Chapter 8

The Second Law of Thermodynamics

One of the books on the reading list is "The Two Cultures", by C.P. Snow. Snow, speaking in 1959, deplored the increasing division of the intellectual world into two cultures, one comprising the pure and the applied scientist, the other the world of the arts, in particular the literary arts. He found this division deplorable because he saw each group being so ignorant of the other that they lack any common language in which dialogue can occur. This ignorance is not entirely symmetrical; most engineers have heard of Dickens, and if they can't name any of his novels, will at least be embarrassed that they can't. The corresponding challenge he proposes for the arts student is to state the Second Law of Thermodynamics, and to explain why it matters. Most artists of his acquaintance failed this simple test.

Unfortunately we were unable to coax any arts students into this course, but I'd nevertheless like to have a go at explaining the Second Law of Thermodynamics to you, so that if you ever find yourself in conversation with an arts student, you can explain it to them. I do not have time to go into the details of who Dickens was, but you should be able to hold your own in most literary discussions by remembering that he's dead, white and male.

Thermodynamics begins, not, as you might expect, with the First Law, but one step *before* that, with the Zeroth Law. Historically, the reason for this is that, after figuring out the First, Second and Third laws, thermodynamicists realised that none of these made sense without a prior law, which they therefore went back and inserted at the beginning. (So future thermodynamicists may find themselves introducing a 'Minus-Oneth Law', and so on back.)

Consider three flasks, each full of a liquid and fitted with a bung and a hollow glass tube penetrating the bung, so that the liquid is visible in the tube above the bung. If a flask gets hotter or colder, the level of liquid in the tube will rise or fall. The Zeroth Law says that, if when I put A and B together, the levels of their tubes stay the same, and when I put B and C together, the levels of their tubes stay the same, then I can confidently predict that, when I put A and C together, the levels of *their* tubes will stay the same. A fancy way of saying this is that the property of thermal equilibrium is transitive.

But even dressed up with such fancy jargon, this doesn't look like much of a law, more like commonsense. How could anyone possibly imagine that *not* being true? Isn't it just saying that, if two things have the same temperature as a third thing, they have the same temperature as each other?... NO, without the Zeroth Law, there's no such thing as temperature.

Well, consider chemical equilibrium instead. Is it true that if two substances are in chemical equilibrium with each other, and the second is in equilibrium with a third, then the first must be in equilibrium with the third? (Hint: let the first be water, the second, oil, and the third, sodium. Wear safety glasses.)

So the Zeroth Law could be false. But it happens to be true. With that out of the way, let's go on to the First Law. Formally, the First Law says that if any system (such as an engine) goes through a cyclic process, the net amount of work crossing the boundary of the system is equal to the net amount of heat crossing the boundary of the system. So, for example, if the system does not exchange any heat with its surroundings, the net amount of work it does must equal the net amount of work that's done on it. Or, to put it more simply still, there's no such thing as a free lunch.

The First Law can also be stated as 'Energy can neither be created nor destroyed'. Although this looks simpler than the formal definition, it turns out that we need the formal definition before we can define energy unambiguously.

Let us temporarily restrict our attention to devices that do not extract heat from their surroundings – for example, the electric motor, the clockwork motor, the fuel cell, and the electric kettle. The First Law tells us that none of these devices can do more work than we put into it.

This has not stopped ingenious inventors from designing perpetual motion machines that purport to produce work out of nothing; a brief search of the Web will turn up a large number of proposed perpetual motion machines. These are now more commonly marketed as 'zero-point energy devices' or 'over-unity machines', several millenia of failure having tarnished the 'perpetual motion' brand name. Curiously, the proponents of these machines are always seeking investors' dollars, though you would think that anyone who could produce free energy would simply sell the energy and get all the income they needed.

Although no perpetual motion machine has ever worked, study of such machines can give us new insights into physics. For example, consider the question, "Can there be a substance that is opaque to magnetism, in the same way that cardboard is opaque to light?" If there were such a substance, we could use it to build a perpetual motion machine: imagine a ball-bearing at the foot of a ramp, and a powerful permanent magnet at the top. We allow the magnet to pull the ball-bearing up the ramp, then interpose a sheet of the hypothetical substance between the ball-bearing and the magnet. The ball-bearing will now roll down the ramp; its energy can be used to do useful work, and we then remove the sheet and repeat the cycle. Such a machine would violate the First Law, so we deduce that no such opaque-to-magnetism substance can exist.

On the other hand, the First Law does not prohibit any machine from being 100% efficient. It's rare for a machine to be 100% efficient in practice – friction tends to turn part of the work output into heat – but this part can be made arbitrarily small. And for an application like the electric kettle, where heat is the output that we want, 100% conversion *can* be attained.

Now consider a device that does extract heat from its surroundings – for example, a device that takes the energy from a beaker of hot water and uses it to raise a weight. The First Law does not rule out the existence of such devices.

If we could build such a device, we would have solved the energy crisis, and also the problem of global warming: we just extract heat from the atmosphere or the oceans and turn it to some form of useful work, such as crushing ore. The work done by the device will eventually end up as heat, due to friction, and that heat can then be recycled through the device. So this machine can run forever, and is therefore known as a perpetual motion machine of the second kind.

The existence of such machines is ruled out by the Second Law of Thermodynamics.

To illustrate the Second Law, consider a particular example of a heat engine, the Stirling Engine. The Stirling can be thought of as a mass of gas contained in a cylinder with a piston. We first heat the gas, then allow it to expand, doing work. We next cool the gas, then compress it, doing some work on it in the process, but not as much as we obtained from its expansion. So the net result of this cycle is that we heat the gas, by putting it into contact with a high-temperature source, obtain some net work from it, then reject the remaining heat by putting the gas in contact with a low-temperature source.

The Second Law says that *any* process that turns heat into work behaves in a similar way to the Stirling engine: it needs *two* fixed temperatures, it adsorbs heat at the higher temperature, turns some into work, and rejects the rest at the lower temperature. And the maximum amount of work you can get out of this process, no matter how you do it, is limited by this equation:

$$Work < HeatIn * (T_h - T_c)/T_h \tag{8.1}$$

The quantity $(T_h - T_c)/T_h$ is known as the *Carnot efficiency*, and the Second Law can be restated as 'No heat engine, operating between reservoirs at T_h and T_c can have higher efficiency than the Carnot efficiency.'

So the first law says, electricity, motion, magnetism, heat, these are all forms of the same thing, energy. The second law says, yes, but not all forms of energy are created equal. To the first law, a swimming pool of lukewarm water looks the same as a gallon of gasoline. To the second law – and to a motorist – one form of energy is useful, the other is not.

You will notice that one surprising result of this analysis is that discovering a new source of coldness would be just as useful as discovering a new source of energy, as it would enable us to turn the low-grade energy of the air and oceans into useful work. And in fact some sources of cold are used in this way: Japan imports refrigerated liquid natural gas, and the importers have investigated using a heat engine to make use of the low temperature of the gas.

There is a way to factor the usefulness of a given source of energy into our analysis. This is simply to multiply the quantity of energy we can get by the Carnot efficiency. We call the resultant quantity the 'availability', or 'exergy'. Thus we can build up a league table.

Notice that chemical energy, such as the energy stored in gasoline, is of the highest quality. But as soon as you burn it, it turns into heat energy. And the temperature you burn it at is limited by the fact that you don't want your engine to melt. The alternative is to convert it to another form of high-quality energy without going through the intermediary of heat, for example, by letting it react in a fuel cell.

Similarly, nuclear energy is of the highest quality. But most existing reactors use that energy to heat water, and that immediately reduces its quality to that characteristic of the water temperature – a few hundred degrees, even with the water being pressurised.

From the point of view of a society trying to balance its energy budget, this looks like bad news. The exergy of a source is always less than, or at best equal to, its energy, so wherever we look, we find we've got less useable energy than we thought. But there is an upside to availability analysis, which is that we don't always need energy of the highest quality. For example, one of our biggest needs is to keep warm... and even if we're addicted to saunas and steam baths, we seldom want to be more than a few tens of degrees hotter than the outside world. How can we best keep warm?

One way is to turn on an electric fire.

This looks pretty good from an energy point of view; 100% of the electrical energy goes into making us warmer, so how can we do better than that? Actually, we can do a lot better than that. We can make use of a heat pump. The heat pump is essentially a Stirling engine run in reverse – we do work to compress a gas, obtaining a *very* hot gas, allow it to reject heat to some hot-temperature sink (for example, the inside of our house), then expand it to obtain a *very* cold gas, cold enough that it will absorb heat even from a cold source, such as the outside of our house during a Yukon winter. Once it's warmed up to Yukon temperature, we compress it again. The net result is that we are expending energy to move heat in the unnatural direction – from cold to hot. The advantage is that we can move several times as many Joules of heat as we expend in Joules of energy.

So for most efficient energy use, in a home or in a province, we need to analyse energy flows using the Second Law. And in analysing supply and demand, we must pay attention both to the amount of energy needed and to the quality needed. From this point of view, one of the inelegancies of our present nuclear power stations is that we're taking high-quality energy – the energy of fission – and immediately degrading it to the temperature of boiling water.

One of the most significant mismatches between supply and demand, from a second-law perspective, is representing the public's need for energy as a need for electrical energy, that is, energy of the best quality. In Canada, the public's biggest energy need is for heating and cooking; both these needs can be satisfied without the use of electricity. For example, most cars in our society are powered by heat engines: we burn the fuel to produce heat, then convert some of the heat to mechanical work. This conversion is necessarily imperfect: we have to reject some energy to a sink below the combustion temperature. This means that, as long as we're using the mechanical energy for driving, we can keep warm for free: we just use the inside of the car as the low-temperature sink. Battery-driven or fuel-cell cars don't suffer from the need to reject part of their energy input; this makes them more efficient, but also means that they need to make separate arrangements for keeping the passengers warm. (*Note: not all engineers would call the internal-combustion engine a heat engine, since after all, to make it run you have to put in gasoline, not heat. On the other hand, the chemical energy of the gasoline does get converted to heat energy as soon as it burns.*)

We have the same relationship in stationary power generation from fossil fuels: the power station must produce a stream of waste heat at some sink temperature, but this heat is often of just the right quality for warming houses and office buildings. There is the drawback that heat cannot be transmitted as effectively as electricity, but, if the power station is adjacent to a town, steam pipes or hot-air ducts can make effective use of the waste heat.

In second-law analysis, we are looking for *irreversibilities*: points at which energy is being converted from a form higher in the league table to a form lower down. Wherever we have an irreversibility, we have a point where efficiency could be increased.

One very common irreversibility is friction, which converts mechanical energy to heat. Ohmic heating converts electrical energy to heat. Any time you mix two fluids at different temperatures you lose availability, since you *could* have run a heat engine between the two temperature extremes.

Thermodynamics is one of the two fields of science developed by engineers. It was developed for practical purposes, before we had our present detailed knowledge of the atomic structure of matter. Suppose we use our knowledge of the structure of matter to take a look at what's going on inside a gas. We see little molecules darting around in all directions. If, without changing their speeds, we were to get them all lined up, we would have a wind. And we could change that wind energy into mechanical energy, via a windmill. The availability of the gas would have increased. The change we've made hasn't affected the energy of the gas molecules, but it has made their motion more orderly. So availability seems to be correlated with *order*. The more ordered a source of energy, the more useful it is.

We call the disorder in a system the *entropy* of the system. The greater the entropy of a system, the less use the energy in the system is. Unlike energy, entropy is not conserved. Instead of a conservation law, we have an inequality: the entropy of a closed system can only increase. This is another way of stating the second law.

There is a strange difference between the First and Second Laws. For the First Law, there is no distinction between past and future; the energy of an isolated system always remains the same. Consider a box containing two gases, separated by a partition. Now remove the partition. The energy of the system remains the same.

Even if we take a microscopic look at the system, we just see perfectly elastic collisions between particles. A film of these collisions could be run backwards, and it would look just the same as it would running forward. But for the Second Law, there is a distinction: the future is the direction of greater entropy, the past is the direction of greater order. We can tell which direction the movie is running in because we see the gases mixing, and we know that gases never unmix spontaneously. The asymmetry in time seems to emerge from the microscopic behaviour, even though that behaviour itself is time-symmetric.

So the Second Law says that, left to themselves, systems are bound to become more disordered – or at best, stay the same. There seems to be one process in nature that does not obey this trend, namely, the process of evolution. This process starts with single-celled creatures like the amoeba, and produces peacocks, orca whales, and us. Isn't this a trend towards greater order, and doesn't it show that there can be exceptions to the Second Law?