Numerical Simulations of Superkick Binary Black Holes

Lydia Nevin Mentor: Dr. Mark Scheel

October 31, 2014

Abstract

Gravitational wave detection by LIGO depends on accurate waveform templates for comparison. The development of these templates is accomplished by numerical relativity simulations of events such as colliding black holes. Two black holes orbiting and merging, called a black hole binary, are described by seven independent parameters: the three components of spin for each black hole, and the ratio of their masses; the total mass of the system can be scaled out. Simulations can take weeks or months, and the possible cases are numerous, so a robust interpolation strategy is desired. We use the SpEC numerical relativity code to investigate cases where the masses are equal and the spins are anti-aligned and in the plane of the orbit, a situation called a Superkick, which results in unusually high radiated linear momentum. We aim to discover how significantly waveforms differ when starting conditions are slightly changed, and therefore how easily interpolation can be used between similar instances of this situation. We also describe improvements in algorithms used to specify the orbital eccentricity of binary black hole simulations.

1 Background

Gravitational waves resulting from the acceleration of massive bodies are predicted by Einstein's theory of General Relativity, but have yet to be measured directly. The Laser Interferometer Gravitational Wave Observatory, known as LIGO, aims to detect these fluctuations in the structure of spacetime. Astronomical events with gravitational wave signals large enough to be potentially detected include neutron star pulsars, supernovae, and binaries involving black holes, neutron stars, or both.

LIGO has two detectors; each is a Michelson interferometer with arm lengths of 4km, located in Hanford, Washington and Livingston, Louisiana. Gravitational waves cause minute changes in the length of the arms, resulting in length differences measurable by examining the phase difference of beams of laser light that travel each of the two arms separately and then are recombined. Numerous noise sources such as seismic activity, magnetic fields, photon shot noise , and thermal fluctuations limit the device's sensitivity. One result of this is that accurate templates are needed to detect a signal out of the noise; additionally LIGO is limited to detections between about 10^1 and 10^4 Hz [4]. LIGO's detectors are currently being upgraded and refined from Initial LIGO to Advanced LIGO, expected to be ten times more sensitive than the original.

Numerical Relativity is one way of generating accurate templates to describe signals that could be detected by LIGO. Numerical simulations of compact binaries, such as two black holes orbiting each other, generate gravitational waveforms which can be compared to collected data in order to make a detection and, if accurate enough, estimate the parameters of the detected event. Einstein's field equations for General Relativity are complicated, nonlinear partial differential equations that can only be solved analytically for very simple configurations. Black hole binaries can be approximated using post-Newtonian methods, but reliable descriptions of the entire event, through inspiral, merger, and ringdown, can only be obtained by numerical simulations.

One major difficulty in establishing a comprehensive collection of binary black hole simulations is the combination of large parameter space and heavy computational demands for each run. A single simulation depends on seven independent parameters: the mass ratio of the black holes, and the 3 components of spin for each black hole. The total mass of the two black holes can be scaled out of the problem and reincorporated once the simulation is complete. A catalog of waveforms based on results of binary black hole simulations performed on various subsets of the possible parameter space is described in [3] and is currently being expanded. A single simulation, with one selection of these parameters, can take weeks to complete. Thus it is essential to construct a faster method of producing a gravitational waveform at an arbitrary point in parameter space; one way of doing so is by interpolating from a set of numerical waveforms computed at a small number of points.

The reduced basis method is a new way of doing this interpolation which is just starting to be explored. Preliminary tests show that reduced basis can produce accurate interpolated waveforms using very few numerical simulations [1]. However, the efficiency of this method relies on the assumption that waveforms vary smoothly as a function of parameters. Thus it is prudent to study regions of parameter space in which the waveforms might change rapidly as a function of parameters. One such region occurs when the black holes have large and opposite spins that lie in the plane of the orbit. This causes the resulting black hole after merger to have strong momentum in a certain direction, known as a superkick [2]. Because the magnitude and direction of the kick vary greatly with a small change in initial parameters, it is thought that the radiated gravitational waves may change significantly even with a small change in initial conditions, making it much more difficult to interpolate under these conditions.

2 Simulation Outcomes

This project aimed to discover how variations in initial superkick parameters would affect the gravitational waveform. We compared waveforms between several simulations in the superkick region of parameter space. We varied both the initial separation of the holes, and the angle that the initial spin of one hole (arbitrarily labeled A) makes with the line segment separating the centers of the two holes (see Table 1). The spin of the other hole (labeled B) is opposite the spin of A. Waveforms are compared by multiplying and integrating normalized frequency-domain waveforms in a particular sky direction; the mismatch is the difference of this integral from one. We found that waveforms changed very little for the superkick case with spin magnitude of 0.5. Comparisons between any two of the final waveforms we obtained yielded a match value of over the 97% threshold required by LIGO to detect an event based on a given template (see Figure 1).

Table 1: Initial Separation and Spin parameters for this project's runs. Each case had mass ratio 1, anti-aligned spins with magnitude 0.5 in the plane of the orbit. Each simulation ran for about 16 orbits. An angle of zero means that the initial spin of hole A points directly away from hole B. Distances are in geometric units.

Initial Separation	Angle of Spin A
14.94	0
14.95	0
14.96	0
14.97	0
14.97	$\pi/2$
14.97	π
14.97	$3\pi/2$

We found the of the magnitude of the kick based on the initial angle to be in agreement with the expected sinusoidal dependence (see Figure 2). The magnitude of the kicks (Table 2) also agrees well with the literature [2]. We predict that with a mass ratio of 1 and spin magnitude of 0.5, the maximum magnitude of the kick will be 1717 m/s. Based on the data in Figure 1, we also predict that even with the two most different of such configurations, those with maximal kick in opposite directions, the observed waveforms will be similar enough in all sky directions that a single template of equal mass, spin 0.5 superkicks would be sufficient to detect any binary fitting this description.

Table 2: Initial Separation and Spin parameters for this project's runs. Each case had mass ratio 1, antialgined spins with magnitude 0.5 in the plane of the orbit. Distances are in geometric units.

Init. Angle of Spin A	Kick (v/c)	Kick (km/s)
0	0.04499	1349
$\pi/2$	0.03546	1063
π	-0.04504	-1350
$3\pi/2$	-.003541	-1061

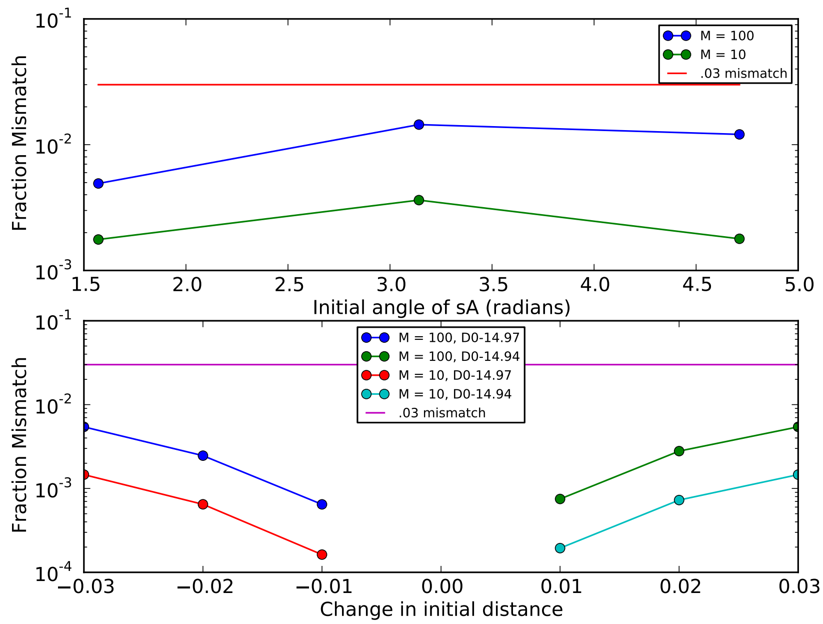

Figure 1: These plots show the degree of mismatch between waveforms among those completed for this project. Notably, all comparisons showed a mismatch of less than 0.03. the maximum allowable mismatch for detection by LIGO. The top graph is a comparison of mismatch between the run with initial separation of 14.97 and the initial spin of hole A pointing directly away from hole B. This shows that the degree of mismatch does correlate with the degree of difference in the kick (see Figure 2). The lower graph shows that the runs with smaller differences in initial distance had a lower degree of mismatch than those with more different initial distances. In all cases, choosing a higher total mass made the mismatch higher, because more of the merger and ringdown phases were visible in the LIGO band. The mismatches shown are the highest found by computing overlaps in many sky directions.

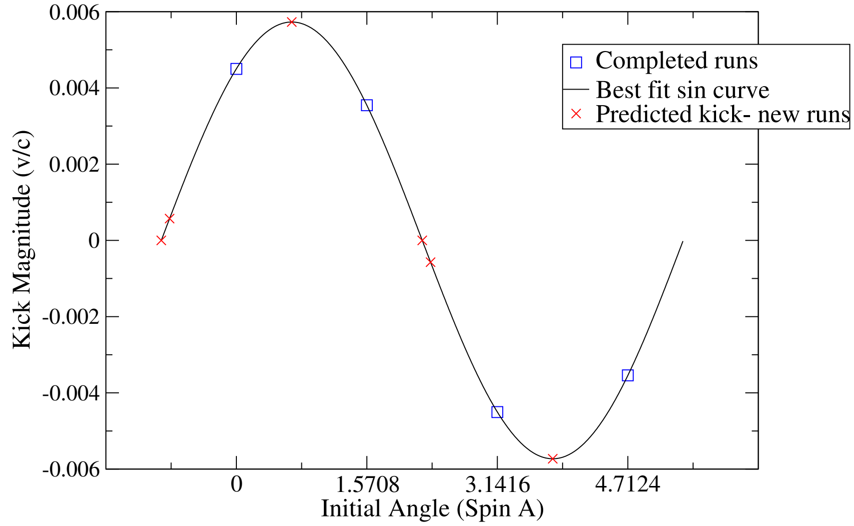

Figure 2: This shows the results of the completed runs with varying initial angle. The magnitude and direction of the kick agrees well with the expected sine wave dependence, and allows more runs to be started at the angles that will cause the maximum kick and maximum derivative.

3 Eccentricity Reduction Modeling

The eccentricity of an orbit quantifies how much it deviates from being perfectly circular. We expect that by the time a black hole binary has progressed enough for its frequency band to be detectable by LIGO, the eccentricity of the orbit will be very low. However, numerical simulations can follow only dozens of binary orbits, which is many orders of magnitude fewer than what would be needed for the eccentricity to sufficiently decrease of its own accord; therefore it is necessary to choose initial data that will yield low eccentricity. There is no way to perfectly predict eccentricity ahead of time, so the simulation starts with a guess for the initial data, runs for about two orbits, examines the trajectories of the black holes. Then adjustments to the initial data are computed that would yield smaller eccentricity, and the simulation starts again with updated initial data. This process is repeated until improvement ceases or the desired level of eccentricity is reached.

The updating scheme for choosing new initial data with the lowest possible eccentricity involves measuring oscillations in the time derivative of the orbital angular velocity of the binary and fitting them to a model. The existing strategy used the same model independent of the eccentricity and independent of how well the model fit the data, unless certain manual options were invoked. A smoothly increasing function accounts for the speeding up of the orbit due to inspiral, and an oscillatory part fits to oscillations in the angular momentum caused by eccentricity. Secondary oscillations due to spin interactions are another component of the model. See Figure 3 for an example showing the inclusion of the spin oscillation terms. The full model is:

$$\begin{aligned} \dot{\Omega} = P_1(t-P_0)^{-11/8} + P_2(t-P_0)^{-13/8} \\ +P_3\cos(P_4t+P_5+P_6t^2) \\ -P_7\sin(\alpha(t)+P_8t) \end{aligned} \quad (1)$$

where $\alpha(t)$ is a function related to the spin vectors of the black holes. We found that using this model, the existing fits for eccentricity reduction were sometimes unreliable, especially when eccentricity was low. Unnecessary inclusion of the P_6 term in the argument of the cosine made estimates of the oscillation frequency unreliable, therefore disrupting eccentricity estimates and parameter updates (Figure 4).

Sometimes the inclusion of P_8, a linear term in the argument of the sine, was needed to achieve a good fit for the spin oscillations (see Figure 5). A rewritten version replicates the logic of the old eccentricity reduction updating scheme, but determines whether to include parameters P_6 and P_8 based on how well the model fits the data. A more complex fit with more parameters is only used if it fits the data significantly better, which we defined as reducing the rms error of the fit by more than 10%.

4 Methods

4.1 Superkick Runs

Because of the limited time of the SURF program, we decided against performing runs with very high spin, or starting from very large separations. We started runs with a range of small changes in initial separation, and with changes of $\pi/2$ in initial angle (see Table 1) . Since the exact correspondence between changes initial angle and initial separation was unknown, as was the required sampling density, we used two different ways of selecting simulations (i.e, varying separation and varying spin angle) to increase the potential to achieve meaningful results.

We were able to reduce the eccentricity of each simulation to about $3*10^{-4}$. Adding terms to Eq. (1) to account for oscillations caused by spin-spin effects was found to improve the fits to $d\Omega/dt$, but did not improve the eccentricity overall. Each simulation was repeated using three choices of numerical resolution, for the purpose of estimating errors. Eccentricity reduction was done at the medium resolution. Waveforms were extracted at finite radii and extrapolated to infinite distance; then overlaps were computed between waveforms with junk radiation removed, and with incorporation of the LIGO noise curve. The minimum match over a selection of sky directions was used to generate the worst-case overlap data shown in Figure 1. Attempts to smoothly hybridize the NR waveforms with post-Newtonian ones for a longer inspiral were unsuccessful; hybridization with precessing systems is an ongoing subject of work.

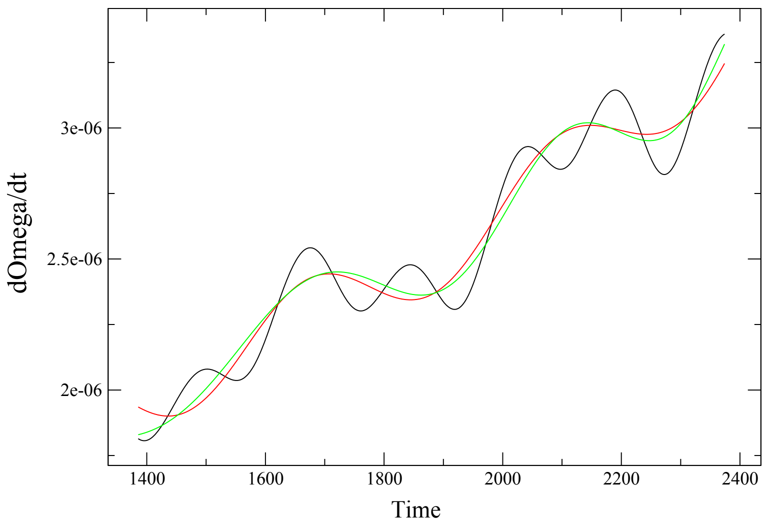

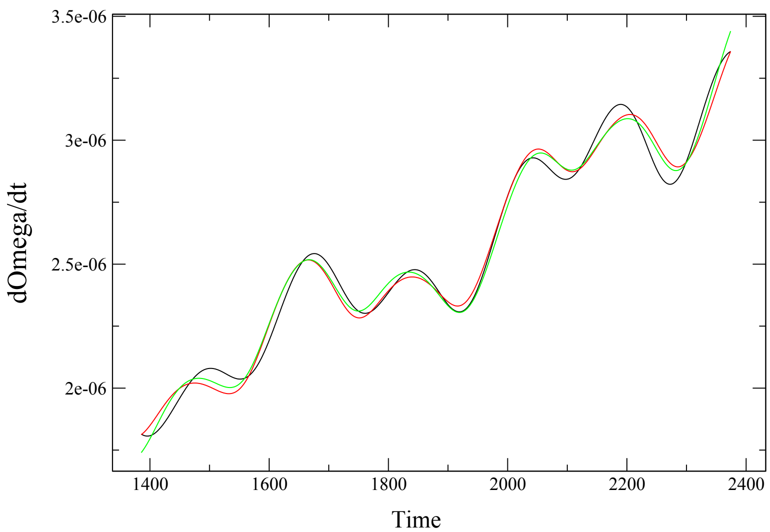

Figure 3: Time derivative of orbital angular velocity versus time. The top panel shows the numerical simulation in black, and the fit from Eq. (1) where P_6 is set to zero (red) and where P_6 is included in the fit (black). The terms accounting for spin-spin oscillations, P_7 and P_8, are not included. The bottom panel shows the numerical data and the same fits, but also including P_7, the basic model for the spin-spin oscillations at twice the frequency of the eccentricity oscillations. It is clear that the inclusion of P_6 makes very little difference in the quality of the fit, but accounting for the spin oscillations made the fits much better in this case.

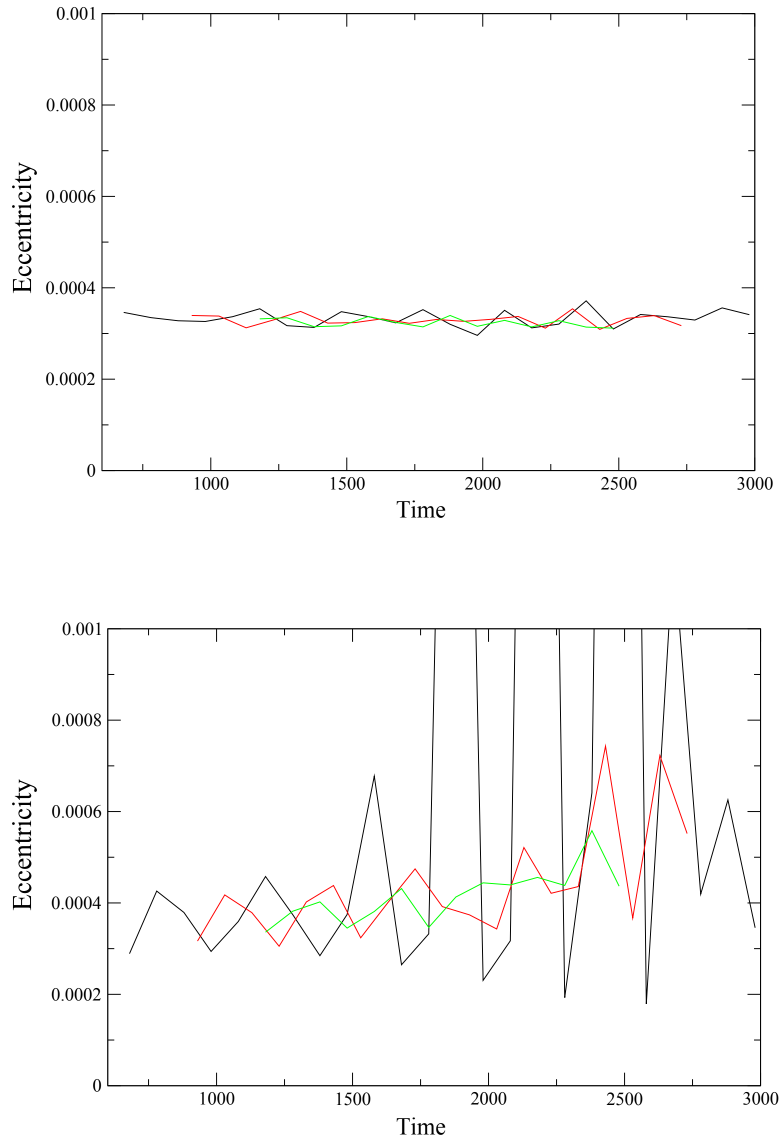

Figure 4: These graphs show the measured eccentricity vs. time for an arbitrarily chosen run. On top is the measured eccentricity when the fit uses a linear argument in the cosine ($P_6 = 0$), and below are the results when the argument in the cosine is quadratic (P_6 determined by fit). The top and bottom panels are plotted on the same scale to emphasize the difference in the variability of results. Each of these fits were performed over a fitting range of Δt = 1000 (black), 1500 (red) and 2000 (green). A wider range of times to fit over improves the data from fits including P_6 dramatically, but the eccentricity estimate is still more consistent when P_6 is excluded.

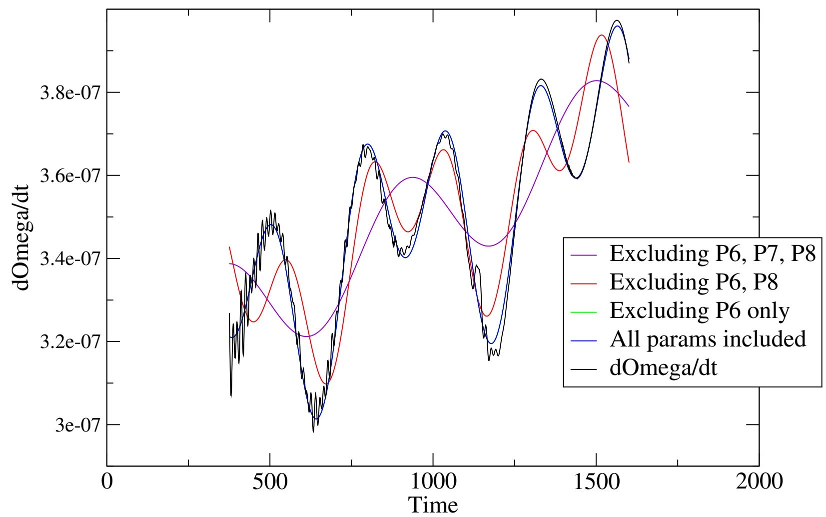

Figure 5: This graph shows the benefit of including P_8 as part of the spin corrections. The green and blue curves fit the data (black) much better than the old model. Very little difference is seen between the models with and without the P_6 term, however. This particular run was chosen as an example because the difference in fits can be clearly seen; however the extra term makes a significant difference on many of the runs examined.

4.2 Programming Eccentricity Fits

Improvement of the eccentricity reduction updating began by using the existing Python implementation to develop a C++ version that would be more compatible with the other code in SpEC. The Gnu Scientific Library fitting was used, to make use of its built-in capabilities to give an error estimate for the fitting parameters. I compared the results of the previous code with my own on runs that were already completed, and the resolution of discrepancies between them led to the discovery of bugs in the existing code. After the new code consistently agreed with the existing code, the significant speed improvement allowed it to be applied over different time windows to obtain trends for eccentricity as a function of time. Using this new ability, we discovered that the eccentricity was very noisy (see the bottom panel of Figure 4). Upon further investigation it was determined that inclusion of the P_6 term in Eq. (1) was making measurement of the frequency of oscillations unreliable, and that excluding it yielded much smoother results (see Figure 4). A bank of 217 runs, part of another project, with their eccentricity removal steps completed was used to compare different fitting models. For some runs, fitting to the spin oscillations was less effective than for other runs, so a model including P_8, a simple first-order term to adjust the frequency, was tested. These terms that improved the fit under some cases only, and have the potential to make the other parameters less reliable, can be included selectively as described above.

5 Results and Future Work

The results of waveform comparisons inspire optimism for the future of the reduced basis method of interpolation. The waveform results made sense and were similar enough between different runs that interpolation over this subset of the parameter space should not be more difficult, or require a much higher density of full simulations, than other cases such as those where the spin is aligned with the orbital angular momentum. Because simulations take so long, it was possible to complete only one round of simulations during the SURF project. Completion of more runs could provide confirmation that the pattern of similarity between waveforms continues for the maximum kick magnitudes. Future research on this subject could investigate whether waveform similarity is applicable with higher spins, and determine the effects of small changes in the mass ratio or spin alignment, moving slightly away from the maximal superkick case.

The updates to eccentricity removal will help this stage of simulations proceed more efficiently and reliably. Further work could allow measurement of eccentricity during a running simulation, rather than by examining a window of data after the simulation has halted. This would allow the simulation to recognize when enough trajectory data has been collected to compute a satisfactory update to initial parameters, and also serve as a sanity check in later stages of the run to ensure that the orbit stays smooth and circular.

Making a dynamic choice between eccentricity fitting models is a strategy that would benefit from further investigation. An understanding of which situations make certain fitting terms useful, and which make them redundant, would allow future versions of the code to be more specific. A more advanced, general model could be developed that uses a greater understanding of the applicable physics; knowing when, for example, the post-newtonian spin corrections benefit significantly from a fitted adjustment in frequency could provide insight into the physics of these situations.

References

[1] J. Blackman, B. Szilagyi, C. R. Galley, and M. Tiglio. A sparse representation of gravitational waves from precessing compact binaries. 2014.

[2] Manuela Campanelli, Carlos O. Lousto, Yosef Zlochower, and David Merritt. Maximum gravitational recoil. *Phys. Rev. Lett.*, 98:231102, Jun 2007.

[3] Abdul H. Mroué, Mark A. Scheel, Béla Szilágyi, Harald P. Pfeiffer, Michael Boyle, Daniel A. Hemberger, Lawrence E. Kidder, Geoffrey Lovelace, Serguei Ossokine, Nicholas W. Taylor, Anı l Zenginoğlu, Luisa T. Buchman, Tony Chu, Evan Foley, Matthew Giesler, Robert Owen, and Saul A. Teukolsky. Catalog of 174 binary black hole simulations for gravitational wave astronomy. *Phys. Rev. Lett.*, 111:241104, Dec 2013.

[4] S. J. Waldman. The advanced ligo gravitational wave detector. 2011.

Acknowledgements

Thanks to:
Dr. Mark Scheel
Dr. Dan Hemberger
LIGO SURF
The NSF
Caltech