Energy transfer and opacities

Perspective on final goal: As a reminder, we've talked about two aspects of the microphysics: energy generation and the equation of state. The third major bit of microphysics is energy transfer, and once we have that we can in principle construct a model of a star in equilibrium.

What are some of the ways in which energy can be transferred? From the particle perspective, it can be transferred via photons, electrons, or neutrinos (why not ions? Because they're too slow). Neutrinos have such a small cross section that they essentially just act as energy sinks most of the time. How do we compute the optical depth to neutrinos from the center to the surface of the Sun? We can make a rough estimate using $\tau \approx \bar{n}\sigma_\nu R$, where σ_ν for electron-neutrino scattering is about 10^{-44} cm^2 at an energy of about 1 MeV. At $\bar{n} \approx 10^{24}$ cm^{-3} and $R \sim 10^{11}$ cm, the optical depth is only $\sim 10^{-9}$, so we can just consider any energy in neutrinos as a lost cause for normal stars. In supernovae, the optical depth can be a lot higher; in the region where the supernova shock tends to stall in numerical simulations, the optical depth to neutrinos is about 10^{-2}, but because in supernovae neutrinos contain $\sim 100\times$ the energy that photons or kinetic energy do, this is a major deal. For ordinary stars, though, we'll just ignore neutrinos.

The next question is: how do we determine, qualitatively, whether photons or electrons are more important for energy transfer? What criterion should we use? Very roughly, the idea is that energy transfer occurs when the particle of interest comes from a high-temperature environment and goes to a low-temperature environment, or the reverse. Thus, there are two things that come in: (1) how great an energy density gradient is typically sampled by the particle, and (2) how rapidly does it carry the energy back and forth? In the interiors of most stars, the cross section for photon scattering is less than for electron interactions, and of course photons move faster than electrons, so photons travel farther, faster, and typically dominate energy transfer there. A caveat is that convection (bulk motion of matter) can be important in parts of many stars, but we'll treat that later. However, does photon transfer dominate energy transfer in our normal environment on the Earth? No. A big reason is that for a given temperature T, the energy density in photons in thermal equilibrium goes like T^4, whereas the energy density in ions and electrons goes more like T, so at low temperatures the radiation energy density is small. Therefore, for cool things radiation transport isn't as important as convective or conductive transport of energy.

In this and the next couple of classes, we'll talk about various interactions, and will phrase them in terms of cross sections and/or opacities. These are all related to energy transfer, so here are some reminders from our previous discussions:

- The *cross section* is the "effective area" of a single particle. Equivalently, it's the number of interactions per second divided by the flux of incoming particles (number per area per second).
- The *opacity* is the "effective area" per gram of material. This is *not* equivalent to cross section. For example, the Thomson cross section for scattering of a photon off a free electron is about 10^{-24} cm^2. However, the Thomson scattering opacity also depends on the number of free electrons; the Thomson opacity of neutral material is zero, even though the Thomson cross section is still the same as it was.
- For a fixed channel of energy transfer, *highest* opacity dominates. For example, suppose we are interested in the opacity to photons of a given frequency. Then photons will typically (but not always) interact first via the highest opacity process. If, e.g., there is a line at that frequency, the photons will interact quickly. This means that for a fixed channel, opacities add *linearly*.
- For independent channels of energy transfer, *lowest* opacity dominates. For example, consider the "average" opacity over a full blackbody spectrum. At energies with low opacity, photons can travel easily. So, photons tend to diffuse into those energies away from energies where the opacity is high. That's why in optically thick material, lines are dark. Another example is conduction versus radiative energy transport. When energy is easily transferred via conduction, that's what determines the overall energy transfer rate. When energy is easily transferred via radiation, that's what dominates. This means that for independent channels, opacities add *harmonically*, that is, $1/\kappa = \sum(1/\kappa_i)$ for different channels i, weighted appropriately.

Those last two concepts (fixed versus independent channels of energy transfer) have proven over the years to be very difficult for students to absorb, for reasons that I haven't quite been able to figure out. Let's try this. Suppose that you have two different sources of opacity. Make one of them infinitely opaque, while you leave the other unchanged. Can energy get through? If it can, then the channels are independent and the opacities add harmonically, like resistors in parallel. If it cannot, then the channels are fixed and the opacities add linearly, like resistors in series.

Consider, for example, a photon at a single energy and polarization. Say that it can interact in various ways (e.g., scattering or absorption). If the scattering or absorption opacity is infinite, then that photon interacts immediately. As a result, the channel is fixed and the opacities add linearly. But now consider energy transfer by radiation and by conduction. If photons are completely blocked from moving anywhere, then the radiative opacity is infinite. But if electrons can still move and carry energy, then the conductive opacity is finite, and energy *does* get through. Thus radiation and conduction are independent channels of energy transfer, and their opacities thus add harmonically.

It may help to realize that energy can be transferred back and forth between channels. For example, radiation and electrons/particles exchange energy with each other. If the photons are stopped dead, they are interacting with electrons and thus their energy can flow to the electrons and be conducted elsewhere. If you consider photons with a single energy, and if you further suppose that at that energy the photons can't go anywhere (e.g., maybe they're in the middle of a strong spectral line), then you can also note that absorption followed by re-emisison, Doppler shifts, and other effects can eventually change the energy of the photon, and if the energy wanders to a less opaque region of the spectrum then it can transport energy more easily.

Radiation in vacuum: For a start, let's think about radiation when there is no matter present. In particular, consider a bundle of rays moving through space. Very generally, Liouville's theorem says that the phase space density, that is, the number per (distance-momentum)^3 (i.e., the distribution function), is conserved. For photons, this turns out to mean that if we define the "specific intensity" I_ν as energy per everything:

$$I_\nu = \frac{dE}{dA\,dt\,d\Omega\,d\nu} \ , \tag{1}$$

then the quantity I_ν/ν^3 is conserved in free space. The source of the possible frequency change could be anything: cosmological expansion, gravitational redshift, Doppler shifts, or whatever. The integral of the specific intensity over frequency, $I = \int I_\nu\,d\nu$, is proportional to ν^4.

Because this phrasing of things can cause confusion, let's put this another way. Suppose that when you get right up to the source, I_ν at some ν_0 is $I_\nu(\nu_0)$. The light at that frequency travels through vacuum, and when it gets to us its frequency is now ν_0'. The specific intensity we measure at ν_0' is now

$$I_\nu'(\nu_0') = I_\nu(\nu_0)(\nu_0'/\nu_0)^3 \ . \tag{2}$$

Thus if there is no net redshift or blueshift, so that $\nu_0' = \nu_0$, then $I_\nu'(\nu_0) = I_\nu(\nu_0)$. If we have the same circumstance but we integrate over the entire spectrum to get I, then we observe $I' = I(\nu_0'/\nu_0)^4$.

One application is to the surface brightness. This is defined as flux per solid angle, so if we use S for the surface brightness, then $S = I$. How does surface brightness depend on distance from the source, if ν is constant? It is independent of distance (we can also show this geometrically). However, how does the surface brightness of a galaxy at a redshift z compare with that of a similar galaxy nearby, assuming no absorption or scattering along the way? The frequency drops by a factor $1+z$, so the surface brightness drops by $(1+z)^4$. Note that in a given waveband, the observed surface brightness also depends on the spectrum, because at a fixed range of observed wavelengths the range of source wavelengths depends on the redshift (the correction for this is called the K-correction, because why should we name things in ways that make sense?).

Another application is to gravitational lensing. Suppose you have a distant galaxy which we would normally measure to have a certain flux. Gravitational lensing, which does not change the frequency of the light, splits the image into two images. One of those images has twice the flux of the unlensed galaxy. Assume no absorption or scattering. How large would that image appear to be compared to the unlensed image? Surface brightness is conserved, meaning that to have twice the flux it must appear twice as large. This is one way that people get more detailed glimpses of distant objects. Lensing magnifies the image, so more structure can be resolved.

This is an *extremely* powerful way to figure out what is happening to light as it goes every which way. The specific intensity is all you need to figure out lots of important things, such as the flux or the surface brightness, and in apparently complicated situations you just follow how the frequency behaves. I have, for example, used this extensively in computations of ray tracing around rotating neutron stars, where in general the spacetime is quite complicated.

Effects of having matter around: what are the effects of matter? Matter can emit and can absorb, so the specific intensity in any given direction can be altered by the presence of matter. For stellar interiors, we're lucky in that the distance a photon (or other energy-transferring particle) can travel is small compared to other length scales, so we can treat the propagation of radiation as diffusion, or as a random walk.

Diffusion and random walks: It would be wonderful to be able to treat the motion of photons as if it were a random walk, independent of frequency. Of course, a problem with this is that photons of different energies have different opacities, so it isn't completely straightforward. For the moment, however, let's ignore this and assume that the opacity is independent of the frequency ("gray atmosphere" approximation).

If the photons were simply to scatter, then their progress would be a random walk. Suppose that in step 1 the photon travels a distance and direction $\mathbf{r}_1$, in step 2 it travels a distance and direction $\mathbf{r}_2$, and so on, in directions that are random and drawn from an isotropic distribution and the distance of each step is drawn from the same distribution (and need not be the same distance each time). The mean location after N scatterings is zero (it has to be, due to the isotropy), but the mean square distance is

$$\langle \mathbf{R}^2 \rangle = \langle \mathbf{r}_1^2 \rangle + \langle \mathbf{r}_2^2 \rangle + \langle \mathbf{r}_3^2 \rangle + \ldots + 2\langle \mathbf{r_1} \cdot \mathbf{r_2} \rangle + 2\langle \mathbf{r_1} \cdot \mathbf{r_3} \rangle + \ldots \quad . \tag{3}$$

The cross terms all average to zero because the directions are uncorrelated, but the squared terms are all the same and are the mean square distance traveled in a single scatter. The net result is that the mean square distance after N steps is Nl^2, where l^2 is the mean square distance for a single scattering. Therefore, we get the result that the average distance from the origin is $\sqrt{N}l$. Similarly, if the optical depth to escape is $\tau \gg 1$, the photon typically undergoes $\sim \tau^2$ scatterings before it escapes. If the medium is optically thin, i.e., $\tau \ll 1$, then of order τ scatterings on average are required (which is to say that most photons don't scatter, a few do, and the average number of scatterings is τ). For most rough estimates you can use just $\max(\tau, \tau^2)$ for the number of scatterings. In stellar interiors, $\tau \gg 1$ so we're in the random walk limit and can treat the process of radiation transfer as diffusion.

For pure *elastic* scattering (Thomson scattering, which applies when the photon energy is much less than the rest mass-energy of the particle off of which it is scattering), the photons don't change energy in the rest frame of the particle (they *can* change energy in a "laboratory frame" in which the particle was initially moving, due to Doppler shifts. Give it some thought...). However, for scattering with recoil (Compton scattering) the photons do change energy even in the rest frame of the particle, and for absorption and emission the photon loses its identity, so we'd like a somewhat more general way to treat things. We also would like to incorporate the fact that the opacity will in general depend on frequency. The most general way to do this would be to calculate the emission and opacity (or their ratio, the source function) at each point in the star as a function of frequency and solve the transfer equation that way. However, there's a simplification that saves a lot of work and usually comes up with remarkably good answers.

Rosseland mean opacities: The idea is that local thermodynamic equilibrium (LTE) is such a good approximation that we can pretty much assume that locally the photons are distributed in blackbodies and the particles in Maxwell-Boltzmann distributions, where the photon and particle temperatures are the same. However, over large distances, the temperature will change. We then want to know an "average" opacity such that if we assume that the opacity at all frequencies is this average, we'll get the radiative transfer about right. The way to do this is to weight the average harmonically by the Planck function:

$$\frac{1}{\kappa_R} = \left[\int_0^\infty \frac{1}{\kappa_\nu}\frac{\partial B_\nu}{\partial T}\,d\nu\right]\left[\int_0^\infty \frac{\partial B_\nu}{\partial T}\,d\nu\right]^{-1} \tag{4}$$

where $B_\nu = (2h\nu^3/c^2)/(e^{h\nu/kT}-1)$ is the blackbody function. This is the Rosseland mean opacity. Why is this harmonically weighted? Because the *low* opacities will dominate the transfer, since photons at those energies will go farther and sample larger gradients in temperature. With this definition, the flux is

$$F(r) = -\frac{c}{3\kappa_R\rho}\frac{d(aT^4)}{dr}\ . \tag{5}$$

Gray atmosphere: We won't go into stellar atmospheres much in this course, but it is useful to know that for many purposes the gray atmosphere solution (using the Rosseland mean opacity) is remarkably good. In this solution, if τ_R represents the Rosseland optical depth from the surface, the temperature as a function of depth is given by

$$T^4 \approx \frac{3}{4} T_{\rm eff}^4 (\tau_R + 2/3) \ , \tag{6}$$

where $T_{\rm eff}$ is the effective temperature, such that the radiated flux from the surface is $\sigma T_{\rm eff}^4$. This isn't perfect, of course, but it's surprisingly accurate even when there are lines and edges and stuff. The more important difference from gray atmospheres that exists in real atmospheres is that the emergent spectrum is usually much different from the Planck spectrum we would actually find from a gray atmosphere.