Binary Sources of Gravitational Radiation

We will now explore binary systems in more detail. These obviously have a large and varying quadrupole moment, and they have the additional advantage that we actually know that gravitational radiation is emitted from them in the expected quantities (based on observations of double neutron star binaries). 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 and then discuss individual types of binaries.

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 then 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 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 are

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

and

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

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

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

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 high frequencies, such as will be detected with LIGO-India, can usually be considered to be effectively circular.

Do we have evidence that these formulae actually work? Yes! Nature has been kind enough to provide us with the perfect test sources: binary neutron stars. Several such systems are known, all of which have binary separations orders of magnitude greater than the size of a neutron star, so the lowest order formulae should work. Indeed, the da/dt predictions have been verified to better than 0.1% in a few cases. The de/dt predictions will be much tougher to verify, though. The reason for the difference is that de/dt has to be measured by determining the eccentricity orbit by orbit, whereas da/dt has a manifestation in the total phase of the binary, so it accumulates quadratically with time. These systems provide really spectacular verification of general relativity in weak gravity. In particular, in late 2003 a double pulsar system was detected, that in addition has the shortest expected time to merger of any known system (only about 80 million years). Having two pulsars means that extra quantities can be measured (such as the relative motion, which gives us the mass ratio), and in fact the system is now dramatically overconstrained (more things measured than there are parameters in the theory). The tests of GR by observations of binary neutron star systems deservedly resulted in the 1993 Nobel Prize in physics going to Hulse and Taylor, who discovered the first such binary.

Estimates of compact binary merger rates

Reference webpages:

Abadie et al. 2010, Class. Quant. Grav., 27, 173001

The only guaranteed source for Advanced LIGO and other ground-based detectors, in the sense that we have seen the precursors for those sources, is binary neutron stars. As a result, estimates of the coalescence rates for binary neutron stars have played a critical role in the acceptance and funding of these detectors. Here we will discuss how the known double neutron star systems in our Galaxy have been used to estimate the relevant rates, and the uncertainties on those rates and on the inferred rates for other double compact object mergers. As a reminder, the measured masses of neutron stars in double NS binaries is in a tight range of $1.25 - 1.44\ M_\odot$, although there are neutron stars up to $2\ M_\odot$.

There are currently six known double neutron star systems in our Galaxy (one in a globular cluster) that will merge within a Hubble time, about ten billion years. There are another few that will not merge within a Hubble time. Given that, **Ask class** how might they use the properties of these systems to get a best estimate for the rates? We'll worry about uncertainties later.

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

$$R = \sum_i^d \frac{V_{\rm Gal}}{V_{\rm max}(i)\tau(i)}\ , \tag{8}$$

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."

Let's contemplate what this means. Consider the lifetime first. Here the total lifetime means the time from the formation of the second neutron star in the binary to the time of merger. That seems pretty reasonable. We can figure out the time to merger from the properties of a current system, by using the Peters equations along with the measured masses, period, and eccentricity for each system. We can estimate the time from formation by assuming (as is common) that the stars spin down by magnetic dipole radiation and that their current periods are much longer than their birth periods. With those assumptions,

$$\tau = P/(2\dot{P}) + \tau_{\rm mrg} \tag{9}$$

where P and $\dot{P}$ are the current period and period derivative, and $\tau_{\rm mrg}$ is the time to merger.

But wait a second! Is this right? **Discuss with class** what if we find a system that is a very short time from merger, say a century, but its total lifetime is ten billion years? Should we give this a very large weighting factor because of the time to merger, or is the low weighting factor appropriate? Although it might be counterintuitive, we shouldn't give it any greater weight than if it were in the middle of its lifetime. Any short period within that lifetime has equal probability, whether it is a century after formation, a century before merger, or a century five billion years after formation. Thus this weighting is correct.

Now focus on the other factor: $V_{\rm Gal}/V_{\rm max}$. **Ask class**: what does this mean? It means that if you could only have seen the system in a small fraction of the Galaxy, it deserves greater weight. **Ask class**: what are ways that a pulsar could have been difficult to see? One way is that if the pulsar is dim, we can only see it out so far. Another way is if the pulsar beam is narrow; then only observers in a fraction of 4π steradians from the pulsar could see it at all. The dimness correction is easy in principle for systems we see: we simply measure the flux from the system and apply a correction factor based on the limiting flux of the particular survey. Of course, as instruments and surveys improve their sensitivity, this means that the correction factors have to be adjusted. If your sensitivity has improved by a factor of ten yet you have seen no new systems, clearly your rate estimates must drop. The comment about "weighted by pulsar density" simply means that if there are far more pulsars in one location than another, you have to take that into account. If almost all pulsars are near the Galactic center but you see a double neutron star system nearby, your system gets greater weight. Similarly, when you estimate $V_{\rm Gal}$ you have to estimate the weighted volume that has pulsars; for example, it would not be appropriate to consider a 200 kpc radius sphere, because very little is past the ~ 15 kpc disk of our Galaxy!

The beaming factor estimates are a little trickier. The standard analogy to a pulsar is a lighthouse, whose beam sweeps by periodically. We can certainly measure the duty cycle of the pulsar as we see it (i.e., the "on" time divided by the total time), but this only gives us a one dimensional cut through the beam. Maybe we didn't cut through the diameter of the beam, or maybe the beam is elongated one way or another. Thus models are needed to estimate the solid angle swept out by the beams. These models are calibrated against the thousands of known pulsars, but we should always keep in mind that because double neutron star systems have to go through particular evolutionary paths, maybe their beam angle distribution is different from that of isolated pulsars. As an example of these potential biases, recall the tight mass distribution of binary neutron stars, versus the much broader possible mass distribution.

The above will give us an estimate, with uncertainties, about the merger rate of double neutron star systems similar to those that we have seen. But what about ones *different* from what we have seen? For example, even for pulsars whose beam shines at us, there will be some that are too dim to see. In fact, it could be that there are a large number of pulsars with *luminosities* smaller than those that have been detected. Those pulsars might not be close enough for us to see, so they are a different population. One's optimism about how many dimmer pulsars there are plays a big role in one's merger rate estimates.

Finally, there is the role of luck. Maybe we are in a galaxy with significantly more, or significantly fewer, double neutron star systems than would be average for a galaxy of our type. Here "significantly" can only mean in the statistical sense, so we use Poisson statistics to quantify how great a range that implies for the rates.

Great, so now we have a range of estimates for the rate per Milky Way Equivalent Galaxy (or MWEG to its friends). How do we translate this into an expected rate for the network of ground-based gravitational wave detectors? This is a comparatively straightforward part of the exercise. The waveforms of binary neutron star systems, in the $\sim 100 - 200$ Hz best sensitivity range of LIGO-India, are expected from theory and from our known cases to be circular (recall our previous discussion about the decrease in eccentricity) and to be at most moderately affected by the spins (which, remember, are small in magnitude). Thus you set a threshold of detection (usually S/N=8), and average properly over the sky to get a maximum distance of ≈ 200 Mpc. Then you multiply this volume by the number density of MWEGs (0.01 Mpc^{-3}) and by the rate per MWEG to get your answer. Abadie et al. 2010 get a range between 0.4 per year and 400 per year for Advanced LIGO at design sensitivity, with a best guess of 40 per year. Now, that 0.4 per year might worry you; if it could be that low, doesn't that mean that LIGO has a pretty good chance of not seeing anything for five years? I need to stress that this is a *very* conservative lower limit. It assumes, for example, that our Galaxy has significantly more double neutron star systems than normal *and* that there are few to no pulsars with intrinsic luminosities lower than the current lower limit.

One more point to make is that B2127+11C is actually in the globular cluster M15. This is pretty remarkable, because only about 10^{-4} of the stellar mass in our Galaxy is in globular clusters. That would make it seem as if we should give B2127+11C a pretty high weight. And yet, analyses of this type usually ignore B2127+11C completely. Outrageous!

Or is it? 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 we don't know any examples of such systems and thus possibly formation channels in globulars dominate.

It is thus worth taking a brief diversion to discuss what might be different about globulars. The main difference is stellar number density; in the Solar vicinity there are roughly 0.05 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).

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.

At this point you may have realized that there are a lot of uncertainties involved in these estimates, even for NS-NS systems where you can point to examples in our Galaxy. Is there another way?

Yes. Along with the extrapolations from known sources that we just discussed, the other method that has been used for rate estimates is population synthesis. Speaking candidly, it is my opinion that at the current time, and for compact binary rate calculations, population synthesis does not give us extra information, and in the worst circumstances it could be misleading. However, when detections from ground-based detectors roll in, it will be important to use a population synthesis framework to intrepret the results.

Let me elaborate on those comments by first defining population synthesis, then giving an example of when it is demonstrably useful, then indicating why that is not yet the case for compact binary rate estimates.

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. A good example of where this *is* helpful is in characterizing the stellar population of a galaxy. You look at a galaxy, and you get spectral or photometric information about its total light. Then you simulate a stellar population, where your variables are things such as age and heavy element fraction (called metallicity in astronomy). You compare the expected lines, colors, etc. from your simulated population with what you actually saw, and iterate until you have best values and uncertainties on the age, metallicity, etc. of the population.

The reason this works is that we *understand* how stars look. Sure, there are always uncertainties at the frontier, but we see so many stars that we can calibrate our models extremely well. Thus we can be confident in our inferences.

Contrast this with models of the evolution of massive binary stars into compact object binaries. There are huge uncertainties involved that are not present in the evolution of single stars. One of the nastiest examples is the dreaded common envelope phase. Before settling to their final fate (white dwarf, neutron star, or black hole), stars go through giant phases. Massive stars go through more than one, in fact. Thus if you have a binary that is close enough to ultimately produce merging compact objects, there will be a phase in which the envelope of the first giant overlaps the second star, and a later phase when the envelope of the second overlaps the first. Too little overlap, and the stars stay far apart and never merge. Too much, and one objects spirals fully into the other; no compact object binary. We don't have observations that help, and different theoretical prescriptions yield monumentally different rates. For example, a few years ago, a new prescription produced predicted BH-BH rates 500 times lower than before!

Common envelopes are far from the only problem. Metallicity, the direction and magnitude of supernova kicks, etc., all produce potentially large effects in unknown directions.

Therefore, unlike in the star case, here we rely critically on the observed systems. But given that the total number of parameters in the models (more than 30 overall, although 7-8 likely produce the dominant effects) is greater than the number of systems, we can't actually check that the model framework is correct.

People barrel ahead anyway, though, with the general procedure being that you pick some parameters, generate many synthetic sources, "observe" your sources, and compare the ones you detect with the ones you actually see. This allows you to constrain the input parameters of your models. From your constrained parameters, you then produce ranges of rates. For NS-NS systems you get no better answers than in our previous method, which is why I said it is unhelpful for now. For NS-BH and BH-BH systems, we don't know of any, but the population synthesis method gives you answers anyway. The reason I characterize this as potentially misleading is that because we don't have enough data points to judge whether the model is right at all, confident predictions of rates are unreasonable in my opinion. They are also strongly dependent on proper interpretation of the observations. For example, two systems were discovered in the last few years that had periodically varying lines, which were analyzed and interpreted as having a black hole and a very massive star stripped of its hydrogen envelope (called a Wolf-Rayet star). If this were true, then the argument went that these would become BH-BH binaries, so we'd be able to use those systems to constrain the BH-BH rate. However, recently it was shown that the lines had been misinterpreted: instead of coming from the companion, they come from a stellar wind. Now, within the uncertainties, the compact objects in those systems might even be neutron stars instead of black holes! We are always looking at indirect data, so we have to be careful.

I've been down on the population synthesis method as it is currently applied, but I do think that it will be necessary to interpret binary results from ground-based detectors. The reason is that if we really detect tens of events per year, then in just the first year we will have multiplied by several our known population of NS-NS binaries, and might have detected currently unknown NS-BH and BH-BH binaries. 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 will therefore play an important role.

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. 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 $1/r^6$ energy 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 int he last 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 equivalent 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, hence 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}.

This process is 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 $\sim 5\%$:

$$\begin{aligned} f_{\rm qnr} &\approx [1-0.63(1-j)^{0.3}](2\pi M)^{-1} \\ Q &\approx 2(1-j)^{-0.45}\ . \end{aligned} \tag{10}$$

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 $E_{\mathrm{inspiral}} \sim \mu$. 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, $E_{\mathrm{inspiral}} \sim 0.06M$ and E_{merger} and E_{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.

Table 1: Double neutron star systems

Source	P(ms)	$P_{\rm orb}$(day)	e	$\log_{10}(t_c/{\rm yr})$	$\log_{10}(t_g/{\rm yr}))$
J0737-3039	22.7/2770	0.102	0.088	8.3/7.7	7.9
B2127+11C	30.5	0.3	0.68	8.0	8.3
J1906+0746	144.1	0.17	0.085	5.1	8.5
B1913+16	59.0	0.3	0.62	8.0	8.5
J1756-2251	28.5	0.32	0.18	8.6	9.3
B1534+12	37.9	0.4	0.27	8.4	9.4
J1829+2456	41.0	1.18	0.14	10.1	10.8
J1518+4904(??)	40.9	8.6	0.25	10.3	12.4
J1811-1736	104.2	18.8	0.83	9.0	13.0
B1820-11(??)	279.8	357.8	0.79	6.5	15.8