

Introduction and Astrophysical Thinking

The era of direct detection of gravitational waves has commenced. With it comes the opening of a new window onto the universe, which will bring unique insights into black holes, neutron stars, stellar dynamics, and possibly categories or properties of unanticipated sources. In the five lectures I will give in this school, I will go through some of the properties of the sources that will likely be seen with gravitational waves, and some of the currently unresolved questions.

But before we get to that stage, it is important to say a bit about astrophysical thinking, and how the most productive approaches differ from those that might be more familiar for normal physics problems. A key difference is that in astrophysics we usually can't experiment on the sources we are studying. This is probably good for humanity in general; we're not ready as a species to be trusted with laboratory black holes and neutron stars! But it does mean that, unlike what is often true in terrestrial labs, we cannot probe systems by changing one variable at a time to see what happens.

The net result is that astrophysicists are routinely confronted by systems that are fundamentally too complicated to understand completely. A star, for example, is spectacularly complex when all its details are considered. Thus to make progress it is standard to simplify until you can get a handle on at least the aspects of the system that you would like to study. As a result, it is extremely valuable for astrophysicists to have the ability to make quick estimates of effects (often called order of magnitude calculations) to see whether something might be important, before settling down to the more involved task of doing the detailed work to pursue an idea. This also means that instead of simply solving a problem, you need to decide in advance what you want to get out of your study before you determine what approach to use.

Let me give two examples. The first is that you are presented with a filled water tank that is 97 cm high, 103 cm wide, and 109 cm long. Suppose you are asked to compute the mass of the water in the tank to the nearest gram. Your approach might initially be to multiply 97 by 103 by 109 to get the volume in cm^3 and then use 1 g cm^{-3} as the density of water to get the mass in grams. But on further thought you should realize that the density of water isn't *exactly* 1 g cm^{-3} in all circumstances; for example, it depends a bit on the temperature of the water and on the local atmospheric pressure. Thus to answer this question about the tank to the desired precision you would need to have extra information and would need to do quite a detailed calculation.

But suppose instead you were asked whether you personally could, using only your own muscles on Earth, lift the tank over your head. Since the volume of the tank is roughly 100 cm by 100 cm by 100 cm, or about a million cubic centimeters, and that volume of water

has a mass of about a ton, the answer is no! You don't need to go into details in that case.

For the second example, let's draw something from astronomy. If I asked you to solve the Sun, you would quite properly look blankly at me, because there are many aspects of our observations of the Sun that aren't yet explained. So I'll be more specific: I ask you to note that the Sun has not been twice its current size, or half its current size, at any point in your life. Can you explain that? First you'd need to understand that there is something to explain; you could come to this understanding by doing a quick free-fall calculation to see that if nothing were holding up the Sun, it would collapse in about an hour. This is much less than your lifetime, so now you realize that you have to determine what holds the Sun up. For this purpose, you might start by assuming that the Sun is a sphere, and guess that gradients of gas pressure hold up the Sun. Of course, the Sun is *not* a perfect sphere, and other things help hold it up, but you'd find that this is a pretty good approximation. Such a calculation would take a few minutes if you knew the direction to pursue.

In contrast, if I ask you to explain all the details of the magnetic structure of the Sun, you (and a few hundred colleagues) could work your whole career on this and still not get a full solution! I hope that these examples make clear that your approach to an astrophysical problem should depend on the level of accuracy you need.

Another general point to make relates to the process of creativity. Linus Pauling said "The way to get good ideas is to get lots of ideas and throw the bad ones away". So what can we do about this in practice? My experience suggests that when you are trying to come up with ideas, it is helpful to separate the process into two steps:

1. Brainstorming.
2. Critical evaluation.

In the brainstorming step, you just think of ideas and put them on a list, without trying to assess at that point whether they are reasonable. In the critical evaluation step, you become your own harshest critic: what are the holes or flaws in each idea you just listed? Can you *cleanly* rule out any of them and thus take them off the list? Are any of the remainder really promising? Using this approach you don't miss genuinely creative solutions but, at the same time, you remain grounded in reality. A trap to avoid, which I have seen claim many astrophysicists (including well-known ones), is to fall deeply, madly in love with your own idea. In such a state, you will ignore counterarguments and evidence against your idea. That's really bad. It is your *responsibility*, as a scientist, to try to think of arguments against your own ideas and to be honest about them. That's the way that science can progress most smoothly; for myself, I can tell you that I have extra respect for scientists who will tell me of possible weaknesses in their ideas, because I feel that I can trust them more and have a productive dialog with them.

You can get some practice in the brainstorming→critical evaluation path by trying to think of the fundamental power source of the Sun. You probably know that it is actually nuclear fusion, but suppose you didn't know that. List some sources; fusion would be on there, as would normal chemical energy (e.g., burning). How many more can you think of? After you have a full list, take a break, then come back to rule them out. Note that when I say above that you want to *cleanly* rule out possibilities, this often implies that you want to make your arguments as simple as possible. For example, you can estimate the total amount of energy that the Sun has produced so far in its lifetime. How does this compare with what you'd get if the Sun were a ball of paper and you got energy by burning the paper? You'll find that paper burning falls short by orders of magnitude, so you can be convinced that even slight variations in your calculations make little difference.

In contrast, if you tried to do something much more involved then the arguments would become more complicated and therefore less convincing. For example, suppose that you wanted to rule out the burning of paper by computing the detailed helioseismic oscillations that the Sun would have if it were made of paper. Maybe you'd eventually get to a contradiction with observations, but it would take a long time and a huge effort, and it wouldn't be followed very easily by others! Simple checks are critical: are the units right? Are the symmetries and limits right? There are published papers that are trivially disproven by these simple checks; try to avoid having yours join that list!

The final point I'd like to make before we discuss the basics of gravitational wave sources has to do with how you can improve your astrophysical intuition. Remember, part of the idea is to be able to make quick assessments of ideas so that you can spend your time on the most promising ones. How do you know what you can discard and what you should keep for further analysis? I'd say that the best approach is to do various calculations for fun in different circumstances to see what you can do. It even helps to do some sample detailed calculations of processes that you know can't be important, so that you get a sense of when you can absolutely rule something out. For example, how strongly and simply can you rule out the idea that the solar wind is largely driven by the scattering of solar neutrinos off of electrons and nuclei?

You also can acquire high-level understanding of some processes and numbers that will allow you to take shortcuts. For example, how fast would something orbit in a circle at a distance of a parsec from the center of a globular cluster with a mass of $10^5 M_\odot$? You could put in factors of G and the like, but if you know that (1) the Earth moves at 30 km s^{-1} in its 1 AU orbit around a $1 M_\odot$ object, (2) there are about 200,000 AU in one parsec, and (3) the orbital speed scales as $(M/r)^{1/2}$, then you can see that the speed is $(10^5/2 \times 10^5)^{1/2}$ times the Earth's speed, or about 20 km s^{-1} . Fast, high-level reasoning such as this also helps in the informal scientific conversations that are essential to your career development.

To sum up this introductory part, astrophysics is tremendously fun to pursue, but for those of you who are used to exact analysis (like I was entering grad school), you need to get used to these techniques!

Overview of Gravitational Radiation

As we contemplate the triumphant direct detection of gravitational radiation, it is useful to consider what such detections will teach us about the universe. The first detection, GW150914, was of course of immediate significance because it was a direct confirmation of a dramatic prediction of general relativity: to paraphrase John Wheeler, that spacetime tells sources how to move, and moving sources tell spacetime how to ripple.

Beyond the initial detections, gravitational wave science will pass into the realm of astronomy, and will give us new observational windows onto some of the most dynamic phenomena in the universe. These include merging neutron stars and black holes, supernova explosions, and possibly echoes from the very early history of the universe as a whole. They have also already provided the cleanest tests of predictions of general relativity in the realm of strong gravity, with much more to come.

However, there are important differences from standard astronomy. In electromagnetic observations, in every waveband there are sources so strong that they can be detected even if you know nothing about the source. You don't need to understand nuclear fusion in order to see the Sun! In contrast, as we will see, most of the expected sources of gravitational radiation are so weak that we expect that usually sophisticated statistical techniques will be required to detect them at all (with occasional happy exceptions such as GW150914, which was so strong that it could be seen by eye after moderate bandpass and notch filtering of the data). A standard technique involves matching templates of expected waveforms against the observed data stream. Maximum sensitivity therefore requires a certain understanding of what the sources look like, and thus of the characteristics of those sources. In addition, it will be important to put each detection into an astrophysical context so that the implications of the discoveries are evident.

As an aside, it is useful to remember that historically the most interesting sources discovered with a new telescope or satellite have often been unexpected. This is also possible with gravitational radiation. However, you can't sell a large project by appealing entirely to the unknown, so we should at least describe what we *can* imagine at this point!

Before discussing types of sources, though, we need to have some general perspective on how gravitational radiation is generated and how strong it is. We will begin by discussing radiation in a general context.

By definition, a radiation field must be able to carry energy to infinity. If the amplitude of the field a distance r from the source in the direction (θ, ϕ) is $A(r, \theta, \phi)$, the flux through a

spherical surface at r is $F(r, \theta, \phi) \propto A^2(r, \theta, \phi)$. If for simplicity we assume that the radiation is spherically symmetric, $A(r, \theta, \phi) = A(r)$, this means that the luminosity at a distance r is $L(r) \propto A^2(r)4\pi r^2$. Note, though, that when one expands the static field of a source in moments, the slowest-decreasing moment (the monopole) decreases like $A(r) \propto 1/r^2$, which implies that $L(r) \propto 1/r^2$ and hence no energy is carried to infinity. This tells us two things, regardless of the nature of the radiation (e.g., electromagnetic or gravitational). First, radiation requires time variation of the source. Second, the amplitude must scale as $1/r$ far from the source.

We can now explore what types of variation will produce radiation. We'll start with electromagnetic radiation, and expand in moments. Suppose that we are far from some distribution of electric charges, which could be in motion. For a charge density $\rho_e(\mathbf{r})$, the monopole moment is $\int \rho_e(\mathbf{r})d^3r$. We assume that the volume over which we perform the integral encompasses the entire system; no charges can enter or leave. As a result, the monopole moment is simply the total charge Q , which cannot vary, so there is no electromagnetic monopolar radiation. The next static moment is the dipole moment, $\int \rho_e(\mathbf{r})\mathbf{r}d^3r$. There is no applicable conservation law, so electric dipole radiation is possible. One can also look at the variation of currents. The lowest order such variation (the "magnetic dipole") is $\int \rho_e(\mathbf{r})\mathbf{r} \times \mathbf{v}(\mathbf{r})d^3r$. Once again this can vary, so magnetic dipole radiation is possible. The lower order moments will typically dominate the field unless their variation is reduced or eliminated by some special symmetry.

Now consider gravitational radiation. Let the mass-energy density be $\rho(\mathbf{r})$. The monopole moment is $\int \rho(\mathbf{r})d^3r$, which is simply the total mass-energy. This is constant, so there cannot be monopolar gravitational radiation. The static dipole moment is $\int \rho(\mathbf{r})\mathbf{r}d^3r$. This, however, is just the center of mass-energy of the system. In the center of mass frame, therefore, this moment does not change, so there cannot be the equivalent of electric dipolar radiation in this frame (or any other, since the existence of radiation is frame-independent). The counterpart to the magnetic dipolar moment is $\int \rho(\mathbf{r})\mathbf{r} \times \mathbf{v}(\mathbf{r})d^3r$. This, however, is simply the total angular momentum of the system, so its conservation means that there is no magnetic dipolar gravitational radiation either. The next static moment is quadrupolar: $I_{ij} = \int \rho(\mathbf{r})r_i r_j d^3r$. This does not have to be conserved, and thus there can be quadrupolar gravitational radiation.

This allows us to draw general conclusions about the type of motion that can generate gravitational radiation. A spherically symmetric variation is only monopolar, so it does not produce radiation. No matter how violent an explosion (even a supernova!) or a collapse (even into a black hole!), no gravitational radiation is emitted if spherical symmetry is maintained. In addition, a rotation that preserves axisymmetry (without contraction or expansion) does not generate gravitational radiation because the quadrupolar and higher moments are unaltered. Therefore, for example, a neutron star can rotate arbitrarily rapidly

without emitting gravitational radiation as long as it maintains stationarity and axisymmetry and rotates around the axis of symmetry.

This immediately allows us to focus on the most promising types of sources for gravitational wave emission. The general categories are: binaries, continuous wave sources (e.g., rotating stars with nonaxisymmetric lumps), bursts (e.g., asymmetric collapses), and stochastic sources (i.e., individually unresolved sources with random phases; the most interesting of these would be a background of gravitational waves from the early universe). We will discuss each of these in subsequent lectures.

We can now make some order of magnitude estimates. What is the approximate expression for the dimensionless amplitude h of a metric perturbation, a distance r from a source? Note, by the way, that because gravitational waves are perturbations in spacetime, h is related to the fractional deviation of the spacetime from the Minkowski (flat) spacetime. Thus h is of the order of the fractional change in length induced by a passing gravitational wave, if the length in question is of order the gravitational wavelength.

We argued that the lowest order radiation has to be quadrupolar, and hence depends on the quadrupole moment I . This moment is $I_{ij} = \int \rho r_i r_j d^3x$, so it has dimensions MR^2 , where M is some mass and R is a characteristic dimension. We also argued that the amplitude is proportional to $1/r$, so we have

$$h \sim MR^2/r. \quad (1)$$

We know that h is dimensionless, so how do we determine what else goes in here? In GR we usually set $G = c = 1$, which means that mass, distance, and time all have the same effective “units”, but we can’t, for example, turn a distance squared into a distance. Our current expression has effective units of distance squared (or mass squared, or time squared). We note that time derivatives have to be involved, since a static system can’t emit anything. Two time derivatives will cancel out the current units, so we now have

$$h \sim \frac{1}{r} \frac{\partial^2(MR^2)}{\partial t^2}. \quad (2)$$

Now what? To get back to physical units we have to restore factors of G and c . It is useful to remember certain conversions: for example, if M is a mass, GM/c^2 has units of distance, and GM/c^3 has units of time. Playing with this for a while gives finally

$$h \sim \frac{G}{c^4} \frac{1}{r} \frac{\partial^2(MR^2)}{\partial t^2}. \quad (3)$$

Since G is small and c is large, the prefactor is *tiny!* That tells us that unless MR^2 is large, the system is changing fast, and r is small, the metric perturbation is minuscule.

Let’s make a very rough estimate for a circular binary. Suppose the total mass is $M = m_1 + m_2$, the reduced mass is $\mu = m_1 m_2 / M$, and the semimajor axis is a , so the orbital

frequency Ω is given by $\Omega^2 a^3 = GM$. Without worrying about precise factors, we say that $\partial^2/\partial t^2 \sim \Omega^2$ and $MR^2 \sim \mu a^2$, so

$$h \sim (G^2/c^4)(\mu/r)(M/a) . \quad (4)$$

This can also be written in terms of orbital periods, and with the correct factors put in we get, for example, for an equal mass system

$$h \approx 10^{-22} \left(\frac{M}{2.8 M_\odot} \right)^{5/3} \left(\frac{0.01 \text{ sec}}{P} \right)^{2/3} \left(\frac{100 \text{ Mpc}}{r} \right) , \quad (5)$$

which is scaled to a double neutron star system. This is really, really, small: it corresponds to less than the radius of an atomic nucleus over a baseline the size of the Earth. That's why it is so challenging to detect these systems!

Remarkably, though, the flux of energy is *not* tiny. To see this, let's calculate the flux given some dimensionless amplitude h . The flux has to be proportional to the square of the amplitude and also the square of the frequency f : $F \sim h^2 f^2$. This currently has units of frequency squared, but the physical units of flux are energy per time per area. Replacing factors of G and c , we find that the flux is

$$F \sim (c^3/G)h^2 f^2 . \quad (6)$$

Now the prefactor is *enormous!* For the double neutron star system above, with $h \sim 10^{-22}$ and $f \sim 100$ Hz, this gives a flux of a few hundredths of an erg cm⁻² s⁻¹. For comparison, the flux from Sirius, the brightest star in the night sky, is about 10⁻⁴ erg cm⁻² s⁻¹! That means that if you could somehow absorb gravitational radiation perfectly with your eyes, you would see hundreds to thousands of events per year brighter than every star except the Sun. To put it another way, the energy per time emitted by the GW150914 event, during the last part of its coalescence, was tens of times greater than the energy per time emitted by every star in the visible universe *combined* during that same time (!!!). What this really implies, of course, is that gravitational radiation interacts *very* weakly with matter, which again means that it is mighty challenging to detect.

Let us conclude with an idea of the frequency range available for a given type of binary. There is obviously no practical lower frequency limit (just increase the semimajor axis as much as you want), but there is a strict upper limit. The two objects in the binary clearly won't produce a signal higher than the frequency at which they touch. If we consider an object of mass M and radius R , the orbital frequency at its surface is $\sim \sqrt{GM/R^3}$. Noting that $M/R^3 \sim \rho$, we can say that the maximum frequency involving an object of density ρ is $f_{\text{max}} \sim (G\rho)^{1/2}$. This is actually more general than just orbital frequencies. For example, a gravitationally bound object can't rotate faster than that, because it would fly apart. In addition, you can convince yourself that the frequency of a sound wave that involves most of

the object can't be greater than $\sim (G\rho)^{1/2}$. Therefore, $\sim (G\rho)^{1/2}$ is a general upper bound on dynamical frequencies.

This tells us, therefore, that binaries involving main sequence stars can't have frequencies greater than $\sim 10^{-3} - 10^{-6}$ Hz, depending on mass, that binaries involving white dwarfs can't have frequencies greater than $\sim 0.1 - \text{few}$ Hz, also depending on mass, that for neutron stars the upper limit is $\sim 1000 - 2000$ Hz, and that for black holes the limit depends inversely on mass (and also spin and orientation of the binary). In particular, for black holes the maximum imaginable frequency is on the order of $10^4(M_\odot/M)$ Hz at the event horizon, but in reality the orbit becomes unstable at lower frequencies (more on that in the next lecture).

The net result is that for ground-based interferometers such as the two US-based LIGO detectors, GEO-600, Virgo, KAGRA, and LIGO-India, which are sensitive to frequencies $\sim 10 - 2000$ Hz, the only individual sources that will be detected are neutron stars and black holes and their creation events (supernovae); some might argue that cuspy cosmic strings might fall into this category, but we'll leave that for a later discussion. In the next two lectures we will therefore lay out the astrophysical basics of black holes and neutron stars.

For a rigorous derivation of the evolution of a well-separated binary under the influence of gravitational radiation, see Peters 1964, *Physical Review*, 136, B1224. This classic paper derives the rates of energy and angular momentum loss, and hence the rate of change of the semimajor axis and angular momentum, for a binary which is imagined to move in a Keplerian way over a full orbit. This paper also does a good job of providing the framework in which other gravitational radiation calculations can be performed (e.g., of a lump on a rotating star).