Stellar Structure and Evolution

For as long as humans have existed, we have looked up at the stars and told ourselves stories about them. We still do that, but now those stories are enriched by the remarkable connection that stars have with physics, and the amazing array of inferences we can make about the universe by studying stars.

In this course we will tell the story of stars. We will do so with a focus on the physics of stars, and with an emphasis on developing physical intuition about the birth, life, death, and afterlife of stars. As appropriate, we will also talk about some topics of current research. At any point during the course it is appropriate to ask "how do we know?" for any particular assertion. Sometimes we might know something from basic physics. Other times it might be inference from observation, although we should realize that inferences are very rarely direct; we have to interpret observations through the lens of current theory. Sometimes "how do we know?" has an answer of "it's our best current guess", and those guesses can obviously change with time. But asking the question is a key to wisdom :)

So let's begin by thinking about how to reason as astrophysicists.

Developing Astrophysical Reasoning Skills

As discussed in detail in the "Hints about doing research in astrophysics" file on the class web page, there's quite a transition between classwork and research. In this course I will encourage development of research-oriented skills. One of these is the ability to size up a problem and determine how best to approach it, given the goal of the research and the needed accuracy. Some things are best solved analytically and some with a computer; some require great accuracy and some are best done with order-of-magnitude estimates; and so on. In all cases, though, you've got to be able to sit back and ask yourself "Does this make sense?" so that a programming bug doesn't convince you that energy isn't conserved!

One aspect of "does this make sense" is that you need to be able to look at a result and determine if it satisfies several "common-sense" criteria, from simple to complex. Does it have the right units? Is it correct in limits that I can check easily? Does it possess the appropriate symmetries? Does it depend on what it should depend on, and no more? Ideally, you should do this before you embark on a calculation, and also afterwards, to check your result. You'd be surprised at how often you can catch errors this way or sharpen your intuition.

Approaches to Creativity in Astrophysics

Creativity can be said to have two steps: (1) coming up with a list of possibilities, (2) going through those and throwing out what doesn't work, to focus effort on the more promising explanations.

Let's say you want to explain a phenomenon. One approach is to simply make a big list of anything you can think of that might explain it (without culling them at this stage), then later go through the list and see if observations or other constraints absolutely rule out some of the proposals. Doing it in a two-step way like this gives you a chance to come up with something really original (by not cutting it down first), but also serves as a check against errors.

To do this successfully, you need to have a wide range of knowledge of physics and astrophysics, both to generate ideas and to test them. In this class we will try to generate a set of tools to approach problems in high-energy astrophysics, so that we can come up with ideas and cull them for the most promising.

As a creativity exercise, let's see if we can figure out what powers the Sun. Yes, we know that it's nuclear fusion, but suppose that we didn't know that; what possibilities might we suggest, and how would we test them?

Stars

What do we know about stars and how do we know?

These are major questions! In detail, we could list a nearly unlimited set of things that we know about stars, from their masses, luminosities, composition, spectra, evolution, fate, etc. We could then list the (often long) chains of reasoning that have led researchers to have various degrees of confidence in any given measurement. In many cases we could point to misunderstandings that often persisted for years to decades (e.g., about the composition of stars) and why (a) the incorrect answer wasn't silly, and (b) how the mistake was discovered. Science is a process, not a final answer.

We will therefore place ourselves in the position of talented physicists who want to work out the structure and evolution of the stars. That means that we will not only have to come up with answers to questions, but we need to come up with the questions themselves. Given the phenomenal inherent complexity of stars, we have to choose carefully what problems we are studying. These have to be problems that are simplified enough to be analyzed, but realistic enough to give us insight and to be compared with observations.

But how can we strive for the goal (to paraphrase Einstein) that everything should be made as simple as possible, but no simpler? First, we have to decide what we can ignore. "Ignore" here means that what we are neglecting is small enough compared to what we are including that the accuracy we need is satisfied. Thus, we need in every case to have some standard against which to compare, and what this means in practice is that the approach we take to a problem has to depend on the level of precision that we need or want to get out of the solution.

To start with a non-astronomical example, suppose that I have an exactly filled water tank that is a parallelepiped with dimensions 107 cm wide by 109 cm long by 103 cm deep. My task is to determine the mass of the water in the tank to the nearest gram. My first step is to compute the volume of the tank ($107 \times 109 \times 103$ cm^3) and multiply by the density of water, which is about 1g cm^{-3}. But "about" won't cut it here; because the volume is about 10^6 cm^3 the mass is about 10^6 g, which means that I need to know the density to a part in 10^6 to get the required one gram precision. But that means I need to know the composition of the water (pure H_2O? Salt water? Any minerals in there?), its temperature, the ambient pressure, the local surface gravity, and so on. To this precision, I have a *very* complicated task ahead of me.

But if I take the same tank and ask whether I could lift it over my head (using just my own muscles, on Earth, no funny business), then I can eyeball the volume and note (1) it's about $100 \times 100 \times 100 \approx 10^6$ cm^3, (2) at roughly 1 g cm^{-3} that's about 10^6 g=10^3 kg, and so (3) no way! No complexity needed.

For a stellar example: we have each lived long enough to notice that the Sun hasn't collapsed in our lifetime. But is this something that we need to explain? One way to get at the answer is to ask ourselves "what if nothing were holding the Sun up against gravity?" If you do the calculation you find that in that case the Sun would collapse in about an hour. That's much shorter than your lifetime, so you can conclude that something *is* holding the Sun up against gravity. Your first try at an explanation should be to assume that the Sun is a spherical, nonrotating ball of gas, without any magnetic fields, and that it is held up by gradients in gas pressure (more about that later, but you should have seen that in ASTR 320). This does an excellent job. But we'd like to see whether our neglect of rotation and magnetic fields is reasonable.

Can we ignore rotation?—The Sun has a rotation period of about a month. The Keplerian period at the surface is $2\pi\sqrt{R^3/GM}$, or about 3 hours. Thus, the rotation rate of the Sun is about 1/250 of Keplerian. Centrifugal acceleration scales as the square of the angular velocity (**Ask class:** is there a symmetry reason why it's not linear?), which means that rotational effects are only at the 10^{-4} level. Thus unless we need that level of precision, or there is some effect that only emerges when there is rotation, we can ignore rotation.

Can we ignore magnetic fields?—The strongest magnetic fields seen in the Sun are in sunspots, where $B \sim 1000$ G. If that were the average in the entire Sun, then the total magnetic energy would be $(B^2/8\pi)\frac{4\pi}{3}R^3 = 6 \times 10^{37}$ erg. It's actually much less, because the bulk average magnetic field in the Sun is significantly lower (**Ask class:** how might we learn about the *interior* magnetic field strength in the Sun?). But in astronomy (and more generally in physics), we need to have context; that seems like a large amount of energy, but with what should we compare it?

Since we're thinking about the gravity of the Sun versus gas pressure gradients, maybe we want to compare the gravitational potential energy in the Sun with the magnetic energy. A first guess about the gravitational potential energy is $GM_\odot^2/R_\odot = 4 \times 10^{48}$ erg, where $M_\odot \approx 2 \times 10^{33}$ g is the mass of the Sun and $R_\odot \approx 7 \times 10^{10}$ cm is the radius of the Sun. Now, the gravitational potential energy could be larger than that, because the Sun has a highly compressed core. But we're already ~ 11 orders of magnitude higher than the magnetic energy, so we don't need to work harder on this part.

This is the type of thinking you should always, at least implicitly, undergo when you analyze a problem. Note that what you keep and what you ignore really does depend on the problem. For example, if you wanted to explain the X-ray emission from the Sun you would find that magnetic fields are crucial, because they are essential to understanding flares and prominences, which produce the X-rays. Thus even the process of deciding "I can neglect this" can be used to sharpen your intuition and stock your toolbox.

Another major question you should ask is "is this system in equilibrium?" Maybe that sound like an easy question, but there are multiple types of equilibrium. Let's consider a few that are relevant to stars.

Dynamic, or hydrostatic, equilibrium.—This means that the star as a whole stays put. Said another way, the forces acting on any given parcel of gas balance each other. First **Ask class:** what would happen if the Sun were far away from this balance? Answer: it would collapse or expand on the dynamic time scale, which to a decent approximation is just the free-fall time scale. That, in turn, is proportional to $1/\sqrt{GM/R^3}$, or about $1/\sqrt{G\rho}$. For the Sun, the average density ρ is about 1 g cm^{-3}, so that's 1 hour, as we said above. Since we don't see dramatic changes in the Sun on 1 hour time scales, and indeed not on scales of millions of years (**Ask class:** how do we know? Fossils and the geologic record tell us that nothing overwhelming has happened on that time scale, although from stellar models the Sun was about 30% less luminous when it was born than it is now.), we know that this overall balance holds to extreme accuracy. This is not a good approximation for supernovae, of course, but even for most pulsating stars (such as Cepheids) the bulk of the star is in hydrostatic equilibrium.

To pursue this further, we need to quantify what is balancing what. For a given parcel of gas, gravity pulls toward the center of mass. **Ask class:** what could oppose gravity? If the pressure *gradient* (not just the pressure; why?) is in the same direction of gravity (i.e., more pressure farther down), then this opposes gravity. Let's say that we have a parcel of gas with area perpendicular to $\hat{r}$ of A and thickness dr. If the density is ρ, then the gravitational force on this is $-GM\rho A dr/r^2$, where the negative sign indicates a downward force. This can be written as $g\rho A dr$, where $g = -GM/r^2$ should really be a vector, and of course M is really M_r, the mass interior to r. The force due to the pressure gradient is

$P(r)A - P(r+dr)A = dr(dP/dr)A$. The sum of the two has to be zero for force balance, so $dr(dP/dr)A + g\rho A dr = 0$, or finally $dP/dr = -\rho g$ (more generally, when spherical symmetry doesn't apply, $\nabla P = -\rho \mathbf{g}$). This is the equation of hydrostatic equilibrium, and is one of the four fundamental equations of stellar structure. **Ask class:** how would this be modified if a star were rotating rapidly? In that case, we would need to include centrifugal terms. In general, force balance is required for dynamic equilibrium. In the case of stars (and most other things in the universe), this translates to gravity vs. everything else, because gravity is universally attractive and hence other forces are needed to balance it. Indeed, for much of astrophysics you can get a lot of insight by asking "what opposes gravity in this case?"

It is often helpful to write such equations in terms of the mass instead of the radius. This formulation, in which we follow the mass, is called the *Lagrangian* formulation. Then, since the mass in a spherical shell is $dM_r = 4\pi r^2 \rho dr$, the equation of hydrostatic equilibrium is $dP/dM_r = -GM_r/4\pi r^4$.

Thermal equilibrium.—Is the Sun in thermal equilibrium? Of course not! It has a photospheric temperature of almost 6000 K, and it is radiating into cold, empty space. So what does it mean to say it is in thermal equilibrium? What about deeper in the Sun? What is the microscopic condition for something to be pretty close to thermal equilibrium? One criterion could be that whatever is carrying the energy can't go so far that it samples temperatures that are dramatically different from where it started. In the case of the Sun, we know that the whole temperature run is about 10^7 K in 10^{11} cm, or about 10^{-4} K cm^{-1}. As we'll find out later, photons carry most of the energy. The cross section is no less than about 10^{-24} cm^2 (the Thomson scattering cross section). At about 1 g cm^{-3}, which is the average for the Sun (1.4 g cm^{-3} for those who like to be precise), that's a number density of 10^{24} cm^{-3}, which implies that the average photon travels roughly 1 cm. It therefore samples an average temperature change of 10^{-4} K. Since the temperature even at the photosphere is 6000 K, this is a tiny fraction and thermal equilibrium is a good approximation for almost all of the Sun, even where the density is much smaller than the average.

Thermal balance.—Ah, but there is another issue. Energy isn't just leaking out of the Sun. It is being generated as well. For each parcel of fluid the energy generated must flow out at the same rate it is generated, or there will be a buildup of heat. Thinking again about a parcel of area A and thickness dr, and assuming that the energy generation rate per mass is ϵ and the flux at r is $F(r)$, we have $F(r+dr)A - F(r)A = \epsilon\rho A dr$. In this expression we can replace A by $4\pi r^2$, the area of the spherical shell at r, and define the luminosity at r as $L(r) = 4\pi r^2 F(r)$. Then $L(r+dr) - L(r) = \epsilon\rho 4\pi r^2 dr$, or $dL/dr = 4\pi r^2 \rho\epsilon$. This is the equation of thermal balance. In the Lagrangian formulation this is $dL/dM_r = \epsilon$.