Compact binary formation: isolated binary evolution

The basic ways that compact binaries can come together break down to two major categories:

1. Evolution of an isolated massive binary. That is, we start with a pair of massive stars that both evolve into black holes, and merge, without any other stars coming close enough to do anything.
2. Dynamical processes. Examples include single-binary interactions, the Kozai-Lidov resonance, direct dynamical capture, and capture in the accretion disk of an active galactic nucleus.

Today we will talk about isolated massive binaries, and tomorrow we will explore some of the dynamical paths.

At first glance, the formation of a compact binary seems easy. We know that black holes and neutron stars evolve from massive stars, so just start with a binary of two massive stars and voila! When we learn that massive stars almost always have binary companions (at the $> 80\%$ or even $> 90\%$ level), and that the companion to a massive star is typically another massive star, it seems even easier. So why even have the discussion?

The answer is that the first glance is very misleading. We can get an initial quantitative sense for this by referring to the best current estimate of the merger rate per volume of double black hole systems, and then comparing that with the rate of formation of massive stars.

We'll start with the estimated double black hole merger rate in the "local" universe (redshift less than a few tenths!): about 20 $\mathrm{Gpc}^{-3}\ \mathrm{yr}^{-1}$. A commonly used concept in gravitational wave astronomy is the "Milky Way Equivalent Galaxy", or MWEG, which is a collection of stars equal to the number in the Milky Way. A typical estimate of the average number density of MWEGs over a large volume in the local universe is 0.01 Mpc^{-3}. Since $1\ \mathrm{Gpc}^3 = 10^9\ \mathrm{Mpc}^3$, the estimated rate of double BH mergers is $2 \times 10^{-8}\ \mathrm{Mpc}^{-3}\ \mathrm{yr}^{-1}$, or $2 \times 10^{-6}\ \mathrm{MWEG}^{-1}\ \mathrm{yr}^{-1}$.

The estimated stellar mass formed per year in the Milky Way is around 5 solar masses per year (of course, each star takes a long time to form, but this is the average). The distribution of the number N of stars in some mass interval M to $M + dM$ is given reasonably well by the Salpeter mass function $dN/dM \propto M^{-2.35}$ above, say, $M = 0.5\ M_\odot$. This means that the mass in stars formed above some mass M_0 is proportional to $M_0^{-0.35}$ (one of the two extra powers of M is because this is mass rather than number, and the other is because we are integrating rather than looking at the differential distribution). If we estimate that we need an initial stellar mass of 25 $M_\odot$ to get a black hole (which is in the right order of magnitude although it's uncertain), this means that a fraction $\sim (25/0.5)^{-0.35} \approx 0.25$ of the mass in stars forms in stars above 25 $M_\odot$. Rounding, we'll say that $5 \times 0.25 \approx 1\ M_\odot\ \mathrm{yr}^{-1}$. Since we are thinking about a binary, we'll say that the binary we need has a total mass of 50 $M_\odot$, which means that we expect a formation rate of $\sim 1/50 = 0.02$ MWEG yr^{-1}.

But this is *huge* compared with the observed rate! $0.02/(2 \times 10^{-6}) = 10^4$. This already tells us that we're in trouble when we make our calculations of the expected rate of double BH mergers coming from binary massive star systems: only a fraction $\sim 10^{-4}$ of the binary systems can result in a double BH coalescence. Whatever the factors that make this such a small fraction, we could imagine that slight adjustments of parameters could make that fraction 10^{-3} or 10^{-5} or something even farther away from 10^{-4}.

Armed with this somewhat disturbingly small fraction, we can now ask about the specific processes that might make it tough for binaries to form coalescing double BH systems. The fine line that must be walked to result in a compact object merger is that the stars must begin far enough apart that they do not merge before both are compact objects, but close enough together that the final double compact object binary can then merge within a few billion years under the influence of gravitational radiation alone. The study of the evolution of massive binaries is particularly difficult because observational evidence is tough to obtain: massive stars are rare and short-lived, and the most critical evolutionary phases for compact object mergers occupy very small fractions of the short lives of these systems.

We are therefore largely dependent on theory to tell us what is likely to happen. From the first lecture, a binary of two $\sim 10\ M_\odot$ black holes needs to get to a semimajor axis around 0.1 au or less to coalesce within the age of the universe. But that's a problem! Massive stars will under most circumstances expand out to be giants after they run out of hydrogen in their cores (an exception might be if they rotate rapidly enough to continue to cycle hydrogen into the core). The size of the giants is usually several au, i.e., tens of times larger than the needed final separation. This does raise an interesting possibility: a pair of massive stars that are initially much too far separated to spiral in via gravitational wave emission can, in the "common envelope" phase (where the envelope of a giant encompasses its companion), be dragged much closer together. If the pair begins too close together, it might merge; if one of the stars was already a compact object, it could then reside in the center of the other star and thus form a hypothesized "Thorne-Żytkow object", but it will not produce a compact binary. Thus binaries need to start their lives far enough apart to avoid merger, but not so far apart that common envelope drag is insufficient to reduce the separation to a tenth of an au or so.

Unfortunately, the common envelope phase is *very* difficult to understand from a purely theoretical point of view and also extremely challenging to observe. Thus the uncertainties are huge. In fact, there have been times when different treatments of common envelopes have given rate estimates (say, for double black hole binaries) that differ by more than two orders of magnitude!

That's not the only problem, either. For example, both neutron stars and black holes are produced by core-collapse supernovae. When we look at neutron stars it is clear that many of them have received kicks (i.e., net linear momentum) because of the core collapse. There is also evidence of supernova kicks for some black holes. However, the origin of these kicks is not known, and neither are the kick direction or the kick magnitudes as a function of the compact object mass (perhaps neutron stars are often kicked at hundreds of km s^{-1}, with some exceptions, and black holes are kicked at tens of km s^{-1}). Neutron star kicks are easily capable of separating the original binary; indeed, most neutron stars that we know are single. It's not as clear for black holes.

Prior to the direct detection of gravitational waves, the best hope was to look at double compact object systems in our Galaxy and tune the parameters of binary evolution models to agree with those systems as well as possible. The problem at this stage is that there are only ~ 20 double neutron star systems known in our Galaxy, and no known binaries in our Galaxy have two black holes or a black hole and a neutron star. Binary evolution models aren't simple, and the interpretation of the aftermath systems is far from easy, so these models are very underdetermined.

But now that we have detected close to 100 double compact coalescences, why should we worry about having too small a number of double compact object binaries in our Galaxy?

It might not be a problem *if* your perspective is that double compact object sources in gravitational waves come from the evolution of massive binaries, and thus that when we observe coalescences we can simply use them to constrain the parameters that go into the models. But it is a problem if we would like to distinguish one formation channel (in this case, from isolated massive binaries) from another (such as dynamical channels). To distinguish channels it is necessary that there be clear predictions. Unfortunately, there are so many parameters involved in the formation channels that as new detections have rolled in, the models have been able to accommodate everything.

One can hope that with the passage of time it will be possible to nail down some of the many uncertainties in the isolated binary formation channel. Some of this can come from observations, e.g., of the distribution of masses, mass ratios, and semimajor axes of massive binaries. But it appears that more will have to come from theory: for example, what can we learn about the common envelope phase, kicks, and so on? We could also lean on the hope that as more detections are made, then either (1) nontrivial features in the mass, mass ratio, or spin distributions will become clear, or (2) certain specific events emerge that can only happen in one channel or another.

I admit that I'm pessimistic because the models have proven to be far more flexible than was apparent prior to the detections. As one example, prior to any detections when binary modelers were asked what their model could *not* do, the most common answer was that they could not produce black holes that would have come from the pair instability gap. So let's discuss what that is, and then the current state of the observations.

As we know, as stars evolve they first fuse hydrogen to helium, and if they are massive enough they proceed to fuse helium to carbon, carbon to oxygen, oxygen to neon, and so on up to iron. Normally we think of the progression to iron, followed by a core collapse and the formation of a neutron star or black hole. But it is expected theoretically that if you have a massive enough star, then by the time that oxygen fusion is needed to support the star against gravity, there is an instability. The basis of the instability is that when the mass is high, the required oxygen fusion rate requires a temperature large enough that at the high-energy tail of the associated blackbody function, photon-photon interactions can produce electron-positron pairs. The problem is that these pairs will have low momentum, which means that they will produce less pressure than the photons, and thus less support against gravity. Thus the oxygen core shrinks to produce a higher rate of fusion, which increases the rate of production of low-pressure pairs. The result is that the oxygen core shrinks quickly, undergoes explosive fusion, and blows the star to bits without leaving any remnant. When the star is *very* massive, then the temperature is so high that heavy nuclei can be split by photons, which takes energy, and therefore the core can't support itself and it collapses into a black hole.

The net outcome is that, prior to any detections, there was a firm prediction that the isolated binary channel could not produce black holes in roughly the 50 $M_\odot$ to 110 $M_\odot$ range. We know that black holes can be produced at lower masses, and it is expected that they could also be produced at higher masses. Thus the prediction was that there would be a mass gap.

And then, lo and behold, black holes were detected in the mass gap! For instance, GW190521 has best-estimated masses of 66 $M_\odot$ and 85 $M_\odot$, which means that *both* components were apparently in the mass gap. Other examples exist as well. But rather than showing definitively that at least GW190521 had to come from some other channel, this event forced people to confront some of the fuzziness of what had previously been thought to be definite predictions. Some of this comes from unavoidable statistical uncertainty; it is difficult to measure both masses accurately rather than individual combinations of the two masses that depend on the mass and frequency range. Thus, for example, it was proposed that one mass was below the mass gap and one was above; you could make it work, although it's not the preferred solution. But other caveats are in the physics, and could have been investigated prior to detection. The detections provided the motivation.

For example, the actual boundaries of the predicted mass gap turn out to depend on details of nuclear physics. In particular, the rate at which carbon-12 combines with alpha-particles to make oxygen and release a photon (or as the cool nuclear kids would represent it, $^{12}\mathrm{C}(\alpha,\gamma)^{16}\mathrm{O}$) makes a difference, within current laboratory uncertainties, of tens of solar masses. People have also talked about rare scenarios that could lead to black holes with higher mass than expected via this channel (e.g., rapid rotation of the star, or the presence of a large hydrogen envelope). Frustrating though this flexibility is, it's not crazy: the rates of such events might be 1% or less of the overall rate, so you're allowed to appeal to unusual circumstances.

Overall, the isolated binary channel seems able to account for at least most of the events observed so far, and possibly all of them. Is there anything that this channel *can't* do? That's something that the modelers will eventually have to confront.