Convection

Having examined energy transport by radiation and by conduction, we now get to look at the final entry on the list: convection. Convection is important in some parts of most stars, and its very basics are simple. Unfortunately, in detail convection is so complicated that not only has no one developed a clear analytical theory, but it is in practice impossible to do everything correctly even with computers! Here, we give the qualitative aspects of convection; please see our textbook for a much more detailed treatment.

The general process of convection has some fluid element rising or falling and dissipating its energy as it does so. Let's begin our consideration of convection by thinking about a spherically symmetric star which is everywhere an ideal gas but with temperature and density that vary with radius. In an ideal gas, the pressure P, number density n, and temperature T are related by $P = nkT$. Now suppose that some element of that gas has slightly higher temperature than the average of its environment. If n is equal to or higher than average as well, then P is higher than average. Higher pressure would mean that the element would expand against its environment, and do so quickly: it would happen at the speed of sound. We will, therefore, assume that even if the temperature is higher than average, the fluid will adjust so that the pressure is the same as average. Thus, since T is high and k is a constant, n must be lower than average.

The lower density of the fluid means that it is buoyant, like a helium balloon in air: basically, the pressure gradients around it would balance the weight of something of average density, so since this fluid element has less than average density and thus less than average weight, the net force is upward. Therefore, the first thing that the fluid element does is to rise.

Now what? If it rises, then it goes to a region of lower pressure, which means that the element expands until it establishes equilibrium with the new pressure. Expansion also cools the element, so both the cooling and the lowering of the density have to be taken into account when we compute its equilibrium state. One possibility is that after reaching pressure equilibrium the element is now denser than the average density at the new height. In that case, the element sinks back down. But if after reaching equilibrium the element has a lower density than the average at the new height, then it keeps rising.

Now let's think about what this means for energy transport. Suppose that the fluid element rises much more slowly than the speed of sound, so that pressure balance is maintained, but much more rapidly than the time necessary to have heat leak out of (or into) the fluid element. Then the total heat in the fluid element is conserved, and if we can ignore viscosity (which we usually can for this purpose), it means that the element moves adiabatically. This means that the entropy is conserved, so that the temperature gradient is fixed for a given pressure gradient. Call this gradient $\nabla T_{\rm ad}$. Thus, for a given fluid element with an initially small perturbation, we know how its temperature will change as it rises.

Given this, what is the condition on the gradient of the temperature of the surrounding medium such that the fluid element will continue to rise once perturbed upwards? Since the surrounding temperature has to continue to be smaller than the temperature of the fluid element (why? Because the element has to be buoyant, i.e., its density has to be lower than that of the surroundings, and if $P = nkT$ is the same as the surroundings then T has to be higher than the surrounding), the temperature of the environment has to drop with increasing height faster than the temperature of the element drops. Therefore, the temperature gradient ∇T must satisfy $\nabla T > \nabla T_{\rm ad}$. This is the Schwarzschild criterion for convection, brought to you by the same person who came up with the Schwarzschild spacetime for uncharged, nonrotating black holes.

If the Schwarzschild criterion is satisfied, it implies that the fluid has large-scale motion. That is, the fluid isn't just individual particles that move for one (tiny!) mean free path, but it can have elements that move large distances. In addition to hot parcels of fluid rising, cold parcels will drop, so there is a net transfer of energy from the hot stuff below to the cold stuff above. Note that the macroscopic motions mean that concepts such as mean free path are a bit dicey, compared to their rather clearer analogues for radiative or conductive energy transport. There are also complicating effects in reality. For example, heat really can leak into or out of the fluid element as it rises or sinks, and compositional gradients can also complicate things. Before addressing issues like that, though, let's hold on a bit and develop our general intuition about convection some more.

People frequently determine whether a given layer of fluid is convectively unstable by calculating its structure without including convection, then applying the Schwarzschild criterion or an alternative. Let's assume that for the layers of interest, all the energy was generated at a much deeper layer. Given that large temperature gradients are needed for convection, what does this imply about other opacities? The amount of energy per time to be transported is fixed by the deeper layers (where energy is generated by nuclear fusion in stars), so large temperature gradients would have to mean high opacities (why is that?). With this in mind, should we expect white dwarfs or neutron stars to have convective layers under ordinary conditions? No, because for these stars conduction is extremely effective, so the temperature gradient is tiny and convection is suppressed. Only when other forms of energy transport are ineffective does convection come to the rescue. This, by the way, is one reason why we can often ignore the diffusion of energy out of our fluid elements: for convection to be important at all, this diffusion needs to be relatively weak. In main sequence stars convection tends to be most important in the somewhat cooler layers, where line and edge opacities are important. Low-mass stars have large convective zones for this reason.

One interesting consequence of this is that because of all the roiling motion of the gas, magnetic field lines get tangled and can be amplified, in a process called a magnetic dynamo. Thus low-mass stars have relatively strong magnetic fields and as a result they have overwhelming magnetic flares. For that reason, life around a low-mass star would be extra challenging because it would have to deal with major surges of energy that come along pretty frequently. Never count out life, but you'll probably want to look elsewhere for a vacation home! In contrast, higher-mass stars have fewer sources of opacity and therefore have relatively weaker fields.

So how do we put this in to a description of energy transport? The most widespread model is mixing length theory (MLT). In mixing length theory the assumption is that the fluid element rises some distance ℓ, then releases its energy (i.e., comes into thermal equilibrium with the new environment). The reverse process can also happen: a cool fluid element, which is denser than average because of pressure balance, sinks some distance ℓ and gains energy from the environment. In both cases, energy is transported outward. Of course, the fluid element can "leak" energy on its way up or down as well.

This is the core of mixing length theory. The amount of energy that it transported depends on ℓ, the buoyant velocity w, the heat content of the star as a function of depth, and how much energy drifts out of the elements as they rise or fall. These quantities can be computed at a single radius, making MLT a "local" theory and therefore useful for computation. But let's not forget that this is a crude approximation rather than being the real thing...

Assumptions of MLT

These assumptions constitute the "Boussinesq" approximation:

1. A readily identifiable fluid element has a dimension comparable to ℓ.

2. The mixing length is much shorter than any other scale length associated with the star (e.g., pressure, temperature scale heights). This assumption is violated for the sun, where the pressure scale height is comparable to what is needed for ℓ to explain the inferred convective transport!

3. The fluid element is always in pressure equilibrium with its surroundings. This means, for example, that ℓ/w is much shorter than ascent or descent times.

4. Acoustical phenomena, shocks, etc., may be ignored.

5. The density and temperature of the fluid element deviate only slightly from the environment.

Together, these assumptions mean that the fluid is almost incompressible and that density and temperature variations are small. Note that assumption 2 is often violated, and that assumption 1 is troublesome by itself: if the fluid element is the same size as ℓ, it is tough to picture the slow drift and leakage of energy!

Another problem is that lab experiments can't simulate the conditions in stars, in the sense that some of the dimensionless numbers that characterize the fluid are a factor of a billion or so different between stars and laboratories. However, note also that this is a common situation in astrophysics. What you have to do in these circumstances is do the best you can, make reasonable assumptions, and solve a simpler problem. It would, of course, make no sense to go wild with this and end up with a finely refined variant of MLT when you know that some of the initial assumptions are suspect!

Suppose now that convection occurs. What does this do to the temperature gradient? It tends to reduce the temperature gradient, by putting hot fluid higher and cold fluid lower. If we were to try to incorporate convection into the same general opacity scheme as we used for radiation and conduction, would we include it linearly or harmonically? Harmonically, because as another channel for energy transport it increases the flux and decreases the effective opacity.

Suppose that convection is extremely efficient. To a first approximation, what is the temperature gradient you'd end up with after convection had changed the gradient? It will be very close to, but slightly larger than, the adiabatic gradient. It has to be larger to continue convection, but if convection is efficient then ∇T doesn't have to be much larger than $\nabla T_{\rm ad}$ to transport lots of energy.

The flux due to efficient convection thus turns out to be

$$F_{\rm conv} = \rho w c_P \Delta T \ . \tag{1}$$

Here a parcel of energy is typically going at a speed w, the typical temperature contrast with the surroundings is ΔT, ρ is the density of fluid and c_P is the specific heat at constant pressure (energy per mass per temperature). How can we test if this is reasonable? We can verify that the units work. It makes sense that if there is no temperature contrast ($\Delta T = 0$) there will be no flux. The higher the speed the larger the flux, which is what the equation says. If the specific heat is low, then little energy can be transported, so F should depend directly on c_P as it does. Finally, if there is more matter (high ρ), more flux is transported.

At this point you might be, and indeed I hope you are, dissatisfied. Mixing-length theory was developed in the 1950s. But we have advanced decades in our understanding, and have computers undreamt of in that time. Why not just do it right!

To address that question, I need to expose you to the full horror of the Navier-Stokes equation:

$$\rho\left[\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v}\cdot\nabla)\mathbf{v}\right] = -\nabla P + \eta\nabla^2\mathbf{v} + \left(\zeta + \frac{1}{3}\eta\right)\nabla(\nabla\cdot\mathbf{v})\ . \tag{2}$$

Feeling better? No? Maybe it will help if I tell you that this is a gussied-up version of $\mathbf{F} = m\mathbf{a}$ in fluids, where ρ is the density, $\mathbf{v}$ is the velocity, P is the pressure, and η and ζ are related to the shear and volume viscosity of the fluid. Didn't help? Sorry about that. In any case, not only do you have to solve this as well as the equations for heat flow into or out of the fluid element, but the biggest issue here is that the Navier-Stokes equation is nonlinear: the $(\mathbf{v}\cdot\nabla)\mathbf{v}$ term is like the square of the velocity.

The reason this is a big deal is that for nonlinear equations you cannot add two solutions together and get another solution. Let's contrast that with a linear system such as Newtonian gravity. Say that you have a test particle and you want to know its gravitational acceleration in a system such as the Solar System, which has many separate gravitating bodies. All you need to do is to figure out the acceleration due to one body (say, the Sun) and add to that the acceleration from the next body (say, Jupiter), and so on. This leads to a lot of nice mathematical properties and makes it relatively easy to solve on a computer (although challenges always exist).

But for a nonlinear system, it doesn't work that way. General relativity, for example, is nonlinear, although for weak gravity you can linearize the system of equations (you have to be able to do this, since Newtonian gravity is the weak-gravity limit of general relativity). If you have a test particle near two black holes, you can't figure out its total acceleration by adding the acceleration that would be due to one black hole by itself, to the acceleration from the other black hole by itself.

General fluid dynamics, which we would have to solve to get a first-principles treatment of convection, is nonlinear. This leads to a lot of very complicated phenomena, including turbulence, which are important in astrophysics but can't be treated exactly. We'll close this lecture with an apocryphal quote from Werner Heisenberg: "When I meet God, I will ask him two questions. Why quantum mechanics? And why turbulence? I really think he'll have an answer for the first."