy to compute photon energy in Joules given the wavelength in microns.[2] Visible light has a wavelength around 0.4–0.7 μm (violet–to–red), so a typical photon energy, at 0.5 μm, is about 4×10^{-19} J. It's a *tiny* number!

2: One micron, or micro-meter, abbreviated μm, is 10^{-6} m.

Example 13.1.1 About how many photons strike a 0.4 m^2 patch of sidewalk per second if the overhead sun is delivering 1,000 W/m^2?

For the visible light characteristic of sunlight, we can use a convenient wavelength of 0.5 μm, amounting to 4×10^{-19} J of energy per photon. The patch of sidewalk we describe receives light energy at a rate of 400 W or 400 J/s.[3] How many 4×10^{-19} J photons does it take to amount to 400 J? Divide[4] to get 10^{21}.

3: . . . 0.4 m^2 times 1,000 W/m^2

4: . . . or try reasoning it out: 10^{19} of them would make 4 J, so we need 100× more

The final form in Definition 13.1.1 relates to the fact that photons frequently interact with electrons as we will see in Section 13.3, making it convenient to convert to another energy unit called the electron-volt, or eV (introduced in Sec. 5.9; p. 78). One electron volt is the energy it takes to move an electron through an electric potential of one Volt. The conversion is 1 eV $= 1.602\times10^{-19}$ J. For instance, the 0.5 μm (blue-green) photon we used in the previous example would have an energy around 2.5 eV.

Why should we care about unthinkably small quantities of light? Three reasons come to mind:

1. Eq. 13.1 elucidates that bluer[5] photons have higher energy than red photons, which is important to know;

5: . . . shorter wavelength

© 2022 T. W. Murphy, Jr.; Creative Commons Attribution-NonCommercial 4.0 International Lic.;
Freely available at: https://escholarship.org/uc/energy_ambitions.

2. Individual photons interact with matter at the microscopic scale and are relevant to understanding solar photovoltaics and photosynthesis;
3. It's how nature really works.

13.2 The Planck Spectrum

We should first understand where photons originate, which will help us understand how solar panels work and their limitations. Until recent technological advances, photons tended to come from *thermal* sources. It's true for the white-hot sun,[6] and true for flame and incandescent light bulb filaments.[7] Likewise, hot coals, electrical heating elements, and lava are all seen to glow. Physics tells us how such hot sources radiate, as covered by the next three equations. The first (with units) is:

$$P = A\sigma T^4 \text{ (W)}. \tag{13.3}$$

We already saw this equation in the context of Earth's energy balance in Sections 1.3 and 9.2. It is called the Stefan–Boltzmann law, describing the total power (in W, or J/s) emitted from a surface whose area is A (in square meters) and temperature, T in Kelvin.[8] The constant, $\sigma \approx 5.67 \times 10^{-8}$ W/m^2/K^4 is called the Stefan–Boltzmann constant, and is easy to remember as 5-6-7-8.[9]

$$B_\lambda = \frac{2\pi hc^2}{\lambda^5} \frac{1}{e^{hc/\lambda k_B T} - 1} \left(\frac{\text{W/m}^2}{\text{m}}\right), \tag{13.4}$$

Eq. 13.4 might look formidable, but only λ and T are variable. It describes the Planck spectrum, otherwise known as the blackbody[10] spectrum. For some temperature, T, this function specifies how much power is emitted at each wavelength, λ. Three fundamental physical constants from key areas of physics make an appearance: $c \approx 3 \times 10^8$ m/s is the familiar speed of light from relativity; $h \approx 6.626 \times 10^{-34}$ J · s is Planck's constant from quantum mechanics, and $k_B \approx 1.38 \times 10^{-33}$ J · K is the Boltzmann constant of thermodynamics.[11]

$$\lambda_{max} \approx \frac{2.898 \times 10^{-3}}{T(\text{in K})} \text{ (m)}. \tag{13.5}$$

Eq. 13.5 is called the Wien law and is a numerical solution identifying the peak of the blackbody spectrum as a function of temperature. Higher temperatures mean higher kinetic energies at a microscopic scale, so that higher-energy (shorter-wavelength) photons can be produced. This is why as objects get hotter, they move from red to white, and eventually to a blue tint.

All this may seem overwhelming, but take a breath, then just look at Figure 13.1. Everything so far in this section is captured by Figure 13.1.

6: . . . and thus stars and even the moon, which is just reflected sunlight

7: Modern lighting like fluorescent and LED sources rely on manipulating energy levels of electrons within atoms and crystals.

8: Recall that temperature in Kelvin is the temperature in Celsius plus 273 (273.15, technically).

9: The Stefan–Boltzmann constant is actually a witch's brew of more fundamental constants h (Planck's constant), c (speed of light), and k_B (the Boltzmann constant) as $\sigma = 2\pi^5 k_B^4/(15c^2h^3)$.

10: The term blackbody effectively means a perfect emitter and absorber of thermal radiation.

11: This last one may be more familiar to students in its chemistry form of the gas constant $R = k_B N_A \approx 8.31$ J/K/mol, where $N_A \approx 6.022 \times 10^{23}$ is Avogadro's number.

© 2022 T. W. Murphy, Jr.; Creative Commons Attribution-NonCommercial 4.0 International Lic.;
Freely available at: https://escholarship.org/uc/energy_ambitions.

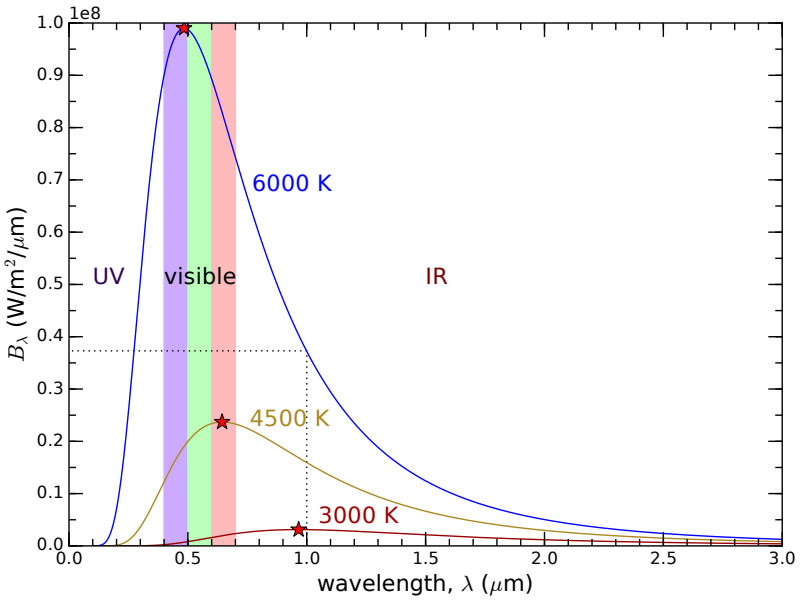


Figure 13.1: Planck spectra, or blackbody spectra for three temperatures, indicating where the ultraviolet, visible, and infrared regions lie. The shapes of the three curves (spectra) are described by Eq. 13.4, the star locations are found by Eq. 13.5, and the total power radiated, per square meter of surface is the area under each curve, as captured in Eq. 13.3. The dotted line relates to Example 13.2.2. Note the 1e8 factor on the vertical axis, meaning that the axis goes from 0 to 1.0×10^8 W/m^2/μm.

The shape of each spectrum is a plot of the function in Eq. 13.4 for three different temperatures. If comparing output of Eq. 13.4 to Figure 13.1, be aware that the units have been manipulated to favor microns over meters.[12]

Let's come at this again with numbers to help us make sense of things. Looking at the curve (spectrum) for 6,000 K, we will verify that each equation makes some sense.

Example 13.2.1 First, Eq. 13.5 says that the wavelength where emission peaks should be about $2.898 \times 10^{-3}/6000 \approx 0.483 \times 10^{-6}$ m, or 0.483 μm.

Now look at the graph to see that the peak of the blue curve is indeed just short of 0.5 μm, denoted by the red star at the top.

Example 13.2.2 Let's now verify a point on the Planck spectrum, picking 6,000 K and 1 μm to see if Eq. 13.4 lands in the same spot as indicated in Figure 13.1.

If we go through the laborious exercise of plugging in numbers to Eq. 13.4 for $T = 6000$ and $\lambda = 1 \times 10^{-6}$ (1 μm), we find[13] the overall outcome is 3.73×10^{13} W/m^2 per *meter* of wavelength. Once we adjust by 10^{-6} for the units on the plot (see earlier margin note), we expect 0.373×10^8 W/m^2 per micron.

Indeed, the blue curve passes through this value at a wavelength of 1 μm, as indicated by the dotted line in Figure 13.1.

12: Eq. 13.4 uses units of meters for λ, but Figure 13.1 uses microns (μm, or 10^{-6} m) for convenience. Also, Eq. 13.4 delivers an answer in units of W/m^2 per *meter* of wavelength, but for the plot we divided by 10^6 so it would be W/m^2 per *micron* of wavelength. By taking care of this detail, the area under each curve in Figure 13.1 should match σT^4 as in Eq. 13.3.

13: Numerically, the numerator is 3.74×10^{-16}, the denominator is 10^{-30}, and the argument in the exponential is 2.4, so that the second fraction is 0.1.

© 2022 T. W. Murphy, Jr.; Creative Commons Attribution-NonCommercial 4.0 International Lic.;
Freely available at: https://escholarship.org/uc/energy_ambitions.

Example 13.2.3 Finally, to assess Eq. 13.3, we can crudely estimate the area under the blue curve by drawing a rectangle that we think has about the same total area. We put the top of the rectangle at the top of the blue curve and ask how wide it would need to be to approximately match the area under the blue curve.

If we make it 0.5 μm wide, it seems too thin: the area is smaller than what's under the blue curve. 1.0 μm wide seems like too much area. So we pick something in the middle like 0.75 μm (Figure 13.2).

This choice makes the area about 1.0×10^8 W/m^2/μm (value at the top of the rectangle) times 0.75 μm, coming to 7.5×10^7 W/m^2. Since this is power per area, we make a minor rearrangement of Eq. 13.3 to $P/A = \sigma T^4$ and evaluate for $T = 6000$ K to find 7.35×10^7 W/m^2. Hey, not bad! So it all hangs together.

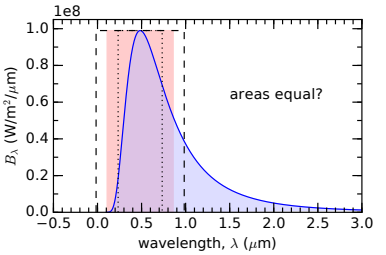


Figure 13.2: Attempts to crudely match the area under the 6000 K Planck spectrum using rectangles of three different widths. The widest (1 μm; dashed line) is too much area; the narrowest (0.5 μm; dotted line) is too small. The middle (0.75 μm; pink area) seems closest, by eye.

Now that we've batted the equations around a little bit, like a cat might do if given a new mouse toy, let's absorb some key takeaways. First, as the source gets hotter, the area under the curve[14] increases *drastically*—as the *fourth-power* of T, according to Eq. 13.3. This is seen in Figure 13.1 in that going from 3,000 K to 6,000 K[15] leads to a *tremendous* increase in area under the curve: 16 times, in fact.

Follow along using Figure 13.1.

14: . . . total power emitted

15: . . . doubling temperature

Second, as an object gets hotter, it emits at shorter wavelengths, going from red-hot to yellow-hot to white-hot. The sun, at 5,800 K, peaks at 0.5 μm, in the blue-green region. We don't see it as blue-green because it emits plenty of red light as well, making a *blend*. Notice how well the 6,000 K spectrum in Figure 13.1 covers the visible colors. A cooler star at 3,000 K has a red tint to it, since the 3,000 K spectrum (Figure 13.3) shows a marked red-heavy tilt: blue is not as well represented as red is. Conversely, a hot star at 10,000 K will have a blue tint to it, since it peaks around 0.3μm and has considerably more flux at the blue end of the spectrum than at the red end.

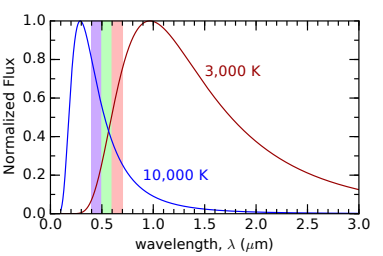


Figure 13.3: A 10,000 K star (or any blackbody) has a spectrum that tilts blue in the visible spectrum, while a cooler star (object) at 3,000 K has a red slant. Spectra are normalized to the same peak for easier comparison.

Finally, it is worth absorbing the overall lesson that photons from a glowing source emerge as a *distribution*, spanning a wide range of wavelengths (colors). This will turn out to be important in understanding why solar panels have the efficiencies that they do.

13.3 Photovoltaics

We are now prepared to dig into how photovoltaic (PV) panels actually work, and what governs panel efficiency.[16] The word "photovoltaic" can be loosely read as: volts from photons, or electricity from light.

16: Solar thermal generation will be covered in Section 13.8.

Various materials are used as the principal component in PV panels, but the vast majority are made of high-purity silicon, so we will speak in these terms alone. The basic physics will be the same for other materials as well. Getting too far into the weeds in describing the semiconductor

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*; *Freely available at*: *https://escholarship.org/uc/energy_ambitions*.

physics is outside the scope of this course, but it is worth painting a general picture—enough to appreciate how much we can expect to get out of a PV panel.

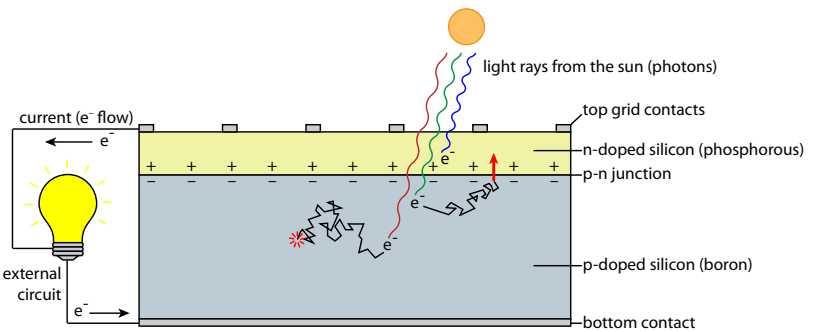


Figure 13.4: PV cell structure and function. A junction between n–doped and p–doped semiconductors sets up an electric field across the junction. If an electron promoted to the conduction band by an incoming photon wanders into the junction, it is swept across (red arrow) and successfully contributes to current. Electrons do not contribute if created above the junction—as is more probable for blue photons that are not likely to penetrate as far. Electrons do not contribute to the external (useful) current if they recombine (fill a hole) before finding the junction (red "poof").

Figure 13.4 illustrates the basic scheme. The underlying idea is that adjoining two slabs of silicon into which small amounts of two kinds of impurities have been deliberately introduced[17] that either contribute extra electrons (n-type, for negative charge donors) or create vacancies for electrons (p-type, for effective positive charge donors). Putting p-doped and n-doped materials together creates a junction[18] exhibiting a *contact potential*. The result is that "donated" electrons right at the junction in the n-doped material decide to relocate across the junction to vacancies in the p-doped material, creating a wall of negative charge on the p-side of the junction and leaving behind "holes" (missing electrons) effectively creating positive charges[19] on the n-side of the junction. In the region right around the junction[20] an electric field is set up between the separated charges. *Any electron wandering into this depletion region will be swept across the junction*, across the contact potential, and will contribute to a current that is then driven around the external circuit to return to its p-side home.[21] Figure 13.4 shows additional salient features that will be pointed out as the story develops below.

17: . . . either during or after the semiconductor crystal growth; a process called "doping"

18: So-called p–n junctions form the basis of diodes and transistors.

19: When a (negatively-charged) electron leaves an otherwise neutral medium, the medium becomes more positively charged.

20: . . . called the "depletion region," as electrons have been depleted from the portion of the n-side adjacent to the junction

21: Once it is "home," the electron will fill a vacancy created by a sun-liberated electron to end the journey.

13.3.1 Theoretical Efficiency of Photovoltaics

We will now follow the fate of one photon as it encounters the photovoltaic material. Doing so will expose the physical process of photovoltaics and simultaneously track losses to elucidate efficiency expectations.

The basic scheme is that a photon knocks an electron away from an atom in the PV cell, and this electron has some chance of being swept across the junction upon which it contributes to a useful current.[22] The goal is to get an electron across the line. It is not unlike some sports where crossing a line is the goal, but many factors are lined up defending against successful attainment of this goal. Efficiency is related to the chance that a photon will produce a "win."

22: Current is just flow of charge, and in this case is just movement of electrons through the external circuit.

© 2022 T. W. Murphy, Jr.; Creative Commons Attribution-NonCommercial 4.0 International Lic.;
Freely available at: https://escholarship.org/uc/energy_ambitions.

A photon leaves the hot solar surface aimed right at a PV panel on Earth. The photon can be any "color," distributed according to the Planck spectrum[23] in Figure 13.1. The most probable wavelength for a 5,800 K blackbody—according to Eq. 13.5—is ~0.5 μm, but it could reasonably be anywhere from 0.2–3 μm. The atmosphere will knock out (absorb or scatter) most of the ultraviolet light before it reaches the panel, and some of the infrared light is absorbed in the atmosphere as well. But almost 75% of the energy[24] makes it to the panel. What happens next depends on the wavelength.

First, we must understand something about the silicon material. The atoms in a typical silicon PV cell are arranged in an orderly lattice, grown as a single crystal. Expensive panels have mono-crystalline silicon, meaning that each 15 cm square cell comprising the panel is a thin slice of one giant crystal. Less expensive poly-crystalline (or multi-crystalline) panels have cells that are a patchwork[25] of randomly-oriented crystals at the millimeter to centimeter scale. But microscopically, both types are orderly crystals. Silicon has four electrons in its valence shell (outermost shell), so that a "happy" silicon atom is home to a four-outer-electron family. These electrons are said to exist in the valence band.[26] But provided a sufficient energy kick, an electron can leave home and enter the conduction band,[27] where it can freely move through the crystal and can potentially contribute to an electric current, if it finds the junction. The threshold energy level to promote an electron from the valence to the conduction band is called the band gap,[28] which for silicon is 1.1 eV (1.8×10^{-19} J).

Infrared photons at a wavelength of $\lambda > 1.1\,\mu$m have an energy of $E < 1.1$ eV,[29] according to Eq. 13.1. The energy falls below the band gap of silicon, and as such is not capable of promoting an electron within the silicon from the valence band to the conduction band. These longer-wavelength photons sail right through the silicon crystal as if it were transparent glass. Since these photons are not absorbed, the part of the incident energy in the infrared beyond 1.1 μm is lost. For the solar spectrum, this amounts to 23%, and is portrayed in Figure 13.5.

For the 77% of sunlight whose photons *are* energetic enough to bump an electron into the conduction band,[30] it's game-on, right? Well, not so fast—literally. Photons whose energy is *higher* than 1.1 eV have more energy than is needed to lift the electron into the conduction band. It only takes 1.1 eV to promote the electron, so a blue-green photon at 0.5 μm (~ 2.5 eV) has an excess of about 1.4 eV. The lucky electron is not just lifted out, but is given a huge boost in the process, rocketing out of the atom. It's going too fast! It knocks into atoms in the crystal and generally shakes things up a bit before settling down. We call this heat, or thermal energy:[31] its excess kinetic energy is transferred to vibrations (randomized kinetic energy of atoms) in the crystal lattice. The blue curve in Figure 13.5 reflects this loss: we get to keep all the energy at 1.1 μm (1.1 eV), thus the blue curve joins the overall black curve here.

23: The spectrum can be thought of as a probability distribution for photon wavelength, if picking out one photon.

24: This is roughly 1,000 W/m^2 out of the 1,360 W/m^2 incident at the top of the atmosphere (the solar constant, which will be derived in Section 13.4).

25: See banner image for this chapter on page 197.

26: The term "band" is used to describe energy levels. The valence band is a lower energy level.

27: . . . higher energy level

28: . . . difference between conduction and valence band energy levels

29: That $\lambda = 1.1\,\mu$m happens to correspond to 1.1 eV is a numerical coincidence, but perhaps convenient, in that remembering 1.1 for silicon covers it from both directions.

30: . . . denoted as e^- in Figure 13.4

31: Solar panels in the sun get pretty hot.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

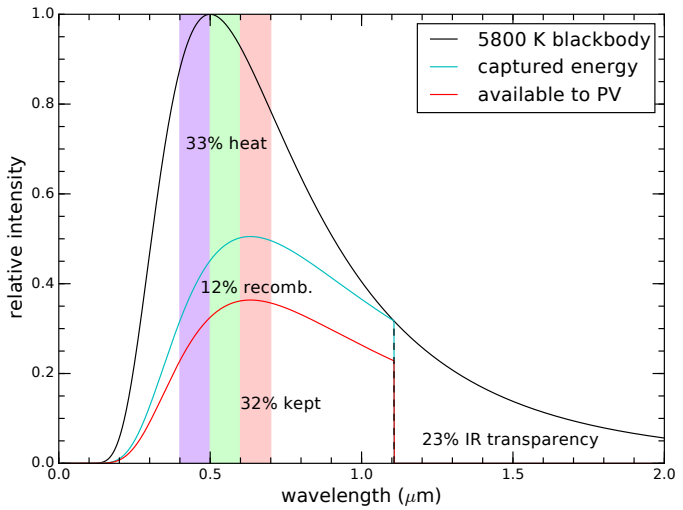


Figure 13.5: Energy budget in silicon PV cell. The areas of the four regions represent the fraction of the total incident energy going to each domain. All light at wavelengths longer than 1.1 μm (infrared; 23%) passes through the silicon without being absorbed. The photons that *are* absorbed give excess kinetic energy to electrons, losing 33% of the incident energy as heat. This effect is progressively more pronounced the shorter the wavelength. Of the remaining 44%, about a quarter disappear as electrons "recombine" with vacancies (holes) in the silicon before getting a chance to contribute to a useful current by crossing the junction, leaving 32% as the maximum theoretical efficiency.

But as the wavelength gets shorter and the energy gets higher, a greater fraction is lost to heat. Overall, 33% of the incident photon energy is lost to heat as the boosted electrons rattle the crystal before being tamed.

Now we're down to 44% of the original incident energy in the form of conduction-promoted electrons that have shaken off their excess kinetic energy. But then here's the rub: electrons are dumb. They don't know which way to go to find the junction, so aimlessly bounce around the lattice, in a motion called a random walk.[32] Some get lucky and wander into the junction, where they are swept across[33] and contribute to external current. Others fall into an electron vacancy (a hole) in a process called recombination: game over.[34] Roughly speaking, about three-quarters[35] of the electrons get lucky by wandering into the junction before being swallowed by a hole. So of the 44% available, we keep 32% (called the Shockley-Queisser limit [86]).

Another significant loss arises because some photons are absorbed in the very top layer above the junction, so that the resulting electrons do not have the opportunity to be swept across the junction to contribute to useful energy. The shorter the wavelength, the shallower the photon is likely to penetrate into the cell.[36] Meanwhile, photons around 1 μm are likely to penetrate deep—well past the junction—making it less likely that the liberated electrons will find the junction before settling into a new home (lattice site) via recombination. Figure 13.4 reflects this color-dependence, and also depicts one electron from the blue photon being generated above the junction, which will not have an opportunity to do useful work by crossing the junction.

In total, the basic physics of a PV cell is such that 20% efficiency is a reasonable expectation for practical implementations.[37] Indeed, commercial silicon-based PV panels tend to be 15–20% efficient, not far from the theoretical maximum. This may seem like a low number, but don't be disappointed! Biology has only managed to achieve 6%

32: . . . sometimes called drunken walk, depicted as meandering paths in Figure 13.4

33: . . . red arrow in Figure 13.4

34: . . . red "poof" in Figure 13.4

35: Naïvely, 50% are lucky and wander up to the junction, and 50% make the wrong choice and go down. But even those initially going down still have a chance to wander back up to the junction before time expires and they recombine, so that effectively 75% make it.

36: Any given photon has a probability distribution of being absorbed as a function of depth. Blue photons *can* penetrate deep, but are more likely to be absorbed near the front surface. A 1 μm photon *can* be absorbed near the front surface, but it is more likely to penetrate deeper into the silicon.

37: Fancy, very expensive multi-junction PV cells may be used for special applications like in space, where size and weight are extremely important and cost is less of a limitation. These devices can reach efficiencies approaching 50% by stacking multiple junctions at different band gaps, better utilizing light across the spectrum.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

efficiency in the best-case photosynthesis (algae). PV technology beats that by a factor of three! And as we'll see in Section 13.6, the only thing higher efficiency really does—besides driving up the price—is it makes the footprint (area occupied) smaller for the same power delivery. But it's already small enough to comfortably fit on most roofs, so efficiency is not a chief limitation at this point.

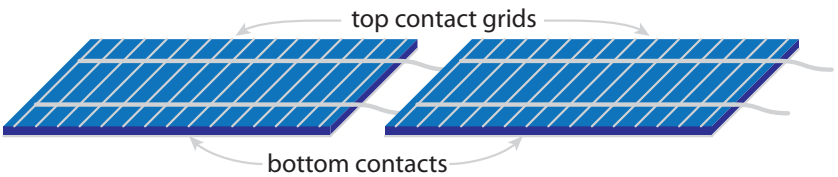


Figure 13.6: PV panels (modules) are constructed of typically 18, 36, 54, or 72 cells in series, two of which are depicted here. Cells are usually ~15 cm squares layered exactly as depicted in Figure 13.4. The bottom sides are covered by a continuous metal contact and the tops host a fine grid of metal contacts that minimize blockage of incident sunlight. Each cell presents ~0.5 V, arranged *in series* to add up to tens of volts per panel. To accomplish this, the top grid of one cell is connected to the bottom contact on the next, all down the line.

PV panels are usually constructed of many individual PV cells wired in series, as depicted in Figure 13.6. Each cell delivers maximum power when it's at a voltage around 0.45 V, and cells are usually arranged in strings of 18, adding to about 8 V. Panels commonly have 2, 3, or 4 such strings of 18—thus 36, 54, or 72 cells total—becoming 16 V, 24 V, or 32 V devices at peak power. Figure 13.7 shows typical performance curves for a PV cell (or whole panel) operating in various light levels. Recalling from Eq. 5.2 (p. 77) that electrical power is current times voltage, the power put out by a PV panel can be found as the area of the rectangle formed from the origin to the operating point somewhere on the curve. The point that maximizes area (power) is shown in Figure 13.7 as the "maximum power point." A battery being charged might hold the panel to a lower voltage, whose corresponding rectangle has a smaller area, thus operating at less than the panel's maximum power.

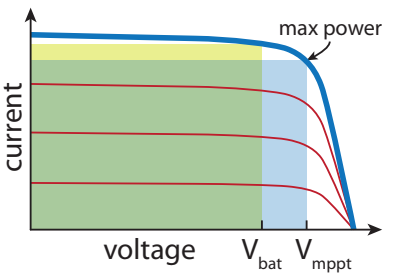


Figure 13.7: Current–voltage (I–V) curve for a PV cell. The cell in full sunlight will operate somewhere on the thick blue curve, and on a lower red curve under weaker illumination. The maximum power point (mppt) is about 0.45 V for silicon, while the nominal design might be for 0.35 V so that a 36-cell panel is sized to charge a 12 V battery. Rectangle area is proportional to power delivered, since $P = IV$.

One serious downside of panels is that because cells are wired in series, partial shading of a *single* cell limits the current the whole panel delivers to that of the *weakest link* in the chain. In other words, 17 of 18 cells in a chain might be in full sun, but if the shadow from a roof vent, chimney, or tree shades one cell and limits its current to 10% of its full value, the whole chain[38] is knocked to 10%. Bypass diodes can isolate problem sections, but usually in chunks of 18–24 cells, so that the panel can still be seriously impaired by partial shading. Connecting panels in series also creates vulnerability to partial shadowing.

38: . . . because in series, each cell shares the same current

This problem is sometimes mitigated via micro-inverters: each panel has an inverter so that higher-voltage outputs are combined in parallel.

Box 13.1: Why Not Parallel?

Given the downsides of series connection of cells, why not connect cells in *parallel*—the only other option for connecting many cells together?

© 2022 T. W. Murphy, Jr.; Creative Commons Attribution-NonCommercial 4.0 International Lic.;
Freely available at: https://escholarship.org/uc/energy_ambitions.

Series combination adds voltages, keeping the same common current. Parallel combination shares a common (low) voltage but adds currents. The same power ($P = IV$) obtains either way. But two problems arise from a parallel combination of cells. First, the ~0.5 V voltage is too small to be useful for most devices. Second, the power lost in connecting wires scales as the square of current, so designing a system with a large current is asking for trouble.[39]

That said, PV installations often combine panels in both series *and* parallel—like 10 panels in series in parallel to another 10 in series. By this time, the voltage is plenty high to offset the losses.

39: Making things worse, the voltage drop in the lines is proportional to current, diminishing an already small voltage to even less by the time it gets to its application.

13.4 Insolation

Let's start our journey from the physics principles we covered in Section 13.2. The sun's surface is a sweltering 5,770 K, meaning that it emits $\sigma T^4 \approx 6.3 \times 10^7$ W/m^2 over its surface. The sun's radius is about 109 times that of the earth's,[40] which itself is 6,378 km at the equator. Multiplying the radiation intensity by the area gives total power output: $4\pi R_\odot^2 \sigma T^4 \approx 3.82 \times 10^{26}$ W. That's one bright bulb!

40: Why this convoluted path? Context. Building from pieces we are more likely to know/remember better engages our *understanding* and ownership of the material.

Sunlight spreads out uniformly into a sphere expanding from the sun. By the time it reaches Earth, the sphere has a radius equal to the Earth–Sun distance, which is $r_{\oplus\odot} = 1.496 \times 10^{11}$ m.[41] Spreading 3.82×10^{26} W over a sphere of area $4\pi r_{\oplus\odot}^2$ computes to 1,360 W/m^2. That's what we call the solar constant [4], and it's a number worth committing to memory.[42]

41: . . . ~150 million kilometers, or 1 AU

42: See: isn't it satisfying to know that the number *comes from somewhere*? It's not just a random fact, but *connects to other pieces*. That's what the earlier margin note meant about context.

See Fig. 9.6 (p. 144) for a visual example.

Earth intercepts sunlight over the *projected* area presented to the sun: a disk of area $\pi R_\oplus^2$. Bright features like clouds and snow reflect the light back to space without being absorbed, and even darker surfaces reflect *some* of the light. In all, 29.3% of the incoming light is reflected, leaving 960 W/m^2 absorbed by the $\pi R_\oplus^2$ projected area of the planet. But now averaging the 960 W/m^2 input over the $4\pi R_\oplus^2$ *surface area* of Earth cuts the number down by a factor of four,[43] to 240 W/m^2.

43: We can understand this factor of four as two separate factor-of-2 effects determining how much solar power a particular location receives: one is simply day vs. night: half the time the sun is not up. The other half relates to the fact that the sun is not always *overhead*, so the amount of light hitting each square meter of land is reduced when the sun's rays are slanting in at an angle.

High latitude sites suffer more from low sun angles, and obviously cloudier locations will receive less sun at the surface. Taking weather into account, a decent number for the average amount of power from sunlight reaching the ground is about 200 W/m^2. This is called insolation[44]—the "sol" part of the word stemming from solar.

44: . . . also called global horizontal irradiance

Solar Flux Context	W/m^2
Arriving at Earth	1,360
Full, overhead sun (no clouds)	~1,000
Absorbed by $\pi R_\oplus^2$	960
Absorbed by $4\pi R_\oplus^2$	240
Typical insolation, includes weather	~200
Typical delivered by 15% efficient PV	30

Table 13.1: Summary of solar power densities. Full overhead sun can be larger than the global absorbed number because the global number includes reflection from clouds, while overhead direct sun corresponds to a local cloud-free condition.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

Table 13.1 summarizes these various power densities, the last line being typical insolation multiplied by 0.15 to represent the yield from a 15% efficient photovoltaic panel lying flat in a location receiving an insolation of 200 W/m^2. Figure 13.8 shows global insolation, variations arising from a combination of latitude and weather.

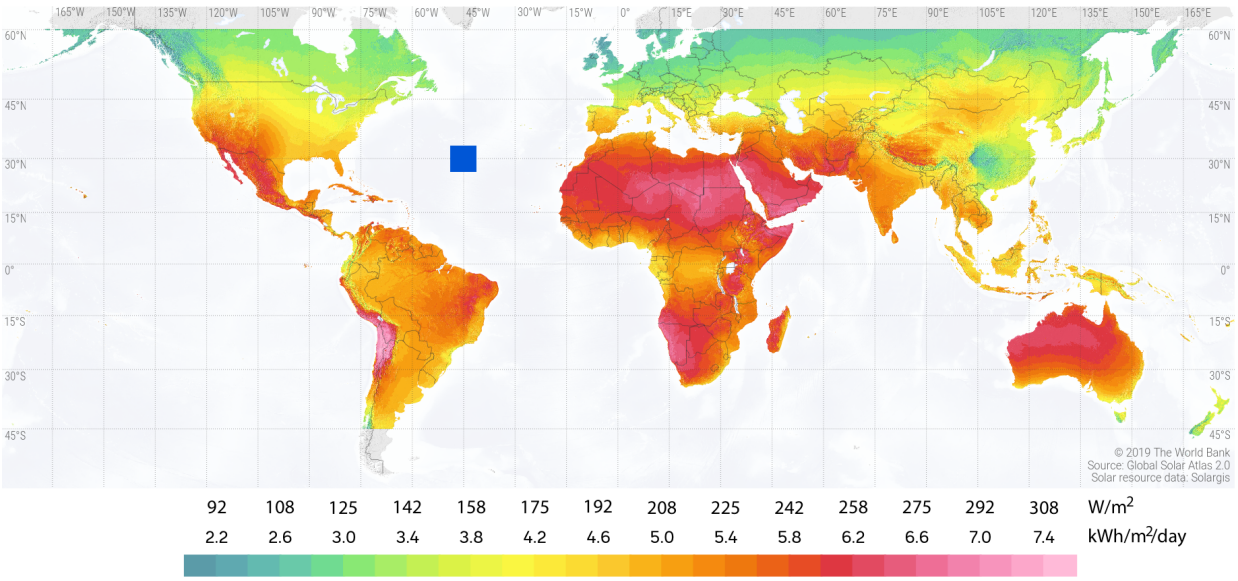


Figure 13.8: Insolation onto locally horizontal surfaces for the world (for flat plates facing directly upward), in units of W/m^2 and kWh/m^2/day. The area of the blue square in the middle of the Atlantic ocean is enough to satisfy current global energy demand, using 15% efficient solar collection (but distributed, of course). Source: The World Bank.

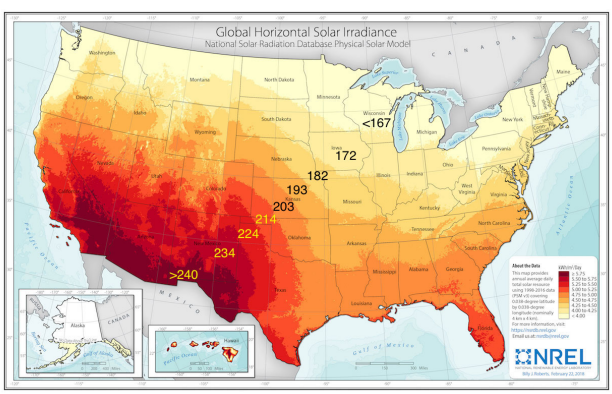


Figure 13.9: Horizontal insolation for the U.S. for a flat plate facing directly upward. Native units for the graphic are in kWh/m^2/day, the break-points between colors running from 4.0 to 5.75 kWh/m^2/day in steps of 0.25. These values can be converted to W/m^2 by multiplying by 1,000 W/kW and dividing by 24 h/day. Annotations are added once in each color band (in black or yellow) to indicate the equivalent measure in W/m^2 [87]. Alaska is not even close to scale. From NREL.

Figure 13.9 shows the variation of insolation across the U.S. The latitude effect is evident, but also weather/clouds make a mark, giving the southwest desert the highest solar potential. Even so, the variation from best to worst locations[45] is *not even a factor of two*.

Figures 13.8 and 13.9 are in the context of a flat surface.[46] For solar panels, it makes sense to tilt them to an angle equaling the site latitude and oriented toward the south.[47] The noon-time sun is always high in

45: . . . e.g., 250 vs. 150 W/m^2

46: . . . as is the definition of insolation

47: . . . toward the south for northern hemisphere locations; a more generally correct way to say it would be "toward the equator"

© 2022 T. W. Murphy, Jr.; Creative Commons Attribution-NonCommercial 4.0 International Lic.;
Freely available at: https://escholarship.org/uc/energy_ambitions.

the sky near the equator, so panels there should lie flat.[48] But at high northern latitudes, the sun is lower toward the southern horizon, so the panels should tilt up to best face the sun. Tilting at an angle equal to the latitude is the best compromise, as Figure 13.10 illustrates.

48: . . . point mostly up

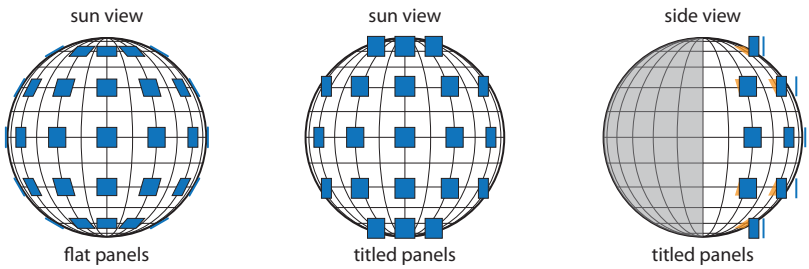


Figure 13.10: The left globe shows the sun's view of 21 panels of the same size sitting flat on the ground at their various sites. In the middle globe, the panels are all tilted up toward the equator. Notice the improvement in how much panel area is visible to the sun by doing this—especially at higher latitudes. At right is the side view, from which it is easier to appreciate why the best tilt angle is equal to the site latitude.

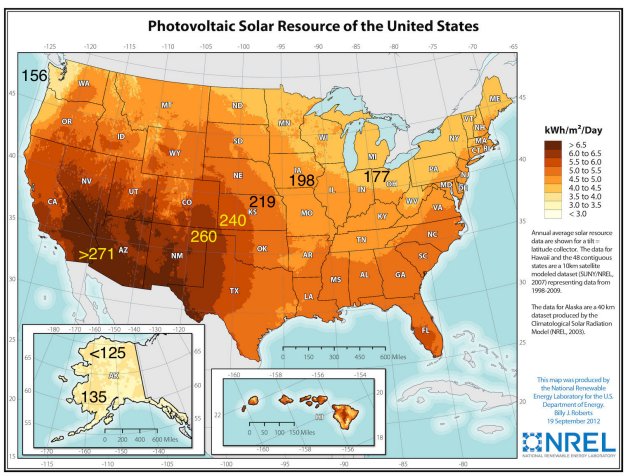


Figure 13.11: Solar potential for flat panels tilted to latitude, oriented south—relevant to PV panel installations. The graphic is presented in units of kWh/m^2/day, the breakpoints between colors running from 3.0 to 6.5 kWh/m^2/day in steps of 0.5. Annotations are added once in each color band (in black or yellow) to indicate the equivalent measure in W/m^2 [87]. From NREL.

Tilting panels toward the equator at an angle equal to site latitude optimizes annual yield, and the results are shown in Figure 13.11. Note that the numbers in Figure 13.11 are not strictly insolations, since that's defined as what reaches *flat* ground. In this case, the area (square meters) is that of the *panel*, not of the *land*.

The fact that the numbers in Figure 13.11 are higher than in Figure 13.9 is not to say that the land offers more solar energy if the panels are tilted: just that an individual panel can get more light. But in this case, panels need to be spaced out to avoid shadowing,[49] as Figure 13.12 illustrates.

Some applications need to track the sun, like those that concentrate solar power, and only work when the sun is not blocked by clouds.[50] This brings us to Figure 13.13, showing the potential per square meter of collector (mirror or lens) used for the concentration (the topic of Section 13.8.2). The same pattern holds, in that the desert southwest dominates. But a look at the numbers indicates that the cloudier regions are not much better than just a flat panel facing upward (as is the case for Figure 13.9). In the southwest, where skies are often cloud-free, the boost can be

49: . . . which can be more devastating than just fractional area blocked, due to series arrangement of cells in panel modules

50: Photovoltaics still produce 10–50% of full capacity under cloudy skies during daylight hours, depending on how thick the clouds are: daylight still means photons.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

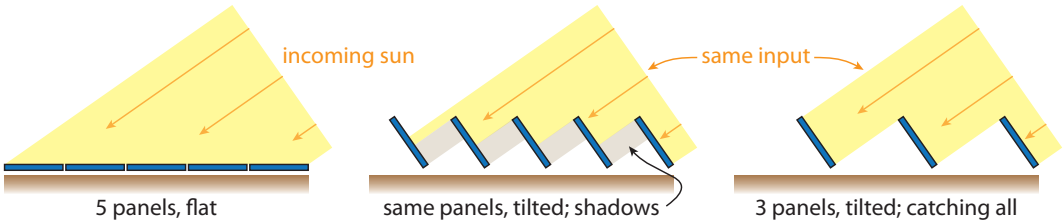


Figure 13.12: On a fixed piece of land receiving a fixed amount of sunshine at a slant angle, the amount of energy received is independent of whether the panels are flat or tilted. Just tilting the flat panels up (middle) results in self-shading. It makes the most sense to tilt and separate panels (right), one benefit being that fewer panels are needed to collect the same incident energy.

about 30% over the flat, upward-facing panel. Concentration schemes make less sense away from such regions.

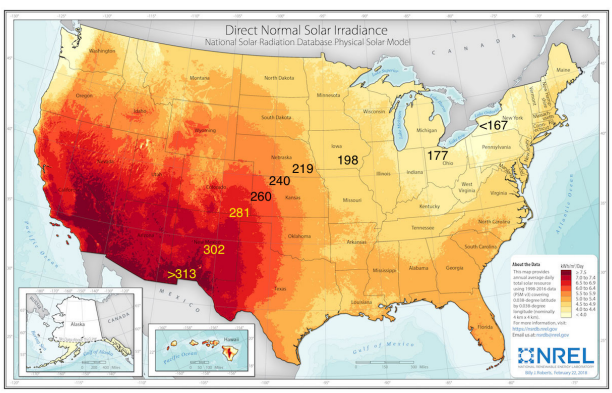


Figure 13.13: Solar potential for tracking panels, facing directly toward the sun's position and requiring a cloud-free view of the sun (concentrating collectors). The graphic is presented in units of kWh/m^2/day, the break-points between colors running from 4.0 to 7.5 kWh/m^2/day in steps of 0.5. Annotations are added once in each color band (in black or yellow) to indicate the equivalent measure in W/m^2 [87]. From NREL.

Stepping back, let's appreciate a few big-picture facets from these maps. First, numbers tend to be in the general neighborhood of 150–300 W/m^2. Burn this range in—it's a useful context. Second, the variation from the most solar-intense places in the contiguous U.S. to the weakest areas[51] is not more than a factor of two on an annual basis. This is astounding. The Mojave desert in California and the rain-forest Olympic Peninsula in Washington would seem to be practically day vs. night with respect to solar illumination. But not so much: only a factor of two.[52] Part of what this means is that if storage over *annual* timescales could be realized, solar power would become practical almost everywhere.[53]

51: . . . ignoring Arctic-leaning Alaska

52: The northwest benefits from long summer days when clouds are also less likely.

53: This would require huge storage capacity: giant batteries, for instance.

Box 13.2: Hours of Full-Sun Equivalent

A useful take-away comes from the native units used in the three maps presented here: kWh/m^2/day, as opposed to our preferred W/m^2. Although they look different at a glance, kWh is a unit of energy, so kWh/day is a power, just like W. Since a kilowatt is 1,000 W and a day is 24 h, 1 kWh/day is 1,000 Wh/24 h = 41.67 W.[54] So we can multiply 6 kWh/m^2/day by 41.67 to get 250 W/m^2.

54: The hours in numerator and denominator cancel, since the kilowatt-hour is kW *times* hours.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

But the main purpose of this box is to point out the following. Full overhead sunshine bathes the ground in about 1,000 W/m^2.[55] So if you could contrive to keep the sun directly overhead for 5 hours, you'd get 5 kWh of solar energy for each square meter on the ground. Therefore, if your site is listed as getting 5 kWh/m^2/day, it's the *equivalent* amount you'd get from 5 hours of direct overhead sun.[56]

What *actually* happens is that the day is longer than 5 hours, but for much of the day the sun is at a lower angle so that the panel is not directly illuminated, and weather can also interfere. This leads to a concept of full-sun-equivalent-hours. A site getting an annual average of 5.4 kWh/m^2/day might be said to get 5.4 hours of full-sun-equivalent each day. It's a pretty useful metric.

55: It's 1,360 W/m^2 at the top of the atmosphere, and the atmosphere blocks/scatters some of the wavelengths outside the visible part of the spectrum.

56: This equivalence relies on the convenient fact that full overhead sun is about 1,000 W/m^2. It would not work otherwise.

Box 13.2 leads to a crucial bit of understanding on characterizing a PV system. Panels are rated on what they would deliver when illuminated by 1,000 W/m^2 at a temperature of 25°C.[57] So the measure in kWh/m^2/day, or full-sun-equivalent hours tells you effectively what fraction of a day the panel will operate at its rated capacity.

Example 13.4.1 A 250 W panel at a location getting 4.8 kWh/m^2/day, or 4.8 full-sun-equivalent hours, is basically operating at 250 W for 4.8 hours out of every 24, or 20% of the time. So the panel delivers an *average* power of 50 W, not 250 W.[58]

The 250 W rating is referred to as "peak" Watts, sometimes denoted 250 W$_p$. Panels are sold this way, and now cost about \$0.50/W$_p$.

57: The 1,000 W/m^2 is reasonable, but a photovoltaic panel in full sun will be about 30–40°C hotter than its surroundings (it gets hot!), so it would have to be very cold outside to meet the specification of 25°C panel temperature. Solar panel performance wanes when hot, and will only reach 85–90% of rated capacity in typical conditions.

58: We would need to apply a de-rating of 0.85 to 0.9 to account for typical PV temperatures in the sun, bringing the panel to about 45 W average power.

A 30-year study by the National Renewable Energy Lab [88] initiated in 1960 characterized solar potential across the U.S. and produced detailed statistics on what each location might expect to collect each month for panels in different orientations. Table 13.2 is a subset of the complete data for St. Louis, Missouri.[59] All cases in Table 13.2 correspond to a panel facing south, at various tilts (including flat, at 0° and vertical at 90°; other tilts are relative to the site latitude of $\theta \approx 39°$). From this, we see that tilting the panel at the site latitude delivers an annual average of 4.8 kWh/m^2/day, matching the graphic expectation from Figure 13.11. Also shown is the monthly breakdown and how different tilts translate to performance. We will visit this table again in Section 13.6 to help us establish an appropriate size for a residential installation.

[88]: National Renewable Energy Lab (1994), *Solar Radiation Data Manual for Flat-Plate and Concentrating Collectors*

59: . . . a fairly typical solar location in the U.S.

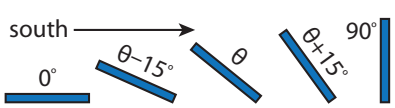


Figure 13.14: Panel tilts for Table 13.2, for $\theta = 39°$.

Angle	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Year
0°	2.2	2.9	3.9	5.0	5.9	6.4	6.4	5.7	4.6	3.5	2.3	1.8	4.2
$\theta - 15°$	3.2	3.8	4.6	5.4	5.9	6.3	6.3	6.0	5.3	4.5	3.2	2.7	4.8
θ	3.6	4.2	4.7	5.3	5.6	5.8	5.9	5.7	5.3	4.8	3.5	3.1	4.8
$\theta + 15°$	3.8	4.3	4.6	4.9	4.9	5.0	5.1	5.2	5.1	4.8	3.7	3.3	4.6
90°	3.5	3.7	3.4	3.1	2.6	2.4	2.6	3.0	3.5	3.8	3.2	3.0	3.2

Table 13.2: Solar exposure (kWh/m^2/day) for a south-facing panel in St. Louis, MO, at various panel tilts (θ is latitude, which happens to be 39° for St. Louis). 0° means a panel lying flat, pointing straight up (like on a flat roof), and 90° means vertical, like on a (south-facing) wall (see Figure 13.14).

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

13.5 The Incredible Solar Potential

The potential of sunlight can be assessed by pieces we have already seen. Multiplying the solar constant of 1,360 W/m^2 by the projected area of Earth ($\pi R_\oplus^2$) and by 0.707 to account for the 29.3% reflection loss, we compute that Earth absorbs solar energy at a rate of 1.23×10^{17} W, or 123,000 TW (1 TW is 10^{12} W). Compared to the 18 TW societal scale, that's *huge*! Notice that the non-reflected entries back in Table 10.2 (p. 168) add[60] to this same value.

Placing solar panels on just 10% of the land (itself 29% of Earth's surface area) capturing the incoming energy[61] at 15% efficiency would produce solar power at ~500 TW: about 25 times today's 18 TW usage rate. This, in turn, means that our current energy demand could be met by covering only 0.4% of land[62] with photovoltaic panels: see Figure 13.8 for a visual representation of how much this is. Solar is the *only* currently-available resource that can come anywhere *close* to satisfying our current energy appetite.[63] And it exceeds our demand by such a huge margin! We therefore have reason to be excited about solar energy: the raw numbers are great news, indeed.

So it seems like a done deal. Solar. Let's get started! Wait, why aren't we there yet?

Naturally, solar has some downsides. First, the cost. Panel cost has dropped to something like $0.50 or less per peak-Watt. To get 10 TW of delivered (average) power at typical locations getting 20% capacity factor[64] would require about 50 TW$_p$ (Example 13.4.1 defines W$_p$), costing 25 trillion dollars. The cost of other necessary components and installation double the cost for utility-scale projects [89], bringing the cost to $50 trillion.[65] The global annual economy is not quite twice this. To outfit the world with the requisite number of panels would take about 60% of the economy for a year, or 6% over 10 years, or 1.5% continuously.[66] For comparison, the world goes through 30 billion barrels of oil each year at a cost of $50/bbl, meaning $1.5 trillion per year. Installing enough panels to fully satisfy demand would cost three-decades-worth of the *entire* global petroleum budget. So it's not going to happen fast. To put things on a personal scale, Americans use power at a rate of 10,000 W. To cover this, we would need about 50 kW$_p$ of panel per person, costing $50k per person.[67] When can we expect your payment?

Another daunting realization is that even though only 0.4% of the land is needed to match today's demand, this is comparable to the amount of area currently covered by roads and buildings. A road trip across the country conveys a sense for how vast (and boring) all that pavement can be. And pavement is a fancy form of dirt. It is true that PV panels are *also* an ultra-pure, ultra-fancy form of dirt. But it's a different level of high-tech. It becomes hard to fathom that much PV around the world.

60: . . . 83,000 TW plus 40,000 TW absorbed by the surface and atmosphere, respectively

61: . . . out of the 123,000 TW total

62: . . . 1/25 of the 10% starting point

63: As we have seen, global wind may be limited to a few TW, and global hydropower would be hard-pressed to reach 2 TW.

64: 20% of 24 hours corresponds to 4.8 full-sun-equivalent hours per day.

65: The cost over the ~40 year lifetime of the panels is already competitive with conventional means, but when fuel cost is zero, *all* the cost is up-front, which presents a significant barrier.

66: . . . based on 40 year effort, at which point the first panels need replacing

67: A subtlety here is that most of the 10 kW Americans use presently is in thermal form (fossil fuels), at ~35% efficiency. For non-heating applications that can use electricity, solar has an advantage. On the other hand, mitigating intermittency via storage requires a larger PV installation by as much as a factor of two. For the purposes of a crude estimate, we'll call it even and say that 10 kW per capita from PV would cover the entire demand 100% of the time.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

A major impediment to solar power is its **intermittency**.[68] Figure 13.15 shows 31 days of solar capture, along with typical state-wide electricity demand. The two do not look very similar: not well matched. Demand is far more constant than the solar input, which is obviously zero at night. Even the peaks do not line up well, since demand is highest in the evening, well after solar input has faded away.

68: Recall that wind has a similar problem (Fig. 12.6; p. 190).

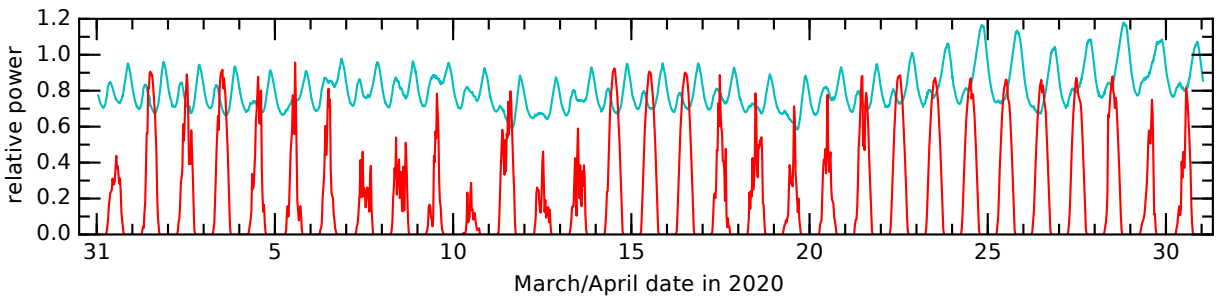


Figure 13.15: Solar input (red) and electricity demand (blue) look nothing alike. Solar data from the author's home begins 31 March 2020, while demand is for California. Tick marks denote the start of each date, at midnight. April 22–27 are essentially perfect cloudless days, while the earlier part of the month had rainy periods. Note that even a very rainy day (April 10) provides *some* solar power (15% as much as a full-sun day). Intermittent clouds cause the "hair" seen on some days. The capacity factor for the month is 19%, while the perfect six days near the end perform at 27% capacity. From this, we infer that weather caused the yield to be 70% what it would have been had every day been cloudless.

Storage is required in order to mitigate the intermittency, allowing the choppy solar input to satisfy the demand curve of Figure 13.15. We don't have good solutions for seasonal storage,[69] so a complete reliance on solar energy would necessitate over-building the system to handle months of low-sun conditions through winter (see the annual variation in Table 13.2), making it cost all that much more.

69: It's not a matter of how long the energy is stored, but a matter of sufficient *capacity* to store *enough* excess energy in summer for use during the darker winter.

Finally, all energy is not equivalent and substitutable. Solar PV cannot power passenger airplanes or power our cars down the road real-time (only via storage).[70] For all its potential, the hangups are serious enough that more than 60 years after the first demonstration of a photovoltaic cell, less than 1% of our energy derives from this ultra-abundant source.

70: It is possible to build a solar-powered aircraft or car, but not airplanes and cars as we know them (see Box 13.3). We can consider such things to be "cute" demonstrations, rather than a viable path to substitution.

Box 13.3: Why no Solar Planes?

Consider that full overhead sun delivers 1,000 W/m^2. The top surface area of a typical commercial airplane (Boeing 737) is about 450 m^2. If outfitted with the most expensive space-worthy multi-junction PV cells getting 50% efficiency, the plane would capture about 500 W/m^2 and a total of 225 kW. Sounds like a lot! The problem is that a Boeing 737 spends about 7 MW while cruising (and more during the climb). We're shy by a factor of about 25, even in optimal[71] conditions! Any solar-powered airplane[72] would be very light and very slow by air travel standards. See also Box 17.1 (p. 290) on the difficulty of battery-powered planes.

71: An airplane will not always have full overhead sun!

72: If your niece or nephew draws a solar plane in crayon, just smile, say it's very nice, put it on the refrigerator, and cry inside.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

Cars have a similar problem: a top area around 10 m^2 equipped with the most expensive PV money can buy would get 5 kW, or less than 7 horsepower. That's about 20 times less powerful than typical cars, so again: light and slow.

13.6 Residential Solar Considerations

Despite these drawbacks, it can still make a lot of sense[73] to invest in solar photovoltaics for the home. We'll explore sizing and cost in this section.

73: The author runs most of his house from an off-grid PV system he built: a solar enthusiast when it comes down to it.

13.6.1 Configurations

A typical household uses much or most of its energy when the sun is not shining: lighting, cooking, evening entertainment, charging an electric vehicle, etc. To get around this, the system would need to have local storage,[74] or be tied to the regional electrical grid so that excess production can be exported in the daytime and electricity produced by the utility used at night or when household demand exceeds solar production. The overwhelming majority of solar installations in the U.S. are grid tied, and very few mess with batteries, which can double the cost of a system and need replacement before the system has paid for itself in electricity bill savings.

74: . . . batteries: expensive, require maintenance, and periodic replacement

Box 13.4: Disappointing Dependence

A disappointing surprise to many who "go solar" is that their house has no power when the electrical service to the house is disrupted—even in the light of day. A grid tied system *needs* the grid to operate. Safety concerns prohibit PV systems from continuing to energize a disabled grid.

Only "off-grid" battery systems continue to work in these scenarios, but then the disappointment shifts to the price tag, maintenance, and replacement of worn-out batteries after a few thousand cycles.

While a description of the components and practical workings are beyond the scope of this book, students might be interested in an article the author wrote after first getting started tinkering with PV systems [90].

[90]: Murphy (2008), "Home photovoltaic systems for physicists"

13.6.2 Sizing and Cost

How large does an installation need to be? If the goal is to cover annual or monthly electricity use in a grid-tied system, the only two pieces of information needed are the typical electricity consumption in the

© 2022 T. W. Murphy, Jr.; Creative Commons Attribution-NonCommercial 4.0 International Lic.; Freely available at: https://escholarship.org/uc/energy_ambitions.

relevant period and the average solar input at that location for the period of interest.

The first can be surmised from electricity bills, usually giving a monthly total usage in kWh. We can get an approximate average scale from Fig. 7.2 (p. 105), which indicates that 42% of residential energy (11.9 qBtu per year) is from electricity. That's 5 qBtu, or 5.3×10^{18} J in one year (3.156×10^7 s), or 167 GW. Distributed among 130 million households in the U.S.,[75] average household electricity consumption is 1,285 W. Applied over 24 hours, this makes for just over 30 kilowatt-hour (kWh) per day for an average household.[76]

The next piece is solar potential at the location of interest. We'll use the excerpted data from [88] for St. Louis, Missouri found in Table 13.2.

Example 13.6.1 Let's design a grid-tied PV system for an average U.S. household in an average[77] U.S. city (St. Louis). We'll orient the panels facing south and tilted to the site latitude (39°) and purchase PV panels at 18% efficiency (pretty typical).

Table 13.2 indicates that for this configuration we can expect an annual average of 4.8 kWh/m^2/day of input. If we're shooting for 30 kWh per day, we would need 6.25 m^2 of panel operating at 100% efficiency.[78]

But 18% panels will require about 35 m^2 of panel,[79] which would be a square array about 6 meters on a side (about 20 feet) or a rectangle 5 by 7 meters, etc. The total area (400 square feet) is much smaller than a typical house footprint, so that's good news.

But panels are not marketed by the square meter. They are sold in terms of peak Watts: what the panel would deliver in 1,000 W/m^2 sunlight (see Example 13.4.1). How do we convert? Two ways are instructive.

Example 13.6.2 In one method, we multiply the 35 square-meter area from Example 13.6.1 by 1,000 W/m^2 and then by the PV efficiency (18% in this example) to get how much would be delivered: 6.3 kW.

Alternatively, we could adopt the interpretation of 4.8 kWh/m^2/day as the equivalent full-sun hours operating at peak output (Box 13.2). To get our target 30 kWh in 4.8 hours of full-sun-equivalent, we would need to produce 6.25 kW for those 4.8 hours.[80]

We get the same answer either way, which is a good check.[81]

We should assume that the panels will not achieve their rated potential due to the facts that:

- The 25°C specification is almost never realized for a PV panel in the sun: PV panels in the sun get hot, and less efficient as a result;
- The panels will get a little dirty;

75: . . . makes sense for a population of 330 million: translates to 2.5 people per household, on average

76: This is another case where students might suggest replacing this whole paragraph with the result. The point is to build connections, context, and tools to apply previous knowledge.

[88]: National Renewable Energy Lab (1994), *Solar Radiation Data Manual for Flat-Plate and Concentrating Collectors*

77: . . . solar-wise

78: . . . Divide 30 kWh/day by 4.8 kWh/m^2/day

79: . . . Divide 6.25 m^2 by 0.18

80: 6.25 kW times 4.8 hours is 30 kWh.

81: The math is actually just the same, but we rearranged the order and interpretation.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

- The equipment that converts panel output to AC electricity is not 100% efficient.

So it's a good idea to bump the number up by 20% or so, and order a 7.5 kW PV system for the case under study. A typical full cost (panels, electrical converters, installation) lately runs just shy of $3 per peak Watt (Figure 13.16), which in this case brings the price tag to roughly $20k. If electricity costs $0.15 per kWh—approximately the national average—each 30 kWh day costs $4.5, accumulating to $20k after 12 years. Federal and state incentives can make the payback time shorter.

What would these numbers become if trying to meet monthly instead of annual demands? December is usually the worst month for PV in the northern hemisphere, when the sun is lowest in the south, and the days are shortest. Table 13.2 backs this up, showing 3.1 kWh/m^2/day for the chosen panel orientation in December. This is about two-thirds the annual average, so we would need to increase the size of the system (and thus cost and payback time) by a factor of 1.5 to produce enough in December.[82]

If sizing for an off-grid system, we need to factor in some inefficiency for battery charge/discharge and design for poorer months, so should increase by another factor of at least 1.5. The cost of batteries can be rather large, too. A good rule of thumb is to have at least three days of storage in the event of no solar input for several days during a stormy period. For our 30 kWh per day target, we would want about 100 kWh of storage. As an easy way to get a cost estimate for storage, the Tesla powerwall 2 is 13.5 kWh[83] and costs about $7k apiece. If we follow along, the cost of the off-grid PV system for 30 kWh/day at an installation cost of $3/W will be 7500 W $\times$ 1.5 $\times$ 1.5 $\times$ $3 $\approx$ $50k for panels/installation plus $56k for batteries.[84] Then the batteries may be in need of replacement every 10–15 years.[85]

If this seems rather alarming, don't worry—there's a trick to making it much more affordable/practical: **don't demand 30 kWh per day**! Even though we picked 30 kWh/day due to the fact that it is the average American electricity demand, it is a worthwhile challenge[86] to seek ways to use *far* less energy than the average. Trying to make solar fit our present expectations may be the wrong approach. Also, expecting to get away from fossil fuels *and* have it be cheaper may be unrealistic.

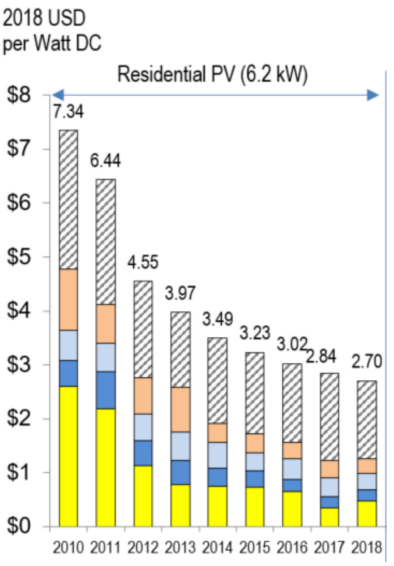


Figure 13.16: Evolving price of PV installation per peak Watt. Yellow is for the panels; the two blues for electronics; peach is labor; and hashed is for the utility hookup, inspection, taxes, profits [89, 91]. From NREL.

82: . . . resulting in over-production in summer

83: . . . so we would need about 8

84: It is often the case that battery cost is comparable to the rest of the system, roughly doubling the total cost.

85: Good batteries generally last a few thousand full charge cycles.

86: See Chapter 20 for examples.

13.7 Photovoltaic Installations

The Energy Information Administration's Electric Power Monthly (EPM) [85] provides detailed statistics on power generation in the U.S. Photovoltaic data is available in the EPM's tables 1.17.B and 6.2.B. In the usual way, we first look at installed capacity, based on the actual average *delivered* power. Figure 13.17 shows the situation in the U.S. California is

[85]: U.S. Energy Inform. Admin. (2020), *Electric power monthly*

© 2022 T. W. Murphy, Jr.; Creative Commons Attribution-NonCommercial 4.0 International Lic.; Freely available at: https://escholarship.org/uc/energy_ambitions.

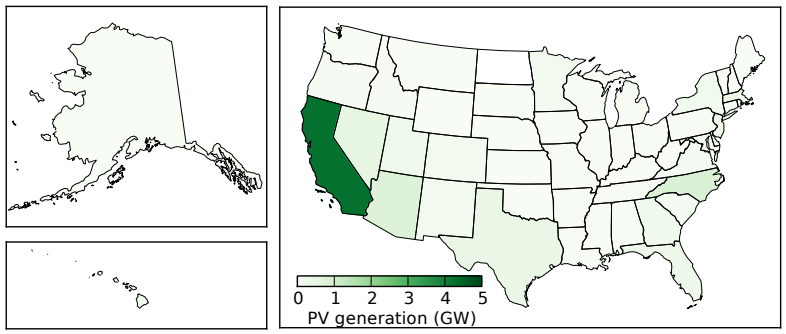


Figure 13.17: Photovoltaic power production by state, in GW, in 2018.

rocking it! The average solar power in California was 4.3 GW in 2018, far ahead of the next biggest: North Carolina at 0.82 GW. For California, this is 13% of its electricity. But electricity production is 38% of all energy in the U.S., so we might say that California gets about 5% of all its energy from solar. This is far ahead of other states.[87] The U.S. as a whole gets about 0.9% of its energy from solar.

87: North Carolina got about 5% of its electricity from solar in 2018, or less than 2% of all its energy.

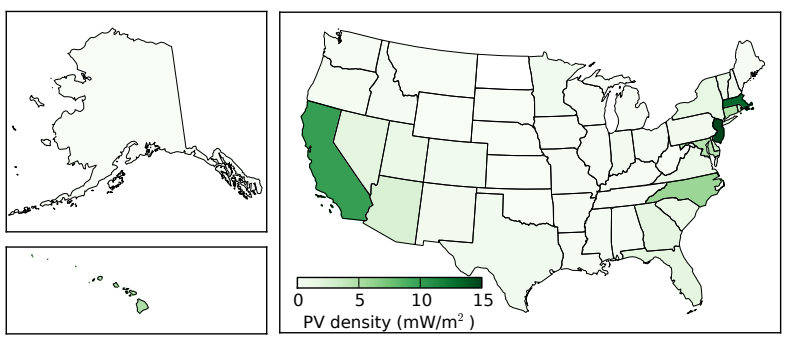


Figure 13.18: Photovoltaic power production areal density by state, in milliwatts per square meter.

Next, we divide by area to get power density from photovoltaic installations. A site having an insolation of 200 W/m^2 and 15% efficient panels has access to 30 W/m^2 of production capability (Table 13.1). Figure 13.18 shows how much we're actually getting. New Jersey has its moment in the sun, here. A few sites (NJ, MA) are pushing[88] 15 mW/m^2, which is a factor of 2,000 lower than the full potential. What this says is that only 1/2,000 of the land (0.05%) is covered by solar panels. This sort-of makes sense, right?

88: Compare to 50 mW/m^2 for hydroelectricity in Washington state (Fig. 11.6; p. 179) and 17 mW/m^2 for wind in Iowa (Fig. 12.9; p. 192).

On a per-population basis (Figure 13.19), Nevada shines brightest, at 180 W per person.[89] The southwestern U.S. is doing well overall, as is North Carolina on this measure.

89: . . . still small compared to the American metric of 10,000 W/person

Finally, we look at capacity factor: how much was generated compared to installed capacity (Figure 13.20). We expect something like 20%, corresponding to 4.8 full-sun-equivalent hours per day. The best states top out at about 0.27, equating to about 6.5 full-sun-equivalent hours per day. States at higher latitude and/or having more clouds will do

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

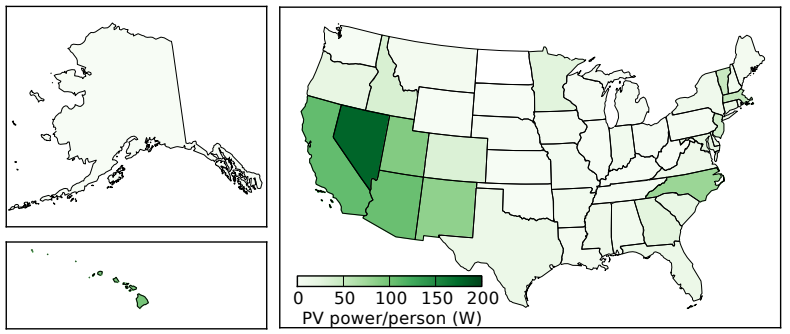


Figure 13.19: Per capita photovoltaic power production by state, in Watts per person.

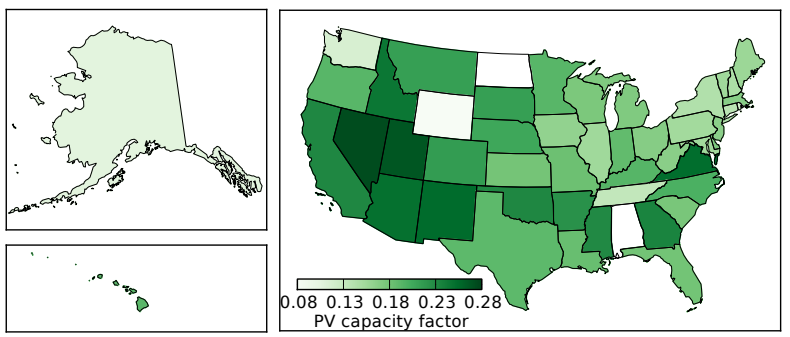


Figure 13.20: Photovoltaic capacity factors by state. While we see lots of darker green, it's because everybody has similarly *low* numbers, due to unavoidable nighttime and low sun angles. Somebody tell Wyoming, North Dakota, and Alabama to get with the program!

more poorly on this measure. Alaska clocks in just over 0.1, mapping to about 2.5 hours per day, on average.

Country	installed (GW_p)	average (GW)	% of all energy	global share (%)
China	175	~18	1.2	27
U.S.	62	10.6	0.9	16
Japan	56	6.5	3.5	10
Germany	46	5.0	3.3	7.5
India	27	4.1	1.5	6
World	510	67	1.5	100

Table 13.3: Top five global producers of PV power in 2018, accounting for two-thirds of the world's total production [92, 93]. The installed PV corresponds to peak watts (W_p), or production in full overhead sun.

Globally, two-thirds of the photovoltaic capacity is represented by five countries, shown in Table 13.3. Note that delivered power is significantly lower than installed capacity because of the low capacity factor for solar.

13.7.1 Pros and Cons of Photovoltaics

Before advancing to solar thermal generation, let's summarize the major advantages and disadvantages of solar photovoltaics. First, the good stuff:

- PV taps into a super-abundant resource—the only renewable that has such a margin;

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

- PV technology has no moving parts or steam; panels are robust and last a long time;
- PV is one of the few resources that can fit on a rooftop and provide self-contained electricity generation;
- PV efficiency is rather good: close to theoretical expectations and much better than biology has managed at getting energy from sunlight;
- PV technology works well, and despite expense has been deployed on rooftops across the world;
- Life-cycle CO_2 emissions are 15 times smaller than that of traditional fossil fuel electricity [68];
- PV is often a good solution when utility electricity is far away.

[68]: (2020), *Life Cycle GHG Emissions*

And now the less attractive aspects:

- PV is intermittent, and not well-matched to energy demand; it would be hard to "balance" the electrical grid if too much of the input came from such an intermittent source, and storage is difficult;
- PV is still expensive[90] relative to prevailing energy resources—especially important in terms of up-front cost;
- Electricity alone is not well-suited to many of our current energy demands, like transportation and industrial heat/processing;
- Stand-alone operation requires batteries, at least doubling the cost and adding maintenance/replacement demands;
- Even partial shading can be disproportionally disruptive;
- PV manufacturing involves environmentally unfriendly chemicals;
- PV deployment can harm habitats if installed in undeveloped areas.

90: Cost *has* been a major barrier, but may cease to be so as prices fall further.

13.8 Solar Thermal

Photovoltaics (Section 13.3) convert sunlight directly into electricity, but this is not the only way to harness energy from the sun. Solar energy can also be used for *heat*. We'll first have a brief look at home heating, then turn to electricity generation from solar heat.

13.8.1 Passive Solar Heat

Full sun delivers something like 1,000 W/m^2 at the earth's surface. Now imagine a window in a house intercepting 1.5 m^2 of sunlight, in effect admitting 1,500 W into the home—like a space heater, and it's free! Depending on window construction, some of the infrared energy may be blocked, so maybe not all 1,000 W/m^2 will make it inside, but a sizable portion will. Clever design has south-facing windows for receiving low-angle winter sun, but an overhang to keep out the high summer sun

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

(Figure 13.21). A dark and massive absorber[91] inside the house capturing the heat can continue to provide warmth through the evening hours. The Passive House designs mentioned in the context of Box 6.1 (p. 87) attempt to maximize solar capture so that little active heating is required.

91: . . . dark rock or brick works well

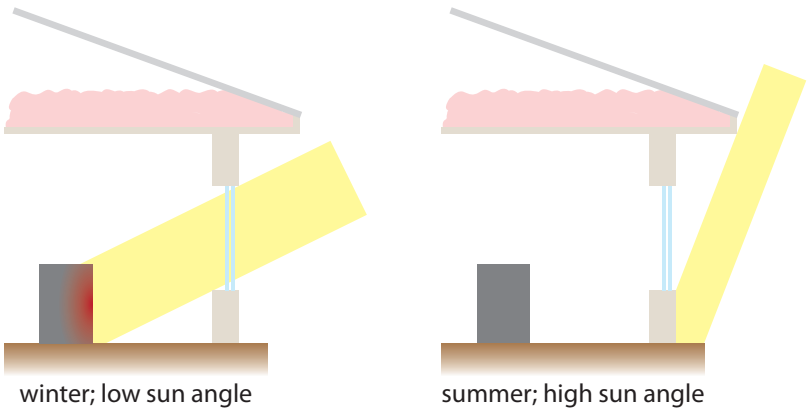


Figure 13.21: A well-designed house has thick walls, thick insulation, and double-paned windows. Even better, it can have south-facing windows that admit sunlight in the winter but not in the summer (the overhang shields the window). A large, dark thermal mass—stone or brick works well—can absorb energy and continue to release heat into the evening.

13.8.2 Solar Thermal Electricity

While 1,000 W/m^2 is nice, the power is too diffuse to get anything very hot and create a large enough ΔT to allow the operation of an efficient heat engine (Sec. 6.4; p. 88). More complex arrangements can *concentrate* solar power—think of a magnifying glass—to heat up a liquid in pipes. Figure 13.22 shows an example of a parabolic reflector that can track the sun to concentrate light onto the energy-absorbing central pipe. This shape can be extruded along a long cylinder—a "trough"—following the pipe.

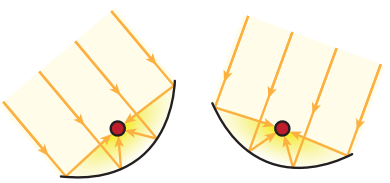


Figure 13.22: Solar trough cross sections showing the focusing of sunlight onto a central pipe. The troughs can be oriented to follow the sun.

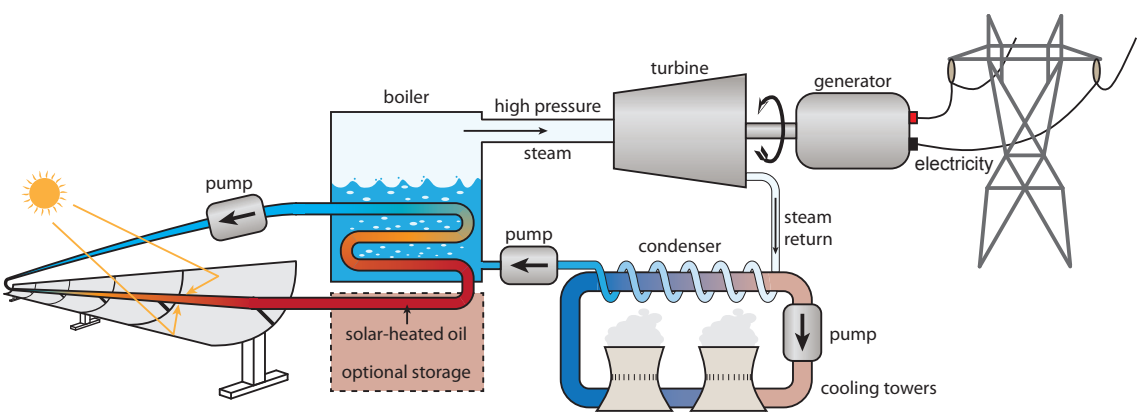


Figure 13.23: A common solar thermal power scheme uses parabolic "trough" reflectors to focus sunlight onto a central pipe, which carries oil that can be heated to very high temperatures for making steam to run a traditional electrical power plant very much like that of Fig. 6.2 (p. 90). Optional thermal storage can save heat for later use.

Figure 13.23 shows a schematic representation of a typical solar thermal (ST) collector, and a picture of one appears in Figure 13.24. A curved reflector tilts to track the sun, concentrating light onto a long pipe in

© 2022 T. W. Murphy, Jr.; Creative Commons Attribution-NonCommercial 4.0 International Lic.;
Freely available at: https://escholarship.org/uc/energy_ambitions.

front of the reflector carrying a fluid (usually oil) that can be heated to a high temperature by the absorbed sunlight. The hot oil pipes can then be run through water to boil it and make steam, thereafter driving a traditional steam power plant. Such ST arrangements are sometimes called concentrated solar power (CSP). Another common variant—called a "power tower"—is shown in Figure 13.25, in which an array of steerable flat mirrors on the ground direct sunlight to the top of a central tower to make steam.



Figure 13.24: Parabolic trough-based ST plant, in which part of the power generation facility is seen in the background. Reflectors must be spaced out to prevent self-shadowing. From U.S. DoE.



Figure 13.25: One of three "power towers" of the Ivanpah facility in California. By Craig Dietrich.

As for efficiency, solar thermal is at face value similar to PV: 15–20% is fairly typical. Broken down, roughly 50–75% of the available energy successfully transfers to the fluid, and then the heat engine delivers about 25–30% efficiency. But these numbers only apply if we count just the area of the reflective *collector*. Because they have to track the sun, and self-shadowing is to be avoided, only a small amount of the *land area* is occupied by the reflectors. Characterization of real facilities indicates that only 3% of the solar energy hitting the patch of land corresponding to the power plant is exported in the form of electrical energy.

But efficiency is not everything. 3% of a gigantic resource like solar energy input can still be tremendously large. It translates to over 6 W/m^2 for a standard insolation of 200 W/m^2, which is about thirty-times better than wind, per land area. While a field of PV panels outperform an ST installation by a factor of 5–6, the technologically simpler solar thermal designs can be more cost effective than PV. Reflectors and oil pipes are low-tech cheap devices, compared to photovoltaic material. The production cost for solar thermal is estimated to be about \$0.06/kWh, which is lower than the typical retail cost of electricity, but still a factor of two higher than fossil fuel electricity production costs.

One disadvantage of solar thermal is that concentration only works when the sun itself is visible in the sky: no obscuring clouds. One way to think of it is: if you can't see your shadow, solar concentration will not work. Meanwhile, PV panels will still produce a meaningful amount of daytime electricity from the bright sky and clouds even if the sun itself is not "out."

Balancing this disadvantage is the fact that solar thermal has some built-in storage capacity, in that the heated oil can be "banked" for some hours[92] and continue to produce electricity even during the passage of a cloud or for a few hours into the evening. In this sense, it can better match the peak of electrical demand (early evening: Figure 13.15) than can PV, which goes to zero once the sun sets.

92: . . . thus the "optional storage" block in Figure 13.23

As seen in Figure 13.13, the desert southwest is the best place in the U.S. for solar thermal electricity generation. It makes sense that deserts would be good spots, since effective concentration requires no interference from clouds. Incidentally, transmitting electricity over intermediate distances (across regions) is fairly efficient: typically better than 90% for distances shorter than ~1,000 km.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

In terms of implementation, solar thermal is a small player. In 2018, only four states produced solar thermal power, 68% from California and 22% from Arizona. Table 13.4 provides some context, comparing ST to PV in each of the four states that have any solar thermal. For the entire U.S., less than 0.1% of electricity derives from solar thermal, and PV is about 25 times bigger on the whole. Globally, ST averages about 1.1 GW[93] (2016), about half in Spain and a third in the U.S.[94]

93: . . . 0.006% of global demand

94: . . . therefore not much left in the rest of the world

Table 13.4: Solar Thermal (ST) generation in the U.S. in 2018, compared to photovoltaic (PV); MW is megawatts.

State	ST MW avg.	ST % elec.	PV MW avg.	PV % elec.	ST/PV %
California	281	1.25	4,285	19.0	6.6
Arizona	89	0.08	765	5.1	11.6
Nevada	35	0.09	552	12.1	6.3
Florida	6	0.002	326	1.2	1.8
U.S. total	410	0.086	10,565	2.2	3.9

13.8.3 Pros and Cons of Solar Thermal

Summarizing the pros and cons for solar thermal (ST), starting with the good aspects:

- ST taps into a super-abundant resource—the only renewable that has such a margin;
- ST technology is low-tech and inexpensive, using well-developed power plant technologies;
- ST has built-in short-term storage capacity for covering evening power demands;
- Life-cycle CO_2 emissions are 20 times smaller than that of traditional fossil fuel electricity [68].

[68]: (2020), *Life Cycle GHG Emissions*

And the less great stuff:

- ST requires direct sunlight; intolerant of clouds;
- ST is only possible at utility-scale, requiring a power plant;
- ST has a lower land-area efficiency than PV panels;
- Some disruption will be imposed on the local environment/habitat.

13.9 Upshot for Solar

Hands down, solar is *the only renewable resource* capable of matching our current societal energy demand. Not only can it reach 18 TW, it can exceed the mark by orders of magnitude. Finding space for panels is not a limitation. The efficiency of PV panels is perfectly respectable based on physics expectations, and beats the best that biology has done by a factor of 3–4. The efficiency is high enough that roof space tends to be more than sufficient to satisfy the demands of individual houses.

Pros and cons are listed separately for PV and ST in Section 13.7.1 and Section 13.8.3, respectively.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

Holding solar back is its intermittency[95] and high up-front cost. Intermittency can be solved by battery storage, but this can double the cost and require maintenance and periodic battery replacement. Additionally—as for many of our renewable options—all of our society's demands[96] are not well met by electricity generation.

95: ... low capacity factor that is weather-dependent

96: ... like transportation and industrial processing

Sizing up a PV installation is fairly straightforward. Having first determined how many kWh per day are to be produced, on average, divide this by the kWh/m^2/day value for the site,[97] which is essentially the number of hours of full-sun[98] equivalent, and tends to be in the 4–6 hour ballpark. This says how many kilowatts the array should produce in full sun (peak Watts). For instance, if only 10 kWh/day are needed,[99] and the region in question gets 5 kWh/m^2/day, the system needs to operate at a peak power of 2 kW$_p$, costing about $6k to purchase and install (grid tied). Inflating by 20% offsets unaccounted losses[100] to better match real conditions.

97: ... either annual or monthly

98: 1,000 W/m^2

99: ... because you are careful about energy expenditures

100: ... hot, dirty panels and conversion efficiencies

Solar thermal energy is another way to run a traditional steam-based power plant, using relatively low-tech mirrors and pipes to concentrate solar energy into a heat-carrying fluid that can later make steam. Effective efficiencies are relatively low,[101] but on the bright side, the low-tech nature makes it fairly cheap, and the technique can accommodate some degree of thermal storage for use some hours into the evening. Anything[102] starting from solar input has the potential to be a major player, given the ~100,000 TW scale of solar energy incident upon the earth.

101: 3% of solar energy hitting the plant area ends up as electricity

102: ... even if the efficiency is modest or low

13.10 Problems

1. If we had two monochromatic (single-wavelength) light sources—a green one at $\lambda = 0.5\,\mu$m and a near-infrared one at $\lambda = 1.0\,\mu$m—each emitting photons at an energy rate of 1 W,[103] how does the number of photons emerging per second from each source compare? Is it the same number for each because both are 1 W sources, or is it a different number—and by what factor, if so?

103: Hint: recall that 1 W is 1 J/s.

2. Overhead sunlight arrives on the surface of the earth at an intensity of about 1,000 W/m^2. How many photons per second strike a solar panel whose area is 1.6 square meters, if the typical wavelength is $\lambda = 0.5\,\mu$m?

3. Using the setup in Problem 2, how many photons enter your pupil every second if you look directly at the sun? When doing so, your pupil restricts to a diameter of about 2 mm.

4. The dimmest stars we can see with our eyes are thirteen orders-of-magnitude[104] dimmer than the intensity of the sun. Building off of Problem 3, how many photons enter your eye per second at this edge of detectability?

104: 10^{-13} times

© 2022 T. W. Murphy, Jr.; Creative Commons Attribution-NonCommercial 4.0 International Lic.;
Freely available at: https://escholarship.org/uc/energy_ambitions.

5. Warm humans at ~300 K and glowing-hot light bulb filaments at ~2,400 K both radiate according to Eq. 13.3. How much more power per unit area (W/m^2) does an incandescent filament emit compared to human skin, roughly?

6. The outcome of Problem 5 indicates that a hot light bulb filament emits thousands of times more power per unit area than human skin. Yet both a human and a light bulb may emit a similar amount of light[105]—both around 100 W. Explain how both things can be true?

105: Humans emit in the infrared, so we don't see it with our eyes.

7. At what wavelength does the Wien Law say the Planck spectrum will peak for a temperature of 4,500 K? Express your answer in microns and compare to Figure 13.1 for confirmation.

8. Human bodies also glow by the same physics as the sun or a light bulb filament, only it is too far out in the infrared for the human eye to see. For familiar objects (and human skin) all in the neighborhood of 300 K, what is the *approximate* wavelength of peak blackbody radiation, in microns?

ⓘ Note that the spectrum is spread over so many microns and we're being approximate about the temperature, so relax your answer to an easy, round number.

9. We might describe the efficiency of a light bulb as the fraction of its total light output that falls within the *visible* range. If using a *thermal* source[106] why would you expect it to be impossible to reach 100% efficiency at any source temperature, based on what Figure 13.1 shows?

106: . . . like the sun or an incandescent light bulb; one that obeys Eq. 13.4

10. Using a technique similar to that in the text, approximate the height and width of a rectangle that has comparable area to the spectrum for $T = 4500$ K in Figure 13.1. Compute the area of your rectangle and compare to the expectation from σT^4.

11. What are two reasons that blue photons are disadvantaged in terms of having their energy contribute to useful current in silicon photovoltaics?

12. Which photons are most responsible for heating up a silicon photovoltaic panel in full sun: blue photons or infrared photons (beyond 1.1 μm)?

13. If a blue photon having 3.3 electron-volt of energy liberates an electron in silicon, whose band gap is 1.1 eV,, what fraction of the photon's energy is "kept" by the electron once it settles down from the excess?[107]

107: The rest going into excess kinetic energy which just heats up the PV cell.

14. If a 2.5 electron-volt photon liberates an electron from silicon with a 1.1 eV band gap, how much kinetic energy does the emerging electron have? Express in both eV and Joules, and then determine the velocity of the electron if the electron mass is 9×10^{-31} kg.

15. Briefly summarize the sequence of events[108] that results in a

108: . . . consistent with the relevant physics, of course

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

successful contribution to PV current, starting with a green photon leaving the sun and ending with an electron crossing the junction.

16. Many people have an instinctive reaction to discount the < 20% PV panel efficiency as disappointingly low—perhaps thinking they should hold out for higher. Present a multi-point argument about why the efficiency is actually pretty good, and why in practice it is plenty good enough to be practical.

17. If aiming for a particular power output[109] from a PV array, describe explicitly/quantitatively how PV panel efficiency interacts with the physical size (area) of the array. For instance, what happens if the efficiency doubles or is cut in half, while keeping the same target output?

109: Make up your own number if it helps.

18. Make the connection between Figure 13.4 and Figure 13.6 by drawing a zoom-in of the bottom left corner of one of the cells in Figure 13.6.

19. Figure 13.7 shows operational curves of a PV cell for different levels of illumination. If the illumination is low and the panel continues to operate at maximum power,[110] which changes the most compared to full-sun operation: the voltage or the current? Why might lower light (fewer photons) directly connect to a lower current based on the physics of PV operation?

110: . . . largest rectangle that fits in curve

20. Replicate the calculation on page 206 (showing work) that starts with the surface of the sun being 5,770 K and finds that we receive 1,360 W/m^2 at Earth.

21. According to Figure 13.8, which continent appears to have the most solar potential? How would you rate China? Do the largest populations and/or largest energy consumers in the world tend to be well-aligned to the best solar resources?

22. Examine Figure 13.9 to determine the insolation at the "four corners" location where Arizona, New Mexico, Utah, and Colorado touch. Express this in both kWh/m^2/day and in W/m^2, showing how to convert from one to the other.

Just using this location as an unambiguous spot on the map.

23. What is a typical value for hours of full-sun-equivalent[111] of solar exposure in the U.S. based on the map and native units in Figure 13.9? Explain how you arrive at this.

111: . . . full sun meaning 1,000 W/m^2

24. A 30 year study by the National Renewable Energy Lab[112] indicates that in San Diego, a typical year delivers an annual average of 5.0 kWh/m^2/day of insolation for a flat panel facing straight up. Convert this to W/m^2.

112: . . . called the Redbook study: [88]

25. The same study mentioned in Problem 24 finds that worst year in San Diego delivered an annual average of 4.7 kWh/m^2/day

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

and the best gave 5.2 kWh/m^2/day for a flat horizontal panel. What, then, is the range of annual average insolation values in units of W/m^2 for San Diego, and what percentage variation is this, roughly (in round numbers)?

26. The study from Problem 24 finds that a flat panel facing south and tilted at various angles[113] relative to the horizontal produce the following annual average yields in units of kWh/m^2/day:

Angle	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Year
0°	3.1	3.9	4.9	6.1	6.3	6.4	6.9	6.5	5.4	4.4	3.4	2.9	5.0
$\theta - 15°$	4.1	4.8	5.6	6.4	6.3	6.3	6.8	6.7	6.0	5.3	4.5	3.9	5.6
θ	4.7	5.3	5.8	6.3	5.9	5.8	6.4	6.5	6.1	5.7	5.1	4.6	5.7
$\theta + 15°$	5.1	5.5	5.7	5.9	5.2	5.1	5.6	5.9	5.8	5.8	5.4	5.0	5.5
90°	4.5	4.3	3.9	3.2	2.4	2.1	2.2	2.9	3.6	4.4	4.6	4.5	3.5

What tilt delivers the best yield for the year, and how much better is this (in percent) than a flat plate facing straight up? What tilt appears to result in minimal seasonal variation?

113: . . . where we use θ to represent the site latitude—32.7° for San Diego

27. What if you adjusted the tilt of panels throughout the season to maximize yield? Reproduce the table above, but only writing in the highest number for each month.[114]. What average does this track produce for the year, and how much improvement (in percent) does this represent compared to the best fixed-tilt performance?

114: The pattern will "graph out" the tilt adjustments over time.

28. If typical insolation is 200 W/m^2, how much land area would be needed for a 15% efficient flat PV array supplying an average of 10 kW of power—which is the U.S. individual share? If arranged in a square, how large is the side-length of this array? Compare its size or area to something familiar.

ⓘ This exercise provides insight into PV at a personal-scale that would cover an American's total *share* of energy demand across all sectors.

29. If typical insolation is 200 W/m^2, how much land area would be needed for a 15% efficient flat PV array supplying an average of 10,000 W for every person in the U.S. (population 330 million). If arranged in a square, how large is the side-length of this array? Draw it on top of a state of your choice, to scale.

ⓘ This exercise provides insight into the total PV area needed to cover America's *total* energy demand across all sectors.

30. Based on what is presented in the text,[115] why is solar power still such a minor player if it is so hugely abundant and the technology has been around for a long time? What are some of the challenges?

115: . . . also fine to bring in prior/outside knowledge

31. You look at a PV panel and mentally estimate it to measure about 0.8 m by 1.5 m. Knowing that PV panels tend to be 15–20% efficient, you guess that it is 18% efficient. How much power would you expect it to deliver in full sun (1,000 W/m^2 incident)?

32. According to the table in Problem 26, San Diego can expect an annual average solar yield of 5.7 kWh/m^2/day when the panel is tilted to the site latitude and facing south.[116] If a household seeks to produce a modest 8 kWh per day using 16% efficient panels,

116: . . . absent any shadows, of course

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

how large will the array need to be? Express as an area in square meters, and in side length for a square of the same area.

It is a good practice to round the final size up a bit to make sure additional inefficiencies do not prevent reaching the goal.

33. Using the parameters from Problem 32, interpreting the solar yield as daily hours of peak-sun-equivalent (at solar exposure of 1 kW/m^2) what should the array size be in terms of peak Watts in order to deliver 8 kWh per day? How much would the system cost to install at \$3.00 per W_p?

34. One way to look at solar payback time time is to note that an installed system will cost something like \$3,000 for each kW_p (peak capacity), and that you'll produce x kWh from that 1 kW_p array if your region gets x hours of full-sun-equivalent on average. Since each kWh of electricity costs something like \$0.15, it becomes straightforward to compute the value per day as \$0.15x, and determine how long to match the \$3k investment. The result is independent of the actual array size, depending only on the cost per W_p, the solar yield at your location, and the cost of electricity. What would the payback time be, in years, if the cost is \$3/W_p,[117] the yield is 6 hours per day of full-sun-equivalent, and electricity in your region costs \$0.15/kWh?

117: . . . consistent with the first sentence

35. From Table 13.3, compute the capacity factors for the countries listed, and for the whole world, based on average vs. installed power. What is a characteristic range of numbers, and why[118] is it so low? Which country does the best, and which does the worst? What clues does Figure 13.8 offer as an explanation?

118: Hint: a large part is insolation vs. overhead sun.

36. How much power would a large window measuring 2 m wide and 1.5 m tall admit if the sun were shining straight[119] into the window from a cloudless sky? How many kilowatt-hours does this translate to over a four hour period, and how much is it worth monetarily compared to electricity at \$0.15/kWh?

119: . . . e.g., what is the most one might expect for direct sun?

37. Solar photovoltaics are practical for individual homes, but solar thermal is only to be found in large utility-scale installations. What is the practical reason why we should not expect solar thermal installations on peoples' rooftops for electricity generation?

38. Solar thermal has a fairly low efficiency in terms of land area of about 3%, compared to 15–20% for PV. Many would shake their heads and say that's too low to be of any use. What is the counter-argument that it may be fine?

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.



14 Biological Energy

14.1 Photosynthesis 227
14.2 Biomass 229
14.3 Biofuels 230
EROEI 231
EROEI of Biofuels 234
14.4 Upshot: Nature's Burn . . . 235
14.5 Problems 236

The renewable energy options discussed thus far have been rather different from the chemically-stored thermal energy provided by fossil fuels. These sources—hydroelectricity, wind, and solar—are good at making electricity, but are intermittent to various degrees and are not directly suitable for transportation, except via bulky batteries.

Biologically-based energy is more similar to fossil fuels in that it is a form of chemical energy burned to create thermal energy. We will focus on two major forms: solid biomass and liquid biofuels. The latter is well-suited to transportation: one of the few renewable energies that can make this claim. In some cases, the same plant can produce either food or bio-energy—depending on whether it is eaten by another biological form or by a machine.

Ultimately, biologically-based energy is a form of solar energy, creating chemical storage by means of photosynthesis. Fossil fuels are also an ancient form of biofuel, deriving from photosynthetic energy captured millions of years ago. So sunlight is the actual energy source, and photosynthesis is the mechanism by which the energy is stored in chemical form.

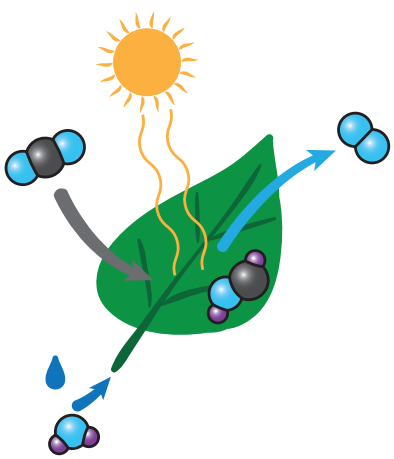


Figure 14.1: Cartoon version of photosynthesis, providing a graphical representation of Eq. 14.1. Water, CO_2, and sunlight are inputs. The leaf "exhales" oxygen and keeps sugar (only part of the final sugar molecule is pictured here).

14.1 Photosynthesis

This textbook will not focus on the complex mechanisms responsible for photosynthesis, but rather will describe the net result and efficiency. Photosynthesis involves the absorption of individual solar photons that ultimately facilitate the movement of electrons in order to change bonding structures, forming sugars, cellulose, and other materials used to construct a plant. The fundamental chemical reaction is depicted in

Photo Credit: Rudy and Peter Skitterians.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

Figure 14.1 and represented as a formula in Eq. 14.1, in which the product is a basic unit of a larger sugar molecule, like glucose ($C_6H_{12}O_6$).

$$CO_2 + H_2O + \text{light} \rightarrow CH_2O + O_2 \tag{14.1}$$

In sentence form: energy from light transforms carbon dioxide and water into a building block of sugar and releases oxygen back to the air.

Box 14.1: Where Do Plants Get Their Mass?

A valuable question to ponder: where do plants get their mass? Is it from the soil? Is it from water? Is it from the air? Take a moment to think about it. We can rule out soil on the observation that massive trees are not sitting in holes of excavated earth. Yes, the roots displace some of the soil, but a fallen tree reveals little root volume compared to the trunk and branches. And while living plant material contains significant water, completely dry plant matter[1] has plenty of mass without water.

Plant matter contains substantial carbon content,[2] and we now connect the knowledge that plants' leaves "breathe in" CO_2 and release O_2, as in Eq. 14.1. Every time this happens, the plant steals a carbon atom from the air, spitting the oxygen right back out. The carbon is stuck into a sugar or other structural molecule and stays in the plant. Thus plants obtain their dry mass out of "thin air."

1: . . . dry wood, for instance

2: . . . which we combust when burning wood

In terms of efficiency, plants tend to convert sunlight into stored chemical energy at a rate of 0.01–6%. The range is rather large due to limitations in water and nutrients. A well-watered and fertilized corn field might reach 1.5% efficiency. Algae tend to top the charts at 5–6%. Dry climates might have ample sunlight, but too little water for efficient use of the available light. Box 14.2 provides an example of how one might estimate what fraction of incident solar energy is turned into chemical storage by a potato plant.

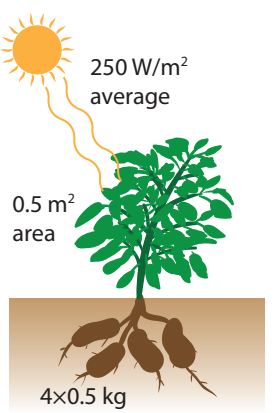


Figure 14.2: The potato plant in Box 14.2.

Box 14.2: Example Photosynthetic Efficiency

Let us consider a potato plant (Figure 14.2) as an example by which to estimate photosynthetic efficiency. The potato plant might have a leafy footprint of 0.5 m^2—a square roughly 0.7 m on a side or a circle about 0.8 m in diameter—and produce four 0.5 kg potatoes, or 2 kg of starchy material. Carbohydrates have an energy density of 4 kcal/g, so the potato plant has stored 8,000 kcal, or about 32 MJ.[3]

If the typical five-month growing-season ($\sim 1.25 \times 10^7$ s) has insolation[4] of 250 W/m^2, the plant collects 125 W[5] times 1.25×10^7 s, or 1.6×10^9 J, while making 32 MJ of spuds. The photosynthetic

3: Recall 4,184 J per kcal.

4: . . . summer, averaging day/night and weather

5: 250 W/m^2 times the plant area

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*; *Freely available at*: *https://escholarship.org/uc/energy_ambitions*.

efficiency is calculated as the ratio of output to input: about 2% in this case.

14.2 Biomass

Biological mass, or biomass, has long been utilized to supplement our energy needs, via controlled use of fire starting hundreds of thousands of years ago. Burning wood or other plant matter, and in some places dried animal feces[6] counts as utilization of biomass. Wood provides about 4 kcal of energy per gram when burned, or about 16 MJ/kg—much like proteins or carbohydrates[7] in our diet. Burning of biomass is most typically used for heating and cooking within individual homes.

6: . . . such as cow dung

7: Fats and hydrocarbons like fossil fuels are 2–3 times more energy dense.

Example 14.2.1 A 10 kg bundle of dry firewood is used to heat a home that requires 4,000 W to stay warm. How long will the wood last?[8]

Every gram of wood has 4 kcal or about 16 kJ of energy. We get 4,000 W by burning one gram every 4 seconds: 16 kJ/4 s is 4 kW. Each kilogram will therefore take about 4,000 s (a little over an hour) and the whole bundle will be gone after 11 hours.

8: . . . assuming a wood stove or other efficient device to prevent most heat from just escaping through the chimney

In the U.S. in 2018, 2.36 qBtu of the 101.25 qBtu total came from burning wood, and an additional 0.5 qBtu came from incinerating waste products [34]. Thus about 2.8% (0.1 TW) of U.S. energy comes from biomass. Out of the 11.41 qBtu of all renewables in 2018, biomass accounted for 25% of the U.S. renewable budget.[9]

[34]: U.S. Energy Inform. Administration (2011), *Annual Energy Review*

9: . . . while 23% was hydro; 22% was wind; 20% were biofuels, 8% was solar, and 2% was geothermal; see Table 10.3 (p. 170)

Globally, biomass use is estimated to be more important, at 6%,[10] constituting more than a third of global renewable resources (Fig. 7.8; p. 109). The high use of biomass in the wider world is a reflection on the difference between developed countries like the U.S. and developing countries that are more likely to rely on more primitive forms of energy like firewood and animal dung. Since most biomass around the world is burned for individual use, emissions controls are essentially non-existent, resulting in high levels of pollution—smoke and harmful chemicals that would be scrubbed out of a power plant's exhaust.[11]

10: . . . down considerably from ~25% in 1950

11: Note that CO_2 is common to both, and is *not* scrubbed out of power plant exhaust, comprising the bulk of the emissions.

Box 14.3: Life is Thin and Precious

The total mass of living organisms on Earth is estimated to be about 2 trillion tons [94]. Having a density similar to that of water, 2×10^{15} kg, if spread uniformly around the planet, would stack to 4 mm high! Put another way, a random line projecting upward from the surface of the earth would go through about 4 mm of living matter, on average. That's a pretty thin shell of life!

If we tried to substitute our 18 TW global power demand by burning

[94]: Elhacham et al. (2020), "Global human-made mass exceeds all living biomass"

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*; *Freely available at*: *https://escholarship.org/uc/energy_ambitions*.

biological matter,[12] we would run through *all* the currently-living mass—land and sea—in a short 15 years!

Can you imagine burning through all of Earth's forests and animals in 15 years? That's the rate at which we use energy today—illustrating the disparity between the biological resources on the planet and our staggering energy appetite. We can't expect to maintain our pace based on biomass and biofuels, and still have a vibrant natural planet.

12: About a quarter of the biomass is "dry" combustible material, at about 4 kcal/g.

14.3 Biofuels

Biofuels deserve their own category because the origins and end uses are different enough to warrant distinction. While the biomass sources from Section 14.2 tend to be in solid form, biofuels—as treated here—are *liquid*. Liquid fuels are instantly a big deal because they have the energy density and versatility to be used in transportation applications. An airplane can't very well fly on firewood, hydroelectricity, solar, wind, ocean currents, geothermal, or nuclear energy.[13] matter. Biofuels therefore occupy a special place in the pantheon of renewable resources as the most obvious viable replacement for petroleum—the dominant fossil fuel responsible for 92%[14] of transportation in the U.S.

13: See, for instance, Box 13.3 (p. 212) and Box 17.1 (p. 290).

14: Another 5% is from biofuels, usually blended into gasoline.

In the U.S. in 2018, 2.28 qBtu (2.3%; 0.08 TW) came from biofuels [34], which is very similar to the amount from biomass (wood, waste). Out of the 11.41 qBtu of all renewables, biofuels account for 20% of the U.S. renewable budget (Table 10.3; p. 170).

[34]: U.S. Energy Inform. Administration (2011), *Annual Energy Review*

Most prominently, ethanol is the chief biofuel, accounting for about 80% of the total. It is an alcohol that can be produced by fermenting the photosynthetically-produced sugars in the plant and then distilling the result.[15] Structurally, ethanol is very similar to ethane[16] except that the terminating hydrogen on one end of the chain is replaced by a hydroxyl group (OH; shown in Figure 14.3).

15: . . . also how "moonshine" alcohol is made

16: . . . C_2H_6: the second in the alkane sequence of methane, ethane, propane, butane, . . . , octane, etc.

Though it is not necessary to fully understand the chemistry,[17] combustion of ethanol—for comparison to the fossil fuel reactions in Eq. 8.1 (p. 121)—goes according to

17: Appendix B provides some background on chemical reactions and associated energy.

$$C_2H_5OH + 2O_2 \rightarrow 2CO_2 + 3H_2O + 29.7\text{kJ/g}. \qquad (14.2)$$

In other words, ethanol combines with oxygen via combustion (burning) producing carbon dioxide and water, also releasing energy. It is almost like the photosynthesis reaction (Eq. 14.1) in reverse.

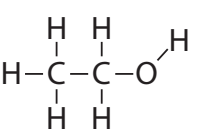


Figure 14.3: Ethanol is similar to ethane, but replacing the hydrogen at the end with hydroxyl (OH).

The energy density works out to 7.1 kcal/g, which is considerably lower than octane (representing gasoline) at 11.5 kcal/g (Table 8.2; p. 121). In terms of CO_2 production, the reaction generates 88 g of CO_2 for each 46 g of ethanol, coming to 1.9 g/g—which is lower than the 3.09 factor

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*; *Freely available at*: *https://escholarship.org/uc/energy_ambitions*.

for octane. In terms of CO_2 energy intensity, ethanol produces 64 g of CO_2 for every 1 MJ of energy: exactly the same as petroleum (Table 8.2). Generally speaking, biofuels—and other forms of biomass—are often considered to be carbon-neutral,[18] as the carbon released upon burning was taken in from the atmosphere in the process of photosynthesis, making it a cycle.

18: In practice, fossil fuels are used in the cultivating of biofuels, so it's not a perfectly clean, closed cycle in present form.

Most of the ethanol in the U.S. is blended into gasoline into E10, E15, or E85 products meaning 10%, 15%, or 85% ethanol. Not all vehicles are equipped to handle the more corrosive ethanol, and those that are ("flex-fuel" vehicles) might expect lower energy performance due to the fact that ethanol has lower energy density than gasoline.

Both the lower energy density and lower carbon mass per input fuel mass can be attributed to the oxygen atom hosted by the ethanol molecule.[19] Ethanol can derive from a number of plants. In the U.S., corn is the most common feedstock. Brazil uses sugar cane, which requires tropical climates.

19: See Sec. B.3 (p. 379) for more insight into chemical reactions and energy content.

14.3.1 EROEI

Before going further, we introduce a crucial metric for evaluating the merit of any energy source: the EROEI.

Definition 14.3.1 ***EROEI****, or Energy Returned on Energy Invested, is a measure of how profitable an energy source is in terms of energy, expressed as a ratio. For instance, a 9:1 EROEI means 9 units were extracted or produced for an investment of 1 unit, leaving a net gain of 8 units of energy. 1:1 is break-even, deriving no net energy benefit*

EROEI (sometimes just EROI) is one level above *efficiency*. Efficiency asks how much of a source is converted into a form useful to us. EROEI asks about what went into *obtaining* that source in the first place. It's still an efficiency, of sorts. Maybe a pre-efficiency.

By and large, energy does not come for free. Oil has to be actively drilled; hydroelectricity requires construction of a dam; solar panels are fabricated in an industrial process requiring energy input. So the question is: how much energy do we *get out* compared to the amount we had to *put in*? If we extract less energy than we invest, we *lose* net energy and probably should not bother.[20] If we only get a little more out, we still may question the investment.

20: In some cases, it may still make sense to pursue EROEI < 1 resources if the resulting form is otherwise hard to obtain, like food energy. As another example, we might use coal to process biofuels, in effect converting a less useful solid to a more useful liquid, even if losing energy in the process. But such desperate measures will not be favored if alternatives exist.

Example 14.3.1 Let's say an oil drilling operation uses petroleum products (like gasoline) as its only energy input for drilling and extracting oil. In one year, the operation pumps 12,000 barrels of oil, and in the effort uses 800 barrels of oil as energy input. What is the EROEI?

In this case, we just have to arrange output to input as 12,000:800, and reduce to 15:1.

Early oil wells were shallow and under pressure, producing "gushers" that exceeded 100:1 in EROEI. To understand what this means, imagine

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

using oil as the energy source for the original exploration, building the equipment, running the drill, and collecting/storing the product. An EROEI of 100:1 means that for every barrel of oil that goes *into* the process, 100 barrels come *out*. That's a good deal. A high EROEI means *nearly* "free" energy: low effort for high reward.

As we progress to more challenging oil resources,[21] the EROEI drops—now around 10–15:1 for conventional oil and as low as 3:1 for tar sands [95]. Table 14.1 provides one set of EROEI estimates for various sources. Note that estimates vary due to difficulties in proper accounting of all energy inputs, so don't take these numbers literally—just as approximate guides.

21: . . . deep water, fracking, tar sands

[95]: Hall et al. (2014), "EROI of different fuels and the implications for society"

Source	EROEI Est.	Source	EROEI Est.
Hydroelectric	40+	Solar PV	6
Wind	20	Soy Biodiesel	5.5
Coal	18	Nuclear Fission	5
Oil	16	Tar Sands	3–5
Sugar Cane Ethanol	9	Heavy Oil (Can., Ven.)	4
Natural Gas	7	Corn Ethanol	1.4

Table 14.1: EROEI estimates for various sources [96]. For example, Wind has an estimated EROEI of 20:1. See Table 7.1 (p. 106) for a refresher on how much energy we get from various sources. Canada and Venezuela tend to have heavy oil deposits.

If life were a video game, we would look at Table 14.1, decide that hydroelectric and wind are "the best," cursor over to them and "plus" those two up until we're getting all our energy from these low-energy-investment sources. But of course the world is constrained, placing real limits to what is possible. We saw in Chapter 11 and Chapter 12 that hydroelectricity and wind cannot be expected to provide more than a few terawatts, leaving a large shortfall. Meanwhile, solar has the largest raw potential. In other words, it is useful to appreciate the EROEI of various resources, but EROEI is not the sole determining factor of what is practical. A low EROEI can be tolerable if abundance makes up for it.

For resources whose energy investment is mostly up-front, before production begins, the resulting EROEI depends critically on *how long* the resource will provide energy. After all, the energy return gets larger the longer the facility can operate, while the investment part may be essentially done and unchanging. It can be difficult to predict how long a resource will last, which is part of why EROEI estimates are just that: approximate guidelines.

Example 14.3.2 Let's say a particular wind turbine achieves a 20:1 EROEI after operating for a 40 year lifetime. How many years-worth of its energy output went into constructing and installing it?

Each year the turbine produces some amount of energy, which we can call E (in Joules, for instance). In this case, it will produce $40E$ Joules in its lifetime. Since EROEI is 20:1, it must have taken $40E/20 = 2E$ Joules of input energy to create.[22] At a rate of E per year, it will produce $2E$

22: . . . because the delivered energy is 20 times the input energy

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*; *Freely available at*: *https://escholarship.org/uc/energy_ambitions*.

in 2 years, therefore taking two years to produce as much as went into its manufacture—paying for itself, energetically.

In a self-supporting sense[23] the *net* energy is $x - 1$ for an EROEI of x:1. In other words, an EROEI of 1.25:1 only "really" produces 0.25 units of exportable energy for every one unit invested, if that one invested unit *comes from* the 1.25 units extracted in a closed system. In this case, for every one unit netted,[24] 4 units went in and 5 came out–only 1 of the 5 free and clear.

23: . . . if the energy extracted is then used as the input to extract more

24: . . . multiplying the 0.25 net by 4

Example 14.3.3 A self-contained operation to produce ethanol manages to only use its own ethanol to run the entire operation of growing, harvesting, and processing the crops to produce ethanol. In one year, the operation produced a total of 250,000 L of ethanol at an EROEI of 1.25:1. How much ethanol were they able to export/sell from the operation?

The 1.25 number is associated with *total* production, which is 250,000 L in this case. Multiplying both sides of the 1.25:1 ratio by 200,000 results in an output:input ratio of 250,000:200,000 meaning that the operation required 200,000 L of input. Thus the operation was able to deliver 50,000 L to market.

Low EROEI cuts into the effective available resource, demanding investment of precious energy. As conventional resources are exhausted, forcing us to lower-EROEI deposits, even if we keep up with energy demand in absolute terms,[25] the net energy available will decline as a greater fraction of the harvest must go back into extraction.

25: . . . e.g., same number of barrels of oil produced each year

Example 14.3.4 What would have happened to an early agricultural society if the EROEI of growing food[26] slipped below 1:1, if all of the energy used to harvest the food came from workers and animals fed by the same food?

26: . . . a form of energy

At 1:1, every unit of energy extracted requires one unit of investment. Then 100% of the energy is spent acquiring energy, leaving no energy for other functions of the society (shelter, defense, etc.). Such a marginal existence could not be maintained, so some minimum exists below which the society becomes non-viable.

Note that many of the entries in Table 14.1 have low numbers, translating to a tough life in which a substantial fraction of the total energy resource is dedicated to continued energy procurement. Biological forms of energy are not superstars in this regard.

Box 14.4: Eating Our Fossil Fuels

Relatedly, and in a familiar context, the food industry in the U.S. today expends about 10 kcal of mostly fossil-fuel energy for every 1 kcal of

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*; *Freely available at*: *https://escholarship.org/uc/energy_ambitions*.

food energy consumed [97]. In a sense, we are *eating*[27] our fossil fuels! It also points to an EROEI of 0.1:1, which is well below break-even. Obviously in times prior to fossil fuels, when we used human and animal labor in our agricultural pursuits, an EROEI less than 1:1 would spell starvation: more energy going in than was recouped from the land. Today, fossil fuels give us a temporary exception, so that we can afford to lose useful energy in the bargain, turning 10 units of fossil fuel energy into one unit that we eat. We might view this as a negative aspect of the Green Revolution.

[97]: Pfeiffer (2006), *Eating Fossil Fuels*

27: . . . or at least subsidizing the energy

14.3.2 EROEI of Biofuels

Various estimates exist for the EROEI for different biofuels. Unfortunately for the U.S., the corn ethanol industry is estimated to have an EROEI of anywhere from 0.8:1 to 1.6:1. The former would mean it's a net loss of energy, and that we would have more energy available if we did not spend any of it trying to get ethanol from corn. Biodiesel (a non-ethanol biofuel produced from vegetable oils or animal fat) is estimated to have an EROEI of 1.3:1 [98]. Sugar cane may be anywhere from 0.8:1 to 10:1 [99] (see Table 14.2).

Table 14.2: Summary: EROEI of biofuels.

Source	EROEI
sugar cane ethanol	0.8–10
soy bean biodiesel	5.5
biodiesel	1.3
corn ethanol	0.8–1.6
algae-derived	0.13–0.71

[98]: Pimentel et al. (2005), "Ethanol production using corn, switchgrass, and wood; biodiesel production using soybean and sunflower"

[99]: Murphy et al. (2011), "Order from Chaos: A Preliminary Protocol for Determining the EROI of Fuels"

To explore an example of how this all plays out, let's say that corn ethanol provides an EROEI of 1.2:1—in the middle of the estimated range. This means that in order to get 1.2 units of energy out, one unit has to go in. Or for every 6 units out, 5 go in. If we use *that same* resource as the energy input—in other words, we use corn ethanol as the energy input to grow, harvest, distill, and distribute corn ethanol—then we get to "keep" one unit for external use out of every 6 units produced. For the U.S. to replace its 37 qBtu/yr oil habit with corn ethanol, it would take six times this much, or 220 qBtu (2.3×10^{20} J) of corn ethanol production each year. If the growing season is 5 months, the solar input is 250 W/m^2 on average, and the corn field is 1.5% efficient at turning sunlight into chemical energy, then each square meter of corn-land produces 4.9×10^7 J of energy[28] and we would therefore need about 5×10^{12} m^2 of land for corn. This is an area 2,200 km on a side (Figure 14.4)! The U.S. does not possess this much arable land (estimated at about 30% of this). About 4×10^{11} m^2 of land in the U.S. is currently used for corn production, which is 8% of what would be needed. And of course we must still feed ourselves. In 2018, 31% of U.S. corn production went into ethanol. We would somehow need to ramp corn ethanol production up by a factor of 40 to derive our current liquid fuels from corn in a self-sufficient way. Don't expect to see this fantasy materialize.

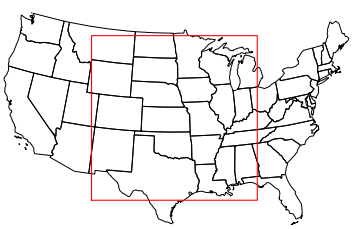


Figure 14.4: Area of corn growth needed to displace U.S. petroleum demand if at EROEI of 1.2:1. This is far larger than agriculturally productive land in the U.S.

28: 150 days times 86,400 seconds per day times 250 W/m^2 times 0.015 gives Joules per square meter produced.

Box 14.5: Why Do Corn Ethanol?

If corn ethanol has such low EROEI, why is it pursued in the U.S.?

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*; *Freely available at*: *https://escholarship.org/uc/energy_ambitions*.

Why do we have mandates to introduce ethanol into fuel blends?

Don't assume that the world is always scientifically rational and run by "adults" whose interests extend beyond personal gain. Many political factors enter: votes from midwestern swing-states,[29] job dependencies, influences from a powerful industry, the appearance of "green" (carbon-neutral) energy[30] all play a role.

29: . . . where corn is grown

30: . . . compromised by low EROEI if using fossil-fuel inputs to run production

A fundamental reason why the EROEI for biofuels tends to be low is that processing the material into ethanol requires a fair amount of energy input in the form of heat. Burning biomass, by contrast, does not have this requirement. Also, burned biomass is often gathered from untended (natural) environments that required little deliberate energy input on the parts of humans. Therefore, low EROEI is more a problem for biofuels than biomass.

14.4 Upshot for Biomass and Biofuels

Wood has always provided a source of heat for people, and will continue to do so. Its use occupied a much higher fraction of energy resources hundreds of years ago before being supplanted by fossil fuels. Still, several percent of U.S. energy comes from wood (and over 5% globally). Wood represents a renewable resource that can often be locally obtained, and will likely continue steady use,[31] potentially assuming a greater fraction again if overall energy expenditure declines.

31: . . . subject to availability in the face of deforestation

Biofuels are special due to their liquid nature, as a potential replacement for oil to drive transportation. Because photosynthesis is not terribly efficient, and the EROEI of biofuels tends to be on the low side, the amount of land needed to replace petroleum is anywhere from daunting to prohibitive. This is even before addressing the crunch an extensive expansion would place on water resources or food supply, or the degradation of arable land that may result from depleting nutrients in the soil. Algae may represent another approach, but so far the process appears to be well below break-even in terms of EROEI (from 0.13–0.71:1 [100]). It is difficult to see a meaningful path forward for wholesale replacement of liquid fuels using biological resources.

[100]: Saad et al. (2019), "Algal Biofuels: Current Status and Key Challenges"

A final perspective is that the total biological scale on the planet is estimated to be 100 TW (Table 10.2; p. 168), which is not outrageously more[32] than the current 18 TW scale of the human endeavor. Can we really imagine commandeering 20% of all life on Earth to serve our energy needs? It would actually need to be substantially more than this, given EROEI limits. It may be that Earth does not possess enough biology to offer a substitute for our current fossil fuel appetite—even if we *tried* to use it *all*.

32: . . . compared to solar or wind budgets, for instance, which are 5,000 times and 50 times our demand–not just 5 times as is the case for all biology

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

We conclude by listing some pros and cons for biologically-derived energy, beginning with the advantageous aspects:

- Biofuels offer a possible liquid fuel substitute to support transportation needs;
- Biological energy relies on dependable solar input, replenished as harvested stocks grow back;
- Biofuels represent a form of storage of solar energy, mitigating intermittency;
- Methods for growing and harvesting crops are well established;
- Burning biomass is low-tech and likely to remain part of our energy portfolio.

And the less savory aspects:

- It is difficult to scale biological energy to meaningful multi-terawatt levels;
- Heavy reliance on biological energy co-opts earth's biology and displaces natural habitat;
- Cultivating biofuels competes with food production for water and land resources;
- Low EROEI for biofuels reduces net energy available;
- Smoke and other pollutants from burning biomass can be problematic.

14.5 Problems

1. A large tree might have a trunk 0.5 m in diameter and be 40 m tall. Even though it branches out many times, pretend all the wood fits into a cylinder maintaining this 0.5 m diameter for the full height of the tree. Wood floats,[33] so let's say it has a density around 800 kg/m^3. How many kilograms of CO_2 did this tree pull out of the atmosphere to get its carbon, if we treat the tree's mass as 50% carbon?

33: ⓘ Water has a density of 1,000 kg/m^3.

Hint: carbon dioxide is 44/12 times the mass of plain carbon.

2. Using the geometry and density of the tree in Problem 1, if the resulting wood has an energy density similar to carbohydrates (4 kcal/g), and the tree spent 50 years accumulating this bulk while receiving an average of 250 W/m^2 of solar input over 5 months each year in a leafy area averaging 200 m^2 to receive sunlight, what is the net photosynthetic efficiency of the tree?

3. The U.S. gets 2.4 qBtu per year of energy from burning biomass (mostly firewood). At an energy density of 4 kcal per gram, and a population of 330 million, how many 5 kg logs per year does this translate to per person?

Now if you could just think of a way to put the answer on a *log* scale, ha ha.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

4. How many logs of firewood per day (whose parameters are specified in Problem 3) would you need to burn to provide 5,000 W of heating to a house?

5. Replicate the conclusion of Box 14.3 by assuming one-quarter of the 2 trillion tons (5×10^14 kg) of mass is combustible at 4kcal/g. How long—in years—would this amount of energy last if burning at 18 TW?

6. Given the energy densities of ethanol vs. octane (gasoline), how many liters of ethanol does it take to replace one liter of gasoline, if the densities are 789 g/L for ethanol and 703 g/L for octane?

7. If 50 kJ of energy are spent to extract 1 MJ of energy content in the form of coal, what is the EROEI?

8. Re-express an EROEI of 1.5:1 in terms of how many total units of energy must be produced in order to extract one unit of *net* energy in a self-supporting operation.

9. Imagine that the extraction of a low-EROEI biofuel is performed using energy derived from that biofuel alone—in other words, a self-contained operation not using any other (external) form of energy. We can think of the situation thusly: for each hectare of land producing fuel for external use, some additional land must be dedicated to raising the energy used to perform the extraction operation—like an overhead. If the EROEI is 1.5:1,[34] how much *total* land area would need to be devoted to the endeavor for every hectare (or any area unit you wish) that contributes to net production? See Problem 8 for a related scenario.

10. It takes a certain energy investment to fabricate a solar panel. Referring to Table 14.1, figure out how many years of the panel's output energy it takes to "pay off" the investment if the EROEI estimate assumed operation for 30 years.

11. Our modern food industry has an EROEI of 0.1:1. In pre-industrial settings, when energy investment for food production was in the form of muscle power (animal and human), why would a 0.1:1 EROEI for food have been untenable? Next, describe the conditions for an exact break-even food EROEI of 1:1. What would this mean in terms of where effort/energy goes? What would this leave for building shelters, cathedrals, or esthetic pursuits?

12. Parallel the development in the text for the land required for corn ethanol if it is self-sufficient[35] in the case for EROEI of 1.5:1—near the optimistic end of the range. How much larger is this than the area now devoted to corn, and how does it compare to total arable land in the U.S.?

13. Using the setup for Problem 12, how large would the required corn

34: You can always multiply both sides by the same factor to make both sides integers, if this is easier to understand.

35: . . . i.e., relying on its own energy to re-invest in its extraction

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

area be in terms of the side length of a square in units of kilometers. Draw this to approximate scale on a crude representation of the contiguous U.S. And this is at the optimistic end.

14. What fraction of the earth's 100 TW biological budget (all life on the planet) do you think is justifiable to use in service of human energy needs? Explain your reasoning. What does this become in TW, and how does it compare to our 18 TW current appetite?

If the EROEI is less than 2:1 (as it is for many biofuels), we would cut your estimate in TW by more than half to account for the diverted energy used for extraction.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.



15 Nuclear Energy

Most of the energy forms discussed thus far derive from sunlight—either contemporary input or fossilized storage. The *fuel* for nuclear energy is truly ancient, predating the solar system. The heavy elements participating in fission were produced by astrophysical cataclysms (most likely merging neutron stars), while the hydrogen building blocks for fusion originated in the Big Bang itself. In brief, fission involves the splitting of heavy nuclei into smaller pieces, while fusion builds larger nuclei from smaller ones.

While only fission has been successfully implemented as a source of societal energy, both types essentially boil down to the same thing: a source of heat to make steam and drive a heat engine. How and why nuclear material generates heat will be a primary focus of this chapter. Many practical concerns surround nuclear power, such as safety, weapons, waste, and proliferation of dangerous material. Self-pride for the impressive accomplishment of mastering nature well enough to implement nuclear power may not adequately justify continued reliance upon it—even if it is not a direct emitter of CO_2.

Understanding nuclear energy requires a longer journey than was needed for hydroelectricity, wind, and solar photovoltaics. We first learn about the nucleus and its many configurations, how nuclei transform from one to another through radioactive decay, the role $E = mc^2$ plays, and finally dig into the workings of fission and fusion.

15.1 The Nucleus 239
15.2 Radioactive Decay 242
15.3 Mass Energy 246
15.4 Fission 249
The Basic Idea 250
Chain Reaction 251
Fission Accounting 251
Practical Implementations . 255
Nuclear Waste 260
Weapons and Proliferation 262
Nuclear Safety 263
Pros and Cons of Fission . . 264
15.5 Fusion 264
Fuel Abundance 266
Fusion Realities 267
Pros and Cons of Fusion . . 268
15.6 Upshot on Nuclear 269
15.7 Problems 269

15.1 The Nucleus

First, what is a nucleus? Every (neutral) atom consists of a positively-charged nucleus surrounded by a cloud of negative electrons (Figure 15.1).

Cooling towers of the decommissioned Satsop nuclear power plant in Washington. Photo credit: Tom Murphy

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

The nucleus is about 100,000 times smaller than the electron cloud,[1] but contains 99.97% of the atom's mass in a super-dense nugget composed of protons (positive charge) and neutrons (no charge). While electromagnetic forces vehemently resist the close congregation of positively-charged protons, the strong nuclear force overpowers this objection and sticks the protons and neutrons together in a stable existence.

1: . . . which defines the size of the atom

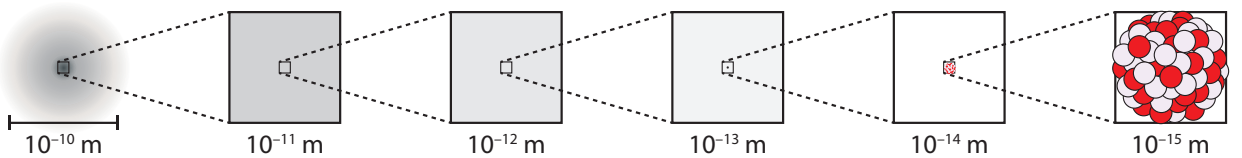


Figure 15.1: Zooming in on an atom in steps of 10×. At left, we see the entire extent of the atom's electron cloud. For a while, no nucleus is visible, being 100,000 times smaller than the atom itself. The nucleus within an atom is like a small dust grain in a bedroom.

By convention, the number of protons is labeled as Z and the number of neutrons as N. The total number of nucleons[2] is called the mass number: $A = Z + N$. It's just counting.

2: . . . a name describing *either* protons or neutrons: any nuclear constituent

Picking carbon as an example, all carbon atoms have $Z = 6$: six protons.[3] Most carbon atoms (98.93%) have $N = 6$, making $A = 12$. But some isotopes carry a different number of neutrons. In natural carbon samples, 1.07% have $N = 7$, making $A = 13$. We label such an isotope as ^{13}C ($A = 13$), or sometimes $^{13}_{6}$C ($A = 13$; $Z = 6$), or even in some cases the fully-described $^{13}_{6}C_{7}$ ($A = 13$; $Z = 6$; $N = 7$). The latter two forms are somewhat redundant—though sometimes appreciated/helpful—because *all* carbon atoms have $Z = 6$, and $N = A - Z$ always. Therefore, ^{13}C says it all, provided you can easily find or remember the Z number for carbon.[4] The general pattern for an isotope of element X is $^{A}_{Z}X_{N}$. Other common designations are, for example, C12, C13, U238, or C-12, C-13, U-238 as alternatives to ^{12}C, ^{13}C, and ^{238}U, respectively.

3: One might say this is what *defines* the carbon atom.

4: . . . just the sequential number labeling boxes in the periodic table: Fig. B.1 (p. 375)

Example 15.1.1 Write down all the various ways of designating the isotope of plutonium (Pu; 94 protons) that has mass number $A = 239$.

First, the math. $A = 239$ and $Z = 94$ so $N = A - Z = 145$. Starting at the simple end and working up, we can label this Pu239, Pu-239, ^{239}Pu, $^{239}_{94}$Pu, and finally $^{239}_{94}Pu_{145}$.

The physicist's version of the periodic table is called the Chart of the Nuclides, and contains a wealth of information. The basic layout idea is introduced in Figure 15.2, for the extreme low-mass end of nuclides.

Definition 15.1.1 *A **nuclide** is any unique combination of nucleons, so that every nucleus is one of the possible nuclides. For example, the ^{12}C nucleus is one nuclide, while ^{13}C is a distinct, different nuclide.*

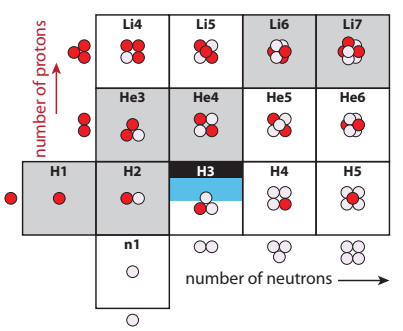


Figure 15.2: Lower left start of the Chart of the Nuclides, shown pictorially in terms of the number of protons (red) and number of neutrons (lavender) in each nuclide. Gray boxes are stable nuclides, and H3 (tritium) is semi-stable for a decade or so.

Figure 15.3 provides a full view of the chart layout: neutron number, N, runs horizontally and proton number, Z, runs vertically. Stable nuclei are indicated by black boxes at some particular integer value of N and

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

Z. Notice how they bend away from the $N = Z$ line, preferring to be neutron-rich. This can be traced to the fact that protons repel each other due to their electric charge, so the nucleus can be more tightly bound if fewer protons than neutrons are present—balanced against another penalty for being too far away from $N = Z$.

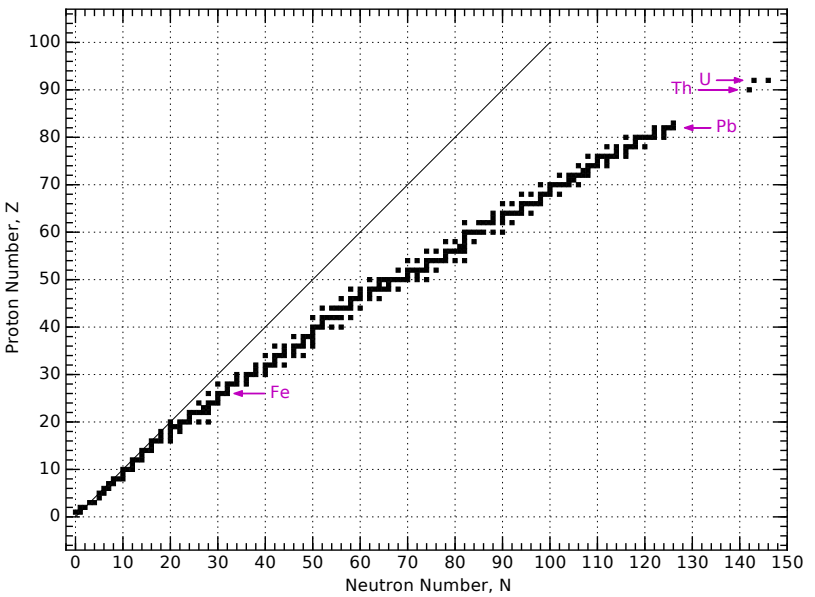


Figure 15.3: Layout of the Chart of the Nuclides, showing positions of naturally-occurring nuclei (stable or long-lived enough to be present on Earth). Stable nuclei tend to have more neutrons than protons—especially for heavier nuclei. This is why the track of stable nuclei bends away from the $N = Z$ diagonal line. Arrows point to important elements of iron, lead, thorium, and uranium at Z values of 26, 82, 90, and 92, respectively.

Figure 15.4 shows the lower-left corner of the chart in much greater detail.[5] For each element (horizontal row), properties of all known isotopes are listed—even those that are radioactive and do not persist for even a small fraction of a second before decaying. Stable isotopes are denoted by gray boxes. The mass of each, in atomic mass units (a.m.u.)—defined so that the neutral ^{12}C atom is exactly 12.0000 a.m.u.—is given, and the natural abundance as found on Earth, in percent. The Chart of the Nuclides lets us peak inside the periodic table in great detail, as Example 15.1.2 suggests.

5: Even this level of detail is short of what can be found in the actual Chart of the Nuclides, which also provides quantitative values for neutron absorption, nuclear spins, excited states, additional decay paths and associated energies.

Example 15.1.2 From the Boron row ($Z = 5$) in Figure 15.4, we can see that 19.9% of boron is found in the form of ^{10}B, while the other 80.1% is ^{11}B.

The weighted composite mass is therefore $0.199 \times 10.0129370 + 0.801 \times 11.0093055$, yielding 10.81103 a.m.u., which is the number presented as the molar mass on the periodic table.[6]

6: . . . and in the summary information in the blue box at the left of each row

Because the Chart of the Nuclides has neutron number, N, increasing from left to right, and proton number, Z, increasing vertically, nuclei having the same mass number, $A = Z + N$, are arranged on diagonals. Notice that in the region shown in Figure 15.4, we never find more than one stable element at each mass number (constant A).

Follow $A = 12$, for instance, from O12 through Be12, crossing through C12 as the only stable element of this mass.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

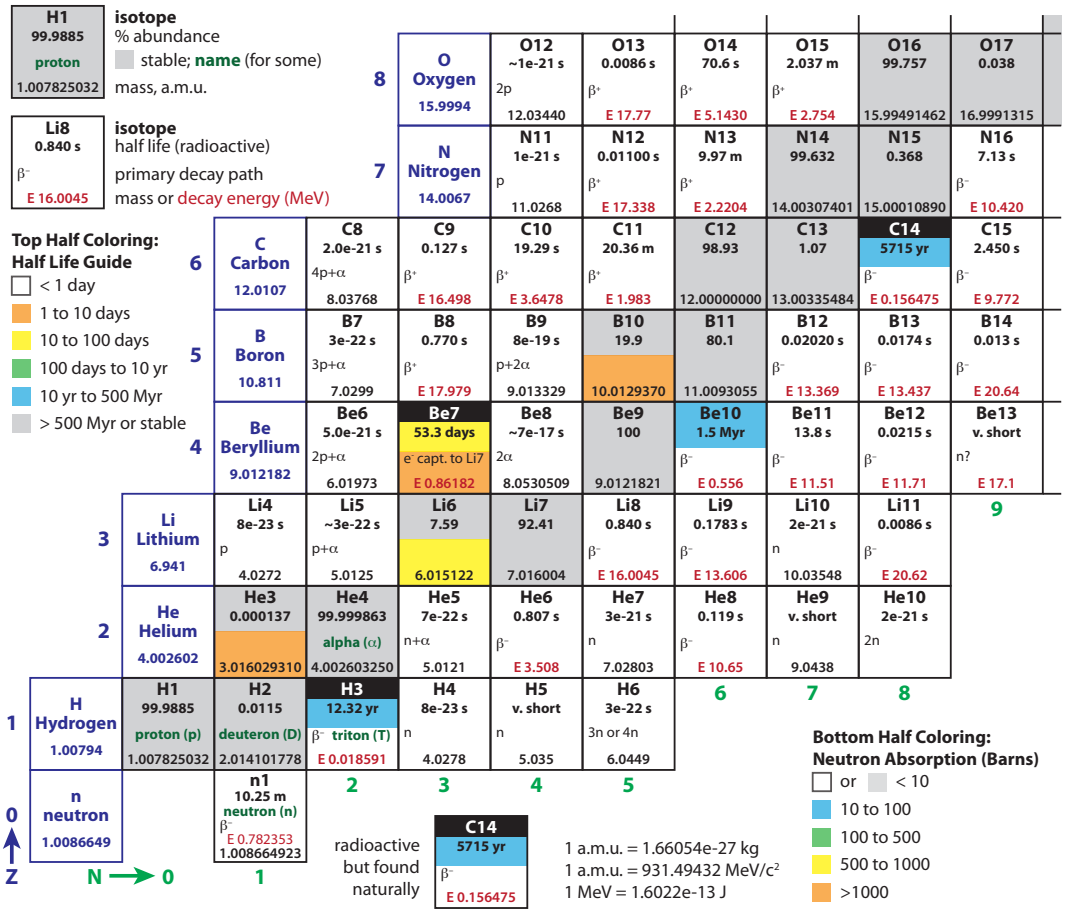


Figure 15.4: Chart of the Nuclides for the low-mass end. Neutron number, N, increases toward the right (green numbering at bottom) and proton number, Z, increases vertically (blue numbering at left). Scientific notation is expressed as, e.g., 8e–23, meaning 8×10^{-23}. A wealth of information is included: spend some time studying the surrounding guides to learn what data each box contains.

15.2 Radioactive Decay

When one nuclide, or isotope changes into another, it does so by the process of radioactive decay. Stable nuclides have no incentive to undergo such decays, but unstable nuclides will seek a more stable configuration through the decay process.

The black squares in Figure 15.3, or gray squares in Figure 15.4 are *stable*,[7] leaving all others as *unstable*,[8] meaning that they will undergo radioactive decay to a different nucleus after some time interval that is characterized by the nuclide's half life.

Definition 15.2.1 *The **half life** of a nuclide is the time at which the probability of decay reaches 50%. A large sample of such nuclides will be reduced to half the original number after one half-life. Each subsequent half-life interval removes another half of what remains.*

7: . . . or long-lived enough to be found in nature

8: Nuclides are unstable if a lower energy (more stable) configuration is within easy reach, better balancing desire for $N = Z$ against the cost of proton repulsion.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

Figure 15.4 lists a half-life[9] for each unstable nuclide. For example, the half-life for a neutron (n1 in Figure 15.4) is 10.25 minutes, meaning that a lone neutron has a 50% chance of surviving this long. The process is statistical, so an individual neutron might only last 3 seconds, or might still be around in 15 or even 60 minutes. The predictive power is sharpened the larger the sample is: half will remain after 10.25 minutes.

9: . . . in units of seconds, minutes, hours, days, or years

Example 15.2.1 If starting with 16 million separate neutrons, we would expect 8 million to still be present after 10.25 minutes, 4 million after 20.5 minutes, 2 million after 30.75 minutes, and down to 1 million neutrons in 41 minutes.

Correspondingly, a single isolated neutron has a 50% chance of still being around in 10.25 minutes, a 25% chance of lasting 20.5 minutes, and a 6.25% chance of surviving 41 minutes. Every half-life interval cuts the probability of survival in half again.

Table 15.1 summarizes these results, adding jumps to 10 and 24 half lives for illustration, ending at one neutron.

Table 15.1: Decay of 16 million (M) neutrons, having a half life of 10.25 minutes, mirroring Example 15.2.1. Time is in minutes. The number remaining at each step is given, as well as the probability of any particular neutron surviving this long. After about four hours, only one would be expected to remain (and not for much longer).

Time (min)	Half Lives	Remain	Prob.
0	0	16 M	100%
10.25	1	8 M	50%
20.5	2	4 M	25%
30.75	3	2 M	12.5%
41.0	4	1 M	6.25%
⋮	⋮	⋮	⋮
102.5	10	15,625	0.1%
⋮	⋮	⋮	⋮
246	24	~1	1/16M

Luckily, radioactive decays don't go just any which way, but stick to a very small menu of possible routes. When a decay happens, the nucleus always spits *something* out, which could be an electron, a positron, a helium nucleus (called an alpha particle), a photon, or more rarely might spit out one or more individual protons or neutrons. Because these particles can emerge at high speed (high energy), they are like little bullets firing at random times and directions into their surroundings. These bullets are potentially damaging to materials and biological tissues—especially DNA, able to cause mutations and/or initiate cancerous growth. The primary decay mechanisms pertaining to the vast majority of decays are listed below and accompanied by Figure 15.5.

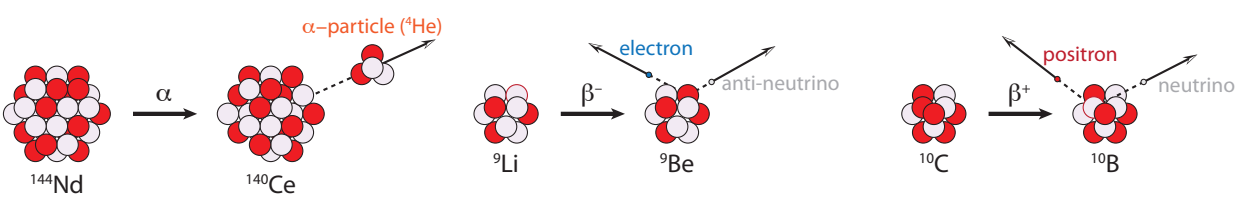


Figure 15.5: Radioactive decay mechanisms for α, β^-, and β^+. Protons are colored red, and neutrons light purple. The total nucleon counts are correct for the two beta decays, but only schematic for the larger ^{144}Nd nucleus used to illustrate alpha decay, which is predominantly seen only in heavier nuclei (aside from ^{5}Li and ^{8}Be). The positron is an anti-electron: a positively-charged antimatter counterpart to the electron. Neutrinos are sometimes called "ghost" particles for their near-complete non-interactivity with ordinary matter.

1. **Alpha decay** (α), in which a foursome of two protons and two neutrons—essentially a ^{4}He nucleus—leaps out.[10] When this happens, the nucleus reduces its N by two, reduces its Z by two, and therefore A by 4. On the chart of the nuclides, it moves two squares left and two squares down (see Figure 15.7). For example, ^{8}Be decays this way, essentially splitting into two ^{4}He nuclei;

10: Helium is found mixed in with natural gas, and derives from alpha particle decay of elements in the earth's interior.

Follow along on Figure 15.4.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

2. **Beta-minus** (β^-) decay is a manifestation of the weak nuclear force, in which a neutron within the nucleus converts to a proton, and in the process spits out an electron (β^- particle, really just e^-) to conserve total electric charge, and a neutrino—which we will ignore.[11] The mass number, A is unchanged, but N goes down one and Z goes up one (gaining a proton and losing a neutron). Thus on the chart of nuclides the motion is one left, one up. It's like a chess move (Figure 15.7);
3. **Beta-plus** (β^+) decay, like β^-, is a manifestation of the weak nuclear force, in which a proton within the nucleus converts to a neutron, emitting a positron (β^+, or e^+, or anti-electron; a form of antimatter) again maintaining charge conservation, and an ignored neutrino. Similar to β^- decay, A is unchanged, but Z is reduced by one and N gains one. On the chart, the move is diagonal: down one and right one (Figure 15.7).
4. **Gamma decay** (γ) happens when a nucleus is in an excited energy state, having been rattled by some other decay or bombardment, and it emits a high-energy photon, called a gamma ray, as it settles into a lower energy state (Figure 15.6). For γ decays, Z, N, and A do not change, so the nucleus does not morph into another flavor, and thus does not move on the Chart of the Nuclides.

11: Perhaps it is fair to ignore neutrinos since they ignore us. Neutrinos interact so infrequently with matter that a neutrino could fly through light years of rocky (Earth-like) material before being expected to hit something (interact). This extreme non-interactivity earns it the title of "ghost" particle.

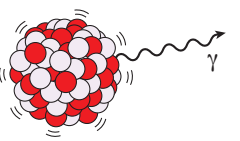


Figure 15.6: Gamma decay of an excited nucleus.

Figure 15.7 demonstrates the motion of each of these decays on the Chart of the Nuclides, and Table 15.2 summarizes the nucleon arithmetic.

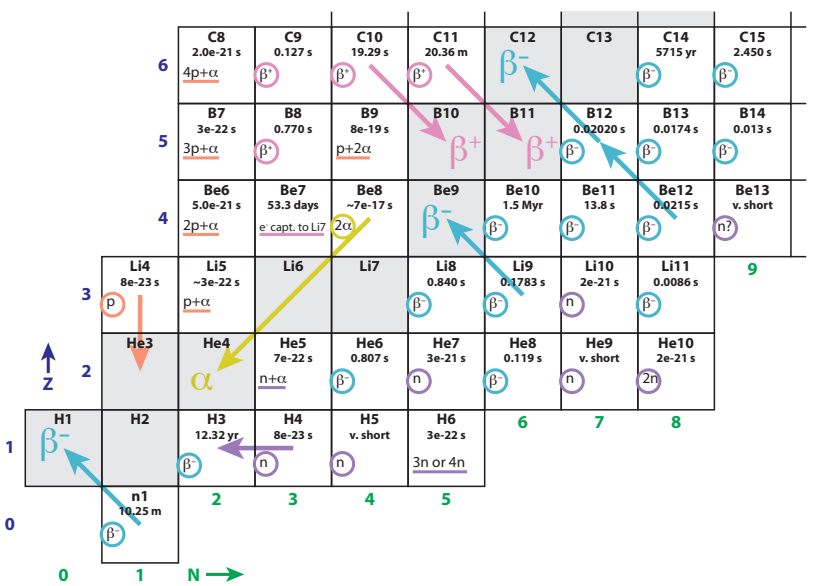


Figure 15.7: Radioactive decays shown as moves on the "chess board" of the Chart of the Nuclides. The different decay types are color-coded to match Figure 15.8, and are only shown in a few representative squares. Decays frequently occur in a series, one after the other (a decay chain), as hinted by the double-sequence starting at ^{12}Be and ending on ^{12}C. Note that the square of every unstable nuclide indicates a decay type, even if arrows are not present.

Table 15.2: Summary of decay math on nucleon counts.

Decay	$Z \rightarrow$	$N \rightarrow$	$A \rightarrow$
α	$Z-2$	$N-2$	$A-4$
β^-	$Z+1$	$N-1$	unchanged
β^+	$Z-1$	$N+1$	unchanged
γ	unchanged	unchanged	unchanged

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

Example 15.2.2 What will the fate of ^{8}He be, according to Figure 15.4?

We can play this chess game! According to the chart, the primary decay mechanism of He8 is β^- with a half-life of about a tenth of a second. It will become ^{8}Li, which hangs around for about a second before undergoing another β^- decay to ^{8}Be. This one lasts almost no time at all ($\sim 10^{-16}$ s) before α decay into two alpha particles (two ^{4}He). Such a sequence is called a decay chain.

As is evident in Figure 15.8, unstable isotopes *above* the stable track in Figure 15.3 tend to undergo β^+ decays to drive toward stable nuclei, while those *below* the track tend to experience β^- decays to drive up toward the stable track. The α decays are more common for heavy nuclei (around uranium), which drive toward the end of the train of stable elements in Figure 15.3, ending up around lead (Pb). We can understand the abundance of lead as a byproduct of heavy-element decay chains.

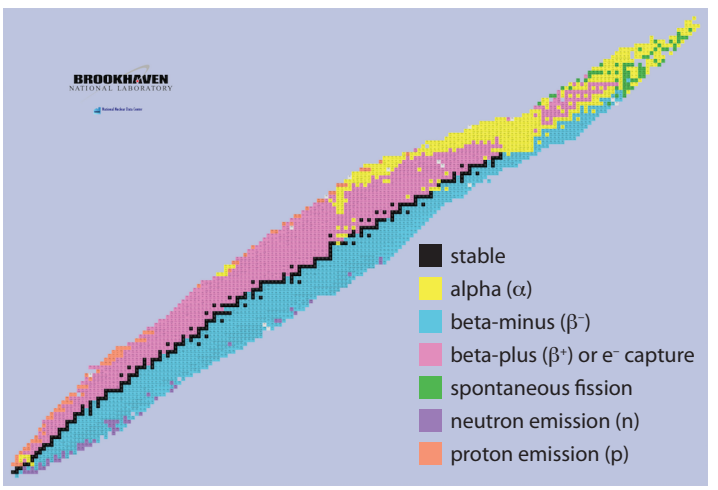


Figure 15.8: Another view of the Chart of the Nuclides, color coded to indicate prevailing decay modes as a function of position on the chart. Note that β^+ sometimes captures an electron rather than emitting a positron, but amounting to the same thing, essentially. From U.S. DoE.

Box 15.1: The Weak Nuclear Force

An aside worth making is that having discussed beta decays, governed by the weak nuclear force, we have now covered all four known forces of nature: gravity, electromagnetism, the weak nuclear force, and the strong nuclear force. That's it: a small menu, really. The latter three are unified into a Standard Model of Physics, but gravity—described by General Relativity—has defied all attempts at "grand unification," or a "theory of everything" trying to unite all four forces under a single theoretical framework. One implication is that known physics offers no other "magic" solutions to our energy needs. No new forces have come to light in more than half-a-century, despite dramatic advances in tools to probe the fundamental nature of physics.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

15.3 Mass Energy

Energy—whatever the form—*has mass* and actually changes the weight of something, although almost imperceptibly. A hot burrito has more mass than the exact same burrito—atom for atom—when it's cold.[12] Most of us are familiar, at least casually, with the famous relation $E = mc^2$. More helpfully, we might express it as

$$\Delta E = \Delta m c^2, \tag{15.1}$$

where the Δ symbols indicate a *change* in energy or mass, and $c \approx 3 \times 10^8$ m/s is the speed of light. Using kilograms for mass results in Joules for energy. Because c^2 is such a large number (nearly 10^{17}), the mass change associated with daily/familiar energy quantities is negligibly small. Box 15.2 explains why $E = mc^2$ is valid for all energy exchanges—not just nuclear ones—but generally results in mass changes too small to measure in non-nuclear contexts. Earlier, we discussed conservation of energy. More correctly, we observe conservation of mass-energy. That is to say, a system *can* actually gain or lose net energy if the mass changes correspondingly. In the case of nuclear energy release, the "new" energy comes at the expense of *reduced mass*.

12: The burrito is also ever-so-slightly more massive if it has kinetic energy, gravitational potential energy, or any form of energy. A battery is more massive when charged, even if no atoms or electrons are added. Incidentally, charging a battery does not mean literally adding electrical charges (adding particles), but amounts to rearranging electrons among the atoms within the battery.

Box 15.2: $E = mc^2$ Everywhere

Physics is not selective about when we might apply $E = mc^2$. It always applies, to every situation. It's just that outside of nuclear reactions it does not result in significant mass differences.

For example, after we eat a 1,000 kcal burrito to fuel our metabolism, we expend the energy[13] and lose mass according to $\Delta m = \Delta E/c^2$. Since $\Delta E \sim 4$ MJ (1,000 kcal), we find the associated mass change is 4.6×10^{-11} kg, which is ten orders-of-magnitude smaller than the mass of the burrito itself.[14] So we'd never notice, even though it's really there.

When we wind up a mechanized toy, coiling a spring, we put energy into the spring and the toy *actually gets more massive*! But for every Joule we put in, the mass only increases by about 10^{-17} kg. Forgive us for not noticing. Only in nuclear contexts are the energies large enough to produce a measurable difference in mass.

13: . . . ultimately given off as thermal energy to our environment

14: This amount of mass corresponds to that of a tiny length of hair that is shorter than it is wide.

Example 15.3.1 Since mass and energy are intimately related, it is common to express masses in *energy* terms. How would we express 12.0 a.m.u. in MeV (a unit of energy; see Sec. 5.9; p. 78)?

1 a.m.u. is equivalent to 1.66×10^{-27} kg (last row of Table 15.4), so 12 a.m.u. is 1.99×10^{-26} kg. To get to energy, apply $E = mc^2$, computing to 1.8×10^{-9} J of energy. Since 1 MeV is 1.6×10^{-13} J, we end up with

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

11,200 MeV corresponding to 12 a.m.u. (1 a.m.u. is 931.5 MeV).

In practice, and perhaps surprisingly, atoms (nuclei) weigh *less* than the sum of their parts due to binding energy. In order to rip a nucleus completely apart and move all the nucleons far from each other, energy must be *put in* (left part of Figure 15.9). And any change in energy is accompanied by a change in mass, via $\Delta E = \Delta mc^2$. All the energy that must be injected to completely dismantle the nucleus *weighs something*! So the mass of the individual pieces after dismantling the nucleus is effectively the mass of the original nucleus *plus* the mass-equivalent of all the energy that was put in to tear it apart (middle panel of Figure 15.9). Therefore, binding energy effectively *reduces* the mass of a nucleus, which we will now explore quantitatively.

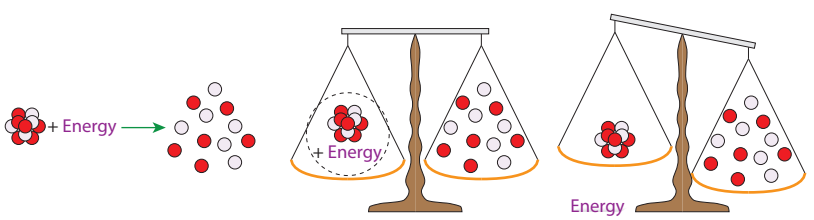


Figure 15.9: One must add energy to overcome nuclear binding energy in order to bust up a nucleus into its constituent nucleons (left). Thus, the collective mass of a nucleus *plus* the mass associated with the energy it takes to break it apart (via $E = mc^2$) must be equal to the sum of the masses of the constituent parts (middle). Therefore, if we compare the mass of the nucleus *alone* (removing the energy's mass from the scale) it must be less than the mass of the loose collection of nucleons (right).

A careful look at Figure 15.4 reveals that lighter stable nuclei (gray-squares) at the lower left of the chart have a mass a little larger than the corresponding mass number, but by the upper right—around oxygen—the mass has edged just lower than A. Table 15.3 shows this trend, confirmable in Figure 15.4 for the first four nuclides in the table. The difference between mass and A is most negative around iron, then turns around and becomes positive again for heavy elements like uranium.

Table 15.3: Example mass progression.

Nuclide	A	mass (a.m.u.)
^{2}H	2	2.014
^{4}He	4	4.003
^{12}C	12	12.000
^{16}O	16	15.995
^{56}Fe	56	55.935
^{235}U	235	235.044

What is going on here? If the mass of a nucleus were just the sum of its parts, we would expect the total mass to just track linearly as we add more pieces. In fact, if we try to build a neutral carbon atom out of 6 protons, 6 neutrons, and 6 electrons, the sum, according to Table 15.4, should be 12.099 a.m.u., not 12.000. The discrepancy is due to nuclear binding energy, as was introduced in Figure 15.9.

Table 15.4: Constituent masses of atomic building blocks, expressing the same basic thing in three common units systems.

Particle	a.m.u.	10^{-27} kg	MeV/c^2
proton	1.0072765	1.6726219	938.2720882
neutron	1.0086649	1.6749275	939.5654205
electron	0.00054858	0.000911	0.510999
(a.m.u.)	1.0000000	1.660539	931.494102

Nuclear binding energy is *incredibly* strong[15] and is able to overpower the natural electric repulsion between positively charged protons and stick them together in an unwilling bunch. The strong nuclear force only acts over a tiny range within about 10^{-15} m:[16] it is very powerful

15: . . . relating to what we call the strong nuclear force

16: The whole atom is around 10^{-10} m in scale

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

on short length scales, but ceases to operate much beyond the confines of the nucleus. Think about binding energy this way: if we tried to pry a proton or a neutron (nucleon, generically) away from a nucleus, we would encounter a very powerful force opposing the action. But let's say we persist, and *do work* in extracting the nucleon by the usual recipe of force times distance. It is *so much* work, in fact, that $\Delta E = \Delta m c^2$ becomes relevant, measurably altering the mass.

Table 15.5 walks through some example calculations, one of which is traced in Example 15.3.2. Because the ^{1}H nuclide is just a lone proton, it has no binding energy.

Grab a calculator and follow Example 15.3.2 yourself!

Nucleus	$\Sigma m_{p,n,e}$	actual m	Δm	Δmc^2 (MeV)	MeV per nucleon
^{1}H	1.007825	1.007825	0	0	0
^{2}H	2.016490	2.014102	0.002388	2.22	1.11
^{4}He	4.032980	4.002603	0.030377	28.29	7.07
^{12}C	12.09894	12.000000	0.098940	92.16	7.68
^{56}Fe	56.46340	55.934942	0.528447	492.25	8.79
^{235}U	236.9590	235.043920	1.915065	1783.85	7.59

Table 15.5: Example nuclear binding energy calculations. The second column is the simple sum of masses of protons, neutrons, and electrons, per Table 15.4. Next is measured mass, then the difference. The difference is re-cast in MeV, representing the *total* binding energy of the nucleus, inexorably rising with the size of the nucleus. The final column divides by the mass number to get binding energy per nucleon, which peaks around iron. See Example 15.3.2 to understand how these numbers are computed.

Example 15.3.2 Following the entry in Table 15.5 for ^{56}Fe, we first multiply the individual proton, neutron, and electron masses from Table 15.4 by the 26 protons, 30 neutrons, and 26 electrons comprising ^{56}Fe to get a sum-of-parts value of 56.46340 a.m.u..[17]

The *actual* mass, as it appears for ^{56}Fe in the Chart of the Nuclides is 55.934942 a.m.u., which is smaller by 0.528447 a.m.u.[18]

Since 1 a.m.u. is 1.660539×10^{-27} kg, we can convert this mass difference into kilograms, then multiply by c^2, where $c = 2.99792458 \times 10^8$ m/s to get the associated energy in units of Joules. Traditionally, nuclear physics adopts a more convenient scale of electron-volts, and in particular, the MeV.[19] To get our mass-energy difference from Joules to MeV, we divide by 1.6022×10^{-13} J/MeV, and this is the 492 MeV number appearing in the Δmc^2 column of Table 15.5.

Finally, we divide by the number of nucleons in the nucleus—$A = 56$ in this case—to determine how much binding energy is present *per nucleon*—the significance of which will soon become clearer.

17: Find this in Table 15.5.

18: These numbers also appear in Table 15.5.

19: 1 MeV is 10^6 eV, and 1 eV is 1.6022×10^{-19} J (Sec. 5.9; p. 78).

Therefore, the difference between the sum-of-parts mass and actual nucleus mass in Table 15.5 provides a measure of how much binding energy holds the nucleus together.[20]

Notice that the first entry in Table 15.5 for the single-proton hydrogen atom has *no* binding energy in the nucleus: the lonely proton has no other nucleon to which it might bind. But deuterium (^{2}H) has a proton and a neutron, held together by 2.2 MeV of binding energy. The binding

20: . . . thus how much energy would need to be supplied to completely *unbind* the entire nucleus, as in Figure 15.9

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

energy per nucleon in the last column of Table 15.5 starts out small, but soon settles to the 7–9 range for most of the entries. It is extremely insightful to plot the binding energy per nucleon as a function of the nucleon mass number, A, which we do in Figure 15.10.

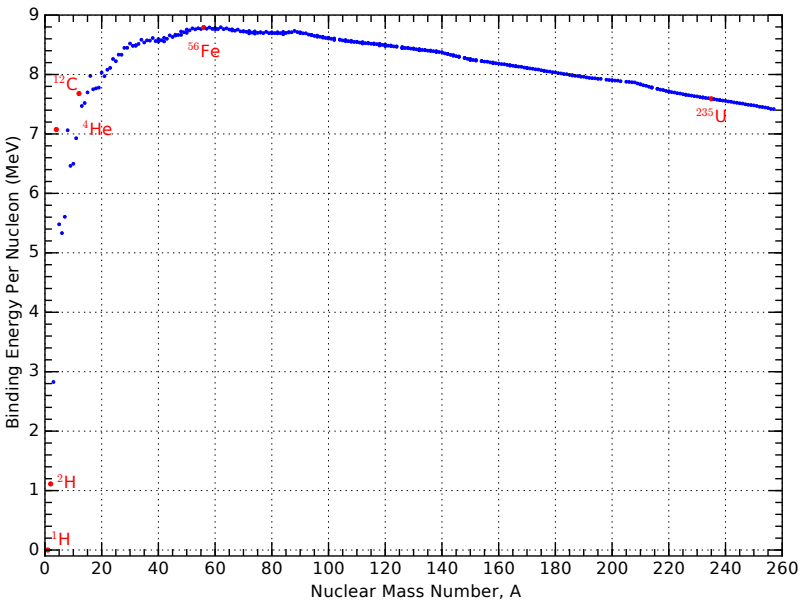


Figure 15.10: Binding energy per nucleon as a function of total mass number, A. The nuclei featured in Table 15.5 are indicated as red points. Note in particular that ^{56}Fe sits at the peak of the curve. Fusion operates from left to right, building larger nuclei, and fission goes from right to left, tearing apart nuclei. Only actions that *climb* this curve are energetically favorable, meaning that fusion is profitable on the the left-hand side, and fission makes sense on the right: each driving toward the peak binding energy per nucleon.

The value of Figure 15.10 is hard to over-emphasize. Key take-aways are:

1. Most nuclei are at around 8 MeV per nucleon, meaning that it would take an average of about 8 MeV of energy to rip out each member (proton or neutron) from a nucleus;
2. The peak is at ^{56}Fe,[21] meaning that this is the most tightly bound nucleus;[22]
3. The slope on the left side is *much* steeper than the slope on the right side, after the peak, which speaks to why fusion (building from small to big) is more potent than fission (tearing apart very massive nuclei);
4. Fusion in stars does not build elements beyond the peak around iron, since to go beyond the peak is not energetically favorable.

21: Actually, ^{62}Ni wins by a hair at 8.795 MeV/nuc, but is somewhat overlooked because it is only 0.006% as abundant as ^{56}Fe, whose binding energy per nucleon is essentially tied for the top at 8.790 MeV/nuc.

22: A peak exists because nucleons initially find advantage in binding together, but ultimately the increasing number of mutually-repelling protons makes the environment less appealing for larger nuclides.

It can be helpful to think of Figure 15.10 upside-down, as in Figure 15.11, turning the iron "peak" into a trough. A ball will roll toward and settle near the bottom of the trough, which is what both fusion and fission do, but from opposite directions.

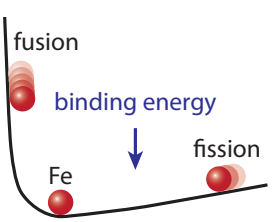


Figure 15.11: Turning the binding energy curve upside-down makes it easier to conceptualize fusion and fission driving toward the most tightly bound point (iron), like a ball might roll.

15.4 Fission

Having covered some fundamentals, we are ready to tackle aspects of nuclear energy. Really it is very simple. Enough nuclear material in a

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

small space will get hot, for reasons detailed below. The heat is used to boil water into high-pressure steam, which then turns a turbine and generator (Figure 15.12). Note that a nuclear fission plant has much in common with a coal-fired power plant, as evidenced by the similarity of Figure 15.12 to Fig. 6.2 (p. 90). Only the source of heat is much different in origin.

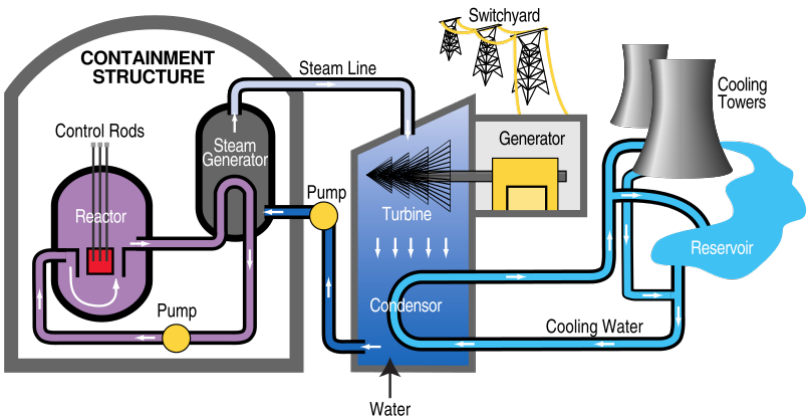


Figure 15.12: Typical nuclear power plant design, bearing much resemblance to the generic scheme from Figure 6.2. Details on the reactor core will follow in Section 15.4.4. Source: TVA.

15.4.1 The Basic Idea

Out of all the nuclides, three are amenable for use in a fission reactor. Two are isotopes of uranium: ^{233}U and ^{235}U; and one is plutonium: ^{239}Pu. Of these, only ^{235}U is found in nature, so we will concentrate on this one, returning later to the other two when we talk about breeder reactors in Section 15.4.4.2.

What makes ^{235}U (and the other two) special is that a slow[23] neutron—one just bumping around at a speed governed by the local temperature, and thus called a thermal neutron—can walk up to and stick[24] to the nucleus and cause it to split into two large chunks—depicted in Figure 15.13. Other nuclei would not break up, just accepting the new neutron and possibly converting a neutron to a proton via β^- decay.

23: This is in contrast to a fast neutron that tends to bounce rather than stick to the nucleus.

24: No forces prevent a neutron from approaching a nucleus. Happening to hit the tiny nucleus is the only barrier.

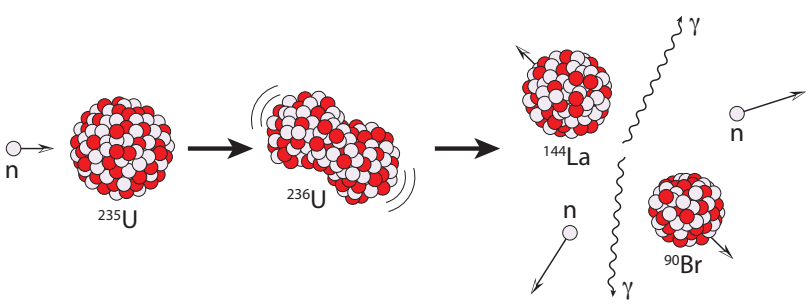


Figure 15.13: Fission schematic for ^{235}U, showing one of many possible outcomes—in this case ^{90}Br and ^{144}La plus two neutrons (an example case treated in detail in the text). The intermediate state, ^{236}U, created when ^{235}U absorbs a neutron, is highly unstable and will spontaneously break into (always) two different-size large fragments ("daughter" nuclei) and perhaps some extra neutrons. Gamma rays and kinetic energy (high-velocity fragments) are also released. Note that at each stage, the total number of nucleons is always 236.

When the nucleus breaks up, the pieces fly out at high speed, carrying kinetic energy that will be deposited in the local material as they bump

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

their way to a halt. Gamma rays[25] are also released. By catching all of this energetic output, the surrounding material gets very hot and can be used to make steam.

25: . . . very high energy photons

15.4.2 Chain Reaction

As we have seen, in order to get fission to happen, we need ^{235}U and some wandering neutrons. Once fission commences, the breakup of the nucleus usually "drips" a few spare neutrons, like crumbs left after cutting a piece of bread. The left-over neutrons provide a replenished source of neutrons ready to initiate more fission events. Now the door is open for a chain reaction, in which the neutrons produced by the fission events are the very things needed to stimulate additional fission events.

When the nucleus splits, any extra neutrons come out "hot" (high speed), which tend to bounce off uranium nuclei without sticking. They need to be slowed down, which is accomplished by a moderator: basically light atoms[26] that can receive the neutron impact as a sort of damping medium. Then the main trick is to prevent a runaway that could occur if *too many* neutrons become available; in which case it's a party that can get out of control. So nuclear plants employ control rods containing materials particularly effective at absorbing (trapping) neutrons. The colors of the lower halves of some squares in the Chart of the Nuclides (Figure 15.4) indicate neutron capture cross section. Boron (^{10}B) is a favorite choice to soak up neutrons and tame (or even halt) the reaction. The goal is to maintain a chain reaction that produces a net balance of **exactly one** unabsorbed slow neutron per fission event, available to attach itself to a waiting ^{235}U nucleus.

26: . . . usually either water or carbon in the form of graphite

15.4.3 Fission Accounting

The nucleus (uranium in the present discussion) always breaks up into two largish pieces, possibly accompanied by a few liberated spare neutrons. Because of the way the track of stable elements curves on the Chart of the Nuclides, the resultant pieces are likely to be neutron rich, to the right of the stable nuclei. To understand this, refer to Figure 15.14 and the associated caption.

The math always has to add up: nucleons are not created or destroyed during a fission event. They just rearrange themselves, so the total number of neutrons stays the same, as does the total number of protons. *After* the split, β^- decays will carry out flavor changes, but we'll deal with that part later.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

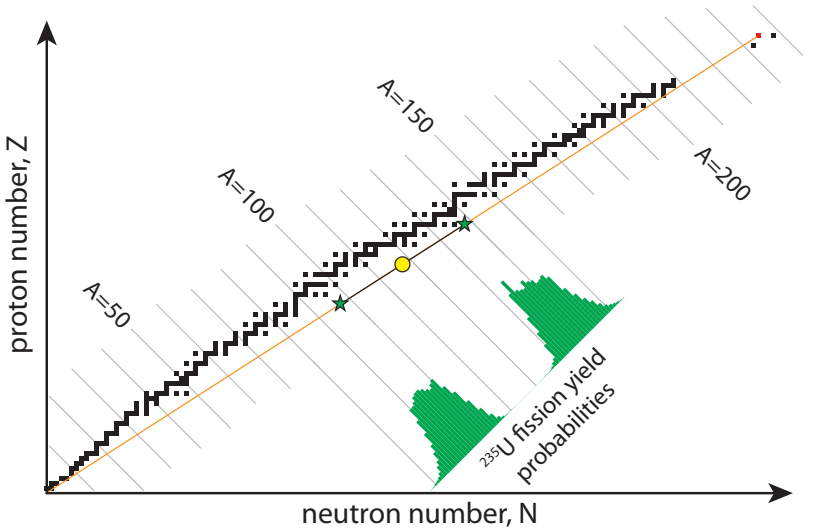


Figure 15.14: Fission of ^{235}U (small red square, upper right) tends to produce two neutron-rich fragments. If it split exactly in two, the result would lie at the midpoint of the orange line connecting ^{235}U to the origin, at the yellow circle. In practice, an equal split is highly unlikely, as one fragment tends to be around $A \sim 95$ and the other around $A \sim 140$, as depicted by the probability histogram in green. The two green stars separated along the orange line represent a more likely outcome for the two fragments. As long as the green stars are located so that the yellow circle is exactly between them, the accounting of proton and neutron number is satisfied. Because the orange line lies to the right of the stable nuclei, the fission products tend to be neutron-rich and undergo a series of radioactive β^- decays before reaching stability, which could take a very long time in some cases.

Example 15.4.1 If one of the two fragments from the fission of a ^{235}U nucleus ($Z = 92$) *after adding* a thermal neutron winds up being ^{90}Br ($Z = 35$), what is the other nucleus going to be?

The other fragment will preserve total proton count, so $Z = 92 - 35 = 57$, and as such is destined to be the element lanthanum. Which isotope of lanthanum is produced depends on how many neutrons escape the split. Table 15.6 summarizes the particle counts of the various players.

If no spare neutrons are left over, the lanthanum must have $N = 144 - 55 = 89$ neutrons,[27] in which case its mass number will be $A = 146$, so ^{146}La. If two neutrons are set free, then the lanthanum will only keep 87 neutrons and be ^{144}La, as depicted in Figure 15.13.

Typically, about 2–3 neutrons are left out of the final fragments, and can go on to promote additional fission events in the chain reaction.

27: ^{235}U had $A - Z = 235 - 92 = 143$ neutrons, plus the thermal neutron addition.

	^{235}U	^{90}Br	^{146}La	^{145}La	^{144}La	^{143}La
A	235	90	146	145	144	143
Z	92	35	57	57	57	57
N	143	55	89	88	87	86
n	1		0	1	2	3

Table 15.6: Possible outcomes for Example 15.4.1 if we set one of the daughter particles to be bromine-90, forcing the other daughter to be lanthanum. Different isotopes of lanthanum will result for differing numbers of spare neutrons left after the break-up (last row).

Being a probabilistic (random) process, each fission can result in a large set of possible "daughter" nuclei—only one set of which was explored in Example 15.4.1. As long as the masses all add up, and the two-hump probability distribution in Figure 15.14 is respected, anything goes. In other words, we have no control over exactly what pieces come out. Figure 15.15 provides a graphic illustration of four different possible pairs of daughter fragments. The counting requirement is satisfied by having the products located diametrically opposite from the ^{235}U midpoint (yellow circle). The positions of the stars will distribute along A-values

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

according to the probability distribution (multi-colored histogram). Note the completely distinct peaks, conveying that virtually *every* fission event results in just two fragments: one bigger and one smaller. At least that aspect of fission is predictable, even if we can't say precisely which nuclei will be left after an individual fission event.

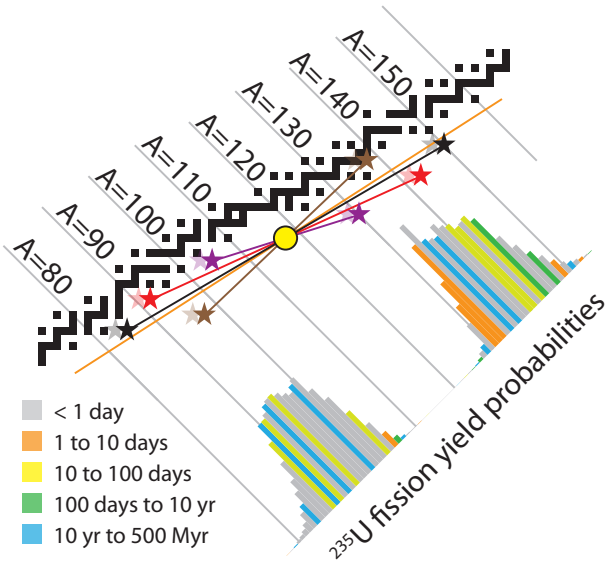


Figure 15.15: Various fission product outcomes are possible, indicated here by four sets of colored star pairs and connecting lines. The average position of each pair is the yellow circle (the stars are diametrically opposite the circle), which guarantees that the total number of neutrons and protons is unchanged from the parent nucleus to the daughter nuclei. To the extent that additional neutrons are left behind like crumbs, the stars will displace to the left of their indicated positions a bit, as hinted by the lighter-shaded "ghost" stars, whose offsets from the nominal star positions will also vary depending on how many neutrons are left out of the two final fragments. The coloring of the histogram indicates radioactive lifetime for the decay chain of a neutron-rich fragment at each mass number, matching the half-life color scheme used in Figure 15.4.

Let us now examine the energetics, using the result from Example 15.4.1, in which ^{235}U breaks into ^{90}Br and ^{144}La, plus two spare neutrons.[28] To be explicit, the reaction we will trace is

$$^{235}\text{U} + \text{n} \rightarrow {}^{90}\text{Br} + {}^{144}\text{La} + 2\text{n}. \tag{15.2}$$

28: . . . also matching the scenario in Figure 15.13 and the penultimate column of Table 15.6

Constituent/Stage	mass (a.m.u.)	mass (MeV/c^2)
^{235}U	235.04392	218,942.0
n	1.00866	939.6
input mass	236.05259	219,881.6
^{90}Br	89.93069	83,769.9
^{144}La	143.91955	134,060.2
2n	2.01733	1,879.1
output mass	235.86757	219,709.3
mass change	0.18502	172.3

Table 15.7: Mass details of Eq. 15.2, tracking before and after masses in both a.m.u. and MeV units. The input mass of around 236 a.m.u. is reduced by about 0.185 a.m.u., or 0.08%.

The masses of each piece, according to the Chart of the Nuclides, appear in Table 15.7. Again, we find that the mass sums don't equal: the final parts are lighter than the inputs. The fission managed to lose 0.185 a.m.u. of mass, corresponding to 172 MeV of energy (via $E = mc^2$; see Example 15.3.1). That's a 0.08% change in the mass, and converts to an energy density of roughly 17 *million* kcal/g, making the process over a million

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

times more energy–dense than our customary ~ 10 kcal/g chemical energy density. See Box 15.3 for an example of how to compute this.

Box 15.3: Nuclear Energy Density

The example corresponding to Table 15.7 is said to correspond to 17 million kcal/g, but how can we get here? The mass change of 0.185 a.m.u. corresponds to a mass in kilograms of 3.07×10^{-28} kg, according to the conversion that 1 a.m.u. is 1.6605×10^{-27} kg (Table 15.4). Multiply this by c^2 to get energy in Joules, yielding 2.76×10^{-11} J.[29] In terms of kcal, we divide by 4,184 J/kcal to find that this fission event yields 6.6×10^{-15} kcal.

We now just need to divide by how many grams of "fuel" we supplied, which is 236.05 a.m.u. (Table 15.7), equating to 3.92×10^{-25} kg, or 3.92×10^{-22} g. Now we divide 6.6×10^{-15} kcal by 3.92×10^{-22} g to get 16.8×10^{6} kcal/g. Blows a burrito out of the water.

29: This result, by the way, is the same as 172.3 MeV in Table 15.7 using the conversion that 1 MeV is 1.6022×10^{-13} J.

Example 15.4.2 Considering that the average American uses energy at a rate of 10,000 W, how much ^{235}U per year is needed to satisfy this demand for one individual?

Since we have just computed the energy density of ^{235}U to be 17×10^{6} kcal/g (Box 15.3), let's first put the total energy in units of Joules, multiplying 10^4 W by 3.155×10^7 seconds in a year and then dividing by 4,184 J/kcal to get kilocalories. The result is 75 million kcal, so that an American's annual energy needs could be met by 4.5 g[30] of ^{235}U. That translates to about a quarter of a cubic centimeter, or a small pebble, at the density of uranium. Pretty amazing!

30: 75 million kcal divided by 17 million kcal/g is 4.5 g.

We can take a graphical shortcut to all of Section 15.4.3, which hopefully will tie things together in an instructive way.

Example 15.4.3 Refer back to Figure 15.10 (and/or Table 15.5) to see that ^{235}U has a binding energy of about 7.6 MeV per nucleon. Where we end up, around $A \approx 95$ and $A \approx 140$, the binding energies per nucleon are around 8.7 and 8.4 MeV/nuc at these locations, respectively.

Multiplying the binding energy per nucleon by the number of nucleons provides a measure of *total* binding energy: in this case 1,790 MeV for ^{235}U, about 825 MeV for the daughter nucleus around $A \approx 95$, and 1,175 MeV for $A \approx 140$.[31] Adding the latter two, we find that the fission products have a total binding energy around 2,000 MeV, which is greater[32] than the ^{235}U binding energy by about 210 MeV—somewhat close to the 172 MeV computed for the particular example in Table 15.7.

31: 7.6×235; 8.7×95; and 8.4×140

32: Binding energy *reduces* mass, so larger binding energy means lighter overall mass.

The graphical method got us pretty close with little work, and hopefully led to a deeper understanding of what is going on. The rest of this

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

paragraph explains the discrepancy, but should be considered non-essential reading. The fission process typically results in a few spare neutrons. Each left-over (unbound) neutron deprives us of *at least* 8 MeV in unrealized binding potential,[33] *and* the subsequent β^- decays from the neutron-rich daughter nuclei to stable nuclei also release energy not accounted in Table 15.7. Both of these contribute to the shortfall in comparing 172 MeV to 210 MeV, but even without this, we got a decent estimate just using the graph in Figure 15.10.

33: Each missing neutron deprives us of more than the standard ~8 MeV per nucleon, as neutrons have no penalty for repulsive electric charge. The 8 MeV per nucleon is an average over protons and neutrons.

15.4.4 Practical Implementations

As we saw above, nuclear fission involves getting fissile nuclei—generally ^{235}U—to split apart by the addition of a neutron. The following criteria must be met:

1. presence of nuclear fuel (^{235}U);
2. presence of neutrons, provided as left-overs from earlier fission events;
3. a moderator to slow down neutrons that emerge from the fission events at high speed;
4. a high enough concentration of nuclear fuel that the slowed-down spare neutrons are likely to find fissile nuclei;
5. neutron absorbers in the form of control rods that can be lowered into the reactor and act as the main "throttle" to set reaction speed (thus power output), and also prevent a runaway chain reaction;
6. a containment vessel to mitigate radioactive particles (gamma rays, high-speed electrons and positrons) from escaping to the environment.

Figure 15.16 shows a typical configuration.

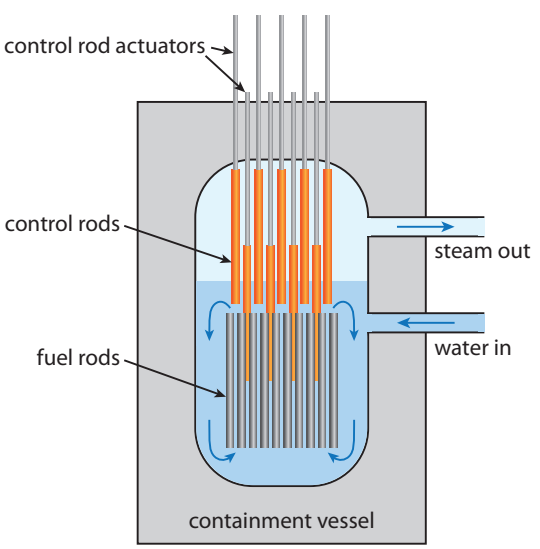


Figure 15.16: Typical boiling water reactor design. A thick-walled containment vessel holds water surrounding ^{235}U fuel rods. The water acts as the moderator to slow neutrons and also circulates around the rods to carry heat away, boiling to form steam that can run a standard power plant. Control rods set the pace of the reaction based on how far they are inserted into the spaces between fuel rods. Extra control rods are poised above the reactor core ready to drop quickly into the core in case of emergency—suddenly bringing the chain reaction to a halt.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

In the design of Figure 15.16, called a boiling water reactor, the water acts as both the neutron moderator and the thermal conveyance medium. Nuclear fuel (uranium) is arranged in fuel rods, providing ample surface area and allowing water to circulate between the rods to slow down neutrons and carry the heat away. Neutron-absorbing control rods—usually containing boron—set the reaction speed by lowering from the top.[34] An emergency set of control rods can be dropped into the core in a big hurry to shut down the reactor instantly if something goes wrong. When the emergency rods are in place, neutrons have little chance of finding a ^{235}U nucleus before being gobbled up by boron.

34: . . . always this direction, so that gravity does the pulling rather then relying on some other drive force

As of 2019, the world has about 455 operating nuclear reactors, amounting to an installed capacity of about 400 GW.[35] The average *produced* power—not all are running all the time—was just short of 300 GW. The thermal equivalent would be approximately three times this, or 1 TW out of the 18 TW we use in the world. So nuclear is a relevant player. See Table 15.8 for a breakdown of the top several countries, Fig. 7.7 (p. 109) for nuclear energy's trend in the world, and Fig. 7.4 (p. 107) for the U.S. trend.

35: From this, we glean that reactors average roughly 1 GW each.

Country	# Plants	GW inst.	GW avg.	% elec.	global share (%)
U.S.	95	97	92	20	31
France	56	61	44	71	15
China	49	47	38	5	13
Russia	38	28	22	20	8
Japan	33	32	8	8	3
S. Korea	24	23	16	26	5
India	22	6	5	3	2
World Total	455	393	295	11	100

Table 15.8: Global nuclear power in 2019 [101], listing number of operational plants, installed capacity, average generation for 2019 (Japan currently has stopped a number of its reactors), percentage of *electricity* (not total energy), and fraction of global production (these 7 countries accounting for over 75%). Notice the close match between number of plants and GW installed for most countries, indicating that most nuclear plants deliver about 1 GW.

Nuclear plants only last about 50–60 years, after which the material comprising the core becomes brittle from exposure to damaging radioactivity and must be decommissioned. The median age of reactors in the U.S. is 40 years, and all but three are over 30 years old. Additional challenges will be addressed in the sections that follow.

When nuclear energy was first being rolled out in the 1950s, the catch phrase was that it would be "too cheap to meter," a sentiment presumably fueled by the stupendous energy density of uranium, requiring very small quantities compared to fossil fuels. The reality has not worked out that way. Today, a 1 GW nuclear power plant may cost $9 billion to build [102]. That's $9 per Watt of output power, which we can compare to the cost of a solar panel, at about $0.50 per W (Fig. 13.16; p. 215), or utility-scale installation at $1 per Watt [89]. While it seems that solar[36] wins by a huge margin, the low capacity factor of solar reduces average power output to 10–20% of the peak rating, depending on location. Meanwhile, nuclear reactors tend to run steadily 90% of the time—the off-time used for maintenance and fuel loading. So nuclear fission costs about $10 per delivered Watt, while solar panels are $2.5–5 per delivered

[102]: Union of Concerned Scientists (2015), *The Cost of Nuclear Power*

36: Recall, for context, that solar is not among the *cheaper* energy resources. Like solar, nuclear power is dominated by up-front costs, rather than fuel cost.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

Watt and installed utility-scale systems are $5–10 per Watt. In short, nuclear power is not an economic slam dunk.

15.4.4.1 Uranium

So far, we have ignored a crucial fact. Only 0.72% of natural uranium on Earth is the fissile ^{235}U flavor. The vast majority, 99.2745%, is the benign ^{238}U.[37] The ratio is about 140:1, so for every ^{235}U atom pulled out of the ground, 140 times this number of uranium atoms must be extracted. The origin of the disparity is a story of astrophysics and eons, covered in Box 15.4.

37: A trace amount, 0.0055%, is in ^{234}U.

Box 15.4: Origin of Uranium

The Big Bang that formed the universe produced only the lightest nuclei. By-and-large, the result was 75% hydrogen (^{1}H) and 25% helium (^{4}He). Deuterium (^{2}H) and ^{3}He were produced at the 0.003% and 0.001% levels, respectively, and then the tiniest trace of lithium. No carbon or oxygen emerged, which must be "cooked up" via fusion in stars.

Fusion in stars does not "climb over" the peak of the binding-energy curve in Figure 15.10, so stops in the vicinity[38] of iron. From where, then, did all of the heavier elements on the periodic table derive? Exploding stars called supernovae and merging neutron stars appear to be the origin of elements beyond zinc.

The relative abundance of ^{235}U and ^{238}U on Earth can be explained by their different half-lives of 0.704 Gyr and 4.47 Gyr, respectively. Even if starting at comparable amounts, most of the ^{235}U will have decayed away by now. Solving backwards[39] to when they would have been present in equal amounts yields about 6 Gyr, which is older than the age of the solar system (4.5 Gyr) and younger than the universe (13.8 Gyr). This is a reasonable result for how old the astrophysical origin might be—allowing a billion years or so for the material to coalesce in our forming solar system.

38: Iron has $Z = 26$; stars tend not to produce elements beyond zinc ($Z = 30$) by fusion.

39: This follows almost the exact same logic and process as carbon-14 radioactive dating, but using much longer half life nuclei to date Earth's building blocks!

Uranium is not particularly abundant. Table 15.9 provides a sense of how prevalent various elements are in the earth's crust. Uranium is more abundant than silver, but the *useful* ^{235}U isotope is four times rarer than silver, and only about 5 times as abundant as gold. Proven reserves of uranium [103] amount to 7.6 million (metric) tons available, and we have used 2.8 million metric tons to date. The implication is that we could continue about 3 times longer than we have gone so far on proven reserves. But nuclear energy has played a much smaller role than fossil fuels, so maybe this isn't so much.

[103]: (2020), *List of Countries by Uranium Reserves*

Evaluating the uranium reserves in energy terms is the most revealing

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*; *Freely available at*: *https://escholarship.org/uc/energy_ambitions*.

Element	Abund.	Element	Abund.	Element	Abund.
silicon	282,000	carbon	200	thorium	9.6
aluminum	82,300	copper	60	uranium	2.7
iron	56,300	lithium	20	silver	0.075
calcium	41,500	lead	14	^{235}U	0.02
titanium	5,650	boron	10	gold	0.004

Table 15.9: Example material abundances in the earth's crust, in parts per million by mass.

approach. First, we take 0.72% of the 7.6 million tons available to represent the portion of uranium in the form of ^{235}U. Enrichment (next section) will not separate *all* of the ^{235}U, and the reactor can't burn all of it away before the fuel rod is essentially useless. So optimistically, we burn half of the mined ^{235}U in the reactor. Multiplying the resulting 27,300 tons of *usable* ^{235}U by the 17 million kcal/g we derived earlier yields a total of 2×10^{21} J. Table 15.10 puts this in context against fossil fuel proven reserves from page 127. We see from this that proven uranium reserves give us only 20% as much energy as our proven oil reserves, and about 5% of our total remaining fossil fuel supply. If we tried to get all 18 TW from this uranium supply, it would last less than 4 years! This does not sound like a salvation.

Table 15.10: Proven reserves, in energy terms.

Fuel	10^{21} J
Coal	20
Oil	10
Gas	8
^{235}U	2

Proven uranium reserves would last 90 years at the *current* rate of use, so really it is in a category fairly similar to that of fossil fuels in terms of finite supply. To be fair, proven reserves are always a conservative lower limit on estimated total resource availability. And since fuel cost is not the limiting factor for nuclear plants, higher uranium prices can make more available, from more difficult deposits. Still, even a factor of two more does not transform the story into one of an ample, worry-free resource.

15.4.4.2 Breeder Reactors

In its native form, ^{235}U is too dilute in natural uranium—overwhelmingly dominated by ^{238}U—to even work in a nuclear reactor. It must be enriched to 3–5% concentration to become viable.[40] Enrichment is difficult to achieve. Chemically, ^{235}U and ^{238}U behave identically. The masses are so close—just 1% different—that mechanical processes have a difficult time differentiating. Centrifuges are commonly used to allow heavier ^{238}U to sink faster[41] than ^{235}U. But it's inefficient and usually requires many iterations to work up higher concentrations. The process is also lossy, in that not all of the ^{235}U finds its way to the enriched pile.[42]

But what if we could use the *bulk* uranium, ^{238}U, in reactors and not only save ourselves the hassle of enrichment, but also gain access to 140 times more material, in effect? Doing so would turn the proven reserves of uranium into about 7 times more energy supply than all of our remaining fossil fuels. Well, it turns out that despite its not being

40: Uranium bombs need at least 20% ^{235}U concentration, but typically aim for 85% to be considered *weapons grade*.

41: . . . in gaseous form

42: Depleted uranium is defined as containing 0.3% or less in the form of ^{235}U, which is not a huge reduction from the 0.72% starting point.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

one of the three fissile nuclei, we can *convert*[43] ^{238}U into the fissile ^{239}Pu the following way.

1. A ^{238}U may absorb a wandering neutron to become ^{239}U.
2. ^{239}U, whose half life is 23.5 minutes, undergoes β^- to become ^{239}Np in short order.
3. ^{239}Np also undergoes β^- with a half life of 2.4 days to become fissile ^{239}Pu.

43: . . . called transmutation

Figure 15.17 highlights this process in a simplified region of the Chart of the Nuclides, while Figure 15.18 shows complete details for the entire region around the fissile materials—the ones with red isotope names—which can be used to track the sequence outlined above.

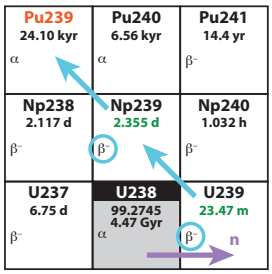


Figure 15.17: Breeder route to ^{239}Pu.

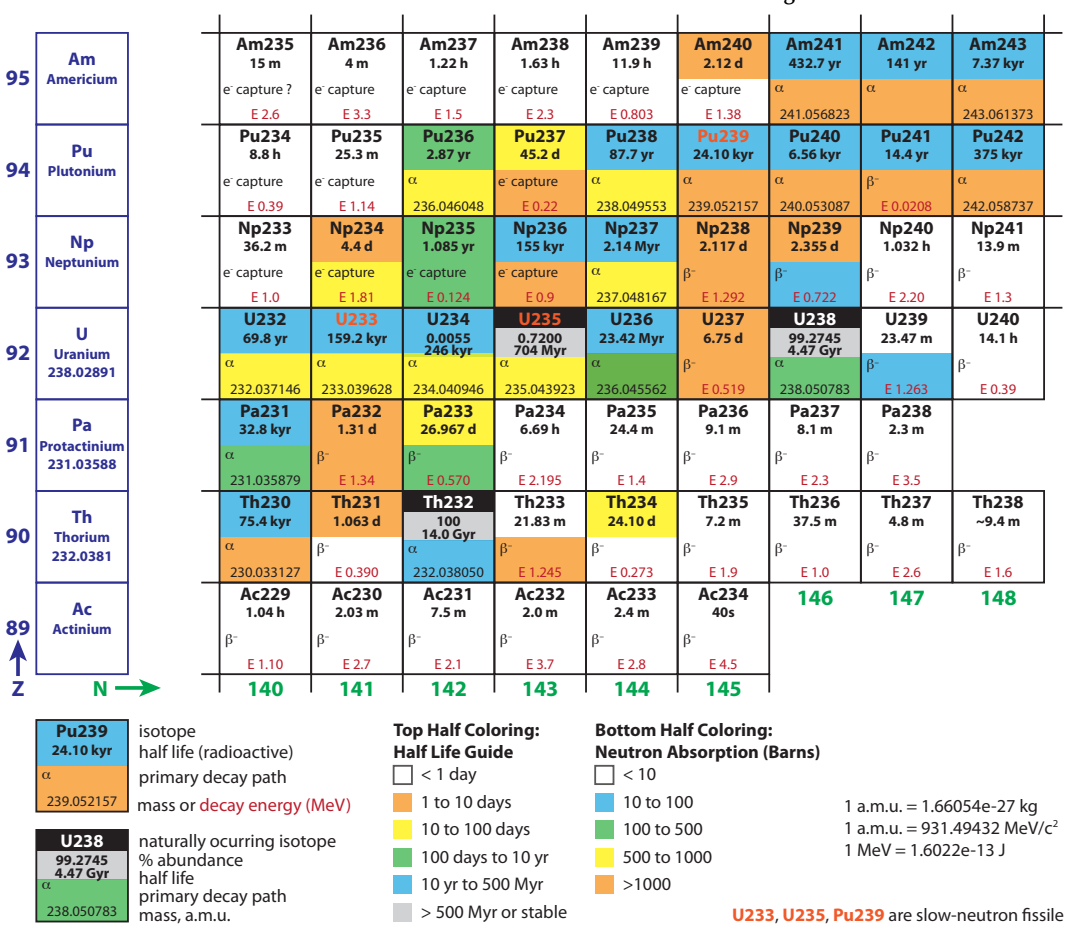


Figure 15.18: Chart of the Nuclides in the fission region. See also Figure 15.4 for the lower-left corner.

The result is that sterile ^{238}U can be turned into fissile ^{239}Pu that can be used in fission reactors. This process of transmuting an inert nucleus into a fissile one is called breeding, and is how we get any plutonium at all.[44] A nuclear reactor is a great place to introduce ^{238}U to neutrons: both are already in attendance. In fact, breeding happens as a matter of course in a nuclear reactor: it is estimated that one-third of the fission

44: . . . e.g., for weapons

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

energy in ordinary nuclear reactors comes from plutonium breeding and subsequent fissioning—without any extra effort. Special reactor designs enhance plutonium production, allowing the fuel rod to be "harvested" for plutonium. Usually, the plutonium is destined for use in weapons, but in principle reactors could be designed to efficiently produce and use plutonium from the ^{238}U feedstock. Downsides will be addressed in Section 15.4.6 on weapons and proliferation.

Box 15.5: Thorium Breeding

Another form of breeding merits mention. Notice that thorium[45] is more abundant than uranium in Table 15.9. But like ^{238}U, it is not fissile. However, applying the breeding trick, the absorption of a neutron by ^{232}Th ends up as ^{233}U—the last of our three fissile nuclei—in about a month's time. This provides an avenue for an even *greater* energy store than exists in ^{238}U via breeding to ^{239}Pu, by virtue of greater abundance. Unlike the plutonium route, thorium breeders are less susceptible to weapons and proliferation concerns.[46] That said, thorium reactors are more complex than uranium reactors, so that technical hurdles have thus far prevented any commercial scale application of the technique, leaving us unclear whether thorium represents a viable nuclear path.

45: ... of which 100% is the desired ^{232}Th isotope

46: ... although, radioactive waste is still problematic

15.4.5 Nuclear Waste

As we saw in our description of the fission process, the fragments distribute over a range of masses in a randomized way (Figure 15.15). The results are generally neutron-rich, and will migrate toward stable elements via β^- decays over the ensuing seconds, hours, days, months, and years. Some will go fast, and some will take ages to settle, depending on half-lives. Radioactive waste is dangerous to be around because the high-energy particles (like sub-atomic "bullets") spewing out in all directions can alter DNA, leading to cancer and birth defects, for instance.

The lighter of the two fission fragments has a 59% chance of landing on a stable nucleus within a day or so. For the heavier fragment, it's a 45% chance. The rest get hung up on some longer half-life nuclide, and could remain radioactive for a matter of weeks or in some cases millions of years. The colors in the fission probability histograms in Figure 15.15 provide a visual guide for the mass numbers that reach stability promptly (gray) vs. those that get hung up for a long time (blue is more than 10 years). For example, the histogram element at $A = 90$ is blue because ^{90}Sr—discussed below—stands in the way of a fast path to stability.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

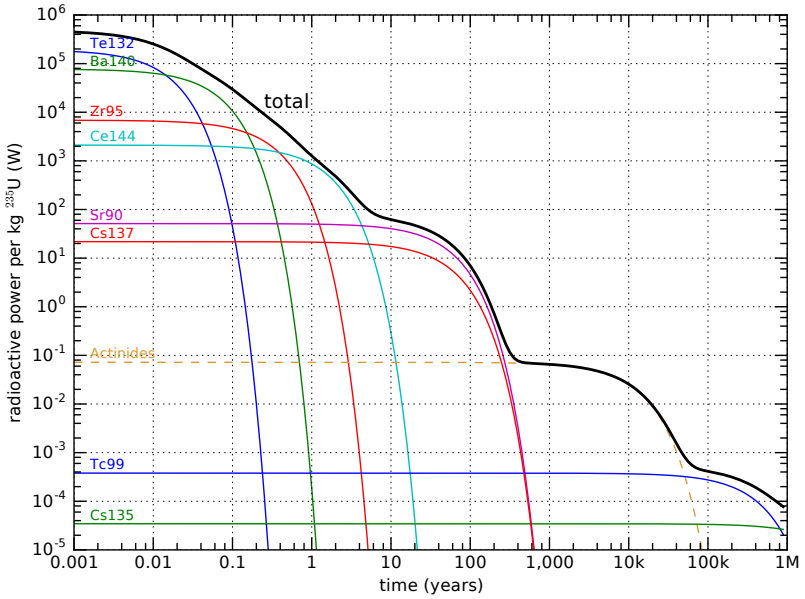


Figure 15.19: Decay activity of fragments from 1 kg of fissioned ^{235}U over time, on a log–log plot. The vertical axis is the power of radioactive emission, in W, for a variety of relevant isotopes—each having their own characteristic half life. The black line at the top is the total activity (sum of all contributions), and some of the key individuals are separated out. The dashed line for actinides is an approximate representative indicator of the role played by heavy nuclides formed in the reactor by uranium absorption of neutrons. Minor tick marks are at multipliers of 2, 4, 6, and 8 for each axis. As a matter of possible interest, the exponential decays of each element on this log–log plot have the functional form of exponential curves drawn upside-down.

Figure 15.19 shows how the fission decays play out over time. For the first month or so out of the reactor, the spent fuel is really "hot" radioactively, but falls quickly as ^{95}Zr and then ^{144}Ce dominate around one year out. At about 5 years, the pair of ^{90}Sr and ^{137}Cs begin to dominate the output for the next few-hundred years. Some of the products survive for millions of years, albeit at low levels of radioactive power. In addition to the daughter fragments, uranium in the presence of neutrons transmutes into neptunium, plutonium, americium, and curium via neutron absorption and subsequent β^- decays, represented approximately and collectively in Figure 15.19 by a dashed curve labeled Actinides.[47]

47: Breeder reactors can "burn" the actinides, reducing some of the long-term waste threat, but will unavoidably still be left with all the radioactive fission products.

The bottom line is that fission leaves a trash heap of radioactive waste that remains at problematic levels for many thousands of years. When nuclear reactors were first built, they were provisioned with holding tanks—deep pools of water—in which to place the waste fuel until a more permanent arrangement could be sorted out (Figure 15.20). We are still waiting for an adequate permanent solution for waste storage, and the "temporary" pools are just accumulating spent fuel. Transporting the spent fuel is hazardous—in part because it could fall into the wrong hands and be used to make "dirty" bombs—and no one wants a nuclear waste facility in their backyard, making the problem politically thorny. On the technical side, it is difficult to identify sites that are geologically stable enough and have little chance of groundwater contamination. Underground salt domes offer an interesting possibility, but political challenges remain daunting.



Figure 15.20: A spent fuel rod being lowered into a storage grid in a pool of water at a nuclear power plant. Source: U.S. DoE.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

15.4.6 Nuclear Weapons and Proliferation

Nuclear bombs are the most destructive weapons we have managed to create. The first bombs from the 1940s were based on either highly enriched ^{235}U or on ^{239}Pu. For uranium bombs, the idea is shockingly simple. Two separate lumps of the bomb material are held apart until detonation is desired, at which point they are slammed together.[48] It's not the collision that creates the explosion, but a runaway process based on having a high concentration of fissile material and no neutron absorbers present to control the resulting chain reaction. The concept is critical mass. The combined lump exceeds the critical mass, and explodes.[49]

48: For plutonium, this process is fouled by the presence of ^{240}Pu, forcing a different approach in which a sphere below critical mass is imploded to create high density.

49: Never stack lumps of fissile material together on a shelf, or a nasty surprise may be in store.

As simple as nuclear weapons are to build, the bottleneck becomes obtaining fissile material. Plutonium does not exist in nature, since its 24,100 yr half-life means nothing is left over from the astrophysical processes that gave us uranium and thorium (Box 15.4). We only still have the latter two thanks to their long half lives. So fissile material has to start with uranium. But as we have seen, natural uranium is only 0.72% fissile (^{235}U). In order to be explosive, the uranium must be enriched to at least 20% ^{235}U, and generally much higher (85%). Reactor fuel, at 3–5% ^{235}U will experience meltdown if the critical mass is exceeded, but will not explode. Enrichment is technically difficult, and attempts to acquire and enrich uranium are monitored closely. Often we hear of countries pursuing uranium enrichment, claiming that they are only interested in domestic energy production—a peaceful purpose. And it is true that the first step in nuclear power generation is also enrichment. So it is very difficult to ascertain true intentions. Once a country has the ability to enrich uranium enough for a nuclear plant, they can in principle keep the process running longer to arrive at weapons-grade ^{235}U.

While we worry about ^{235}U falling into the wrong hands, perhaps more disturbing is ^{239}Pu. Having a much shorter half-life than ^{235}U (24 kyr vs. 704 Myr), it is more dangerous to handle.[50] But plutonium is otherwise easy to deal with, since it requires no enrichment and can be chemically separated to achieve purity. It is the material of choice for nuclear weapons.

50: . . . much higher rate of radioactive decay

Serious pursuit of breeder reactors effectively means manufacturing lots of plutonium, leading to proliferation of nuclear materials: it becomes harder to track and keep away from mal-intentioned groups. The world becomes more dangerous under a breeder program. Thorium breeding (Box 15.5) is less risky in this regard because the ^{233}U prize is mixed with a ridiculously dangerous ^{232}U isotope that puts plutonium to shame, so working with it is pretty deadly, which may deter would-be pursuit of this material by rogue groups.

A related concern involves proliferation of the abundant radioactive waste from fission plants, which could be mixed into conventional explosives[51] to radioactively contaminate a city or local region—poisoning

51: . . . called a "dirty bomb"

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

water, food, and air. In short, nuclear fission carries many perils on a number of fronts.

15.4.7 Nuclear Safety

A properly operating nuclear facility actually emits less radioactivity than does a traditional coal-fired power plant! As is true for many materials mined from the ground, coal contains some small amount of radioactive elements found in the earth's crust: principally thorium, uranium, and potassium. Lacking any shielding or protection, the exhaust from a coal plant distributes these products into the atmosphere. Nuclear plants, by contrast, have no exhaust,[52] and carefully control the exposure to radioactivity.

However, things can go wrong. The U.S. had a scare in 1979 when a six-month-old nuclear plant at Three Mile Island in Pennsylvania (Figure 15.21) suffered a loss-of-cooling incident that resulted in severe damage to (meltdown of) the core. But the containment vessel held and no significant radioactivity was released to the environment. Workers at the plant received a dose equivalent to an extra 100 days of natural[53] exposure. So we dodged a bullet.

Chernobyl was not so lucky in April 1986 when an ill-conceived test went sideways and resulted in an actual explosion of the core. This scenario was previously thought to be impossible, but it was a steam explosion, not a nuclear blast—so more like a "dirty bomb" that scattered radioactive material across the region. Thirty-one people died in the immediate aftermath, and about 200 people got acute radiation sickness. It is estimated that in the long term, 25,000 to 50,000 additional cancer cases will result, but this number is controversial and it is hard to tease Chernobyl-caused cancer/deaths apart from the much larger number of background cancer cases. The town of Chernobyl is still abandoned and only recently has begun to allow strictly limited incursions.

The most recent major accident was the Fukushima Daiichi plant in Japan following the Sendai earthquake in March 2011, resulting in the evacuation of 200,000 people and agricultural loss. The earthquake caused the three operating reactors to shut down (safely), while diesel-fueled generators ran to power pumps maintaining cooling flow over the hot fuel rods. The core of a reactor is still very hot after fission stops and continues to generate heat as daughter nuclei decay, so cooling flow must be maintained or the core can melt. The ensuing tsunami[54] ruined the plan to keep the cores cool, as the generator rooms flooded, causing the cooling flow to fail. The cores of all three reactors melted down and hydrogen gas explosions created a major release of radioactivity. Perhaps in contrast to the Chernobyl plant, Fukushima was designed by General Electric and operated by a well-educated high-tech society. No one is exempt from risk when it comes to nuclear reactors.

52: Note that cooling towers often have a plume of water vapor above them, but this is the result of evaporative cooling, and not exhaust in the usual sense.

53: We are unavoidably exposed to radiation in our daily lives from air, water, food, Earth, and the cosmos.



Figure 15.21: Three Mile Island nuclear plant in Pennsylvania. The two reactor cores are in the foreground of the larger cooling towers behind. Source: U.S. DoE.

54: . . . within 10 minutes of the earthquake

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

15.4.8 Pros and Cons of Fission

Collecting the advantages and disadvantages of fission, we start with the positive aspects:

- Nuclear fuel has extraordinary energy density, about a million times better than chemical energy density;
- Nuclear fission is proven technology providing a substantial fraction of electrical energy at present;
- Life-cycle CO_2 emissions for nuclear fission is only 2% that of traditional fossil fuel electricity [68];
- Breeder reactors could provide thousands of years of fuel, by way of uranium and thorium (undeveloped as yet).

[68]: (2020), *Life Cycle GHG Emissions*

And for the downsides:

- Radioactive waste is dangerous for thousands of years, and no clear solution to its disposal or long-term storage has emerged.
- Conventional uranium fission has limited fuel supply,[55] measuring in decades;
- Breeder reactors exacerbate the waste issue and promote proliferation of nuclear materials;
- Development of nuclear energy technology prepares an easy step to immensely destructive nuclear weapons;
- Accidents happen even to the best-managed reactors, the consequences often being severe for a region.

55: . . . in the absence of breeder reactor implementation

Nuclear fission is a complex topic that has compelling advantages and worrisome faults. Not surprisingly, attitudes are highly mixed. One survey [104] indicates that adults in the U.S. oppose building more nuclear plant by a slim 51% to 45%, while scientists overall favor advancing nuclear plants by a 2:1 margin,[56] and physicists surveyed favored nuclear by 4:1. Scientists are much more likely to view climate change as a serious threat than the U.S. population as a whole, and therefore are likely to be attracted to energy resources that do not emit CO_2. Of the physicists surveyed, it would be a mistake to assume that even the majority know the topic as thoroughly as it is covered in this chapter—given the degree of specialization within the field. Among those who understand the topic thoroughly[57] it is almost certain you'd find a healthy split: those for whom the perils outweigh advantages, and those who are concerned enough about climate change to accept the "lesser of two evils," and/or who are enthusiastic about the technology as a glowing example of our mastery over nature's hidden secrets.

[104]: Pew Research (2015), "Elaborating on the Views of AAAS Scientists, Issue by Issue"

56: . . . a clear, but not overwhelming, result

57: We might also acknowledge an intrinsic psychological appeal for complex topics that have been mastered: a sort of pride in the privileged comprehension that might transfer to warm feelings for the subject.

15.5 Fusion

Given that fission has problems of finite uranium supply, radioactive waste, proliferation and weapons, and safety issues, its future is uncertain.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

Fusion, on the other hand, is not plagued by most of these issues. It's main problem is that it is incredibly difficult and has been in the research stage for 70 years. Other than that, it has many (virtual) virtues. To be clear, the world does not have and never has had an operational fusion power plant. It *may* belong to the future, but is not guaranteed to ever become practical.

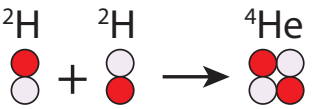


Figure 15.22: Fusion concept: helium from deuterium.

First, the basics. We have alluded to the fact that fusion builds from the small to the big. Putting four ^{1}H nuclei together, at 1.007825 a.m.u. each and forming ^{4}He at 4.0026033 a.m.u. leaves a difference of 0.0287 a.m.u.—0.7% of the total mass—which amounts to 153 million kcal/g.[58] This is *almost ten times* as large as the amount for fission (17 million kcal/g; Box 15.3), making it ten-million times more potent than chemical reactions. Recall that fusion's better performance can be related to the steepness of the left-hand-side of the binding-energy-per-nucleon curve of Figure 15.10.

What makes fusion so difficult is that getting protons to stick together is incredibly hard. Their electric repulsion is so strong that they need to be approaching each other at a significant fraction of the speed of light (about 7%) in order to get within reach of the strong nuclear force that takes over at distances smaller than about 10^{-15} m. The corresponding temperature is a billion degrees.[59] Even the center of the sun is "only" 16 million degrees. The sun has the advantage of being enormous, though. So even at a comparatively chilly 16 million degrees, some rare protons by chance will be going extra fast and have enough oomph to overcome the repulsion and stick together. It's like winning the lottery against very long odds, but the sun is large enough to buy ample tickets so the process still happens often enough.[60] We don't have such a luxury in a terrestrial laboratory setting, so we need higher temperatures than what exists in the center of the sun!

Using ^{2}H nuclei (deuterons, labeled D) instead of ^{1}H (protons) in what is called a D–D fusion reactor, allows operation at 100 million degrees instead of 1 billion. And colliding one deuteron with a triton[61] (^{3}H nucleus, labeled T; 12.3 year half-life), *only* requires 45 million degrees for a D–T fusion reactor. For this reason, only D–T fusion is currently pursued.

For all three types, the relevant reactions[62] are:

$$\begin{aligned} \text{p}-\text{p}: &\quad {}^1\text{H} + {}^1\text{H} + {}^1\text{H} + {}^1\text{H} \rightarrow {}^4\text{He} + 26.7\,\text{MeV} \\ \text{D}-\text{D}: &\quad {}^2\text{H} + {}^2\text{H} \rightarrow {}^4\text{He} + 23.8\,\text{MeV} \\ \text{D}-\text{T}: &\quad {}^2\text{H} + {}^3\text{H} \rightarrow {}^4\text{He} + \text{n} + 17.6\,\text{MeV} \end{aligned} \qquad (15.3)$$

But the 45 million degrees required for D–T fusion is still frightfully hard to achieve. No containers will withstand temperatures beyond a few thousand degrees. Containment—or confinement—is the big challenge then. The multi-million degree plasma[63] cannot be permitted to touch the

58: The calculation is that 0.0287 a.m.u. corresponds to $\Delta m = 4.8 \times 10^{-29}$ kg, or $E = \Delta mc^2 = 4.2 \times 10^{-12}$ J (26.7 MeV). We convert the Joules to kcal by dividing by 4,184, and then divide by the input mass in grams (4.03 a.m.u. times 1.6605×10^{-24} g/a.m.u.) to get 153 kcal/g. Starting with two deuterium nuclei reduces energy yield a bit to to 137 kcal/g, and for deuterium-tritium reactions it's down to 81 million kcal/g.

59: For temperatures this high, it does not matter whether we specify Kelvin or Celsius, as the 273 degree difference is nothing compared to a billion degrees. The scales are therefore essentially identical here.

60: This is no accident: if the center were too cool, the sun would contract in the absence of radiation pressure until the center heated up from the compression and nuclear fusion ignited—just enough to hold off further contraction. It finds its own equilibrium right at the edge of fusion. In the case of the sun, all it takes is one out of every 10^{26} collisions to stick in order to keep the lights on.

61: If only the UCSD mascot were named after *this* triton...

62: ...allowing beta decays to change protons to neutrons in the process

63: Plasma is a hot ionized gas where electrons are stripped off the nuclei. The sun qualifies as a plasma.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

walls of the chamber, despite its constituents zipping around at speeds around 1,000 km/s! This feat can be sort-of managed via magnetic fields bending the paths of the fast-moving charged particles into circles. But turbulence in the plasma plagues attempts to confine the D–T mixture at temperatures high enough to produce fusion yield.

Box 15.6: Successful Fusion

Note that besides stars as an example of successful fusion, we *have* managed to create artificial fusion in a net-energy-positive manner in the form of the **hydrogen bomb**. This is indeed a fusion device, but we could not call it *controlled* fusion. It actually takes a fission bomb (plutonium) right next to the D–T mixture in a hydrogen bomb to heat up the D–T enough to undergo fusion. It's neat (and awful) that it works and is demonstrated, but it's no way to run a power plant.

If a 45 million degree plasma could be confined in a stable fashion, the heat generated by the reactions[64] could be used to make steam and run a traditional power plant—replacing the flame symbol in Fig. 6.2 (p. 90) with something much fancier. The scheme, therefore, requires first heating a plasma to unbelievable temperatures in order for the plasma to self-generate enough *additional* heat through fusion that the game shifts to one of keeping the plasma *cool* enough to produce a steady rate of fusion without blowing itself out. In this scenario, the heat extracted from the cooling flow makes steam. It's the most elaborate[65] possible source of heat to boil water. It may be a bit like working hard to develop a light saber whose only use will be as a letter opener.

64: . . . in the form of radioactive release back to the plasma

65: Should we be proud if we succeed, or embarrassed at the lengths we had to go to?

15.5.1 Fuel Abundance

Deuterium—an isotope of hydrogen—is found in 0.0115% of hydrogen,[66] which means that the occasional H_2O molecule is actually HDO.[67] Therefore sea water is chock-full of deuterium. The global 18 TW appetite would need 3×10^{32} deuterium atoms per year for D–D or 2×10^{32} each of deuterium and tritium atoms per year for D–T. Running with this latter number for the comparatively easier D–T reaction, we would need to process 9×10^{35} water molecules each year to find the requisite deuterium. This corresponds to 26 million tons of water, which is a cubic volume about 300 m on a side. Yes, that's large, but the ocean is larger. Also, it corresponds to a volume of 0.16 billion barrels per year, which is about 200 times smaller than our annual oil consumption. Thus, the volume required should be not at all challenging.[68] The ocean volume is 60 billion times larger than our 300-m-sided cube, implying that we have enough deuterium for 60 billion years. The sun will not live that long, so let's say that we have sufficient deuterium on Earth.

66: See the Chart of the Nuclides abundance information in Figure 15.4.

67: . . . one ^{1}H, one ^{2}H and one oxygen

68: Ocean water is *far* easier to access than underground oil deposits, after all.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

Tritium, however, is essentially nowhere to be found, as it has a half-life of 12.3 years. We can generate tritium by adding a neutron to lithium and stimulating an α decay. So the question moves to how much lithium we have. Proven reserves are at about 15 million tons, currently produced at about 30,000 tons per year.[69] We would need 2,300 tons[70] of lithium per year to meet our 2×10^{32} tritium atom target (for 18 TW). In the absence of competition[71] for lithium resources, the associated R/P ratio timescale is 6,500 years. Yes, that is a comfortably long time, but not eons. The thought is that this would buy time to solve the D–D challenge.

69: Most lithium is used in batteries; the R/P ratio in this case is 500 years.

70: . . . only 8% of current annual production

71: Otherwise, we're still looking at the 500 year R/P ratio.

15.5.2 Fusion Realities

It is clear why people get excited by fusion. It seems like an unlimited supply that can last thousands if not billions of years at today's rate of energy demand. For some perspective, think about what else we know that lasts billions of years. We already have a giant fusion reactor parked 150 million kilometers away that requires no mining, servicing, or any attention whatsoever. In this sense, the sun is essentially as inexhaustible as fusion promises to be, but already working and free of charge. Photovoltaic panels plus batteries work *today* and have already shown a possible path to eternal energy. The author built his own off-grid solar setup on a budget that's tiny compared to the fusion enterprise.

As for the fusion enterprise, an effort called ITER (Figure 15.23) in southern France is an international effort currently constructing a plasma confinement machine that aims to commence experimental D–T fusion by the year 2035 via occasional 8-minute pulses of 0.5 GW thermal power. This machine is a stepping stone that is not designed to produce electricity. Estimates for construction cost range from $22 billion to $65 billion. By comparison, a nuclear fission plant costs $6–9 billion to build. Admittedly, the first experimental facility is going to cost more, but it is hard to imagine fusion ever being a real steal, financially. Even if the fuel is free, *so what*? Solar is the same.

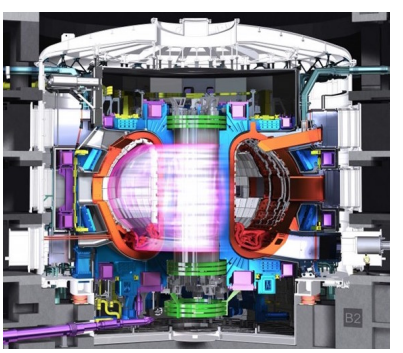


Figure 15.23: ITER tokamak cut-away where the plasma would be created. The white outer chamber is the size of a six-story building. From the ITER Organization.

An effort in the U.S. called the nuclear ignition facility (NIF) is pursuing a different approach to fusion research: attempting to implode a tiny sphere of D–T mixture by blasting it with 192 converging laser beams, crushing it to enormous pressure exceeding that in a star's interior, leading to an explosive release of heat. The building, mostly taken up by gigantic lasers, is the size of three football fields and has so far cost something to the tune of $10 billion. Again, this experimental facility is not provisioned to harness any net energy gain[72] to create electricity.

72: . . . the prospects for which are dubious

Let's say that by the year 2050, we will have mastered the art and can build a 1 GW electrical-output[73] fusion plant for $15 billion. That's $15 per Watt of output, which we can compare to a present-day solar utility-scale installation cost of $1 per peak Watt [89]. Applying typical capacity factors[74] puts fusion at twice what solar costs *already*, today.

73: . . . thus ~3 GW thermal, given typical heat engine efficiency

74: . . . 10–20% for PV and perhaps 90% for fusion?

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

Fusion is therefore a complicated and not particularly cheap way to generate electricity. Meanwhile, we are not running terribly short on renewable ways to produce electricity: solar; wind; hydroelectric; geothermal; tidal. Liquid fuels for transportation represent a greater and more pressing challenge, and fusion does not directly address this aspect any better than other options for electrical production. Fusion is by far the most complex power generation scheme we have ever attempted, evidenced by the 70 year effort to bring it to fruition that is still underway. How many physics PhDs will it take to keep a fusion plant running? Sometimes, we get stuck pursuing a flawed vision of the future, and have trouble reevaluating our options. Imagine being a middle-aged physicist or engineer in the 1950s. In your lifetime, you would have seen the advent of the car, airplane, radio, television, nuclear fission, among a blur of other technology advances. The next frontier was obviously fusion, so let's crack that one! At this point, 70 years later, maybe we should ask: why?

And let's point out that fusion is not without its waste challenges. It is still a radioactive environment, albeit not one that produces dangerous direct products (^{4}He is okay!). It *does* involve a radioactive fuel source (tritium), and it *does* embed the containment vessel with high energy particles and neutrons that over time compromise the integrity of the vessel so that it must be discarded as a radioactively-charged hunk of metal.[75] By comparison, solar, wind, and other renewable sources based on the sun have no such problems. All of the nastiness is created in the sun, and stays in the sun.

75: Transmutation of the nuclei in the material will create radioactivity.

15.5.3 Pros and Cons of Fusion

Collecting the advantages and disadvantages of fusion, we start with the positive attributes:

- Fusion would enjoy an inexhaustible supply of deuterium, easily accessed, outlasting the sun;
- The fusion reactor would serve as a heat source for tried-and-true steam-driven power plant technology.

And now the not-so-good aspects:

- Stable plasmas are exceedingly hard to generate at the requisite temperatures;
- 70 years of effort have not yet borne fruit as an energy supply;
- Tritium is not available, and must be fabricated from a limited supply of lithium;
- Fusion still contends with radioactive fuel (tritium) and a containment vessel that is radioactively contaminated.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

The smaller number of positive points is not in itself an indicator of imbalance, since the first point is huge. One elephant can balance dozens of kids on a playground see-saw.

15.6 Upshot on Nuclear

Nuclear fission is a real thing: it can and does produce a significant fraction the world's power. A number of substantive challenges stand in the way of scaling up significantly.[76] For conventional nuclear fission as it has been practiced thus far, the proven reserves of uranium only last 90 years at today's rate of use, and less than 4 years if we tried to get all 18 TW from fission. Radioactive waste is an unsolved problem that persists for hundreds to thousands of years. Breeder programs can extend the resource by large factors (into the 500 or 1,000 year range under an 18 TW nuclear-breeder effort). But proliferation and bomb dangers become more pronounced—not to mention an even more pressing waste issue and greater accident rates given the profusion of operating reactors. It can be difficult to get excited about a nuclear future. It is very cool that we figured out how to do it. But just because we *can* do something does not mean it is a good idea to scale it up.

76: See [105] for a short article summarizing the various challenges.

Pros and cons are listed separately for PV and ST in Section 15.4.8 and Section 15.5.3, respectively.

Fusion is a harder prospect to pin down. At present, it is not on the table, having never been demonstrated in a viable reactor capable of producing commercial-scale electricity. But even if we did manage it, how could it compete economically, as complex as it is? Even if the fuel itself is free,[77] it may turn out to be the most expensive form of electricity we could muster. Fusion is not without radioactivity concerns, and placed side-by-side, solar can look a lot better—intermittency being the crippling drawback, necessitating storage.

77: . . . as it also is for solar power, which does not mean solar power is cheap

Nuclear options cause us to grapple with the question: who are we? What is our identity? What are our aims, and where do we see ourselves going? Are we plotting a course for a Star Trek future, in which case it seems we have little choice but to adopt the highest-tech solutions. Or are we aiming for a more modest future more aligned with natural ecosystems on Earth? So even if we *can* do something, does it mean we're obligated to? Sometimes the costs may be too high.

15.7 Problems

1. If an atom were scaled up to be comparable to the extent of a mid-sized campus, how large would the nucleus be, and what sort of familiar object would be similar?

2. In parallel to Example 15.1.1, what are all the ways to label the radioactive isotope carbon-14?

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

3. How many neutrons does the isotope ^{56}Fe contain?

4. Use the information in the boxes for ^{12}C and ^{13}C in Figure 15.4 to determine the weighted composite mass of a natural blend of carbon—showing work—and compare this to the number in the left-most box for carbon in the same figure.

5. In Figure 15.4, what are the only mass numbers, A, for which no stable nuclei exist?

6. What are the only three long-lived radioactive isotopes in the portion of the Chart of the Nuclides appearing in Figure 15.4, and which one lives the longest (how long)?

7. Cosmic rays impinging on our atmosphere generate radioactive ^{14}C from ^{14}N nuclei.[78] These ^{14}C atoms soon team up with oxygen to form CO_2, so that plants absorbing CO_2 from the air will have about *one in a trillion* of their carbon atoms in this form. Animals eating these plants[79] will also have this fraction of carbon in their bodies, until they die and stop cycling carbon into their bodies. At this point, the fraction of carbon atoms in the form of ^{14}C in the body declines, with a half life of 5,715 years. If you dig up a human skull, and discover that only one-eighth of the usual one-trillionth of carbon atoms are ^{14}C, how old do you deem the skull to be?

78: Nitrogen is the principal constituent in Earth's atmosphere.

79: . . . and/or eating the animals that eat these plants

The wording is long because without context, it's just math. The real learning is in the *application* of math to the world.

8. If a friend creates a nucleus whose half-life is 4 hours and gives it to you at noon, what is the probability that it will *not* have decayed by noon the following day?

9. In close analog to the half-lives of ^{235}U and ^{238}U, let's say two elements have half lives of 4.5 billion years and 750 million years.[80] If we start out having the same number of each (1:1 ratio), what will the ratio be after 4.5 billion years? Express as x:1, where x is the larger of the two.

80: . . . a factor of 6 different

10. Control rods in nuclear reactors tend to contain ^{10}B, which has a high neutron absorption cross section.[81] What happens to this nucleus when it absorbs a neutron, and is the result stable? If not, track the decay chain until it lands on a stable nucleus.

81: . . . as indicated by the orange lower-half of the corresponding box in Figure 15.4

11. If someone managed to create a ^{14}B nucleus, what would its fate be? Track the decay chain on Figure 15.4—indicating the type of decay at each step—until it reaches stability, and indicate how long each step is likely to take.

12. A particular nuclide is found to have lost 3 neutrons and 1 proton after a decay chain. What combination of α and β decays could account for this result?

13. How would you qualitatively describe the overall sense from Figure 15.8 in terms of where[82] on the chart one is likely to see α

82: Region descriptions can include references to the mass range (e.g., low mass or high mass), above or below the stable elements (proton-rich or neutron-rich).

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

decay, β^- decay, β^+ decay, and spontaneous fission?

14. In a year, an average American uses about 3×10^{11} J of energy. How much mass does this translate to via $E = mc^2$? Rock has a density approximately 3 times that of water, translating to about 3 mg per cubic millimeter. So roughly how big would a chunk of rock material be to provide a year's worth of energy if converted to pure energy? Is it more like dust, a grain of sand, a pebble, a rock, a boulder, a hill, a mountain?

15. The world uses energy at a rate of 18 TW, amounting to almost 6×10^{20} J per year. What is the mass-equivalent[83] of this amount of annual energy? What context can you provide for this amount of mass?

83: ⓘ This is how much mass would have to "disappear" each year to satisfy current human demand.

16. How much mass does a nuclear plant convert into energy if running uninterrupted for a year at 2.5 GW (thermal)?

17. A large boulder whose mass is 1,000 kg having a specific heat capacity of 1,000 J/kg/°C is heated from 0°C to a glowing 1,800°C. How much more massive is it, assuming no atoms have been added or subtracted?

18. Replicate the computations in Table 15.5 for ^{4}He, paralleling the ^{56}Fe case in Example 15.3.2. Along the way, report the Δm in kg and the corresponding ΔE in Joules, which are not in the table.

19. To illustrate the principle, let's say we start with a nucleus whose mass is 200.000 a.m.u. and inject 1,600 MeV of energy to completely dismantle the nucleus into its constituent parts. How much mass would the final collection of parts have?

    a) the exact same: 200.000 a.m.u.
    b) less than 200.000 a.m.u.
    c) more than 200.000 a.m.u.

20. Using the setup from Problem 19, compute the mass of the final configuration in a.m.u., after adding energy to disassemble the nucleus.

Hint: convert MeV to Joules, then kg, then a.m.u.

21. Referring to Figure 15.10, what is the *total* binding energy (in MeV) of a nucleus whose mass number is $A = 180$?

Hint: Fig. 15.10 is binding energy *per nucleon*.

22. Explain in some detail what happens if control rods are too effective at absorbing neutrons so that each fission event produces too few unabsorbed neutrons.

23. Which of the following is true about the fragments from a ^{235}U fission event?

    a) any number of fragments (2 through 235) can be produced
    b) a small number of fragments will emerge (2 to 5)
    c) two nearly identical fragments will emerge

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

d) two fragments of distinctly different size will emerge
e) the fission is an alpha decay: a small piece having $A = 4$ is emitted

24. A particular fission of ^{235}U + n (total $A = 236$) breaks up. One fragment has $Z = 54$ and $N = 86$, making it ^{140}Xe. If no extra neutrons are produced in this event, what must the other fragment be, so all numbers add up? Refer to a periodic table (e.g., Fig. B.1; p. 375) to learn which element has the corresponding Z value, and express the result in the notation ^{A}X.

25. Follow the same scenario as in Problem 24, except this time *two* neutrons are left out of the final fragments. What is the smaller fragment this time, if the larger one is still ^{140}Xe?

26. Provide three examples of probable fragment size pairs (mass numbers, A) from the fission of ^{235}U + n, making up your own random outcome while respecting the distribution of Figure 15.15 in determining A values. For the sake of this exercise, assume no extra neutrons escape the fragments.

Hint: no need to identify elements; just settle on pairs of A values that add up correctly.

27. Paralleling the graphical approach in Example 15.4.3 using Figure 15.10, what total energy would you expect to be released in a fusion process going from two deuterium (^{2}H) nuclei to ^{4}He, in MeV?

Hint: Don't forget to count both ^{2}H.

28. Both nuclear and coal electric power plants are heat engines. What is the fundamental difference between these two, comparing Fig. 6.2 (p. 90) to Figure 15.12?

29. If a nuclear plant is built for $10 billion and operates for 50 years under an operating cost of $100 million per year, what is the cost to produce electricity, in $/kWh assuming that the plant delivers power at a steady rate of 1 GW for the whole time?

Hint: express the plant power in kW and multiply by hours in 50 years to get kWh produced.

30. Since each nuclear plant delivers ~1 GW of electrical power, at ~40% thermodynamic efficiency this means a *thermal* generation rate of 2.5 GW. How many nuclear plants would we need to supply all 18 TW of our current energy demand? Since a typical lifetime is 50 years before decommissioning, how many days, on average would it be between new plants coming online (while old ones are retired) in a steady state?

Hint: how many days will one plant live, then how many plants per day?

31. Extending Problem 16 toward what *actually* happens, we know from Table 15.7 that the change in mass (which was close to 1 kg in Prob. 16) is only 0.08% of the ^{235}U mass.[84] Furthermore, a fresh fuel rod is only 5% ^{235}U—the rest being ^{238}U. So how much total uranium[85] must be loaded into the reactor each year, if all the ^{235}U is used up?[86]

84: 0.185 out of 235 a.m.u.

85: Treat the two isotopes as having the same mass: the rod has 20 times more uranium than just the ^{235}U part.

86: It's not, actually, so this answer is a lower limit on the actual mass that has to be loaded in. So much for the ~1 kg answer from Problem 16.

32. Problem 15 indicated that we need the mass-equivalent of fewer than 10 tons[87] of material to support the world's annual energy

87: One ton is 1,000 kg.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

needs. But given realities that only 0.08% of mass is converted to energy in nuclear reactions, that only 0.72% of natural uranium is fissile ^{235}U, and that only half of the ^{235}U is retrievable[88] and "burned" in reactors, how many tons of uranium must be mined per year to support 18 TW via conventional fission, assuming for the sake of this problem that 5 tons of mass need to convert to energy via $E = mc^2$?

88: . . . given enrichment inefficiency

33. Based on the abundance of ^{235}U in the earth's crust (Table 15.9), how many kilograms of typical crust would need to be excavated and processed per year to provide the ~ 0.005 kg of ^{235}U you need for your personal energy (as in Example 15.4.2)?

ⓘ Of course mining does not work this way, instead seeking concentrations.

34. In crude terms, proven uranium reserves could go another 90 years at the present rate of use. But the world gets only about a tenth of its electricity from nuclear. What does this imply about the timescale for the uranium supply if the world got *all* of its electricity from conventional (non-breeding) nuclear fission?

35. Replicate the calculation and show the work that if we have 2×10^{21} J of proven uranium reserves under conventional fission, we would exhaust our supply in less than 4 years if using this source to support the entire 18 TW global energy appetite.

36. Use Figure 15.18 to reconstruct the breeder route from ^{232}Th to ^{233}U by describing the associated nuclei and decays (and half-lives) involved.

Hint: start by adding a neutron to ^{232}Th

37. For spent nuclear fuel a few decades old, what isotopes are responsible for most of the radioactivity, according to Figure 15.19?

38. Let's say that spent fuel rods are pulled out of the holding pool at the nuclear facility ten years after they came out of the core. Based on the total radioactive power from waste products (black line on Figure 15.19), *approximately* how long will you have to wait until the radioactivity level is down by another factor of 1,000 from where it is at the time of extraction?

39. Operating approximately 450 nuclear plants over about 60 years at a total thermal level of 1 TW, we have had two major radioactive releases into the environment. If we went completely down the nuclear road and get all 18 TW[89] this way, what rate of accidents might we expect, if the rate just scales with usage levels?

89: . . . also a thermal measure

40. On balance, considering the benefits and downsides of conventional nuclear fission, where do you come down in terms of support for either terminating, continuing, or expanding our use of this technology? Should we pursue breeder reactors at a large scale? Please justify your conclusion based on the things you consider to be most important.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

41. The sun is a fusion power plant producing 3.8×10^{26} W of power. How many kilograms of mass does it lose in a year trough pure energy conversion? How does this compare to the mass of a spherical asteroid 50 km in diameter whose density is 2,000 kg/m^3?

Hint: the volume of a sphere is $4\pi R^3/3$.

42. Based on the fractional mass loss associated with turning four hydrogen atoms into a helium atom, what fraction of the sun's mass would it lose over its lifetime by converting all its hydrogen into helium, under the simplifying assumption that it starts its life as 100% hydrogen?

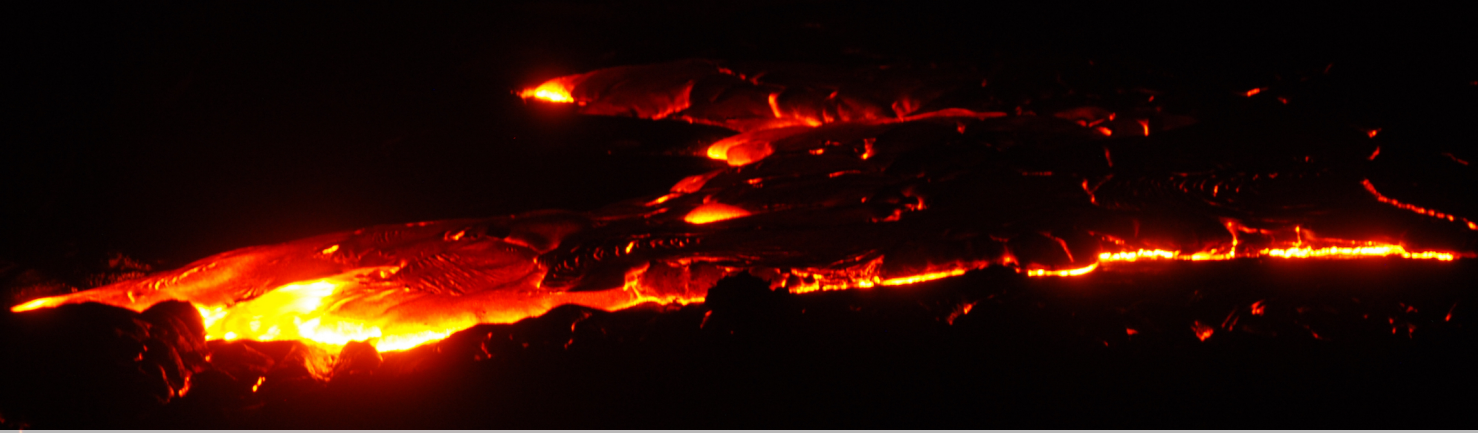
43. The three fusion forms in Eq. 15.3 each have different energy outputs. Looking at Figure 15.10,[90] how would you *qualitatively* describe why the three reactions differ in this way?

90: Tritium is not labeled, but visible just below 3 MeV on the left side.

44. Based on the calculation that 18 TW would require an annual cube of seawater 300 m on a side to provide enough deuterium, what is your personal share as one of 8 billion people on earth, in liters? Could you lift this yourself? One cubic meter is 1,000 L.

45. What are your thoughts about fusion? Are you excited, skeptical, confused, all of the above? Please offer your thoughtful assessment of the role you imagine fusion playing in our future—your best guess.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.



16 Small Players

16.1 Geothermal Energy 275
Sustainable Harvest 276
Geothermal Depletion . . . 277
Geothermal Reality 278
16.2 Tidal Capture 280
16.3 Ocean Currents 281
16.4 Wave Energy 282
16.5 Hydrogen 286
16.6 Upshot on Small Players . . 286
16.7 Problems 287

This short chapter serves to round out the menu of renewable energy options. While all of the entries are viable, at various levels of demonstration and implementation, none of them can scale up to be an important contributor to global power at the relevant level of many terawatts. In this sense, we might call the remaining resources "cute." Because of this, items are kept brief—not warranting unjustified attention. Likewise for graphical adornment: the chapter is a little bare, in part because the subjects do not deserve much in the way of undue promotion.

Skipping this chapter will not sacrifice much in the way of knowledge vital to our future. It simply fills in the gaps and addresses some of the "but what about *insert-scheme*?" questions that may arise. The next chapter, Chapter 17, highlights the pros and cons of the full range of alternative energy resources, so this installment largely exists to provide a basis for the conclusions of that more important capstone.

16.1 Geothermal Energy

The interior of Earth is still hot from the initial collapse of mass in the formation of the planet. We can think of this as the thermal conversion of gravitational potential energy as pieces fell into the gravitational field of the earth, converting first to kinetic energy and then into thermal energy (heat) after collision. But that's only half of the story. The other half is radioactive decay of long-lived unstable elements within the earth [106]. Of the radioactive decay part, 40% is from uranium (^{238}U), 40% is from thorium (^{232}Th), and 20% is from potassium (^{40}K). As was indicated in Table 10.2 (p. 168), a total of 44 TW of geothermal power flows through the earth's crust. Divided by total surface area, this amounts to less than

[106]: Johnston (2011), *Radioactive decay accounts for half of Earth's heat*

Lava seeping along east of Pu'u 'Ō'o crater in Hawaii as an evident display of geothermal energy. Photo Credit: Tom Murphy

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

0.1 W/m^2, making it less than one-two-thousandth of the solar input average.[1]

1: Recall average insolation is around 200 W/m^2.

16.1.1 Sustainable Harvest

Let's evaluate what it would take to tap into the steady 44 TW geothermal heat flow in order to provide some or all of today's 18 TW energy demand. In this way, we could be assured of steady flow for billion-year timescales, as the half lives of the radioactive elements measure in the billions of years and the residual heat of formation is still slowly leaking out after 4.5 billion years. On the face of it, 44 TW compared to 18 TW would suggest that geothermal has margin to spare.

In order to utilize geothermal energy for most of the things we do, we really need to turn the heat into a versatile form of energy like electricity. The only exception would be for heating water and air spaces. Otherwise, it is not hot enough to perform most industrial process demands, like melting metals. We know how to turn a temperature difference into electricity, via a heat engine (Sec. 6.4; p. 88). We saw that the theoretical efficiency of a heat engine is

$$\varepsilon_{\max} = \frac{T_h - T_c}{T_h} = \frac{\Delta T}{T_h}, \tag{16.1}$$

⚠ Temperatures must be in Kelvin.

where the hot and cold thermal reservoirs are at temperatures T_h and T_c, and ΔT is the difference.

A typical temperature gradient in the earth's crust is 25°C per kilometer, meaning that the temperature rises by another 25°C for each kilometer of depth. A heat engine constructed to operate between the surface (288 K average) and 1 km down (313 K) could expect a *maximum* thermodynamic efficiency of 8%. Now imagine an ambitious deployment across 100% of land (29% of the globe): 8% of 29% of the original 44 TW flow is down to a paltry 1 TW.[2] And that's if achieving the *theoretical maximum* efficiency limit[3] and somehow accessing the *entire* flow of heat 1 km down through a network of pipes across *every* bit of land area on Earth. Can you imagine the scale of this effort? And for all that, it would fall well short of 1 TW, considering practical limitations. Don't hold your breath for this long-term-sustainable form of geothermal energy to provide an answer to our 18 TW appetite.[4]

2: 44 TW × 0.08 × 0.29 ≈ 1 TW.

3: Typically, we fall short by a factor of two.

4: In fact, the continental share of the steady geothermal flow—29% of 44 TW—is only 13 TW and not enough to satisfy demand even at the impossible-to-realize efficiency of 100%.

Box 16.1: Deeper and Oceans, Too

While the ambitious global-scale 1 km deep geothermal network was disappointing in its output, the mathematical solution is to "just"[5] use the *entire* globe and dig deeper, allowing a larger thermal gradient and therefore higher efficiency, according to Eq. 16.1. Allowing for

5: ⚠ Beware of the word "just."

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*; *Freely available at*: *https://escholarship.org/uc/energy_ambitions*.

engineering challenges that may limit us to half the theoretical efficiency, we would have to capture 36 TW of theoretical flow to end up at 18 TW. Now we need a theoretical efficiency of 82%,[6] translating to T_h of 1,600 K, which would be about 50 km down: deeper than the earth's crust is thick.

For context, the deepest mine is less than 4 km deep, and the deepest drill hole is about 12 km.[7] So outfitting 100% of Earth's surface—including under the oceans—with a dense thermal collection grid 50 km down sounds like pure fantasy.

6: 36 out of 44

7: Drilling stopped because technical challenges prevented going deeper. The project goal was 15 km.

16.1.2 Geothermal Depletion

The previous section was framed in the context of accessing the 44 TW steady geothermal flow, sustainable for billions of years—finding that we cannot expect to satisfy demand by that route. But when did we ever exhibit collective concern for long-term sustainable solutions? The human way is more about exploiting a resource fully, not worrying about consequences even decades down the line. In that sense, geothermal energy has more to offer—at least on paper.

A one-time extraction of thermal energy under out feet—not worrying about replenishment—amounts to *mining* thermal energy, in much the same way that we mine copper, or fossil fuels. Using a rock density of 2,500 kg/m^3 and a specific heat capacity of 1,000 J/kg/°C (Sec. 6.2; p. 85), each cubic meter of rock has an extra 60 MJ of thermal energy for each kilometer deeper we go—based on a gradient of 25 °C/km, as before. Is that a lot? It's about the same as the energy in 2 L of gasoline. The energy density works out to 0.006 kcal/g, to put in familiar units (see Table 16.1).

Table 16.1: Energy densities of familiar energy substances. For hydroelectricity, a 50 m dam is assumed, and for geothermal, the depth is 1 km.

Substance	kcal/g
Gasoline	11
Fat (food)	9
Carbohydrates	4
TNT explosive	1
Li-ion battery	0.15
Alkaline battery	0.11
Lead-acid battery	0.03
Geothermal (1 km)	0.006
Hydroelectric (50 m)	0.0001

So it's no screaming-good deal, but it's still energy, and the earth's crust has a heck of a lot more rock than it does oil. To appreciate the scale, the land area of the lower-48 states is approximately 10^{13} m^2. A 1-meter-thick slice of earth under the U.S. at a depth of 1 km therefore contains 60 MJ/m^3 times 10^{13} m^3, or 6×10^{20} J of energy. It's a big number, but recall that 1 qBtu is about 10^{18} J, so we're talking about ~600 qBtu. The U.S. uses about 100 qBtu per year of energy, but at an average efficiency of 35% in heat engines, so that we seek about 35 qBtu of *useful* energy. As we saw, the geothermal resource, at lower temperature, is less potent in terms of efficiency. If achieving half of the theoretical 8% efficiency for the 1 km ΔT of 25°C, a one-meter-thick slice would provide about 24 qBtu of useful work.[8] Reaching the 35 qBtu goal would require a slice about 1.5 m thick, at 24 qBtu per meter.

8: . . . 4% efficiency times 600 qBtu thermal resource for one meter

To summarize, we would need to *completely* remove all the heat from all the rock 1 km below our feet in a 1.5 m-thick layer every year. Once we cool the underground rock, it will take a long time for the surrounding

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

heat to diffuse in, so we'd have to keep moving down, year after year. This would mean completely replacing the collection network (whatever comprises it—this is all fantasy) every few years. It's ludicrous to imagine we would endeavor to go to such extremes, plowing through the deep earth to remove every scrap of thermal energy in a massive never-ending effort.[9] Such a colossal scheme makes oil drilling seem like child's play.

9: It would not be surprising if the EROEI is abysmally low or even a net energy drain.

Granted, as one drills deeper, the thermal energy goes up *and* the efficiency increases as ΔT climbs. The result is *quadratic*, in that the energy yield 3 km down is 9 times that at 1 km down.[10] At the same time, drilling gets more challenging and at some point exceeds the current state of the art. By the time the temperature reaches 150°C—which is thought of as a minimum viable temperature for traditional geothermal ventures—drilling technology runs into limitations.

10: . . . because the hotter rock contains 3× more thermal energy *and* that energy can be converted to electricity 3× more efficiently

Box 16.2: There Will Be Hype

As long as an idea is not outright impossible, the world is big enough and competitive enough that enterprising individuals will be able to generate interest and investment in ideas that seem viable and can be touted to have great promise. Whether that idea is truly capable of benefiting humanity as a "good" idea is not fully evaluated. Instead, if it can make money in the short term,[11] then it may get a green light.

So be wary of claims by people or companies whose financial interests lie in the perception of success and promise. Even media coverage that *should* be objective is often quantitatively sloppy,[12] and has a much easier time finding enthusiasts willing to devote time and quotes than non-enthusiast experts who are too busy pursuing their own projects to waste time poking holes in shoddy ideas.

11: . . . or at least attract investment

12: . . . lacking a staff physicist

16.1.3 Geothermal Reality

Enough, then, about "pipe dreams" of massive geothermal exploitation on a planetary scale. Geothermal energy is not *all* fantasy, as some places are able to capture significant energy from this resource.

In a few locations, hot magma is brought near the surface, offering rare cracks of access to high temperatures. An electrical power plant, as depicted in Fig. 6.2 (p. 90), does not care particularly where the thermal energy derives, as long as it's hot enough to make steam. The ideal site has:

- magma near the surface—volcanic regions, for instance;
- fractured rock above the magma in which water can flow;
- water temperatures in excess of 180°C (under pressure);
- a caprock above the fractured rock, able to trap pressurized steam.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

The most common scheme—labeled "hydrothermal"—is to drill two holes into the ground near each other, injecting water into one and collecting pressurized steam from the other. Fractures in the rock permit water and/or steam to flow between the two holes. Alternatively, but far less common, a fluid[13] can be run through a closed loop that passes through the hot medium. By either direct use of the steam in the hydrothermal case, or generating steam from the hot fluid in the closed-loop case, the resulting steam can be used to run a turbine and generator in the usual way.

13: . . . not necessarily water now

A newer form, called "binary" geothermal uses two fluids: water in the ground as in other schemes, but a second fluid having a much lower boiling point to make a steam analog at lower temperatures. This opens additional power generation possibilities at temperatures below 100°C, but of course will suffer the inevitable efficiency hit when T_h is lower, according to Eq. 16.1.

Globally, roughly 10 GW of electricity is produced from geothermal energy [107], and an estimated additional 28 GW of direct heating is obtained from this source [108]. Together, these account for 0.4% of the 18 TW global energy budget, after a thermal equivalent adjustment.

[107]: Intern'l Renewable Energy Agency (2018), *Renewable Energy Statistics 2018*

[108]: (2020), *Geothermal Heating*

Country	GW installed	GW produced	% elec.
U.S.	3.5	1.9	0.4
Philippines	1.9	1.3	27
Indonesia	1.5	1.2	4
New Zealand	1.0	0.85	15
Mexico	1.0	0.7	3
Italy	0.9	0.7	1.5
Iceland	0.7	0.6	30
World Total	12.6	9.4	0.4

Table 16.2: Global geothermal electricity production in 2016 [85, 107, 109]. Note that the percentage is the fraction of *electricity*, not total energy, contributed by geothermal. The capacity factor tends to be relatively high for this non-intermittent resource (see Problem 6).

Table 16.2 lists the top 7 producers of geothermal electricity, capturing 72% of the global total. Note that many are on the Pacific Rim, sometimes called the "ring of fire" for its volcanic activity. Iceland gets 30% of its electricity[14] from geothermal sources. But in absolute terms, it is a small amount of energy. Considering that a single nuclear plant puts out about 1 GW, the countries in Table 16.2 have the equivalent of 1–2 nuclear plants in the form of geothermal (compare to Table 15.8; p. 256).

14: And electricity is only about one third of Iceland's energy demand.

The U.S. gets an average of 1.9 GW of electrical production[15] from geothermal sources [85]. 72% of this is produced in California—almost all at a site called The Geysers in the northern part of the state—accounting for ~6% of the state's electricity. Another 22% of U.S. geothermal electricity is produced in in Nevada. The rest is in Utah, Hawaii, Oregon, Idaho, and New Mexico, in that order (7 states total).

15: . . . ~0.4% of total electricity

[85]: U.S. Energy Inform. Admin. (2020), *Electric power monthly*

Geothermal is just a small player. The fact that a country like Iceland can produce a large fraction of its electricity this way mostly tells us that Iceland is on a geological hot-spot and is not very populated. We should

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*; *Freely available at*: *https://escholarship.org/uc/energy_ambitions*.

not expect geothermal to assume a large role in the energy landscape of tomorrow.

16.2 Tidal Capture

Table 10.2 (p. 168) indicates that the earth currently receives 3 TW of power in the form of tidal energy. Gravity from the moon (and the sun[16]) tug on the earth, pulling slightly harder on the nearest side, and less enthusiastically on the far side. This deforms the otherwise spherical Earth[17] into a *prolate ellipsoid*, which is a fancy way to say oval—somewhat like an egg. The bulges that form on either side (near and far; see Figure 16.1) as a consequence are more pronounced in the ocean than the land, but both indeed deform.

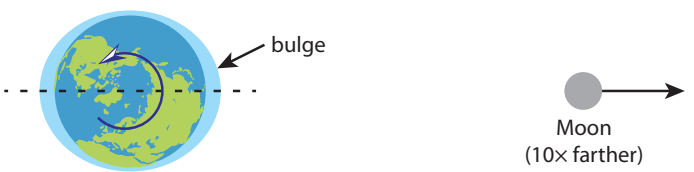


Figure 16.1: The moon (pictured closer here to fit on page) raises bulges on the near and far side of the earth, while the earth rotates underneath, causing two high tides and two low tides each day.

The earth rotates "under" these bulges, since the bulges point along the Earth–Moon axis and take a month to make a complete revolution in space, while the earth rotates once per day. At a fixed position on Earth, then, the experience is two high tides and two low tides each day.[18] As the tide comes in and flows out, friction between land and water results in energy dissipation, to the tune of 3 TW.[19]

The idea for capturing this natural flow of energy is to allow the tide to enter a bay or inlet—raising the water level by perhaps several meters—then closing the exit, trapping the water behind a wall. At this point, the situation is very much like a hydroelectric dam (Sec. 11.2; p. 175), but at a much lower water height than is typical for hydroelectric dams. All the same, exiting water can be forced through turbines turning generators to create electricity. Draining through turbines over a 6-hour period allows the process to begin again at low tide, opening the gates to refill the inlet as the tide comes back in.

The amount of energy depends on the area of the captured body of water and the height of water trapped behind the wall. We use the familiar gravitational potential energy: $E = mgh$ to calculate the energy involved (Sec. 11.1; p. 173). If the height of water trapped behind the wall at high tide is called h, the height of water behind the wall smoothly transitions from $h \rightarrow 0$ as the water drains out, so that the *average* height of water behind the wall is $h/2$, and the gravitational potential energy available really looks like $mgh/2$. The mass is density (ρ =1,000 kg/m^3) times volume, and volume is the area of the captured water surface, A, in square meters, times the initial water height, h. To get power, we

16: The tidal influence from the sun is subdominant, at 45% the strength of lunar tides, so for simplicity we will describe just lunar tides from here on.

17: Actually, an even larger (21 km) effect is the spinning motion of the earth flattening it into an *oblate ellipsoid*.

18: The sun also contributes, sometimes adding to the peaks and troughs (near new and full moon, when Earth, Moon, and Sun are arranged along a line), and sometimes moderating the tides by filing in the gaps a bit (at quarter moon phases).

19: An aggressive global-scale tidal capture enterprise could actually increase the total tidal energy budget, which would have the side effect of pushing the moon further away, slowly, as is described in Sec. D.4 (p. 402).

© 2022 T. W. Murphy, Jr.; Creative Commons Attribution-NonCommercial 4.0 International Lic.;
Freely available at: https://escholarship.org/uc/energy_ambitions.

divide the energy available by the time over which we let out the water: nominally 6 hours.[20] Collecting it all, we have:

$$P = \varepsilon \frac{\Delta E}{\Delta t} = \frac{\varepsilon m g \frac{h}{2}}{\Delta t} = \frac{\varepsilon \rho A g h^2}{2\Delta t}, \tag{16.2}$$

where $g \approx 10\,\mathrm{m/s^2}$ is the acceleration due to gravity and ε is the efficiency of converting gravitational energy into electrical energy.

20: Six hours is a typical time between high tide and the next low tide.

Example 16.2.1 The Rance tidal capture station in France has a capture area of $A = 22.5\ \mathrm{km^2}$, and a height capacity of 8 m. If operating off of a 7 m high tide and draining for 6 hours, what is the expected power delivered at an efficiency of 90% (as is typical for hydroelectric power)?

Expressing A in square meters yields $22.5 \times 10^6\ \mathrm{m^2}$. The time is 21,600 seconds, leading to a calculated value of 230 MW.

Only two large tidal facilities operate in the world today: Rance in France (Example 16.2.1) has a rated capacity of 240 MW, and produces an average of 57 MW. Thus the capacity factor is about 24% due to the fact that it can only generate tidal power half the time[21] and not all high tides are at the full design height. The h^2 in Eq. 16.2 indicates particular sensitivity to height—due to the double-whammy that lower height means less trapped mass *and* lower pressure head.

21: It has to spend half the time letting water back in as the tide flows.

The other facility is Sihwa Lake in South Korea, a 254 MW facility that came online in 2011. Much like Rance, and for the same reasons, its capacity factor is 25%, averaging 63 MW. Its basin is 30 $\mathrm{km^2}$ and has similar operating height as the French installation. The Sihwa Lake facility cost $560 million to build, making it $9 per average Watt of delivered power. This puts it roughly in line with the cost of nuclear power (page 256), and a little higher than utility-scale PV, currently.

Two other large tidal stations in the 300–400 MW capacity range are in the works for the UK and Scotland. But it's not something that works well everywhere: best suited for areas that have large tidal swings and large inlets having narrow mouths that are easy to dam. It's a niche player now and always will be. After all, the 3 TW global budget for tidal energy suggests that it is not an energy jackpot.

16.3 Ocean Currents

Steady currents throughout the volume of the ocean[22] are estimated to total 5 TW (see Table 10.2; p. 168). This is not much more than total tidal dissipation on the planet, at 3 TW. Already, we don't hold out much hope for an energy bonanza. Wind in the atmosphere, by contrast, has

22: . . . contrasted to the oscillating currents from tides

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

almost 200 times as much power. And we saw that getting even a few terawatts out of wind is a challenge, globally.

The physics is very similar to that of wind (Chapter 12). Kinetic energy from the motion of the water is transferred to rotational motion in a turbine to turn a generator to make electricity. Echoing Eq. 12.2 (p. 187), we have

$$P_{\text{current}} = \frac{1}{2}\varepsilon\rho_{\text{water}}\pi R^2 v^3. \tag{16.3}$$

The efficiency, ε is limited by the same physics as for wind, so 50% is a reasonable maximum. Recall that $\rho_{\text{water}} = 1,000\,\text{kg/m}^3$.

So we still have the quadratic dependence on rotor radius and the crucial cubic power of velocity. Ocean currents tend to be far slower than wind velocities, but the density of water is about 800 times the density of air, so that for the same rotor area, water moving at 1 m/s has similar power as air moving at 9 m/s.

One problem is that the giant land-based wind turbines whose rotors move at high speeds relative to the wind would not be practical in water, so individual units would more likely be on the ten-meter scale and not a 100-meter scale. Individual rotors would therefore likely generate a few hundred kilowatts apiece, requiring thousands of rotors to make up a typical 1 GW power plant output.

Example 16.3.1 A 10 m diameter rotor sits in a brisk 2 m/s current. At an efficiency of 50% (similar to wind), how much power would the turbine deliver?

All the pieces[23] fall right into Eq. 16.3 to yield about 160 kW.

23: Just be careful about diameter vs. radius.

Among the advantages, marine currents are very steady and dependable, so that capacity factors would be high. This is in contrast to solar, wind, and even hydroelectric and tidal resources.

However, marine environments are not kind to underwater structures,[24] making it hard to imagine the headache it would be to maintain the moving parts in good condition—all in a difficult operating environment for human workers. Additionally, marine life might not be so happy to get whacked by a rotor blade.

24: Look at pier pilings and sunken ships to get an idea.

In the end, we should not expect this sector to contribute substantively to global energy, given the small scale of total energy available and the practical difficulties associated with harnessing any sizeable piece of it.

16.4 Wave Energy

Waves present an interesting case to study, being composed of a continuous interplay between potential and kinetic energies—water particles

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

executing circular paths as they gain and lose gravitational potential energy. Meanwhile, the energy in a packet of waves travels efficiently onward, until reaching a barrier like a shoreline, where they break in a display of kinetic energy. It is hard to stare at a coastline being pounded by surf without admiring the power of nature. What if we harnessed that power?

As a means of estimation, let's imagine sinusoidal[25] waves 1 m in amplitude (peak-to-trough) arriving every 6 seconds, traveling at about 3 m/s. Crests must then be 18 m apart, since the distance between crests is their rate (3 m/s) times time (6 seconds). We call this the wavelength, symbolized by λ (Greek lambda). See Figure 16.2 for the layout.

25: . . . shaped like a sine wave

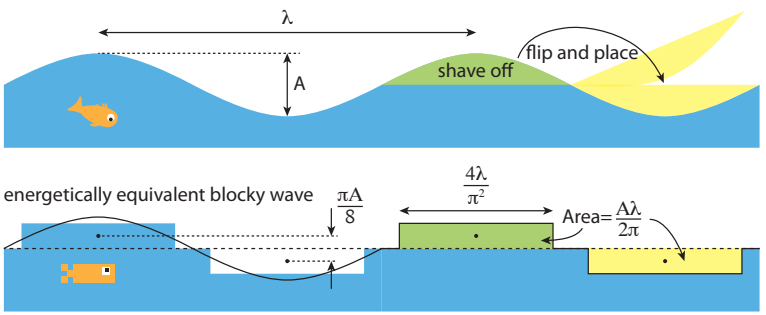


Figure 16.2: Sinusoidal waves have amplitude A, and wavelength λ. In terms of harvesting the gravitational potential energy, we can think of it as lopping off the crest and flipping it over onto the trough to level the water surface. In doing so, we move some mass, m, down a height h to get mgh of energy. The block-equivalent is shown below, where the area and average height of the sinusoidal trough/crest has been faithfully captured by rectangles of height $\sim 0.39A$ and length $\sim 0.41\lambda$. From these, it is possible to figure out the potential energy associated with the wave.

In order to figure out the energy involved, we need the mass of water raised and a height to which it is raised. Notice that in Figure 16.2, the potential energy in the wave can be extracted by making the water flat again, which is equivalent to taking all the water from the crest and putting it into the trough. We just need to know how much water we're moving, and through what height. Figure 16.2 has done the fancy math already and redrawn the wave as rectangular chunks that have equivalent area as the sinusoidal crest and trough and also the same average (midpoint) height relative to the average surface height (dashed line). From this, we learn that the wave crest has area $A\lambda/2\pi$ and the height of the displacement is $\pi A/8$, where A is the wave amplitude from the top of the crest to the bottom of the trough.

Author's note: we'll be going through a bit of a derivation here—not because wave power is particularly important, but because physicists sometimes can't help it. Feel free to skip past all the equations, as this is meant as a sort of entertainment. What can I say: the introduction had a warning about skipping the whole chapter, yet here you are.

To assess potential energy as mgh, we need three pieces. We already know $g \approx 10\,\mathrm{m/s^2}$, and we now know that $h = \pi A/8$. The mass is a density, ρ, times a volume. We already have the area of the crest cross section as $A\lambda/2\pi$. To get a volume, we need a length *along* the wave, which we are free to make up as a variable we'll call ℓ. The mass, m, of our block of water is then $\rho A\lambda\ell/2\pi$. Putting this together, the gravitational potential energy (GPE) associated with putting the water back to a flat state along a length, ℓ, of one wave is

$$E_{\mathrm{GPE}} = mgh = \frac{\rho A\lambda\ell}{2\pi} g \frac{\pi A}{8} = \frac{\rho\lambda\ell g A^2}{16}. \tag{16.4}$$

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

It's not an equation to remember, just a way to keep track of the physics as we build toward a final/useful result. Next, we want to understand the power delivered as the waves come and come. We get a new one every $\Delta t = 6$ s. And we also reasoned before, in slightly different form, that $\Delta t = \lambda/v$, where v is the velocity of the wave: 3 m/s in our example. The power looks like:

$$P_{\text{GPE}} = \frac{E_{\text{GPE}}}{\Delta t} = \frac{\rho\lambda\ell g A^2}{16\Delta t} = \frac{\rho\ell g A^2 v}{16}. \tag{16.5}$$

This is how much power the potential energy part of waves contributes as the waves pile onto the shore. But waves also have kinetic energy. It turns out that kinetic energy and potential energy are balanced in a wave—which is perpetually sloshing back and forth between potential and kinetic forms, much as happens in the motion of a pendulum. So the total power is just double P_{GPE}, or

$$P_{\text{tot}} = \frac{\rho\ell g A^2 v}{8}. \tag{16.6}$$

It is a little awkward to have to specify the length of the wave, but we needed it to make sense of the mass involved. At this point, let's switch to expressing the *power per unit length* of the wave, or P/ℓ.

$$\frac{P_{\text{tot}}}{\ell} = \frac{\rho g A^2 v}{8}. \tag{16.7}$$

Notice that this expression does not actually depend on the wavelength, in the end. The only measures of the wave that enter are the amplitude and velocity.[26]

26: The velocity of near-shore waves is set only by the depth, d, of the water ($v = \sqrt{gd}$).

For our example of 1 m amplitude and 3 m/s velocity, we compute a power per unit length of 3,750 W/m. Okay, this is a new unit, and it looks vaguely encouraging. Blow dryers, toaster ovens, space heaters, or similar power-hungry appliances consume about 1,800 W of power when running full blast, so 3,750 W/m is roughly equivalent to having *two* such appliances plugged in and running for every meter of length along the wave, or coastline. It seems like a bonanza: Our collective hair will be dry in no time! Take a moment to picture a beach cluttered with a power-hog appliance plugged in every 0.5 m all down the beach, all running at full power. That's what the waves can support, and it seems pretty impressive.

⚠ We will use this 3,750 W/m figure from here on for our rough analysis, but it should be borne in mind that larger wave amplitude has a quadratic effect on power, and waves are not all 1 m peak-to-trough!

But what we care about, in the end, is how much *total* power the waves can deliver: how many terawatts? So we need to multiply the wave P/ℓ value by a length along the wave, or a shoreline length.

Example 16.4.1 How much wave power arrives on the U.S. Pacific coast if the whole coastline is experiencing 1 m amplitude waves at a

© 2022 T. W. Murphy, Jr.; Creative Commons Attribution-NonCommercial 4.0 International Lic.;
Freely available at: https://escholarship.org/uc/energy_ambitions.

wave speed of 3 m/s?

The foregoing text already worked out that P/ℓ for these wave parameters is 3,750 W/m. Now we just need to multiply by a coastline length. The Pacific coast of the U.S. is approximately 2,000 km long. Multiplying 3,750 W/m by 2,000,000 m yields 7.5 GW.

Getting 7.5 GW on the Pacific coast, and maybe a similar amount on the Atlantic coast for a total of 15 GW is nothing to sneeze at. But consider that the U.S. electricity demand is about 450 GW, and various alternatives already top the upper limit of wave potential, as shown in Table 16.3.[27] Also, to get 15 GW from waves would require extracting *all* the wave energy from the U.S. coasts. Sorry surfers. Sorry marine life who depend on the waves for stirring nutrients and other functions a physicist can only guess. The point is that when a *fully developed* wave energy resource only provides a few percent of demand, while a promising thing like solar is already roughly matching it and has ample room to grow, we can be pretty confident that wave power will not become an important player.

Table 16.3: Alternative electricity scales in the U.S.

Resource	GW average
Solar	11
Wind	31
Hydro	33
Nuclear	92
All Elec.	450

27: Note that wind and solar are growing year by year, so their ultimate numbers will be significantly higher, still.

But let's make a quick estimate of global potential to compare to the 18 TW demand. The question becomes one of how much coastline[28] receives wave energy.

28: Keep in mind that many coastlines are protected from waves by their orientation or lack of exposure to long stretches of open water.

We can play a crude trick by recognizing that two big clumps of connected land occupy the eastern and western hemispheres, each having an east coast and a west coast, running roughly pole to pole. So the exposed coastlines roughly ring the globe twice. Our familiar unit of length, the meter, was actually chosen to yield approximately 10,000 km from equator to pole, so it takes 40,000 km to circle the earth once, and we therefore estimate 80,000 km of wave-receiving coastline. Multiplying 3.75 kW/m by 80,000 km yields 0.3 TW, or less than 2% of global demand, if *fully* developed.

3,750 W/m seemed so promising at first, especially compared to insolation numbers we saw in Chapter 13 that tend to be around 200 W/m^2. But a huge difference here is that wave density is a *linear* measure (Watts per *meter* of coastline) versus an *areal* measure in the case of insolation (Watts per *square* meter of land). A country has *far* more square meters than linear meters. So wave power, even if fully developed around all the world's coastlines, could not amount to very much, and is therefore placed in the "boutique" category, along with the other occupants of this chapter.

Work this out for yourself, maybe by thinking of something small like a house first, to understand why.

We have not even mentioned the technologies that can convert wave power to electricity—because what's the point? Suffice it to say that the same relative motion of magnets past coils of wire that create electrical energy in a generator can work in applications that do not spin all the way around, as is usually the case. A back-and-forth motion like one cylinder inside another, or a joint bending back and forth can be configured to generate electrical power as well.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*; *Freely available at*: *https://escholarship.org/uc/energy_ambitions*.

16.5 Hydrogen

Combustion of hydrogen gas produces energy, and indeed is sometimes used as rocket fuel and in some demonstration vehicles. It can also be combined with oxygen in a fuel cell (not combustion) to generate electricity. Twenty years ago, references to a *hydrogen economy* replacing fossil fuels were common. So why has it been left out of this book, shoved unceremoniously into the last part of a chapter on miscellaneous small players? One key reason: hydrogen **is not a source** of energy. Pockets of hydrogen gas are not found underground.[29]

29: It is too reactive to survive on its own, instead bound in H_2O and CH_4, for instance.

Hydrogen can be produced from water given enough energy input, typically via *electrolysis*, which splits H_2O into hydrogen and oxygen. Electrolysis efficiency is typically 65–80% in terms of capturing input energy in the form of stored hydrogen. So hydrogen should be thought of as energy **storage**—like a chemical battery—rather than as an energy *source*. Combustion of hydrogen to drive a car, for instance, would suffer the usual heat engine inefficiency of ~25%. Combined with electrolysis inefficiency, total efficiency is down to 15–20%. A standard battery does far better in round-trip[30] storage efficiency: 60–90%. Fuel cells are far more efficient than combustion—around 65%. Combined with the electrolysis step, the overall efficiency is still well short of battery storage, at around 50% overall.

30: . . . meaning charge and discharge energy cycle

Hydrogen in gaseous form is bulky and hard to package—especially in mobile applications like cars. The alternative to gaseous storage is as a cryogenic liquid, adding other complications. Hydrogen is more dangerous than gasoline as an explosion hazard. Fuel cells are not particularly robust and have difficulty in cold weather. On balance, hydrogen is unlikely to come roaring onto the scene as a fossil fuel replacement. An article [110] summarizes the pros and cons well.

[110]: Volkswagen (2019), *Hydrogen or battery? A clear case, until further notice*

16.6 Upshot on Small Players

The topics discussed in this chapter are presented more out of duty to completeness than as set of hopeful candidates for the energy of tomorrow. The first four are physically viable *sources* of real energy. Geothermal and tidal (to a lesser extent) contribute to today's energy mix, and new development is underway to expand both.

Geothermal energy is only practical in a small number of places that provide proximate access to high-temperature magma. Plans to extract geothermal energy from "normal" locations have not yet materialized, and the challenges are substantial. We should not expect geothermal energy to contribute at the terawatt scale, in practice.

Tidal energy, like geothermal, is very location-dependent—requiring large inlets with large tidal amplitude that can be closed off easily. But

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

some locations *are* well-suited to supplement local power generation from this resource. The possibilities for terawatt-scale extraction are essentially non-existent, so tidal will remain a small player.

Ocean currents and wave power are sometimes explored in small demonstration projects, but never yet at scales that add significantly to the total energy mix. Terawatt-scale implementation is not to be expected.

Hydrogen is not a source of energy, but a possible means of energy storage. Various challenges make it less than exciting as a fuel of the future.

As much as anything, this chapter may go to show that physics can be used to explore all sorts of ideas, operating within a finite set of rules and options. Humans have had a very long time to explore the world and identify possible sources of energy.[31] The options on the table are now well known, and any new ideas can be assessed quantitatively—usually only to show that the resource is small compared to present-day demand.

31: . . . which has been understood to be a pressing concern for at least half-a-century

16.7 Problems

1. A typical college campus is probably about one square kilometer in area. How much power out of Earth's 44 TW geothermal budget passes through the campus area, assuming uniform distribution across the 5.1×10^8 square kilometers of earth's surface? How does this compare to a typical college electrical demand of about 20 MW?

2. If it were possible to achieve 60% of the *theoretical* maximum heat engine efficiency for a geothermal plant, how deep would it have to drill to access high enough temperatures to match the 35% efficiency of fossil fuel power plants if the thermal gradient is 25 °C/km? Have we drilled this far before?

   Hint: using $T_c = 288$ K, solve for T_h at the target theoretical maximum efficiency, then figure out ΔT.

3. At 4 km down, we expect the ΔT to be about 100°C, and each cubic meter of rock would contain about 250 MJ of thermal energy. If 50% of the maximum theoretical efficiency were achieved from an ambient environment at 288 K, how much rock thickness would have to be depleted in a year to satisfy a 1 km^2 campus whose output electricity demand is 20 MW?

   Hint: breathe. Realized efficiency; amount of *thermal* power needed; thermal energy in a year; how many cubic meters; how thick for the campus area?

4. Your friend just visited a geothermal power plant, and is excited by the facts that it is environmentally clean, not intermittent like solar or wind, can last ages, totals more than twice our 18 TW budget, and it really works—seen with their very own eyes. What are the key points you might offer to temper their unvarnished enthusiasm?

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

5. What are your thoughts on whether we should "mine" geothermal heat in a way that could theoretically last hundreds or thousands of years, but not tens of thousands of years before depleting the resource? Such activity would not be strictly sustainable in the long haul, but could seem abundant for many generations. Should we care about this? Where do you land?

6. Work out the capacity factors of geothermal installations for various countries and the world as a whole from the data in Table 16.2.

7. The average global production, in GW, for hydroelectricity, wind, solar, and nuclear are 477, 125, 67, and 393, respectively (Tables 11.2, 12.2, 13.3, and 15.8). Add geothermal production from this chapter to the list and express each of the five as percentages[32] and draw an approximate pie chart to help put these alternative electricity resources in context.

8. The Sihwa tidal power plant has a reservoir area of 30 km^2 and is rated to generate 254 MW of electrical output. If generation efficiency is 90%, what initial water height does the power rating correspond to, assuming a 6 hour discharge time?

9. On the basis of fluid power scaling as the cube of velocity, show the supporting math for the claim in the text that a water current at 1 m/s delivers the equivalent power (per rotor area) as a wind speed of about 9 m/s.

10. If placed in a steady current of 1 m/s, and at a generation efficiency of 40%, how large would an ocean current rotor be (diameter) to satisfy the 10,000 W demand of the average American? Put this scale in the context of some familiar object or space.

11. The running example for waves in the text delivers 3,750 W/m of energy along a coastline. It sounds like a lot. But if we put the 40 million residents of California along its 1350 km coastline, how many meters does each person have, and how much power from waves?

12. Let's imagine waves hitting the entire 2,000 km Pacific coastline of the U.S. that are different from those evaluated in the text. This time, the waves have 50 m wavelength, arriving every 10 seconds, and 2 m crest-to-trough amplitude. How much power does the coast receive under these more active conditions?

32: . . . so that the five add to 100%

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.



17 Comparison of Alternatives

17.1 Alternative Energy Matrix . 289
17.2 Individual Tally 294
17.3 Student Rankings 299
17.4 Upshot: Fossil Fuel Gap . . 300
17.5 Problems 301

This book has explored the energy surge of the last ~ 150 years, largely casting it as a story of fossil fuels. Combining the finite nature of fossil fuels with their environmental cost makes it clear that we cannot expect fossil fuels to carry us into the indefinite future—and the sooner we transition away, the better. We have explored aspects of the chief alternatives in the preceding chapters. This chapter gathers key results.

One striking realization is that no fundamentally new energy technologies have emerged in the last half-century. Hydroelectricity, nuclear fission, wind, and solar photovoltaics, had all been invented. Fusion has been in the research phase the entire time. It would be unwise to expect a miraculous new entry into the field of energy—a knight in shining armor.[1] It is not as though scientists and engineers have not prioritized energy research: energy has been recognized as a lynchpin of modern life all this time. The peak of U.S. oil production in 1970 followed by two global oil crises that decade served as a wake-up call. Alarm bells have been ringing over climate change almost the entire time. Interest has been keen, and many ideas[2] have been entertained. Scratching our heads to generate more ideas is not so promising at this stage.

Perhaps, then, the table is set. We know the actors. We know the pros and cons, having many decades of experience with each of the technologies. Where does this leave us? This chapter[3] aims to pull it all together to take stock. The chapters to follow will begin to address how we might move forward in light of our predicament.

1: Extensive understanding of fundamental physics offers no hopeful revolutionary energy source—especially on relevant timescales.

2: . . . most untenable

3: The content of this chapter is an edited reprisal of the author's contribution to *State of the World 2013: Is Sustainability Possible?* by The Worldwatch Institute [111]. Copyright © 2013 Worldwatch Institute. Reproduced by permission of Island Press, Washington, DC. Originally, the content appeared in a *Do the Math* blog post called The Alternative Energy Matrix.

17.1 The Alternative Energy Matrix

In exploring potential replacements for fossil energy, it soon becomes apparent that fossil fuels are unparalleled in many respects. Even though

Imagery from the groundbreaking movie *The Matrix*, by J. Rue.

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.

viewed as a *source* of energy from the ground, fossil fuels are perhaps more aptly described as nearly perfect energy *storage* media, at energy densities that are orders of magnitude higher than anything achieved thus far in the best available battery technology. The storage is nearly perfect because it is reasonably safe, not especially corrosive, easy to transport,[4] lightweight yet dense enough to work in airplanes (see Box 17.1), and indefinitely storable—indeed, for millions of years—without loss of energy. No alternative storage technique can boast all the same benefits, be it batteries, flywheels, hydrogen, or ethanol.

4: . . . via pipelines, often

Box 17.1: Electric Airplanes?

Box 13.3 (p. 212) demonstrated that an airplane receiving solar energy in ideal conditions could gather only 4% of the power typical of a cruising airliner, and would not be able to get off the ground on direct solar input. But what about battery storage?

The best lithium-ion batteries store 0.17 kcal/g, which is 65 times less energy-dense than gasoline, at ~11 kcal/g. A Boeing 737—the workhorse of the airline industry—has an empty weight around 35 tons, and accommodates about 15 tons of fuel and 15 tons of passengers/cargo. A comparable energy storage in the form of batteries—even allowing a factor of three difference in thermal efficiency versus electric[5] efficiency—would take 300 tons of battery: *far* in excess of the entire plane's maximum takeoff weight. Or the airplane could just accept a reduced range by a factor of 20: down to 200 km. Recharge time could easily exceed flight time. It's not the same. See Sec. D.3 (p. 397) for additional analysis.

5: This may be too generous, since performance will depend on the propulsion mechanism: propellers?

Electric flight is not *impossible*. Model airplanes, drones, and other cute demonstrations may get airborne. But this is not at all the same thing as practical air travel!

This chapter aims to build a relatively complete list of ways in which we might tap into Earth's various energy flows and stores (as in Table 10.2; p. 168), most of which have been addressed in some detail in Chapters 11–16.

In order to make comparisons, it is helpful to create a matrix of energy source properties so that the relative strengths and weaknesses of each are obvious at a glance. (See Figures 17.1 and 17.2.) The matrix is presented as a color-coded figure based on 10 different criteria.[6] Blue, yellow, and red can be loosely interpreted as good, neutral, and deficient, respectively. Yellow boxes are often accompanied by brief reasons for their neutrality—the reasons for blue or red extremes often being obvious. While some criteria are quantitative, many are subjective. The following 10 properties are assessed for this comparison.

6: The criteria are defined by the author: this method of comparison is not common practice—even if something like it perhaps should be.

- **Abundance**. Not all ideas, however clever or practical, can scale to meet the needs of modern society. Hydroelectric power cannot expand beyond about 10% of current global demand, while the solar potential reaching Earth's surface is easily calculated to exceed this benchmark by a factor of about 5,000. Abundant

© 2022 *T. W. Murphy, Jr.*; *Creative Commons Attribution-NonCommercial 4.0 International Lic.*;
Freely available at: *https://escholarship.org/uc/energy_ambitions*.