LIGO/Virgo results so far

We're almost ready to assess the results from the LIGO/Virgo/KAGRA collaboration in terms of what they mean for astrophysics and gravitational physics. Our final task prior to that is to get some understanding about what information we get from gravitational waves, which will allow us to understand the limitations of current analysis.

As we learned in the first lecture, for binaries that have at least several cycles in the LVK frequency band, the primary mass combination which can be inferred is the chirp mass, $M_{\rm ch} = \mu^{3/5} M^{2/5}$, where for a binary of masses m_1 and m_2, $M = m_1 + m_2$ and $\mu = m_1 m_2 / M$. To higher order, and primarily at frequencies closer to the merger, the "symmetric mass ratio" $\eta \equiv \mu/M$ affects the waveform as well. If $M_{\rm ch}$ and η can both be measured, then it is possible to estimate the two masses independently. But $\eta = m_1 m_2/(m_1 + m_2)$ has a maximum when $m_1 = m_2$. This means that near this maximum, a small change in η implies a large change in m_1/m_2. For example, $m_1 = m_2$ gives $\eta = 0.25$; $m_1 = 1.5 m_2$ gives $\eta = 0.24$. This means that when $m_1 \approx m_2$ (which is common but not universal in the detected sources), there is a large uncertainty in m_1/m_2.

In a similar way, although the magnitude and direction of the spins of both compact objects influence the gravitational waveform, there is a particular combination that has the largest effect. To a reasonable approximation, this is the "effective spin"

$$\chi_{\rm eff} = \frac{\mathbf{S}_1/m_1 + \mathbf{S}_2/m_2}{m_1 + m_2} \cdot \hat{\mathbf{L}} \ , \tag{1}$$

that is, this is the mass-weighted projection of the two angular momenta $\mathbf{S}_1$ and $\mathbf{S}_2$ onto the orbital angular momentum, which has a unit vector $\hat{\mathbf{L}}$. Other components can come in: for example, if the spin angular momenta have components that are perpendicular to the orbit, then the orbit will precess and leave a signature in the waveform. But this effect is weaker than the orbit-aligned effect.

Finally, if one or both compact objects in a binary can be tidally deformed by an external gravitational field (i.e., if at least one of the objects is not a black hole), then that tidal deformation affects the waveform. We can understand this heuristically by thinking about energy: it takes energy to deform a star away from its natural spherical state, and that energy has to come from the orbit, which means that tidal deformation causes a binary to spiral inward faster than it would without the deformation (a weaker but related contribution is that a deformed object stirs spacetime differently than a spherical object, but this contributes only $\sim 20\%$ of the total effect). Again, in a binary, there is a combination of the two individual deformabilities that has the greatest effect. If we define the dimensionless tidal deformability of each object as

$$\Lambda_{1,2} = \frac{2}{3} k_2 \left(\frac{R_{1,2} c^2}{G m_{1,2}} \right)^5 , \tag{2}$$

where k_2 is the tidal Love number and $R_{1,2}$ are the radii of the two objects, then the combination which leaves the strongest imprint on the waveform is the "binary tidal deformability"

$$\tilde{\Lambda} = \frac{16}{13} \frac{(12q+1)\Lambda_1 + (12+q)q^4 \Lambda_2}{(1+q)^5} \tag{3}$$

where $q = m_2/m_1 \leq 1$ is the binary mass ratio.

Okay, *now* we're ready to think about the results!

Results of the first three LIGO/Virgo runs

Since LIGO began its O1 run in the fall of 2015 (with Virgo joining partway through the O2 run in summer 2017, and KAGRA starting to take data in O3), nearly 100 binary coalescences have been detected with confidence. No other category of sources (continuous, burst, or stochastic) has been seen, and as of now the effects of gravitational waves have not been seen in other frequency ranges. We will therefore focus on the LVK results for binaries. The highlights are:

1. Almost all of the events involve two black holes. Two events probably involve two neutron stars (one of which will get a major highlight below), maybe a few more have a black hole and a neutron star, and the rest are double black hole coalescences. Through O3, the 5th to 95th percentile credible rate for BH-BH mergers is $18 - 44$ Gpc^{-3} yr^{-1}, for NS-BH mergers it is $8 - 140$ Gpc^{-3} yr^{-1}, and for NS-NS mergers it is $10 - 1700$ Gpc^{-3} yr^{-1}. The rate per volume for NS-NS and NS-BH mergers can be comparable to or larger than the rate for BH-BH mergers despite their smaller number of detections, because neutron stars are much lower-mass than black holes. This means that at a given frequency their amplitude is less, so they can't be seen in as large a volume, and thus a given number of detections implies a higher rate per volume.

2. The black holes tend to have much higher masses than the black holes that we see in our Galaxy. As with the NS/BH issue, higher-mass black holes can be seen in a larger volume, so this is an observational bias. But some papers have asked the question of whether the intrinsic BH mass distribution from gravitational waves actually is different from what we see in our Galaxy. There isn't strong evidence for this at this time; heavier BH are easier to see in gravitational waves, and tougher to see in our Galaxy (because heavier black holes can't get as far out of the Galactic plane, so we see them through a

lot of gas and dust and can't estimate their masses as easily as we can for the higher-latitude lower-mass black holes). There has been speculation that it could be easier to produce high-mass stellar black holes if the metallicity is low, because then winds are weaker and less mass is lost prior to the core collapse. But there are ways of getting around that, too, and thus we can't draw strong conclusions.

3. There isn't an obvious deficit of black holes in the pair instability mass gap. The masses there are tough to measure, but we can't yet see a dip. There might be some nontrivial features in the mass distribution (e.g., it's possible that there is a drop at low masses and a bump or two), but they are subtle and we'll have to see whether those features survive when we have a lot more events.

4. Most double black hole mergers have black holes of comparable mass to each other. A few are definitely *not* mass ratios of unity, but they're rare.

5. The spin parameters of black holes detected using gravitational waves are close to zero, and do not appear to be symmetrically distributed around zero (that is, to the degree that the spins are nonzero, they appear to be more likely to be prograde with respect to the orbital axis rather than retrograde). The magnitudes of the spin parameters is a significant contrast with what we infer from electromagnetic observations of black holes in the Galaxy, where the spin parameters seem to go from roughly 0 to close to the maximum allowed value, of 1. This *might* indicate different formation channels. It is worth remembering that only a tiny, tiny fraction of stellar-origin black holes are detectable by any means at all, so even if there are some differences we can't draw a conclusion about the full population of all black holes that form from stars.

6. From the standpoint of general relativity, the agreement of the waveforms with the GR predictions indicates that even at roughly the scale of the event horizon, gravity does not deviate detectably from GR. Note that deviations from GR are expected to be more significant when the radius of curvature is smaller, so these black holes should provide more stringent tests of GR than the supermassive black holes in the centers of galaxies.

7. On 17 August 2017 (thus: GW170817), LIGO and Virgo detected a double neutron star merger. This was arguably the most intensely observed astronomical event in human history, and it inaugurated the era of multimessenger astronomy (a term not everyone likes; here it means gravitational waves plus some other messenger such as photons, neutrinos, or cosmic rays). We devote the next section to this event.

GW170817

Shortly after Virgo began taking data, in August 2017, and just three days after a joint LIGO-Virgo detection of a double black hole event proved that Virgo worked as hoped, a very different event was seen. This was an inspiral that lasted more than a minute, culminating in a merger at 12:41:04.4 UTC on 17 August 2017. The duration of the event made it clear that this was a low-mass event; the chirp mass was consistent with the merger between two 1.36 $M_\odot$ objects. 1.74 seconds after the merger, a gamma-ray burst was seen. 11 hours after the merger, an optical transient was observed in the galaxy NGC 4993, which at just 40 Mpc distance means that this was the closest gamma-ray burst ever seen (by a factor of ~ 10 compared with other GRBs with known distances). The flux from the burst was not exceptional, and this was quickly understood to be because (unlike for most gamma-ray bursts) we saw this one at tens of degrees from the axis of the jet, so the flux was lower than might have been expected at that distance if we had observed the burst jet-on.

The later development of the electromagnetic counterpart was consistent with predictions about "kilonovae", which are afterglows from short gamma-ray bursts (now confirmed in at least one case to be due to a neutron star merger), and which are powered by the production of elements heavier than iron (so-called "r-process" nucleosynthesis, where neutrons are captured by nuclei faster than they can decay into equilibrium). Radio observations are still being conducted.

From the gravitational wave standpoint, although tidal deformation was not definitively seen, an upper limit on tidal deformation corresponded roughly to (at the 90% credible level) the radii of these stars being less than about 13.5 km. Also, the close coincidence of time of the merger and the gamma ray signal demonstrated that gravity travels at the same speed as light in a vacuum, to at least a part in 10^{15} accuracy.

So to summarize, GW170817:

1. Confirmed, at least in one case, that short gamma-ray bursts can be produced by the merger of two neutron stars.
2. Confirmed kilonova predictions and supported the idea that neutron star mergers can be a substantial, or even the dominant, source of heavy elements in the universe.
3. Provided a rough upper limit on the radius of a $\sim 1.36\ M_\odot$ neutron star, which therefore constrains the properties of the dense matter in neutron star cores.
4. Ruled out alternate theories of gravity in which gravitational waves travel at a speed more than a part in 10^{15} different from the speed of light in a vacuum.

That's pretty good!

What are our conclusions about the origin of binary black holes?

Disappointingly, we can't really say a lot. This is because, as we've indicated in the previous two lectures, the models have been far more flexible than was really understood prior to detections. Do you want high-mass black holes, on the order of $30-40\ M_\odot$? No problem. Low mass? Sure. Comparable-mass mergers? Sure. Mergers with widely different mass ratios? Yes; that's tougher, but it's a minority of events so we can do a little special pleading. Aligned high-spin systems? They were expected and they haven't been seen, but after the first detections didn't find evidence for much spin, models appeared indicating that, after all, we don't expect much spin (and those models can't apply to many of the electromagnetically observed black holes). A preference for alignment rather than counteralignment? That's not really expected in dynamical scenarios (where you might expect random orientations), but the evidence isn't particularly strong yet. Masses in the pair instability gap? Sure, why not. And so on.

I suspect that almost everyone, prior to the first detection, would have been confident that by the time we got to ~ 100 detections we would know which channel (binary evolution, dynamics, AGN disks, even primordial black hole binaries) dominates. But that hasn't happened.

The next obvious question is: what *would* give us clarity? One hope is that continued accumulation of data will solidify features in the mass spectrum, or the spin distribution, or whatever, perhaps as a function of redshift. Maybe this will be easier to explain for one channel than for others. For example, if ultimately there is a statistically significant excess of AGN in the (huge!) localization volumes of gravitational wave events, that could point to the AGN disk channel. Or if there turns out to be a major excess of prograde over retrograde spin orientations, then the dynamical modelers will have to work extra hard. Another possibility is that a single event could be unambiguously from one channel, or unambiguously *not* from a particular channel. For example, if an $80\ M_\odot$ black hole is involved in a coalescence and the uncertainties on the mass are very small, then it would definitively be in the pair instability mass gap regardless of uncertainties about the $^{12}\mathrm{C}(\alpha,\gamma)^{16}\mathrm{O}$ rate. This *should* rule out an isolated binary origin. Or if we see an event with an eccentricity of $e = 0.5$ at a frequency > 50 Hz (to be well in-band), it will point pretty strongly to a dynamical origin.

But given the adaptability of models demonstrated in response to the first ~ 100 events, I wouldn't be willing to bet that the next 100, or next 1000, will lead to resolution. With that in mind, it seems to me that it will be critical to improve our understanding of the models themselves; what can they do and, importantly what *can't* they do? In that respect I'm more hopeful for improved understanding of the isolated binary channel than of the others. Stellar cluster dynamics and, especially, captures in the accretion disks of AGN have tremendous freedom in their initial conditions. That makes clear predictions difficult. But at least for the isolated binary channel we have plenty of actual systems that we can observe, and at least some parts of stellar physics is known well and is becoming better known. We'll see. Maybe some of you can provide critical advances!