

Entropy

Michael Fowler 7/7/07

A Clue from Carnot

We've seen over the last few lectures how the simple kinetic picture of a gas leads to the Gas Law, $PV = nRT$, with the macroscopic variable T now seen as proportional to the average energy of a molecule. Putting heat into a gas simply means increasing the molecules' average kinetic energy (plus rotational energy, etc., for more complicated molecules) and the work done by an expanding gas is just a transfer of energy from this internal energy to whatever the gas is pushing. The First Law of Thermodynamics is nothing but conservation of total energy, now that we know heat is really microscopic kinetic energy for an ideal gas (and rotational energy, etc.). By applying these findings to the Carnot Cycle, we discovered that this reversible ideal engine had efficiency $1 - T_C / T_H$. Furthermore, we found that no heat engine could be more efficient, at least if we granted the truth of the Second Law of Thermodynamics, which states that it is impossible to construct an engine that just takes heat energy out of the air, say, and converts it into work, without any available "cold reservoir" to dump waste heat into.

The incredible thing is that Carnot, in the 1820's, who believed heat was a conserved caloric fluid, and therefore didn't know the First Law (that heat is just microscopic energy, which is *not* conserved by itself, but can be transformed into other forms of energy) correctly found the efficiency of his cycle to be $1 - T_C / T_H$, and then argued—also correctly—that this set the absolute limit on heat engine efficiency! How did he do that? It's worth a brief examination of his ideas, because they give a clue about something we haven't mentioned so far—entropy.


Remember, Carnot saw his cycle as a water wheel, the caloric fluid being the water, the temperature difference being the height difference from the top to the bottom of the waterfall. But he did know one difference from a waterfall: he knew there was an absolute zero, a level below which nothing could "fall". (This was from the well known results on the contraction of a gas with cooling. He took the absolute zero to be -267 Celsius, only a few degrees off the correct value). This meant that one could assign to the caloric fluid an *absolute* value of "potential energy" in this imaginary water wheel scenario: at temperature T , an amount of fluid F would have "potential energy" FT . This is the total possible amount of heat energy in the fluid, and to extract it all you would have to "drop" the fluid all the way down the temperature scale to the absolute zero. But in our real (if idealized) heat engine, it can only drop as far as the lowest temperature available, that of the cold reservoir T_C . Hence the efficiency factor $1 - T_C / T_H$: it's just the ratio of how far the caloric fluid is able to fall in our engine to how far it would have to fall (down from T_H to absolute zero) to give up all its heat energy.

In this picture, the heat energy delivered from the hot reservoir $Q_H = FT_H$, so the amount of actual "caloric fluid" must be $F = Q_H / T_H$. But the fluid is *conserved*, in Carnot's picture, so the waste heat is $Q_C = FT_C$, with the same F , and therefore $Q_H / T_H = Q_C / T_C$. This equation is the important one: even though the *argument* he used is incorrect—there is no conserved caloric fluid—the *equation* is right, and relates directly to the efficiency of his engine: the heat utilized


is $Q_H - Q_C = Q_H (1 - T_C / T_H)$. But how did he know no engine could do better? Ironically, although he didn't know the First Law of Thermodynamics, he *had* figured out the Second Law, and he presented an argument exactly like that given in the last lecture to demonstrate that no engine could outperform his reversible one.

Tracking Heat Changes on Paths in the (P , V) Plane

Let's go back to the Carnot cycle and consider a bit more how the gas gains heat from a reservoir as it follows a path in the (P , V) plane. We begin with the first half of the cycle, from a to c :




This does *not* mean that we can say the gas at (P_c, V_c) has Q_H more heat than the gas at (P_a, V_a) . Why not? Because we could equally well have gone from a to c by a route which is the second half of the Carnot cycle traveled backwards:



This is a perfectly well defined reversible route, ending at the same place, but with quite a different amount of heat supplied!

So we cannot say that a gas at a given (P, V) contains a definite amount of heat. It does of course have a definite internal energy, but that energy can be increased by adding a mix of external work and supplied heat, and the two different routes from a to c have the same total energy supplied to the gas, but with more heat and less work along the top route.

However, notice that one thing *is* the same over the two routes in the diagrams above: the ratio of the heat supplied to the temperature at which it was delivered: $Q_H/T_H = Q_C/T_C$. Of course, we've chosen two particular reversible routes from a to c , but it turns out that for *any* reversible route from a to c , the integral of increments of heat supplied divided by the temperature of delivery, $\int_a^c dQ/T$, is the same! To see how this can be, consider cutting a corner in the previous route:



If we follow the path $ae fgc$ instead of adc , where ef is an isotherm and fg an adiabat, how does that affect $\int_a^c dQ/T$? The answer is it doesn't: look at the little Carnot cycle $efgd$. We've just changed from the bottom route to the top route around this cycle from e to g , so from the previous argument about the original big Carnot cycle, $\int_e^g dQ/T$ is the same.

But we can now cut corners on the corners: any zigzag route from a to c , with the zigs isotherms and the zags adiabats, in other words, any reversible route, can be constructed by adding little Carnot cycles to the original route. In fact, any path you can draw in the plane from a to c can be approximated arbitrarily well by a reversible route made up of little bits of isotherms and adiabats.

The value of $\int_a^c dQ/T$ is evidently the same along any reversible path from a to c : in contrast to $\int_a^c dQ$, the heat change, which was different for different reversible routes.

We can define a *new state variable*, S , such that the difference

$$S(P_2, V_2) = S(P_1, V_1) + \int_{(P_1, V_1)}^{(P_2, V_2)} \frac{dQ}{T}$$

where the integral is understood to be along a reversible path. Given the value of S at a single point (P_1, V_1) , we can find its value everywhere.

The Third Law

Notice the argument above only tells us entropy *difference* between two points, it's a bit like potential energy. Actually, though, there *is* a natural base point: a system at absolute zero temperature has zero entropy. This is sometimes called the *Third Law of Thermodynamics*, or Nernst's Postulate, and can only be really understood with quantum mechanics. We don't need it much for what we're doing here, we only work with entropy differences, but it makes things convenient because we can now write $S(P, V)$ without ambiguity.

Finding the Entropy Difference for an Ideal Gas

In fact, for the ideal gas, we can find the entropy difference exactly! Recall that the internal energy of a monatomic gas, the total kinetic energy of the molecules, is $\frac{3}{2}nRT$ for n moles at temperature T . For a diatomic gas like oxygen, the molecules also have rotational energy because they're spinning, and the total internal energy is then $\frac{5}{2}nRT$ (this is well-confirmed experimentally). The standard way to write this is

$$\text{total internal energy} = C_V T$$

because if the gas is held at constant volume, all ingoing heat becomes internal energy. Note that C_V is for the amount of gas in question, it is *not* specific heat per mole.

Suppose now we add heat dQ , but allow volume variation dV at the same time. Then, for n moles of gas,

$$\begin{aligned} dQ &= C_V dT + PdV \\ &= C_V dT + nRT \frac{dV}{V} \end{aligned}$$

so

$$\int_{(P_1, V_1)}^{(P_2, V_2)} \frac{dQ}{T} = C_V \int_{(P_1, V_1)}^{(P_2, V_2)} \frac{dT}{T} + nR \int_{(P_1, V_1)}^{(P_2, V_2)} \frac{dV}{V}$$

and doing the integrals

$$S(P_2, V_2) - S(P_1, V_1) = C_V \ln \frac{T_2}{T_1} + nR \ln \frac{V_2}{V_1}.$$

So the entropy change depends only on the final T and V , in other words, the final P and V since we always have $PV = nRT$. This is just restating what we've already established: ***S is a state variable***, the (macroscopic) state of the ideal gas is fully determined by P , V (or P , T), and therefore so is S .

(We should perhaps mention that this formula ignores quantum effects, and consequently is incorrect in the absolute limit of $T_2 \rightarrow 0K$, where for example C_V also goes to zero. But these deviations occur at *extremely* low temperatures, and are irrelevant to our considerations here.)

The bottom line is that we can get to the final state any way we want, *the entropy change is the same*, we don't have to go by a reversible route!

The Divided Box: Entropy in Irreversible Change

Now consider the following scenario: we have a box of volume $2V$, which is two cubes of volume V with a face in common, and that face is a thin partition in the box.



At the beginning, n moles of gas are in the left-hand half of the box, a vacuum in the other half. Now, suddenly, the partition is removed. What happens? The molecules will fly into the vacuum, and in short order fill the whole box.

Question: what happens to the temperature of the gas during this expansion?

Answer: Nothing! There's no mechanism for the molecules to lose speed as they fly into the new space.

Question: what happens to the pressure? Can you explain this?

Question: What is the entropy change? It is $nR \ln 2$. But no heat flowed in on this route! If we had followed a *reversible* route, for example moving slowly along an isothermal, letting the partition gradually retreat to one end of the box like a piston in a cylinder, we would have had to supply heat. That supplied energy would have all been used in pushing the cylinder, the gas itself ending up in the same state as just removing the partition. But we would have had a nonzero $\int dQ/T$. Just removing the partition quickly there is no heat transfer, and this action doesn't correspond to *any* path in the P , V plane, since each point in that plane represents a gas in

equilibrium at that P, V . Nevertheless, the initial and final states are well defined (after the partition-removed state has reached equilibrium, which will happen very quickly) and we *have* moved to a state of different entropy, since the entropy depends only on P, V .

Relating Entropy to the “Space” Occupied by the System: Phase Space

So what does the increase in entropy mean in the molecular picture of the gas? The total number of molecules in n moles is nN_A , where N_A is Avogadro’s number, and remember $R = N_A k$ where k is Boltzmann’s constant. So that total entropy increase of $nR \ln 2$ on doubling the volume available to the gas molecules means an entropy increase of $k \ln 2$ per molecule. *So part of the entropy per molecule is evidently Boltzmann’s constant k times the logarithm of how much space the gas is occupying.*

Question: but you could also increase the entropy by keeping the volume fixed, and increasing the temperature—in fact, according to our equation above, for a monatomic gas, so $C_V = \frac{3}{2} nR$, doubling the temperature will increase the entropy by $\frac{3}{2} k \ln 2$ per molecule! How does *that* relate to “more space”?

Remember the lecture on the kinetic theory of gases. To specify the state of the system at a molecular level at some instant of time, we need to say where in the box every molecule is located, *plus* giving its velocity at that instant. We can imagine the gas filling up *two* kinds of space: the ordinary space of the box, but also the gas is, roughly speaking, occupying a sphere in velocity space, really a sphere with a fuzzy edge, defined by the Maxwell distribution, which is spherically symmetrical, dying away as the radius increases. The radius of this sphere in velocity space is the root mean square velocity $v_{rms} = \sqrt{v^2}$. Therefore the volume of the sphere is $\frac{4}{3} \pi v_{rms}^3$, we’ll take this as a measure of how much velocity space the gas is occupying.

Now, if we heat up the gas, the molecules speed up—which means they spread out to take up more volume in the velocity space. If we double the temperature, since $\frac{1}{2} m v_{rms}^2 = \frac{3}{2} kT$, v_{rms}^2 is doubled and the volume of velocity space the gas occupies, measured by the sphere, must go up by a factor of $2^{3/2}$. But $\ln 2^{3/2} = \frac{3}{2} \ln 2$, so the entropy increase $\frac{3}{2} k \ln 2$ on raising the temperature is *exactly* a measure of the increase in velocity space occupied by the gas, that is, *the increase in entropy per molecule is evidently Boltzmann’s constant k times the logarithm of the increase in the space the gas is occupying*, exactly analogous to the simple volume increase for the box.

The bottom line, then, is that *entropy is a logarithmic measure of how much space a system is spread over, where space is both ordinary space and velocity space.*

This total space – ordinary space plus velocity space – is called “**phase space**”.

The *phase space* gets even bigger for diatomic molecules, because we would have to include all their speeds and directions of *rotation* in giving a full account of the state of the gas, but the same approach works here as well.

A More Dynamic Picture

We have pictured the gas at one instant in time in terms of where the molecules are at that instant, both in the box (position space) and in velocity space. But that's only one instant. If we could follow in this detail as time progresses, we will see constant changing around on this microscopic level, the gas will be passing through *immense* numbers of different states. On a macroscopic level we will be oblivious to all this, the gas will appear to fill the box smoothly, and if we could measure the distribution in velocity space, we would find Maxwell's distribution, also smoothly filled. But the entropy is a logarithmic measure of the space, real space and velocity space, the gas is occupying, and therefore a measure of the numbers of different states it is constantly passing through. Over time, we might imagine the gas eventually goes through *all* states consistent with its total energy. (This is actually a very subtle point—we're certainly not proving it here, we just hope it sounds reasonable.) But this raises a question about the gas in the box after the partition is removed. It didn't lose any energy. Therefore, the initial state, the state with all the molecules in the left-hand half, is an allowed state of the gas in the whole box. So, why don't we ever see all the gas piling up in that half, even if only for an instant, if we just keep watching?

The point is, that of all the possible states the gas has at that temperature, and which it's moving through, only a *minute* fraction correspond to all the molecules being in the left-hand half. This is best appreciated by beginning with a very small number of molecules, say ten. Then if we pick a state of the gas at random, it's 50 – 50 that the first molecule is in the left hand half. In fact, we could just pick a state by flipping a coin, and putting the molecule somewhere in the left-hand half if it comes up heads, somewhere on the right if it's tails. The chances that we get ten heads in a row are $(1/2)^{10}$, or 1 in 1,000 (actually 1024, we'll round it off). For a gas of 20 molecules, it's a one in a million chance they'll all be in one half, for 40, it's one in a trillion. For 100 molecules, it's 1 in 10^{30} .

Question: What is it for a million molecules?

But a liter of air has of order 10^{22} molecules! This is never going to happen. So the states having the gas all in the left-hand half of the box are a very tiny fraction of all the states available in the whole box—and the entropy is proportional to the logarithm of the number of available states, so it increases when the partition is removed.

Entropy and “Disorder”

You have probably seen it stated that entropy is a measure of “disorder”. We have defined it above in terms of the amount of phase space the system is moving around in. These are not really different concepts. If we pinned a little label on the molecules, *l* or *r* depending on whether it was in the left or the right hand side of the box, the original state had all the labels *l*. The final state has a random selection of *l*'s and *r*'s, very close to equal in numbers.

Another way to say the same thing is to visualize an initial state before removing the partition where the left hand half has “red” molecules, the right hand side “green” molecules (we're making up these names for visualization purposes, they could be any two gases). If the gases are ideal (or close enough) after removal of the partition both will have twice the volume available, so for each gas the entropy will increase just as for the gas expanding into a vacuum. Thinking

about what the state will look like a little later, it is clear that it will be completely disordered, in the sense that the two will be thoroughly mixed, and they are not going to spontaneously separate out to the original state.

So, when a system previously restricted to a space is allowed to move into a larger space, the entropy increases in an irreversible way, irreversible meaning that the system will never return to its previous restricted configuration. Statements like this are always to be understood as referring to systems with a large number of particles. Thermodynamics is a science about large numbers of particles. Irreversibility isn't true if there are only ten molecules. And, for a large number, by "never" we mean it's extremely unlikely to happen even over a *very* long period of time, for example a trillion times the age of the Universe.

Everyday Examples of Irreversible Processes

Almost everything that happens in everyday life is, thermodynamically speaking, an irreversible process, and increases total entropy. Consider heat ΔQ flowing by direct thermal contact from a hot body at T_H to a cold one at T_C . The hot body loses entropy $\Delta Q/T_H$, the cold one gains entropy $\Delta Q/T_C$, so the net entropy increase as a result is $\Delta Q/T_C - \Delta Q/T_H$, a positive quantity. The heat flow in the Carnot cycle did *not* increase total entropy, since it was isothermal—the reservoir supplying the heat was at the same temperature as the gas being heated (an idealization—we do need a slight temperature difference for the heat to flow).

Friction, or an inelastic collision such as dropping a ball of clay on the floor, converts kinetic energy to heat—so there is a nonzero ΔQ of heat entering some body, entropy is being generated. Since heat involves random motions of the particles involved, this is sometimes stated as a transformation of ordered energy (a falling ball, with the molecules all moving in the same direction) into disordered motion. Yet another phrasing is that potentially useful energy is being degraded into less useful energy: the falling ball's kinetic energy could have been utilized in principle by having it caught by a lever arm which could then do work. The slight heat in the floor is far less useful energy, some of it could theoretically be used by a heat engine but the temperature difference to the rest of the floor will be very small, so the engine will be very inefficient!

What about Quantum Mechanics?

The whole analysis of the ideal gas above is based on specifying the position and velocity of each molecule at some instant of time. But if you've taken any quantum mechanics, you know this can't be done—the Uncertainty Principle says there are limits to how precisely these can be known. So are all our arguments on a dubious foundation? The answer is yes and no. Although it is quite true that we can't specify position and velocity to arbitrary precision, quantum mechanics only allows each particle to be in certain quantum states, and this turns out to give a much cleaner basis for describing entropy: it is k times the logarithm of the total number of quantum states of the whole system the macroscopic gas state corresponds to. At absolute zero, any quantum system goes into its ground state, a single quantum state, so this formulation naturally gives the entropy of any system to be zero at absolute zero temperature. Furthermore, for ordinary gases at room temperature, the uncertainty principle imposes rather weak requirements, so our analysis above is fine. And, quantum mechanics explains why the diatomic molecule cannot spin around its own axis, but only has the two degrees of rotational freedom

corresponding to the two orthogonal directions. That's because quantum mechanics allows angular momentum to only have certain values (multiples of Planck's constant) so a direction with a small moment of inertia would have to have very high energy to get one unit of angular momentum, and the energy isn't available at ordinary temperatures.