Binaries and implications of the LIGO detections

For our final lecture we will explore binary systems in general and the Advanced LIGO detections in particular. Binaries obviously have a large and varying quadrupole moment, and they have the additional advantage that even before the detections we knew that gravitational radiation is emitted from them in the expected quantities (based on observations of double neutron star binaries in the pre-LIGO era). The characteristics of the gravitational waves from binaries, and what we could learn from them, depend on the nature of the objects in those binaries. We will therefore start with some general concepts, then discuss individual types of binaries, and then wrap up with the specifics about the LIGO detections.

Suppose that the binary is well-separated, so that the components can be treated as points and we only need take the lowest order contributions to gravitational radiation. Temporarily restricting our attention to circular binaries, how will their frequency and amplitude evolve with time?

Let the masses be m_1 and m_2, and the orbital separation be R. We argued in the first lecture that the amplitude a distance $r \gg R$ from this source is $h \sim (\mu/r)(M/R)$, where $M \equiv m_1 + m_2$ is the total mass and $\mu \equiv m_1 m_2/M$ is the reduced mass. We can rewrite the amplitude using $f \sim (M/R^3)^{1/2}$, to read

$$\begin{aligned} h &\sim \mu M^{2/3} f^{2/3}/r \\ &\sim M_{\rm ch}^{5/3} f^{2/3}/r \end{aligned} \tag{1}$$

where $M_{\rm ch}$ is the "chirp mass", defined by $M_{\rm ch}^{5/3} = \mu M^{2/3}$. The chirp mass is named that because it is this combination of μ and M that determines how fast the binary sweeps, or chirps, through a frequency band. When the constants are put in, the dimensionless gravitational wave strain amplitude (which, remember, is roughly the fractional amount by which a separation changes as a wave goes by) measured a distance r from a circular binary of masses M and m with a binary orbital frequency $f_{\rm bin}$ is (Schutz 1997)

$$h = 2(4\pi)^{1/3} \frac{G^{5/3}}{c^4} f_{\rm GW}^{2/3} M_{\rm ch}^{5/3} \frac{1}{r}\ , \tag{2}$$

where $f_{\rm GW}$ is the gravitational wave frequency. Redshifts have not been included in this formula.

The luminosity in gravitational radiation is then

$$\begin{aligned} L &\sim 4\pi r^2 f^2 h^2 \\ &\sim M_{\rm ch}^{10/3} f^{10/3} \\ &\sim \mu^2 M^3/R^5\ . \end{aligned} \tag{3}$$

The total energy of a circular binary of radius R is $E_{\rm tot} = -G\mu M/(2R)$, so we have

$$\begin{aligned} dE/dt &\sim \mu^2 M^3/R^5 \\ \mu M/(2R^2)(dR/dt) &\sim \mu^2 M^3/R^5 \\ dR/dt &\sim \mu M^2/R^3\ . \end{aligned} \tag{4}$$

What if the binary orbit is eccentric? The formulae are then more complicated, because one must average properly over the orbit. This was done first to lowest order by Peters and Matthews (1963) and Peters (1964) by calculating the energy and angular momentum radiated at lowest (quadrupolar) order, and then determining the change in orbital elements that would occur if the binary completed a full Keplerian ellipse in its orbit. That is, they assumed that to lowest order, they could have the binary move as if it experienced only Newtonian gravity, and integrate the losses along that path.

Before quoting the results, we can understand one qualitative aspect of the radiation when the orbits are elliptical. From our derivation for circular orbits, we see that the radiation is emitted much more strongly when the separation is small, because $L \sim R^{-5}$. Consider what this would mean for a very eccentric orbit $(1-e) \ll 1$. Most of the radiation would be emitted at pericenter, so this would have the character of an impulsive force. With such a force, the orbit will return to where the impulse was imparted. That means that the pericenter distance will remain roughly constant, while the energy losses decrease the apocenter distance. As a consequence, the eccentricity decreases. In general, gravitational radiation will decrease the eccentricity of an orbit.

The Peters formulae bear this out. If the orbit has semimajor axis a and eccentricity e, their lowest-order rates of change of the orbital parameters are

$$\left\langle \frac{da}{dt} \right\rangle = -\frac{64}{5} \frac{G^3 \mu M^2}{c^5 a^3 (1-e^2)^{7/2}} \left(1 + \frac{73}{24} e^2 + \frac{37}{96} e^4\right) \tag{5}$$

and

$$\left\langle \frac{de}{dt} \right\rangle = -\frac{304}{15} e \frac{G^3 \mu M^2}{c^5 a^4 (1-e^2)^{5/2}} \left(1 + \frac{121}{304} e^2\right)\ , \tag{6}$$

where the angle brackets indicate an average over an orbit. One can show that these rates imply that the quantity

$$a e^{-12/19} (1-e^2) \left(1 + \frac{121}{304} e^2\right)^{-870/2299} \tag{7}$$

is constant throughout the inspiral. If we ignore the final factor (which is always between 0.88 and 1), we can write this as $a(1-e)(1+e)e^{-12/19} \approx$ const. For high eccentricities such that $1-e \ll 1$, $1+e$ and $e^{-12/19}$ are roughly constant, so $a(1-e) = r_p \approx$ const, which means that the pericenter distance r_p is roughly constant as promised. For low eccentricities such that $1-e^2 \approx 1$, we get $ae^{-12/19} \approx$ const. The orbital frequency (which is half the dominant gravitational wave frequency when $e \ll 1$) is $f \propto a^{-3/2}$, which means that $f \propto e^{-18/19}$, or roughly $e \propto f^{-1}$. Thus for low eccentricities, the eccentricity roughly scales as the reciprocal of the frequency. This means that binary sources at the high frequencies detectable using LIGO can usually be considered to be effectively circular.

How compact binaries can merge

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, and direct dynamical capture.

Let's talk first about isolated massive binaries. 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. As we now discuss, 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. 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). Thus it is possible that 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 few tenths of an astronomical unit (which is needed for the inspiral to take a few billion years or less).

Unfortunately, the common envelope phase is *very* difficult to understand from a purely theoretical point of view, and given that no binary has ever been seen *in* a common envelope state, the uncertainties are huge. In fact, over the last decade plus 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}).

The best that can be done is to look at the final systems of two compact objects 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 ~ 10 double neutron star systems known, and no known binaries in our Galaxy with 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.

An interesting side point is that when double neutron star merger rates are computed or used to calibrate model parameters, one such system, which will merge in a few billion years, is not included in the calculations. That poor, unappreciated system is ignored because it is in a globular cluster. I love globular clusters, so I am outraged, outraged I tell you, at this blatant discrimination!

But should I be? The estimated rate of mergers in the body of our Galaxy is about 10–100 per million years, with large uncertainties. What rate might we expect from globulars? Suppose that every one of the ~ 100 globular clusters around our Galaxy had, initially, 200 neutron stars (probably a large overestimate) and they all merge within 10 billion years (certainly overly optimistic!). Then there are $100 \times 100 = 10^4$ mergers in 10^{10} years, for a rate of 1 per million years. We can ignore this because the disk contribution is much greater. This is *not* true of estimates of double black hole or BH-NS mergers, because as indicated above we don't know any examples of such systems in our Galaxy and thus possibly formation channels in globulars dominate.

It is thus worth taking a brief diversion to discuss what might be different about globulars, and indeed this will lead us into the second general path to mergers: dynamical processes.

The main difference between the main part of our Galaxy and globular clusters (or systems such as nuclear star clusters) is stellar number density; in the Solar vicinity there are roughly 0.15 stars per cubic parsec, but in the center of the densest globulars the density can be 10^6 per cubic parsec. This still isn't enough to have stars collide directly with each other very often, but it does mean that binary systems, which act as if their collision cross sections are the sizes of the orbits, can have collisionless three- and four-body encounters. That can be significant. For example, in a standard semi-rich globular with a velocity dispersion of 10 km s^{-1}, stars pass within 1 AU of each other (and hence binaries with radii of 1 AU have strong encounters) once per few hundred million years on average in a cluster of number density 10^5 per cubic parsec. Thus over the 10^{10} year lifetimes of these clusters, binaries can undergo tens of such interactions.

The interactions are chaotic, but computer simulations show that when a binary and single interact, the binary that emerges from the interaction tends to contain the two most massive of the three original objects. Thus neutron stars and black holes, which are more massive than the average star in a globular, can swap into binaries and eventually find compact objects as companions. Therefore, per stellar mass, globulars are expected to have far more of these systems than the low number density bulk of their host galaxies. For the same reason, there is a high rate per mass in globulars of low-mass X-ray binaries and millisecond pulsars (which are thought to have been spun up by accretion from a companion).

Another popular process that can happen with multiple stars is the Kozai (or Kozai-Lidov) resonance. Kozai and Lidov discovered independently in the early 1960s that if a binary is orbited by a third object in a hierarchical triple (such that the system has long term stability), and the binary orbital axis is strongly tilted with respect to the orbit of the tertiary, then over many orbits of both the binary and the tertiary, the relative inclination of the binary to the tertiary cycles between low and high values, while conserving the semimajor axis. Most importantly for merger possibilities, when the inclination goes down the eccentricity goes up and vice versa. Thus in the right range of orientations the binary could be driven to such a high eccentricity that gravitational radiation grinds it down to merger. Careful observations of massive binaries in our Galaxy suggest that 10% or more of them could actually be triples, so in principle such systems could evolve naturally to high-eccentricity states (although if the system is susceptible to such evolution it is likely that it would be driven to collisions on the main sequence or giant branch rather than when the objects become compact). In dense stellar systems such as globular clusters, hierarchical Kozai-susceptible triples can be created *after* evolution to compact objects, for example as an outcome of binary-binary interactions.

Yet another possibility, although one with much smaller probability, is that two initially unbound compact objects could pass close enough to each other that the gravitational radiation they emit during their closest passage carries away enough energy to bind the objects together. They would then coalesce quickly. The reason that this is a very low-probability event is that the objects would have to come very close to each other to radiate the required energy, and thus the cross section for the process is tiny. If this happens, it seems most likely to happen during the many chaotic interactions that occur during a binary-single interaction than during a random encounter between two single objects. In any case, high stellar densities are obviously important for this mechanism to have any chance.

All this means that globulars, and the nuclear star clusters in the centers of galaxies, could be ripe breeding grounds for the BH-BH and BH-NS systems that we have not yet seen in our Galaxy.

Binaries in strong gravity

When two masses are close enough to each other, the Peters formulae do not quite describe their motion. Instead, there are additional terms corresponding to higher order moments of the mass and current distributions: the octupole, hexadecapole, and so on. This is often expressed in terms of equations of motion that include the Newtonian acceleration and a series of "post-Newtonian" (PN) terms. The order of a term is labeled by the number of factors of M/r by which it differs from Newtonian: for example, the 1PN term is proportional to M/r times the Newtonian acceleration. Since $v^2 \sim M/r$ in a binary orbit, there can also be half-power terms. For point masses (which thus do not have any spin) the first several corrections are at the 1PN, 2PN, 2.5PN (this is where gravitational radiation losses first enter), 3PN, and 3.5 PN orders.

The equations of motion have been fully, rigorously established up to 3PN order, but the algebra is daunting. We note that, fortunately, tidal effects only enter at the 5PN order, which one can justify by realizing that tidal couples have a $1/r^6$ dependence, or five powers of r greater than the Newtonian potential. Therefore, for many purposes, tidal effects can be neglected. A purpose for which they can *not* be neglected is the calculation of tidal effects on waveforms to determine neutron star radii (really their tidal deformability), as we discussed in the third lecture. The post-Newtonian approach is useful, but problematic because succeeding terms are not much smaller than the terms before them. Another way to put this is that the Newtonian acceleration is overwhelmingly dominant for an extremely wide range of separations (out to infinity, in fact), but the range in which the 1PN term is necessary but the 2PN term is negligible is small, and this becomes even more true for the higher order terms. One can therefore often make good progress by taking the lowest-order term, and since the 2.5PN term is the lowest-order that involves energy and angular momentum loss, one can use the Newtonian plus 2.5PN terms. However, more terms turn out to be necessary to get sufficiently accurate waveforms for analysis of future gravitational wave data streams.

Various clever attempts have been made to recast the expansions into forms that converge faster than Taylor series. For example, a path adopted by Damour and Buonanno is to pursue effective one-body (EOB) spacetimes in which an effective test particle moves, and to then graft on the effects of gravitational radiation losses. They also use Padé resummation, in which the terms are ratios of polynomials, in the hopes that this can more naturally model the singularity of black hole spacetimes. When this method is calibrated with the results of numerical simulations it matches those simulations extremely well. Even when tidal deformability is taken into account, the EOB approach appears at this time to be able to explain the waveforms with just that single extra parameter. This is highly promising, but as always more simulations and comparisons are being performed.

One interesting effect that emerges from the higher-order studies of binary inspirals is that gravitational radiation carries away net linear momentum, and thus the center of mass of the system moves in an ever-widening spiral. We can understand this as follows (following an idea of Alan Wiseman). In an unequal-mass binary, the lower-mass object moves faster. As the speed in orbit becomes relativistic, the gravitational radiation from each object becomes beamed, with the lower-mass object producing more beaming because it moves faster. Therefore, at any given instant, there is a net kick against the direction of motion of the lower-mass object. If the binary were forced to move in a perfect circle, the center of mass of the system would simply go in a circle as well. However, because in reality the orbit is a tight and diminishing spiral, the recoil becomes stronger with time and the center of mass moves in an expanding spiral. Note that by symmetry, equal-mass nonspinning black holes can never produce a linear momentum kick, and that if the mass ratio is gigantic the fractional energy release is small and therefore so is the kick. For nonspinning holes, the optimal ratio for a kick is about 2.6, and the maximum kick speed is a bit below 200 km s^{-1}.

It is important to realize that the preceding argument by Wiseman is only heuristic: it gives us an idea of the direction of the recoil and some of the dependences of the magnitude of the effect. There is, however, one critical respect in which this argument is quite misleading. We talked about "the gravitational radiation from each object" as if the situation were similar to what it would be if each object emits independent gravitational radiation. But they don't. The wavelength of the gravitational radiation is of order the speed of light divided by the orbital frequency. Because the objects orbit slower that the speed of light, this guarantees that the gravitational waves have a *longer* wavelength than the size of the system. Thus you can't actually localize gravitational waves as coming from one object or the other; it really is the system as a whole. We encounter such situations frequently in astrophysics: we can produce an argument that gives us some insight, but we have to be careful not to push heuristic arguments farther than their domain of validity.

Kicks are potentially important astrophysically because if the final merged remnant of a black hole inspiral is moving very rapidly, it could be kicked out of its host stellar system, with possibly interesting implications for supermassive black holes and hierarchical merging. There have therefore been a number of calculations of the expected kick. The net result is that spins can increase the kick a *lot* if they are not parallel or antiparallel to the orbital axis. Work by Lousto et al. shows that an optimal kick configuration for two maximally rotating black holes that approach each other in an initially quasicircular orbit can reach almost 5000 km s^{-1}! That would be enough to eject the remnant out of any galaxy in the universe. We have argued (Bogdanović, Reynolds, and Miller 2007; Miller & Krolik 2013) that when significant gas is present the torques from the gas on binary supermassive black holes tend to align the spin axes with the orbital axis, which then lead to much lower kicks (again, < 200 km s^{-1}) and are therefore more consistent with the lack of definitive candidates for recoil. Studies are ongoing.

Generically, if two black holes coalesce, how does it happen? In this field it is standard (and reasonable) to divide the whole process up into three stages. The first stage is inspiral, which follows the binary from large separations to when the binary has reached the stage of dynamical instability. That is, inspiral is roughly where the binary is outside the innermost stable circular orbit, so the motion is mostly azimuthal. Inside the ISCO, the motion becomes a plunge, and this happens on a dynamical time scale. As the event horizons disturb each other and finally overlap, the spacetime becomes extremely complicated and must be treated numerically. This is called the merger phase.

Ultimately, of course, the "no hair" theorem guarantees that the system must settle into a Kerr spacetime. It does this by radiating away its bumpiness as a set of quasinormal modes. The lowest-order, and longest-lived, of the modes is the $l = 2,\ m = 2$ mode. When all but this mode have essentially died away, the system has entered the period of ringdown. With only a single mode left, the ringdown phase can be treated numerically. The result is that the frequency $f_{\rm qnr}$ of the gravitational radiation, as well as the quality factor $Q \equiv \pi f_{\rm qnr}\tau$ (where τ is the characteristic duration of the mode; this measures how many cycles the ringing lasts) depend on the effective spin $j \equiv cJ/GM^2$ of the final black hole (sometimes $\hat{a}$

is used instead of j). Echeverria (1989) gives fitting formulae valid to $\sim 5\%$:

$$\begin{aligned} f_{\rm qnr} &\approx [1 - 0.63(1-j)^{0.3}](2\pi M)^{-1} \\ Q &\approx 2(1-j)^{-0.45}\ . \end{aligned} \tag{8}$$

Thus more rapidly spinning remnants have higher frequencies and last for more cycles. This could allow identification of the spin based on the character of the ringdown.

We can make rough estimates of the energy released in each phase as a function of the reduced mass μ and total mass M of the system. Since the inspiral phase goes from infinity to the ISCO, the energy released is simply μ times the specific binding energy at the ISCO, so $E_{\rm inspiral} \sim \mu$. What about the merger and ringdown phases? We know that the strain amplitude is $h \sim (\mu/r)(M/R)$, where r is the distance to the observer and R is the dimension of the system. For the merger and ringdown phases, $R \sim M$, so $h \sim \mu/r$. We also know that the luminosity is $L \sim r^2 h^2 f^2$, so $L \sim \mu^2 f^2$, and if the phase lasts a time τ then the total energy released is $E \sim \mu^2 f^2 \tau$. But the characteristic frequency is $f \sim 1/M$ and the characteristic time is $\tau \sim M$, so we have finally $E \sim \mu^2/M$, or a factor $\sim \mu/M$ times the energy released in the inspiral. The exact values for a particular mass ratio are somewhat in dispute, but for an equal-mass nonspinning black hole binary, $E_{\rm inspiral} \sim 0.06M$ and $E_{\rm merger}$ and $E_{\rm ringdown}$ are probably $\sim 0.01M$. Note that for highly unequal mass binaries ($\mu \ll M$), the inspiral produces much greater total energy than the merger or ringdown. This is one reason why analyses of extreme mass ratio inspirals have ignored the merger and ringdown phases.

Estimates of compact binary merger rates

The two basic approaches have been to (1) use the systems directly, in the sense that we estimate rates based on what we see and then extrapolate to the number of similar systems that exist in the Galaxy, (2) use the systems indirectly, to calibrate population synthesis models. We will discuss (1) first, then briefly define and discuss population synthesis models.

Ask class: how would they use the observed population directly? One way is to figure out how probable, in some sense, a given system would be to be observed. Then you use this to estimate how many similar systems are out there that are *not* observed because we're looking at the wrong time or we are in the wrong place. The more improbable a system is, the greater weight it has, because you figure that for each one you see there are many you don't see.

A classic early reference for this kind of analysis is Phinney 1991, ApJ, 380, L17. In this work, my graduate advisor noted:

"If d binary neutron star systems i, each of total lifetime $\tau(i)$, are detected in surveys j which could have detected pulsar i in a volume $V_{\rm max}(i) = \sum_j V_{j,\rm max}(i)$, the merger rate in the Galaxy can be estimated as

$$R = \sum_i^d \frac{V_{\rm Gal}}{V_{\rm max}(i)\tau(i)} , \tag{9}$$

where $V_{\rm Gal}$ is the volume of the Galaxy. If the pulsars are not uniformly distributed, $V_{\rm max}(i)$ and $V_{\rm Gal}$ are to be weighted by pulsar density."

You can ponder the reasonableness of these factors. The translation is that if a binary neutron star system is difficult to see (e.g., the system is dim, so that $V_{\rm max}(i)$ is small, or it is short-lived, so that $\tau(i)$ is small) then that system gets a weighting factor that is larger than if a system is easy to see. Does that make sense? Phinney argued that $\tau(i)$ should be the *total* lifetime of the system; does *that* make sense, or should there be an extra weighting factor if the system is very close to merger, or very close to its formation time, or in some other special phase of its existence?

Given that we now have detections of two, and possibly three, double black hole systems, you might ask how you would estimate merging rates directly from the observations. A logical analog to the formula above is that the best estimate of the rate per volume per time given d detections should be

$$R_{\rm LIGO} = \sum_i^d \frac{1}{V_{\rm max}(i)T} , \tag{10}$$

where $V_{\rm max}(i)$ is again the volume in which event i could have been detected, but T is the *total* observation time. Is this reasonable, or are there other considerations that should enter? In the problem set you'll have an opportunity to make your own estimates using the properties of the LIGO detections.

Another approach to take regarding rates (of NS-NS or NS-BH mergers, which we haven't yet seen, and of BH-BH mergers, which we have seen) is population synthesis. Population synthesis refers to the practice of simulating a large number of objects/systems of interest, simulating observations of them, comparing that to what you do see, and inferring something about the proper inputs to your model. This can be very productive when we have an excellent and validated understanding of the relevant theory, such as in the luminosities and spectra of stars. Then population synthesis can be used to characterize the stellar population of a galaxy based on the overall luminosity and spectrum.

In my opinion we are farther away from population synthesis being helpful in rate estimates for compact object mergers. The reason is that massive binary evolution models are extremely complex, and we simply don't have enough systems to observe that we can constrain the parameters and test the models, at least at this point. However, I do think that population synthesis models will be helpful to interpret binary results from ground-based detectors. The reason is that if we get to a point where we are detecting tens of events per year, then a framework is necessary for these data; in a vacuum, data don't mean anything, so we have to determine what it all implies. At that stage, population synthesis could therefore play an important role.

The LIGO 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, 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. A subsequent definite event on 26 December 2015 (thus GW151226) was reported on 15 June 2016, and a likely but not definite event was also recoreded on 12 October 2016 (called LVT151012). 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 Mário Schenberg spherical detector, the 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 Advanced LIGO, the changing arm lengths result in shifting interference patterns that are seen at the output port. Either way, the motions to be measured are miniscule. 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 Advanced LIGO changed by roughly 1/200 of the radius of a proton! The weaker event GW151226 reached only about a third of that amplitude. 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 two definite events and one candidate event 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. Such theories include any that deviate from GR only in the presence of matter; for those, confirmation of a NS-NS or NS-BH GR waveform will be critical. Interestingly, because the GW151226 event involved significantly lower-mass black holes than GW150914 and thus had more cycles in the Advanced LIGO sensitivity band, GW151226 provides stronger constraints on deviations from GR in the inspiral than did the stronger GW150914 event.

- The discovery of the two heaviest (and then the single heaviest!) 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. For GW151226 the masses (best guesses around 8 $M_\odot$ and 14 $M_\odot$ with large uncertainties on everything except the chirp mass $M_{\rm ch} = (m_1 m_2)^{3/5}/(m_1 + m_2)^{1/5} = 8.9 \pm 0.3\ M_\odot$) are in the range seen previously from black holes in our Galaxy. The masses for the black holes in LVT151012 was in between these two (but recall that it is not certain that this was real). As more detections roll in over the next several years, we will be able to build up a mass function.

- The rate of double black hole mergers in the local universe is being tightened rapidly. Prior to the detections, the lack of any known BH-BH binaries in our Galaxy and the enormous uncertainties in models meant that the 90% credible rate of double BH mergers was quoted as $0.1 - 300\ {\rm Gpc}^{-3}\ {\rm yr}^{-1}$. Combining the evidence from GW150914, GW151226, and LVT151012 has brought the latest 90% credible rate to $9 - 240\ {\rm Gpc}^{-3}\ {\rm yr}^{-1}$. That's a factor of 90 increase in the lower limit! Obviously, as the detectors become more sensitive and the discoveries mount, the rate will continue to be refined rapidly.

The origin of these events is still being debated. Whether it came from an isolated binary origin or a dynamical origin, the consensus at this stage is that it is easier to produce the high masses of GW150914 if the stars were in a low-metallicity environment (i.e., the fraction of elements heavier than helium was significantly less than the fraction for the Sun or for the disk of our Galaxy in general). The reason is that stellar winds are largely driven by the interaction of photons with atomic lines, and those lines are obviously more common when there is an abundance of atoms with many electrons. Thus in low-metallicity environments, massive stars can hold onto a greater fraction of their original mass, which makes it easier to form more massive black holes. It will be necessary to accumulate more detections and to look at properties such as the masses, spins, and eccentricities (although these are expected to be undetectable using Advanced LIGO) of binaries to discern their origin.

The future of gravitational wave astrophysics is remarkably bright. Advanced LIGO will start its second science run this summer, and it is hoped that the Advanced Virgo detector will join them towards the end of the run. The third run will be next year, and then runs will commence at a regular rate. 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. In addition, other interferometers such as Advanced Virgo, KAGRA, and LIGO-India, and resonant mass detectors such as Mário Schenberg and Mini-GRAIL, will join the hunt. The space-based detector LISA will launch in 2034 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. 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. 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!