Radiation Part 3: Equilibrium and Detailed Balance

In these notes we take advantage of our experience thus far in the class, and our current discussions of radiation, to explore concepts that are powerful, general, but surprisingly complex: equilibrium and detailed balance. I hope that when you read this set of notes you'll have a better appreciation for why equilibrium is such a useful concept when it can be used, but also why you need to check various types of equilibrium before you make any assumptions.

1. Equilibrium

As I hope to convey in these notes, if you can assume that a system is in equilibrium in some way it really makes your life a lot easier! But first we have to establish that equilibrium is a reasonable assumption in a given circumstance. We'll start with an example.

Is a rock sitting on a table in equilibrium?

As you ponder this question you might realize that the question isn't well-posed. Equilibrium with respect to what? Also, I hope that you've come far enough in the class to realize that even if we specify precisely what we mean (say, thermal equilibrium), the question we might really want to ask is "how far away from ideal equilibrium is my rock and does it matter for the analysis or observation that I'm doing?" So we'll refine our question:

Is a rock sitting on a table close enough to thermal equilibrium with its surroundings that for our purposes we can treat it as being in perfect thermal equilibrium?

Okay, now maybe we can make some progress! First, as a review, in a closed and isolated system (i.e., one that is self-contained and has no contact with anything else), given enough time everything in the system will come to the same temperature. So let's specify our problem further: the room containing the rock on the table is closed and isolated. No molecules and no energy enter or leave this system. Then, if the rock is initially hotter than the table, air, room, etc. it will cool down until it approaches equilibrium, and if the rock is initially cooler it will warm up.

Stated this way it could occur to you that equilibrium is a matter of time, and that's right. If we pluck the rock from a blast furnace and put it on the table (poor table!), one second later it certainly won't have reached the same temperature as the room (assuming that the room is at room temperature). But if we let everything sit for a week and then measure it, we'd expect the rock to be at pretty close to the same temperature as its environment. Similarly, if the rock had been immersed in liquid nitrogen just prior to our measurement on the abused table, it wouldn't have had time to warm up, but after a week it should be pretty close.

This gives us an idea for how we could at least formulate a better statement about equilibrium: "*A system is within some tolerance of equilibrium if it has been isolated for long enough to get to that tolerance.*"

But what do we mean here by “tolerance”? Going back to our rock temperature example, suppose that for our purposes it’s close enough to thermal equilibrium if its temperature differs from the environmental temperature by no more than 1 K. Then, we could ask about the detailed properties of the system: how rapidly does heat flow from hot to cool regions, and how long has the system been standing in isolation? If the required time is much less than the actual time, then we can say that for our purposes the system is indeed in equilibrium. That’s what I did when I compared temperature measurements one second after the rock had been in a furnace or liquid nitrogen (when that time of one second is certainly much less than the time needed to get within 1 K of the environmental temperature) to temperature measurements one week later (when that time of one week is much larger than the needed time).

I hope that it’s clear that whether a system is “close enough” to equilibrium depends on our purpose! If 1 K is enough then a week should be plenty. If 10^{-9} K is needed then it would be a very long time indeed (and at that level the random fluctuations of air molecules might not allow the rock to all have a temperature within 10^{-9} K of the average). But this, at least, is the general algorithm for determining whether your system is close enough to equilibrium, in some sense, for your purposes:

1. Determine how close you need to be to equilibrium to assume equilibrium for your purposes.
2. Determine how long that will take.
3. Determine how long you have; if much longer than the needed time, the system is close enough to equilibrium. If much shorter, the system might not be in equilibrium. Here it’s only “might not” because maybe by great good luck your system started out close enough to equilibrium for your purposes.

This last point allows us to answer a question that you might have at this stage: given all of these caveats, why should we ever assume equilibrium? Why not just treat the system (e.g., the rock) for real, and determine how it evolves?

The answer is that if you can assume that a system is in equilibrium, then at least for the property we care about (temperature in this case) we don’t need to worry about the history of the system or any details. For our rock this means that if we wait a week its state will not depend on whether it started hot, or cold, or half hot and half cold, or any such combination. This independence of history allows us to make profound simplifications, but you can’t assume it; you need to make sure that equilibrium really holds to the required degree.

Still, you might wonder: given that we required all of this “closed and isolated” nonsense and no system in the universe is closed and isolated to a mathematically rigorous level, why are we ever justified in assuming equilibrium?

Here, again, it’s a question about the level of the approximation. The gas deep inside the Sun is pretty close to equilibrium. That statement may seem odd given that the center of the Sun is at about 15 million Kelvin and the photosphere of the Sun is at about 6000 Kelvin. Obviously that's a big difference! True thermal equilibrium would have to mean that the whole Sun is at the same temperature, and clearly it isn't. So how can we justify the approximation that gas at some point inside the Sun is close to equilibrium?

Here we can realize that the photons in the gas can't go very far before they scatter or absorb on something, so the portion of the gas and photons that really interacts with a given bit of gas is quite limited. In round numbers, say that the Sun is 10^9 m in radius and that the total temperature difference is about 10^7 K. That's an average gradient of about 10^{-2} K m^{-1}. Now, somewhere average in the Sun maybe a photon travels about 10^{-2} m on average before it scatters or absorbs. Over that distance, the temperature changes by only about 10^{-4} K. Even near the photosphere, the temperature is around 10^4 K (remember that we're using round numbers). This means that over a volume of gas and photons that is in close contact, the temperature only changes by a fractional amount $\sim 10^{-4}$ K$/10^4$ K $= 10^{-8}$. That's pretty small, so for some purposes we can ignore it. But for other purposes we can't: for example, it is this small temperature gradient that allows energy to flow from the interior of the Sun to the photosphere and from there to the Earth and the rest of the universe. Thus if we are interested in how energy is transported then this tiny change in temperature is critical; but if we're interested in the distribution of speeds of atoms and electrons, we can easily ignore the slight deviation from equilibrium.

The same is true for approximate equilibrium over time. Consider the case of Jupiter. Jupiter radiates a few times as much energy per time as it gets from the Sun; as we discussed earlier, this is because it is contracting gradually and is therefore radiating gravitational potential energy. As it contracts, the state of matter in its interior changes to accommodate the contraction. For example, the pressure at Jupiter's center is increasing because Jupiter is becoming more compact. As another example of a long-term change, Jupiter's mass is increasing with time because of impacts with comets and dust. These changes happen over billions of years, so over a single year the fractional changes are tiny. If, therefore, you care only about short time scales then you can ignore the long-term changes.

Finally, we note that there are many types of equilibrium. Returning to our rock, we could also ask: is the rock in gravitational equilibrium? If it's on a table, the answer is no: the rock would have lower gravitational energy on the floor, so its position on the table means that it is not in gravitational equilibrium. But if the rock isn't too heavy then it won't punch through the table to rest on the floor. Thus as long as the rock is in local equilibrium (e.g., balanced stably and not on a point) then we can effectively consider that the rock is in gravitational equilibrium.

Similarly, the rock is not in nuclear equilibrium. If it were, it would be made entirely of Iron-56, which has the greatest binding energy per neutron or proton of any isotope. Instead, it is made of elements such as silicon and oxygen. In principle, if we could just get those elements to fuse we could reach nuclear equilibrium. But that would require enormous temperatures, which we don't have in a typical room! Thus although the rock is not in nuclear equilibrium, at least its state is steady.

To sum up this part: if you can assume equilibrium you can make some tremendous simplifications that will ease your life and give you great insight. Many times you *can* make that assumption. But you should always check!

2. Detailed Balance

Our discussion in the previous section focused on a macroscopic thing: a rock. A colloquial definition of equilibrium in that case is that the property you care about (thermal, chemical, nuclear, ...) doesn't change by more than you care about on a time you care about :). So if the rock maintains effectively the same temperature for a second, an hour, a week, or whatever, you can say that it's at thermal equilibrium.

But there is another scale of interest where you have to take a different approach. It doesn't make sense to say that a single atom is in equilibrium. Even if the overall system is somehow equilibrated, the atom you care about might be ionized, or might recombine. In that case we need to think about the statistical distribution of many atoms and require that their overall, average characteristics remain the same over some time.

In fact, you can even go beyond that and use a concept developed by Boltzmann, Maxwell, Einstein, and others: the concept of *detailed balance.* The definition is:

Detailed balance: in equilibrium, all microscopic processes and their inverses are in balance with each other.

As our example, let's consider a hydrogen atom. It has a proton and an electron. In a given system, it could be that some photons have enough energy to kick the electron out of the atom, so that now we have a lone proton and some lone electron flying away. But the inverse process also exists. It is possible that a lone electron can go past a lone proton, radiate a photon, and be captured by the proton to form a hydrogen atom. The principle of detailed balance means that in equilibrium the rate at which electrons are kicked out of hydrogen atoms due to the absorption of a photon is exactly the same as the rate at which electrons are captured by protons to make hydrogen atoms, emitting photons as a result.

This statement is much more powerful than simple equilibrium. Say that you have two processes A and B whose inverses are $\bar{A}$ and $\bar{B}$. Detailed balance says that A balances $\bar{A}$ and B balances $\bar{B}$. It isn't immediately obvious that this has to be true. You might imagine, for example, that A and $\bar{A}$ don't balance each other but that A, B, $\bar{A}$ and $\bar{B}$ together all form some sort of balance. For example, another way to kick an electron out of a hydrogen atom is for the atom to bump into another atom with sufficient energy, so maybe that *plus* photons kicking electrons out of atoms are in balance with electrons and protons combining? Nope. In equilibrium, each of these processes are balanced with their own inverses.

Let's be clear about what we mean by "in balance". Suppose we have a gas in which 99 out of every 100 atoms is neutral, i.e., has a proton and an electron. Thus in 1 out of every 100 cases you have a separate proton and a separate electron. Over some stretch in time, detailed balance means that a *given* neutral atom has 1/99th of the probability of having an electron kicked out, as a proton and an electron have of combining. We can see that by imagining that we have 100 total systems, 99 neutral and 1 separate. If the one separate system combines, that removes 1 separate system and adds 1 neutral system. If just 1 of the 99 neutral systems separates, that adds 1 separate system and removes 1 neutral system. Thus detailed balance does *not* mean that over a given time the probability of separating a neutral atom *per neutral atom* must be equal to the probability of combining a separate proton and electron *per specific proton-electron pair*. The balance in the whole system has to be maintained.

Another application of this in the real world is to a glass of liquid water in air, with both the water and air in a sealed container. At any given moment, some fraction of the water molecules will evaporate from the liquid to go into the air. If the air initially had no water molecules at all, then this evaporation would just represent a net loss to the water. But as the evaporation proceeds, the air starts to have some concentration of water molecules. Those water molecules have some probability of rejoining the liquid water. Detailed balance would be achieved when the rate at which water molecules evaporate from the liquid is balanced by the rate at which water vapor rejoins the liquid.

This principle has been used in many applications in physics, from statistical physics and thermodynamics to the theory of gases to Einstein's development of the ideas that led to lasers(!). It also has been used in many statistical applications. It's good to understand this as a microscopic basis for many types of equilibrium.

Practice problems

1. The excess luminosity radiated by Jupiter, beyond what it absorbs for the Sun, is about $L_{\rm excess} \approx 4.5 \times 10^{17}$ Watts. If we assume that this energy comes from gravitational binding energy, and that for Jupiter that energy is roughly $E_{\rm grav} \approx \frac{GM_J^2}{R_J}$ (where $M_J = 1.9 \times 10^{27}$ kg is Jupiter's mass and $R_J = 7.15 \times 10^7$ m is Jupiter's radius), then how long, in years, would it take Jupiter to radiate away a total energy equal to its current binding energy? What is a reasonable time to which we should compare our answer to determine whether Jupiter is roughly in thermal equilibrium?

Answer: Putting in the numbers we get $E_{\rm grav} \approx 3.4 \times 10^{36}$ J. The time needed for Jupiter to radiate this energy is $E_{\rm grav}/L_{\rm excess} \approx 7.6 \times 10^{18}$ seconds. There are about 3.16×10^7 seconds in a year, so this turns out to be about 2.4×10^{11} years, or about 240 billion years. A reasonable time to which we can compare this is the age of the Solar System, or about 4.6 billion years. Because 240 billion is much larger than 4.6 billion we can say that Jupiter is roughly in thermal equilibrium even over the age of the Solar System. Over shorter times, of course, equilibrium is an even better assumption.

2. Do the same calculation for the Sun: $L_\odot \approx 3.8 \times 10^{26}$ Watts, $R_\odot \approx 7 \times 10^8$ m, $M_\odot \approx 2 \times 10^{30}$ kg. Interpret your answer; what does it mean for the mechanism by which most of the Sun's energy is produced?

3. You have 1000 fair coins. Initially, 200 are heads and 800 are tails. Show explicitly that if you flip all the heads separately (and each flip results, randomly, in half becoming heads and half becoming tails), and flip all the tails, that you *do not* expect to maintain the initial distribution. That is, the initial setup was not in detailed balance.

4. You have 1000 fair coins. Initially, 500 are heads and 500 are tails. Show explicitly that if you flip all the heads separately (and each flip results, randomly, in half becoming heads and half becoming tails), and flip all the tails, that you *do* expect to maintain the initial distribution. That is, the initial setup *was* in detailed balance.

5. Use the general ideas of detailed balance to explain why on a hot, dry day you cool off more easily, and dehydrate more quickly, than on an equally hot but humid day.