The LIGO/Virgo detections

The dreams of a generation of gravitational-wave physicists and astrophysicists came true on 14 September 2015 at 09:50:45 UTC when a strong signal appeared in the data stream of Advanced LIGO, a few days *before* the official beginning of its first science run. This event, dubbed GW150914 (for 2015 September 14), not only launched the era of gravitational wave astrophysics but also presented us with new ways to test the properties of extreme gravity and informed us of the existence of stellar-mass black holes twice as massive as we had ever seen before. Since that time, roughly 100 more BH-BH events have been seen, including ones that confounded expectations by probably being in the "pair instability gap" (more below) along with two NS-NS events and two probable NS-BH events. Gravitational wave astrophysics has begun.

Recall that the passage of a gravitational wave stretches and shrinks distances. When a wave goes past a resonant detector such as the ones Joe Weber constructed, or the Mário Schenberg spherical detector in Brazil, the bar or sphere is driven at the frequency of the wave. If the wave has close to the detector's resonance frequency then the amplitude to which the detector is driven can be enough to pick up via sensors on the surface. When a wave goes past a laser interferometric detector such as LIGO, Virgo, or KAGRA (the Japanese detector), the changing arm lengths result in shifting interference patterns that are seen at the output port. Either way, the motions to be measured are minuscule. For example, at its peak the strong event GW150914 reached a strain amplitude, which is the fractional distance change, of only $\approx 10^{-21}$, which means that the 4 km arm lengths of LIGO changed by roughly 1/200 of the radius of a proton! Decades of technological development and instrumental innovation were required for the detection. Resonant mass detectors tend to be most sensitive at ~ 3000 Hz and ground-based interferometers quote a sensitivity range of $\sim 15 - 3000$ Hz.

The implications of the detected events are profound. They include:

- The first clean tests of the predictions of general relativity in extreme gravity. When the best-fit waveform is subtracted from the GW150914 data, less than 4% of the signal remains. This suggests that there is not substantial room for significant deviations from general relativity. On the other hand, because alternate theories of gravity have not been pursued to the extent that their merger waveforms can be determined numerically, there are still non-GR theories that survive. However, such theories tend to be those that have an adjustable parameter that, effectively, goes to zero when GR is correct, and thus are generalizations of rather than genuine alternatives to GR.

- The discovery of theheaviest stellar-mass black holes known. Prior to GW150914 the highest mass established for a stellar-mass black hole was $\sim 15\ M_\odot$, but the best estimate for the components of GW150914 were $\sim 29\ M_\odot$ and $\sim 36\ M_\odot$. When they merged, the final mass was $\sim 62\ M_\odot$, which is less than the sum of the original masses because $\approx 3\ M_\odot c^2$ in energy was radiated in gravitational waves. But then GW190521 really surprised people, with estimated masses of $\sim 66\ M_\odot$ and $\sim 85\ M_\odot$. This leads to...

- ...lots of black hole astrophysics. At the population level, we're learning a lot about the rate of such events, and are finding that the black holes involved in these coalescences typically spin slowly compared with the maximum possible spin (in contrast to the black holes we observe in our Galaxy). The higher observed masses of LIGO/Virgo black holes compared with what we see in our Galaxy might just be due to selection effects. Higher masses are easier to see in gravitational waves because they produce stronger signals. And lower masses are easier to see in the Galaxy because supernovae can kick lower-mass black holes farther out of the Galactic plane, which makes them easier to study carefully because they aren't as obscured by dust.

- But GW190521 was surprising for another reason. Prior to its detection, most modelers who believe that BH-BH mergers come from the evolution of massive binaries did not think that it would be possible to have black holes between roughly $50\ M_\odot$ and $110\ M_\odot$. The reason has to do with something called the "pair instability" expected in very massive stars. Think about the nuclear evolution of a massive star. First, hydrogen fuses to helium in the core. Then helium into carbon, carbon into oxygen, and so on. The heavier the star is, the hotter the core has to be, because with a heavier overburden of matter the rate of fusion needs to be greater to support the matter. For stars that would otherwise produce black holes in the $50 - 100\ M_\odot$ range, the idea is that in the phase where oxygen is being fused to produce neon, the temperature is so high that some of the photons produce electron-positron pairs. The problem is that these pairs have very low kinetic energy and thus low pressure. As a result, to support the weight of the matter above, the temperature has to go up, which means that a greater fraction of photons can produce pairs, et cetera. This produces runaway fusion that blows the whole star apart and leaves no remnant.

  Or, at least, that was the idea. GW190521 seems to have two black holes that are both inside the gap. Ideas since have been various: (1) this system formed in a different way, e.g., dynamical (see below), (2) the measurements are imprecise so maybe actually one component was below the gap and the other was above, (3) oops, we forgot this or that effect that actually does make it possible for black holes with such masses to form in massive binary evolution, despite previously being absolutely convinced that it was impossible. To be a little more fair to proponents of (3), only a small fraction of events have masses in the "forbidden" zone, so maybe rare things do happen on occasion.

- The crown jewel of all of the observations: a double neutron star merger (GW170817) that (1) had an associated off-axis short gamma-ray burst and thus dramatically confirmed a long-standing idea for short GRB formation, (2) had a "kilonova" which was powered by outflow of matter from the merger, and which had properties that aligned closely with predictions from years before, (3) therefore provided strong support that at least a major fraction of the elements heavier than iron are produced in NS-NS mergers, and (4) provided a constraint (but not quite a measurement) of the tidal deformability of the two neutron stars in the merger, which thus constrains the properties of their core matter. Pretty amazing, and incidentally when this event happened it had what was arguably the most intense worldwide electromagnetic follow-up of any astronomical event ever. A second probable NS-NS merger was seen later, and two probable BH-NS mergers, but none of them were strong enough or localized enough for serious follow-up study.

Origins of double compact object systems

The origin of these events is still being debated. The two basic possibilities are:

1. Massive binary evolution. This one seems straightforward: we already know that massive stars are almost always in binaries, and that their binary companions are typically of comparable mass to the primary star. So, what could be simpler? Start with two stars, each of which are massive enough to evolve into a neutron star or black hole. Wait a while. They do so evolve, and then you have a compact binary! But as we discussed in the previous lecture the problem is that you need to evolve into a double compact binary that is close enough to merge, via emission of gravitational radiation, within several billion years. For two 10 $M_\odot$ black holes, that's about 0.1 au separation, and it's less than that for neutron stars because neutron stars are less massive. But in its natural evolution, even a star such as the Sun, let alone a more massive star, will become much larger than 0.1 au. Thus you'd expect the envelope of one star to encompass the other (twice, once for each evolution); this is called a common envelope phase, as we discussed before. Can you have the two stars be dragged close enough to merge in some billions of years, but not so close that they form a Thorne-Żytkow object? Are the rates reasonable? The masses? The spin distributions? We can't say, because the models have so many parameters that you can get pretty much anything you like.

2. Dynamical evolution. Again, we discussed this somewhat in the last lecture. It makes sense that binary-single interactions of various types could lead to an enhanced production of merging compact binaries. But can this account for the properties of the observed population? One hint might be in the spins. Dynamical interactions should be random in direction, so we'd expect that the black hole spins would have random orientations relative to their orbital axes. There are hints that it is more likely that they will be pointed in the direction of the orbital axis rather than in the opposite direction. It's tough to tell because the spin parameters are small, but if this holds up it would be a strike against the dynamical models.

3. Something that got people excited but now is fading away is the possibility that some of the events come from primordial black holes, formed before Big Bang nucleosynthesis. The thrill here would be that primordial black holes might be dark matter. That would be awesome, and might still be true, but what seems to be the case is that if a large fraction of dark matter is primordial black holes of tens of solar masses, then we'd have major contradictions with various lines of evidence (e.g., the LIGO/Virgo rate would be higher than it is, and so would various types of gravitational lensing events). Still, hope springs eternal.

The future of gravitational wave astrophysics is remarkably bright. LIGO/Virgo/KAGRA are now preparing for their next science run, dubbed O4, which is planned to start in March 2023. The design sensitivity of Advanced LIGO is about 2.5 times greater than what it had in its initial run, which means that the volume of sources probed will be about 15 times larger than it was initially. The space-based detector LISA will launch in 2037 and will probe frequencies of $\sim 10^{-4} - 10^{-1}$ Hz, where it will find merging $10^4 - 10^7\ M_\odot$ black holes, extreme mass ratio inspirals of $\sim 10\ M_\odot$ black holes into supermassive black holes, and double white dwarf binaries that might be the precursors of Type Ia supernovae. Pulsar timing arrays are already operating in North America, Europe, and Australia, and they are sensitive to a (likely) stochastic background of orbiting (but not coalescing) $10^8 - 10^{10}\ M_\odot$ black holes in the $10^{-9} - 10^{-6}$ Hz range. There are already intriguing hints of a signal but, unlike with the ground-based detectors, it's a matter of steady accumulation of data rather than waiting for that really bright event. Finally, several sensitive cosmic microwave background polarization experiments are operating or planned that will look for the characteristic B-mode polarization of gravitational waves at $10^{-17} - 10^{-15}$ Hz (a detection was claimed in 2014 but it turned out to be dust). Overall, gravitational wave science will cover 20 orders of magnitude in frequency, which is essentially the same coverage as electromagnetic astronomy. Insight and surprises await!

For more on my thoughts regarding the first LIGO detection, please read http://www.astro.umd.edu/∼miller/reprints/miller16d.pdf (Miller 2016, GRG, 48, 95).