

Final spin and radiated energy in numerical simulations of binary black holes with equal masses and equal, aligned or anti-aligned spins

Daniel A. Hemberger,¹ Geoffrey Lovelace,^{2,3} Thomas J. Loredo,¹ Lawrence E. Kidder,¹
Mark A. Scheel,³ Béla Szilágyi,³ Nicholas W. Taylor,³ and Saul A. Teukolsky¹

¹*Center for Radiophysics and Space Research, Cornell University, Ithaca, NY, 14853, USA*

²*Gravitational Wave Physics and Astronomy Center, California State University Fullerton, Fullerton, CA 92831, USA*

³*Theoretical Astrophysics 350-17, California Institute of Technology, Pasadena, CA 91125, USA*

(Dated: September 24, 2013)

The behavior of merging black holes (including the emitted gravitational waves and the properties of the remnant) can currently be computed only by numerical simulations. This paper introduces ten numerical relativity simulations of binary black holes with equal masses and equal spins aligned or anti-aligned with the orbital angular momentum. The initial spin magnitudes have $|\chi_i| \lesssim 0.95$ and are more concentrated in the aligned direction because of the greater astrophysical interest of this case. We combine these data with five previously reported simulations of the same configuration, but with different spin magnitudes, including the highest spin simulated to date, $\chi_i \approx 0.97$. This data set is sufficiently accurate to enable us to offer improved analytic fitting formulae for the final spin and for the energy radiated by gravitational waves as a function of initial spin. The improved fitting formulae can help to improve our understanding of the properties of binary black hole merger remnants and can be used to enhance future approximate waveforms for gravitational wave searches, such as Effective-One-Body waveforms.

PACS numbers: 04.25.dg, 04.25.Nx, 04.30.Tv

I. INTRODUCTION

Binary black holes are an important source for gravitational-wave detectors such as the Laser Interferometer Gravitational-Wave Observatory (LIGO), GEO, and Virgo [1–3]. Searches for gravitational-wave signals have been able to constrain the event rate for binary black hole mergers, but a direct detection of gravitational waves has not yet been made [4, 5]. These searches require predictions (“templates”) of the expected gravitational waves; so far, only non-spinning templates have been included [5]. However, there is evidence that spin is relevant in astrophysical black holes, from both theoretical predictions [6–8] and observational data [9–11].

Therefore, LIGO and other gravitational-wave detectors need to include spin as a parameter in their template waveforms; otherwise, the search space (and thus the detection rate) is reduced because of an insensitivity to spinning sources [12, 13]. Accurate simulations of spinning binary black hole mergers are also needed to infer the properties (e.g. masses and spins) of binaries from the detected waveforms (“parameter estimation”) [14].

For both detection and parameter estimation, numerical simulations are too computationally expensive to generate waveforms for the entire parameter space of binary black hole mergers. Instead, numerical simulations are used to calibrate and validate the approximate, analytic models that are actually used to generate template waveforms. For instance, the Effective-One-Body (EOB) model, calibrated using numerical simulations that include merging black holes with spins aligned or anti-aligned with the orbital angular momentum [15], is used by the LIGO Collaboration to estimate how sensitive their search is to waveforms from spinning systems [5].

However, Ref. [15] has shown that the EOB model poorly predicts configurations with large aligned spins, and that more numerical relativity simulations are needed in this region of spin parameter space to improve the calibration of the model.

Binary black holes whose spins are aligned (or anti-aligned) with the orbital angular momentum involve far fewer parameters than generic binaries with arbitrary spin directions, but nevertheless they can be used to construct templates capable of detecting a sizeable fraction of precessing binaries [12]. Furthermore, aligned-spin systems are astrophysically motivated by studies including observations of the micro-quasar XTE J1550-564 [16], models of gas-rich galaxy mergers [17], and population synthesis models [18].

In this paper, we introduce ten new simulations of binary black holes with equal masses and equal spins aligned or anti-aligned with the orbital angular momentum. We use the notation $S_{|\chi|}^{\pm\pm}$ to refer to specific cases, where the subscript is approximately the dimensionless spin magnitude at $t = 0$, and the superscripts indicate whether each black hole has the aligned (+) or anti-aligned (−) spin orientation. The new simulations are $S_{0.95}^{++}$, $S_{0.9}^{++}$, $S_{0.85}^{++}$, $S_{0.8}^{++}$, $S_{0.6}^{++}$, $S_{0.2}^{++}$, $S_{0.2}^{--}$, $S_{0.6}^{--}$, $S_{0.8}^{--}$, and $S_{0.9}^{--}$.

To more fully cover the aligned-spin space, we include data previously reported for $S_{0.97}^{++}$ [19], $S_{0.95}^{--}$ [20], $S_{0.0}^{--}$ [21], $S_{0.44}^{--}$ [22], and $S_{0.44}^{++}$ [23] in our analysis. The $S_{0.95}^{++}$ case joins the two simulations from Refs. [19, 20] as the only simulations to date of merging black holes with spin magnitudes above $\chi \approx 0.93$ (the “Bowen-York limit”) [24–26]. We use this combined dataset to improve on prior phenomenological fitting formulae for the final spin of the merger remnant [27–29] and the radiated en-

ergy from inspiral through ringdown [29–31]. These improved formulae can be used to reduce EOB waveform phase errors in the ringdown (see Eq. (19) and surrounding text in Ref. [15]) and therefore provide more accurate templates for gravitational-wave searches [31].

The remainder of this paper is organized as follows. In Sec. II, we discuss the numerical methods that we employ in our simulations. In Sec. III, we report on the values and convergence of the constraint violations, masses, and spins. In Sec. IV, we use the horizon data to improve the phenomenological fitting formulae for final spin and radiated energy as a function of initial spin. Section V contains our conclusions, and Appendix A details our method for constructing the fitting formulae.

II. SIMULATION METHODS

All simulations used in this paper were generated with the Spectral Einstein Code (SpEC) [32]. In this section, we describe the methods for the ten new simulations. For detailed methods of the previously reported SpEC simulations, see Refs. [19–23] and references therein. Throughout this paper, we use units where $G = c = 1$, and we report lengths and times in units of M , the total Christodoulou mass in the initial data.

To produce initial data, we solve the extended conformal thin-sandwich equations with quasi-equilibrium boundary conditions [33–38]. We adopt free data based on a weighted superposition of two Kerr-Schild black holes, which enables us to construct initial data containing black holes with nearly extremal spins [39, 40]. The constraint equations are solved using a spectral elliptic solver [41], and the free parameters are iterated until the target masses and spins are achieved to within some tolerance.

We evolve the initial data on a “cut-spheres” domain [42] using spectral adaptive mesh refinement, which will be detailed in a forthcoming paper. On a timescale of $50M$, we change smoothly from the initial data gauge to damped harmonic gauge [43–45], which helps prevent coordinate singularities. We use a fifth-order Dormand-Prince dense adaptive time stepper.

To reduce eccentricity, we first evolve each system for 2.5 orbits beyond the time when the spurious “junk” radiation is sufficiently far from the black holes so as to have a negligible effect on the black hole trajectories. Then we fit the time derivative of the orbital frequency $\dot{\Omega}$ to find improved initial angular and radial velocities (Ω_0 and \dot{d}_0/d_0) [46]. We iterate this procedure until an eccentricity below 10^{-3} is achieved.

We use the dual-frames technique to do spectral excision of the singularities [47]. As described in other papers reporting high spin simulations using SpEC [19, 20], the most important aspect of this excision technique is careful control of the excision boundary. This must accomplish three tasks. First, it must distort the shape of the boundary so that it matches the shape of the ap-

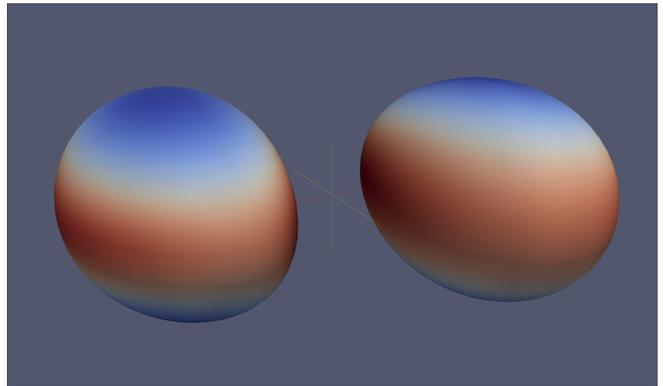


FIG. 1: Effect of spin on horizon geometry. This image shows the intrinsic Ricci scalar on the apparent horizons in simulation $S_{0.85}^{++}$. The proper separation of the horizons along the line connecting their centers is about $1.7M$. Both spin effects (gradients as a function of polar angle) and tidal bulges (dark red regions near the intersection with the line connecting the horizon centers) can be seen.

parent horizon. Second, it must regulate the fractional separation between the excision surface and the apparent horizon — if the separation is too small, then the horizon falls out of the computational domain, but if the separation is too large, then the excision surface falls far inside the horizon, where large gradients are computationally expensive to resolve. Third, it must keep all characteristic speeds on the excision surface positive; i.e., the excision surface must be a pure outflow boundary. This is because we do not impose boundary conditions on the excision surface. Instead, we monitor the characteristic speeds; if they ever become negative, then our evolution system becomes ill-posed, and we terminate the simulation. These three tasks are challenging for high spin systems in part because of the additional distortion of the horizons (see Fig. 1), and they are especially challenging for large aligned spins because such systems spend more time in the dynamic regime before merger.

Using the fast-flow method described in Ref. [48], we find the apparent horizons as an expansion in spherical harmonics, truncated at a given maximum order ℓ . As the system evolves, we adaptively change ℓ to satisfy accuracy criteria for the resolution of the horizon. After a common horizon is found during merger, the evolution continues on a new domain with a single excision surface that subsumes the two individual excision regions [45].

We measure the quasi-local spin, S , on each horizon using the approximate Killing vector method described in Ref. [39]. The dimensionless spin is then

$$\chi = \frac{S}{M_{\text{Ch}}^2}, \quad (1)$$

where M_{Ch} is the Christodoulou mass,

$$M_{\text{Ch}}^2 = M_{\text{irr}}^2 + \frac{S^2}{4M_{\text{irr}}^2}, \quad (2)$$

χ_0	$M\dot{d}_0/d_0 \times 10^4$	$M\Omega_0$	N_{orbits}
0.95	7.26420673	0.01395360	25.4
0.9	5.48222492	0.01419573	24.9
0.85	4.33347923	0.01437107	24.7
0.8	3.54332917	0.01450430	24.2
0.6	1.65215665	0.01487274	22.8
0.2	0.09507527	0.01525060	19.9
-0.2	-0.69081937	0.01538827	17.2
-0.6	-1.95883097	0.01527384	14.6
-0.8	-3.60252091	0.01501397	13.3
-0.9	-5.42657163	0.01474328	12.8

TABLE I: Initial data parameters (radial velocity \dot{d}_0 and angular velocity Ω_0) at separation $d_0 = 15.366M$ for the ten new simulations with target spin, χ_0 . Also included is the approximate number of orbits until merger. Here M is the sum of the Christodoulou masses at time $t = 0$.

and M_{irr} is the irreducible mass, which is a function of the horizon area,

$$M_{\text{irr}} = \sqrt{\frac{A_{\text{AH}}}{16\pi}}. \quad (3)$$

With these definitions, $\chi = 1$ represents an extremal black hole [49].

We choose an integer k to characterize the resolution of each simulation. We call k the resolution level (or ‘‘Lev’’). It sets the resolution by defining the target maximum truncation error for the adaptive mesh refinement and adaptive horizon finding as

$$\epsilon_{\text{max}} = 10^{-4} e^{-k}. \quad (4)$$

Around the excision boundary, where the most resolution is required, we reduce the target maximum truncation error for adaptive mesh refinement by a factor of 10^2 .

III. SIMULATIONS

There are ten new simulations presented in this paper: $S_{0.95}^{++}$, $S_{0.9}^{++}$, $S_{0.85}^{++}$, $S_{0.8}^{++}$, $S_{0.6}^{++}$, $S_{0.2}^{++}$, $S_{0.2}^{--}$, $S_{0.6}^{--}$, $S_{0.8}^{--}$, and $S_{0.9}^{--}$. Initial data were generated with a target Christodoulou mass for each hole $M_0 = 0.5$, target spin for each hole χ_0 , and target ADM linear momentum $p_0^i = 0$. We fix the initial separation at $d_0 = 15.366M$, and then we iterate as summarized in Sec. II to obtain the initial radial velocity \dot{d}_0/d_0 and angular velocity Ω_0 . The targets are met to within an absolute error of $\mathcal{O}(10^{-8})$, and the resulting initial data parameters are reported in Table I. We construct our initial data with a target total Christodoulou mass of $M = 1$ so that our evolution code units are essentially interchangeable with units of M .

At least three different resolutions were evolved for each case to check convergence. Figure 2 shows convergence of the (normalized) volume-averaged L^2 -norm of the generalized harmonic constraint energy [50] for a

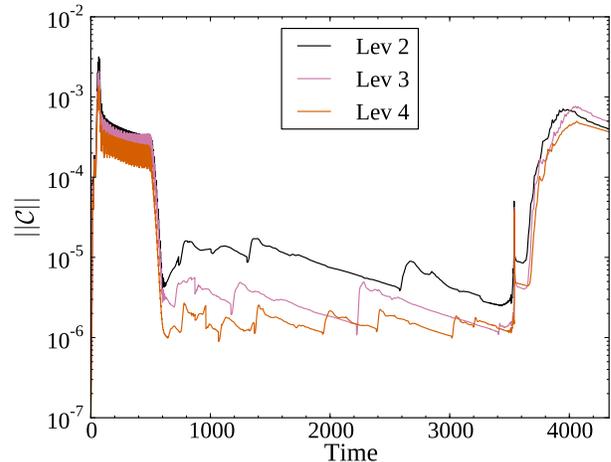


FIG. 2: Normalized constraint violations for $S_{0.9}^{--}$. For each resolution level, k , we plot $\|C\|$, the volume-averaged L^2 -norm of the generalized harmonic constraint energy divided by the volume-averaged L^2 -norm of the dynamical field gradients. This measure is defined in Eq. (71) of Ref. [50]. As the resolution level increases, the constraints decrease. Jumps in the constraints are attributed to changes in the domain structure, and the spike around $t \sim 3500M$ corresponds to the merger.

representative case¹. Each time the domain structure is changed to alleviate grid compression, the constraints jump because of interpolation errors, but then slowly decay back to their baseline levels.

Additional simulations are used in our analysis of masses and spins in Sec. IV: $S_{0.97}^{++}$ [19], $S_{0.44}^{++}$ [23], $S_{0.0}^{--}$ [21], $S_{0.44}^{--}$ [22], and $S_{0.95}^{--}$ [20]. Although these have been previously reported, we include them below for completeness. It should be noted that $S_{0.44}^{++}$ and $S_{0.44}^{--}$ are older simulations and therefore used different initial data and evolution machinery than described in Sec. II. Simulations $S_{0.95}^{--}$, $S_{0.0}^{--}$, and $S_{0.97}^{++}$ used the initial data methods of Sec. II, but earlier implementations of the evolution methods.

A. Mass and spin

We define the initial spin, χ_i , to be the spin after the system has relaxed from the initial data and the junk radiation at the apparent horizon has become negligible. The spin before this time is not physically relevant

¹ Note that we observe poor convergence of the constraints in the late ringdown, which indicates that the simulations may need to be re-run to provide accurate waveforms. However, this issue occurs in the wave zone well after the black hole has settled down to its final mass and spin, and so it does not affect the results of this paper.

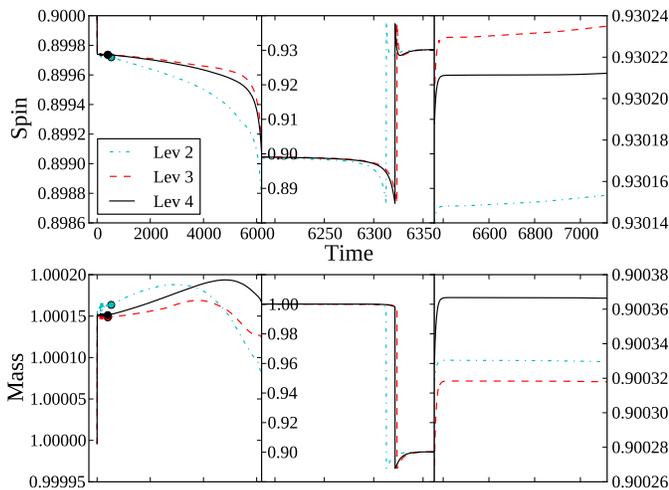


FIG. 3: Plots of the apparent horizon quantities as a function of time for a representative case, $S_{0.9}^{++}$. The top panels display the dimensionless spin and the bottom panels display the Christodoulou mass. From left to right, the panels display the inspiral, merger, and ringdown. We normalize the y -scales separately so that the differences between each resolution can be clearly seen. The discontinuity in the middle panel indicates where we begin to measure the mass and spin on the common horizon. The dots in the early inspiral identify our choice of t_i for each resolution level.

to the rest of the evolution. There are subtle issues to consider when choosing the time to measure χ_i . If we choose too early a time, then junk radiation effects will still be present. If we are overly cautious and choose too late a time, then the system will have emitted enough gravitational radiation to significantly change the spin.

We use a histogram method to determine χ_i . Let $\{\chi_I\}$ be the set of spin measurements during the inspiral. The range of $\{\chi_I\}$ is split uniformly into N bins, where N is the size of $\{\chi_I\}$, and then each element of $\{\chi_I\}$ is put into the appropriate bin.² We choose $\chi_i = \chi(t_i)$, where t_i is the latest time when the spin is in the bin containing the most measurements. In the initial relaxation, the spin is oscillating, and during the inspiral, the spin changes more rapidly as the holes approach each other. Under these conditions, this method selects the spin just after the junk radiation, when the spin is nearly constant for a long interval. Figure 3 shows the evolution of the mass and spin as a function of time for a representative case, and identifies our choice of t_i by the dots in the early inspiral.

We compute t_i from χ because the behavior of the mass is not as simple during the inspiral. The histogram

² If the time interval between spin measurements is not equally spaced, we weight each measurement by the average of the two adjacent time intervals.

Case	$ \chi_i $	$ \chi_f $
$S_{0.95}^-$	0.949053(-30)	0.37567(-18)
$S_{0.9}^-$	0.899569(-11)	0.392748(-12)
$S_{0.8}^-$	0.7997602(59)	0.4268932(30)
$S_{0.6}^-$	0.59993163(71)	0.4942327(-31)
$S_{0.44}^-$	0.437568970(-10)	0.547851(20)
$S_{0.2}^-$	0.1999802(-40)	0.6242202(-61)
$S_{0.0}^-$	$64(-29) \times 10^{-8}$	0.686445(-52)
$S_{0.2}^{++}$	0.200035(-19)	0.7464314(-96)
$S_{0.44}^{++}$	0.4365505(95)	0.8140(10)
$S_{0.6}^{++}$	0.5999635(14)	0.857808(15)
$S_{0.8}^{++}$	0.7998737(-44)	0.907526(14)
$S_{0.85}^{++}$	0.849826(15)	0.919088(30)
$S_{0.9}^{++}$	0.8997371(-15)	0.930212(23)
$S_{0.95}^{++}$	0.9495863(-25)	0.940852(29)
$S_{0.97}^{++}$	0.969504(13)	0.944964(11)

TABLE II: Dimensionless spin measurements. For each case, we provide the initial spin magnitude of each hole and the final spin magnitude of the remnant at the highest resolution. Note that the number in parentheses is not an error estimate, but provides the value at the next highest resolution when added to the last two significant digits.

method applied to the mass will pick out the local maximum late in the inspiral that is present in most of our cases (see Fig. 3). We define the initial mass to be $M_i = M(t_i)$, the sum of the Christodoulou masses at time t_i .

The final spin and Christodoulou mass, χ_f and M_f , are measured at the last observation time, when the merger remnant is in quasi-equilibrium and approximates a Kerr black hole. We report the initial and final spins in Table II, and the initial and final Christodoulou masses in Table III.

From the initial and final Christodoulou masses, we can infer the fraction of the black hole energy that is radiated in gravitational waves during the evolution:

$$E_{\text{rad}} = 1 - \frac{M_f}{M_i}. \quad (5)$$

We expect mass and spin measurements at higher resolutions to be more accurate. However, as illustrated by the comparisons in Fig. 4, these quantities are not strictly convergent in a number of cases. For this reason, we include measurements from all resolutions in our analysis in Sec. IV and weight the uncertainty in the error assigned to a particular measurement by a function of resolution level k .

As described in Sec. II, the most stringent resolution requirements occur in the vicinity of the apparent horizons, but the accuracy may be dominated by short, under-resolved segments of the evolution. The initial masses and spins appear to be randomly perturbed by the junk radiation as the initial data relax. The final masses and spins, which are already limited by the accuracy of the initial masses and spins, appear to also be affected by the details of the coalescence, where we see a

Case	M_i	M_f	$E_{\text{rad}}(\%)$
$S_{0.95}^{--}$	0.999856(68)	0.968134(33)	3.1727(33)
$S_{0.9}^{--}$	1.00016197(73)	0.967909(-27)	3.2248(28)
$S_{0.8}^{--}$	1.0000859(-11)	0.9665941(-16)	3.348894(50)
$S_{0.6}^{--}$	1.00002292(-78)	0.963769(-14)	3.6253(13)
$S_{0.44}^{--}$	2.2470608(-22)	2.159561(-49)	3.8940(21)
$S_{0.2}^{--}$	0.999956(26)	0.9564388(84)	4.3519(17)
$S_{0.0}^{--}$	0.999971(-43)	0.9516182(-74)	4.83791(33)
$S_{0.2}^{++}$	0.999961(22)	0.945471(16)	5.44923(46)
$S_{0.44}^{++}$	2.2451548(28)	2.10099(-44)	6.421(20)
$S_{0.6}^{++}$	1.00001907(-96)	0.926868(-19)	7.3149(18)
$S_{0.8}^{++}$	1.0000765(-22)	0.911275(-28)	8.8794(26)
$S_{0.85}^{++}$	1.000108(-12)	0.906168(-73)	9.3931(62)
$S_{0.9}^{++}$	1.0001513(-29)	0.900366(-48)	9.9770(46)
$S_{0.95}^{++}$	1.00021743(77)	0.893703(-65)	10.6492(66)
$S_{0.97}^{++}$	1.0002384(-94)	0.890691(-22)	10.9521(14)

TABLE III: Christodoulou mass measurements. For each case, we provide the total initial mass of the black holes, the final mass of the remnant, and the radiated energy computed from Eq. (5) at the highest resolution. Note that the number in parentheses is not an error estimate, but provides the value at the next highest resolution when added to the last two significant digits.

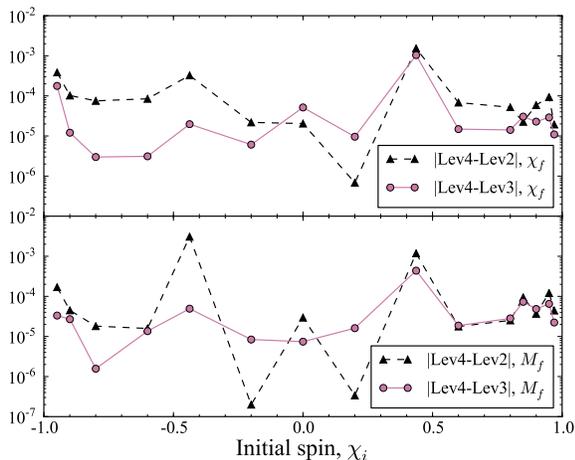


FIG. 4: Differences in the final masses and spins between resolution levels. For each case, we compare χ_f and M_f of the highest resolution to the two lower resolutions. Note that, except for the older $S_{0.44}^{\pm\pm}$ simulations, all differences are $\lesssim 10^{-4}$. Differences in the initial masses and spins behave similarly.

spike in constraint violations (Fig. 2). Apart from these under-resolved segments, we do see convergence in the time derivatives of the masses and spins, but the absolute values remain offset from one another.

We have investigated other potential sources of uncertainty, but found them to lie below the resolution level uncertainty. For example, one source of uncertainty in the masses and spins is the resolution of the surface of the horizon. In Fig. 5, we show a representative plot

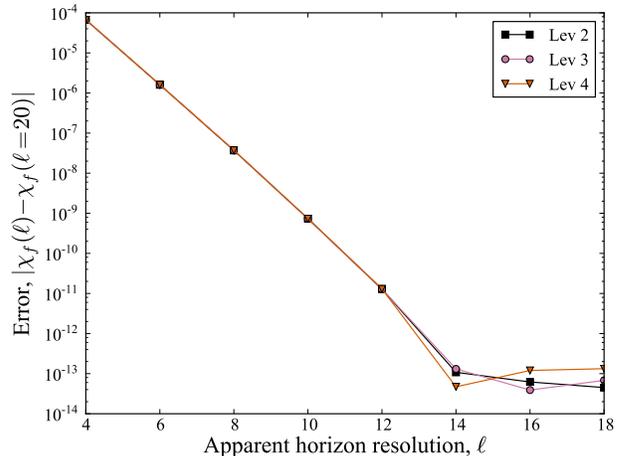


FIG. 5: Convergence of the final spin for $S_{0.9}^{--}$ as a function of the ℓ of the horizon finder. We plot the difference between $\chi_f(\ell)$ and $\chi_f(\ell=20)$ for each resolution. The adaptive horizon finder for this case chose $\ell=8$ at the final time of the simulation.

of error in final spin as a function of ℓ of the horizon finder. Let $\Delta\chi_k$ be the resolution level error between the two highest resolutions, and let $\Delta\chi_\ell$ be the resolution error between the ℓ chosen by the adaptive horizon finder and the ℓ for which the horizon is fully resolved. At the final time of the simulation, we find that, in all cases, $\Delta\chi_k > \Delta\chi_\ell$ by several orders of magnitude.

A source of uncertainty in the radiated energy is the energy that would have been radiated by the binary as it proceeds from infinite separation to the separation d_0 at which we start the simulation. As discussed in Ref. [19], Alvi’s formula [51] estimates that the energy radiated from $d = \infty$ to $d = d_0$ is one part in 10^6 . Since this is smaller than our resolution level uncertainty, it is safe to ignore this difference and we can think of E_{rad} as the total radiated energy from infinite separation through ring-down.

IV. RESULTS

Much effort has been put into constructing phenomenological formulae for the final spin [27–29, 52] and radiated energy [29–31, 52] as a function of initial spin. Because the SpEC code has the capability to generate and evolve initial data of black holes with spins above the “Bowen-York limit” of $\chi \approx 0.93$ [26], we are able to provide new data points to test and improve these formulae.

We use a Bayesian nonlinear measurement error model (described in Appendix A) to fit and compare new parametric formulae. This approach (1) accounts for uncertainties in both the initial spin data and the output data

(i.e., final spin or radiated energy); (2) accounts for the expected improvement in accuracy of results as the resolution level increases; and (3) includes a simple systematic error component quantifying misfit between a chosen formula and the curve the data are converging toward.³ The framework lets us predict an output as a function of initial spin, with prediction uncertainties that account for the uncertainties in the parameters of the chosen formula (including correlated uncertainties) and the typical scale of the systematic error.

The new fitting formulae that we provide here are only applicable to equal mass binary black hole configurations with equal spins aligned or anti-aligned with the orbital angular momentum. More general formulae exist (see e.g. Refs. [28, 29, 53, 54]), but they are less accurate at high spins because of the scarcity of simulations with both unequal masses and high spins in random orientations.

A. Final spin

Using the data from Table II, we construct a new fitting formula for the final spin as a function of initial spin. We fit to a fourth-order polynomial,

$$\tilde{\chi}_f = a_0 + a_1\chi_i + a_2\chi_i^2 + a_3\chi_i^3 + a_4\chi_i^4. \quad (6)$$

The best fit to our data has the parameters a_n and associated covariance Σ_a :

$$a = \begin{pmatrix} a_0 \\ a_1 \\ a_2 \\ a_3 \\ a_4 \end{pmatrix} = \begin{pmatrix} 0.686402(60) \\ 0.30660(14) \\ -0.02684(33) \\ -0.00980(19) \\ -0.00499(35) \end{pmatrix} \quad (7)$$

$$\Sigma_a = \begin{pmatrix} 3.6 & 0.31 & -14 & -0.45 & 11 \\ 0.31 & 21 & -4.8 & -26 & 6.0 \\ -14 & -4.8 & 110 & 7.1 & -110 \\ -0.45 & -26 & 7.1 & 36 & -9.5 \\ 11 & 6.0 & -110 & -9.5 & 120 \end{pmatrix} \times 10^{-9} \quad (8)$$

The uncertainty in a_n , given in parentheses in Eq. (7), is estimated by $\sqrt{\Sigma_a^{nn}}$. However, the parameter estimates are highly correlated; therefore, the full covariance matrix is used in the computation of the fit uncertainty σ_f in Eq. (A12). In Fig. 6, we show the fit and residuals using Eq. (6) with the parameters from Eq. (7).

We fit to a fourth-order polynomial because the high accuracy of our dataset enables us to identify significant third- and fourth-order trends in the residuals of a fit to a second-order polynomial, which is the fitting function used in Refs. [27–29]. The difference between the

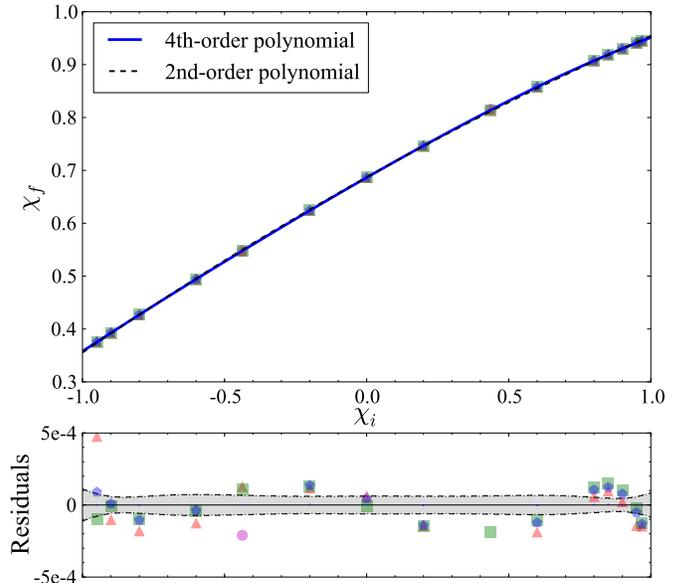


FIG. 6: In the top panel, we plot our preferred fitting formula (solid line), the fourth-order polynomial in Eq. (6), and a comparison with a second-order polynomial (dashed line) for χ_f as a function of χ_i . Our data points are plotted as polygons, where more sides indicates higher resolution level. In the bottom panel, we plot the residuals of the fourth-order polynomial. We indicate our fit parameter (dotted line) and total prediction (dashed line) uncertainties (defined in Appendix A), which in this case are nearly identical. Note that the residuals for the two lower resolution runs for $S_{0.44}^{++}$ are too large to fit in this panel.

logarithm of the marginalized likelihood function (LML) for the best-fit fourth-order and second-order polynomials is ~ 40 , indicating that the fourth-order polynomial provides a significantly better fit. If the two additional degrees of freedom were fitting noise, rather than some underlying structure in the data, then we would only expect a change in maximum LML of $\mathcal{O}(1)$.⁴

The estimated systematic error magnitude $\hat{\sigma}_\Delta$ (defined in Appendix A) for the fourth-order polynomial formula is negligibly small, suggesting that the significant behavior is captured as well as could be expected. However, the residuals, especially at large aligned spins, display trends suggesting that there is additional structure not captured by the fourth-order polynomial (such trends are ignored by our simple systematic error model). This encouraged

³ To keep the calculations analytically tractable, the systematic error component accounts only for the typical magnitude of misfit (essentially, the root-mean-square of the residuals), and does not account for correlations or patterns in the residuals.

⁴ The leading-order term in the maximum LML is proportional to a chi-squared-like quantity, so for nested models, such as the second- and fourth-order polynomials, the change in the maximum LML should roughly mimic the asymptotic statistics of likelihood ratio tests, as given by Wilks's theorem [55]. Two models are said to be nested if the simpler one is a special case of the more complicated one.

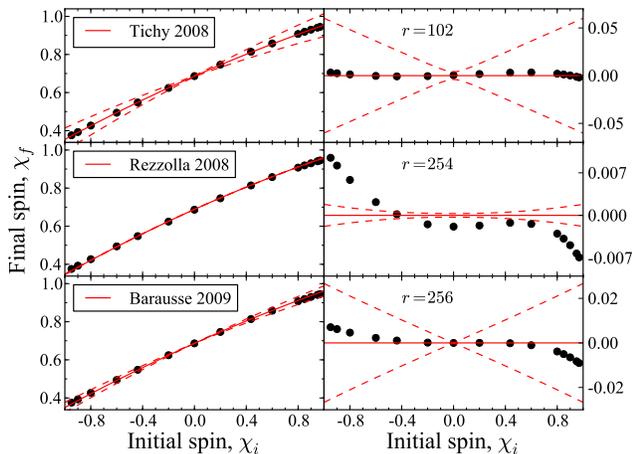


FIG. 7: Final spin as a function of initial spin. In the left panels we plot our data (circles) along with the fitting formula (red line) with error estimates (dashed) from several other studies. The top panel is from Ref. [29], the middle is from Ref. [27], and the bottom is from Ref. [28]. In the right panels we plot the difference between our data and the corresponding fitting formula on the left. The value r quantifies the size of the systematic error compared to the fourth-order polynomial.

us to explore a fifth-order polynomial formula, but it did not reduce the residuals enough to justify the additional degree of freedom. This does not rule out the possibility that a different formula could capture the behavior even more accurately.

We compare our data to existing fitting formulae for the final spin in Fig. 7. The χ_f data corroborate the existing fitting formulae, but indicate deviations at large spins (especially in the aligned direction). This is an expected consequence of the scarcity of high spin numerical relativity data heretofore. In the figure, we provide a quantity for each fit, r , that measures how much larger its systematic error is than that of our fourth-order polynomial. This is essentially a ratio of the root-mean-square residuals (more precisely, r is the ratio of the $\hat{\sigma}_\Delta$ values). The previously reported formulae have roughly 100 to 250 times as much systematic error as our fourth-order polynomial fit. While the formula in Tichy 2008 [29] performs best, we note that it has a large uncertainty.

B. Radiated energy

Following the procedure in Sec IV A, we use the data from Table III to construct a new fitting formula for the radiated energy fraction, E_{rad} , as a function of initial spin. We fit to a hyperbolic function,

$$\tilde{E}_{\text{rad}} = b_0 + \frac{b_1}{b_2 + \chi_i}. \quad (9)$$

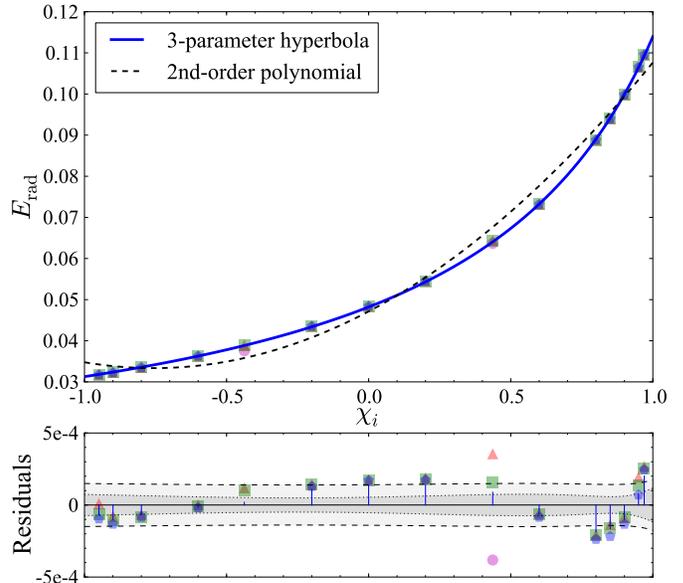


FIG. 8: In the top panel, we plot our preferred fitting formula (solid line), the hyperbolic function in Eq. (9), and a comparison with a second-order polynomial (dashed line) for E_{rad} as a function of χ_i . Our data points are plotted as polygons, where more sides indicates higher resolution level. In the bottom panel, we plot the residuals of the hyperbolic function. We indicate our fit parameter (dotted line) and total prediction (dashed line) uncertainties (defined in Appendix A).

The best fit to our data has the parameters b_n and associated covariance Σ_b :

$$b = \begin{pmatrix} b_0 \\ b_1 \\ b_2 \end{pmatrix} = \begin{pmatrix} 0.00258(29) \\ -0.07730(79) \\ -1.6939(59) \end{pmatrix} \quad (10)$$

$$\Sigma_b = \begin{pmatrix} 0.83 & 2.2 & 16 \\ 2.2 & 6.2 & 46 \\ 16 & 46 & 350 \end{pmatrix} \times 10^{-7} \quad (11)$$

The uncertainty in b_n , given in parentheses in Eq. (10), is estimated by $\sqrt{\Sigma_b^{nn}}$. In Fig. 8, we show the fit and residuals using Eq. (9) with the parameters from Eq. (10).

We use a hyperbolic fitting function instead of a second-order polynomial (as in Refs. [29, 30]) or a constrained second-order polynomial, e.g. $\tilde{E}_{\text{rad}} = c_0 + c_1\chi_i + (c_1/4)\chi_i^2$ (as in Ref. [31]). Parabolic fits show visible offsets in various regions of the initial spin space, which can be seen in plots of the residuals in [31] and in the comparison plot in the top panel of Fig. 8. The difference between the maximum LML for the 3-parameter hyperbola and the second-order polynomial is ~ 36 , indicating that the hyperbola is a dramatically better fit to the data.

In Fig. 9, we compare our data to existing fitting formulae for E_{rad} . All previous formulae suffer from the same systematic deficiencies as the best second-order polynomial fit to our data shown in Fig. 8. The ratio of

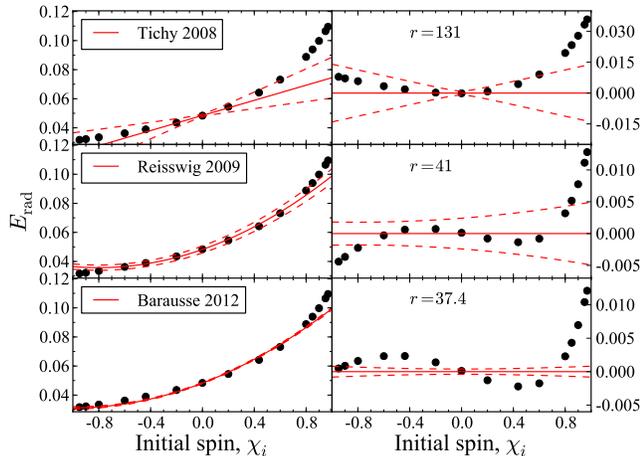


FIG. 9: E_{rad} as a function of initial spin. In the left panels we plot our data (circles) along with the fitting formula (red line) with error estimates (dashed) from several other studies. The top panel is from Ref. [29], the middle is from Ref. [30], and the bottom is from Ref. [31]. In the right panels we plot the difference between our data and the corresponding fitting formula on the left. The value r quantifies the size of the systematic error compared to the 3-parameter hyperbola.

the systematic error magnitude in these formulae to its magnitude in our 3-parameter hyperbolic fit, r , is shown in the figure and ranges from roughly 40 to 130. Note that it is not meaningful to compare these r values to those shown in Fig. 7, because we have not added any additional degrees of freedom in our E_{rad} model compared to a second-order polynomial (unlike in our final spin model, which adds two degrees of freedom).

C. Extremality

An important aspect of these fitting formulae is their ability to predict remnant properties for nearly extremal initial spins. How much of the initial mass can be radiated as gravitational waves, and how fast can the remnant hole spin? Our prediction for the radiated energy and final spin for an extremal initial spin configuration, $\chi_i = 1$, is

$$\tilde{E}_{\text{rad}}(1) = 0.11397(18), \quad (12)$$

$$\tilde{\chi}_f(1) = 0.951383(85), \quad (13)$$

where the uncertainty (in parentheses) is σ_{tot} , the total prediction uncertainty defined in Eq. (A13), evaluated at $\chi_i = 1$. The highest radiated energy predicted by any of the formulae we compare against in this paper is $E_{\text{rad}}(1) = 0.0995(8)$.

Previous estimates of E_{rad} underestimated the mass loss for large, aligned initial spins. The most extreme

data point in this paper, $S_{0.97}^{++}$, was identified as a potential outlier [31] and E_{rad} was expected to be $\lesssim 10\%$ for an extremal, aligned configuration inspiraling from infinity. Additional data presented here, most notably $S_{0.95}^{++}$ and $S_{0.9}^{++}$, suggest that $S_{0.97}^{++}$ is not an outlier. Furthermore, these cases indicate that even a $\chi_i = 0.9$ inspiral is capable of radiating $\gtrsim 10\%$ of its initial mass.

Simulations with $\chi_i > 0.93$ are an important factor in our fitting formulae. This high-spin regime is not accessible with the most popular initial data methods for binary black hole evolutions, which assume conformal flatness (cf. Ref. [39] and references therein). To assess the impact of the high spin simulations, we compare our best fits to fits of a subset of the data, omitting cases $S_{0.95}^-$, $S_{0.95}^{++}$, and $S_{0.97}^{++}$. We identify several key results.

For the χ_f formula, we find that $\Delta\tilde{\chi}_f(1) \gtrsim 2.5\sigma_{\text{tot}}^{\text{sub}}(\chi_f, 1)$. That is, the prediction of the final spin with the full dataset differs from the prediction with the subset by more than 2.5 times the total prediction uncertainty in the fit to the subset. The parameter uncertainty in the full dataset is smaller by a factor $\sigma_f/\sigma_f^{\text{sub}} \approx 0.6$, which is much smaller than would be expected from adding 3 data points randomly distributed in initial spin (the expected improvement based on the root- N rate would be $\sqrt{12/15} \approx 0.9$). In a random sampling context, one would typically have to more than double the size of the dataset to get such a reduction in parameter uncertainties. Of course, we have not chosen the subset randomly. Note that because the systematic error magnitude is negligible, $\hat{\sigma}_\Delta \approx 0$, the total prediction uncertainty of Eq. (A13) has the same behavior as the parameter uncertainty.

For the E_{rad} formula, we find that $\Delta\tilde{E}_{\text{rad}}(1) \gtrsim 3.5\sigma_{\text{tot}}^{\text{sub}}(E_{\text{rad}}, 1)$; the lower spin subset poorly predicts the extremal E_{rad} . Parameter uncertainties decrease only slightly faster than the expected root- N rate for adding 3 randomly placed data points, $\sigma_f/\sigma_f^{\text{sub}} \approx 0.85$. However, the total prediction uncertainty at $\chi_i = 1$ increases, $\sigma_{\text{tot}}/\sigma_{\text{tot}}^{\text{sub}} \approx 1.15$, because the additional high spin data deviate most from the fitting formula that is based on lower spin data. That is, E_{rad} for $\chi_i > 0.93$ is unanticipated by the fit to the lower spin subset, causing the systematic error magnitude $\hat{\sigma}_\Delta$ to increase. While this highlights the importance of the high spin data in assessing the predictive power of the fitting formula for near-extremal initial spins, it also suggests that we are unlikely to capture the behavior of E_{rad} much better with our simple fitting formula. Furthermore, neither manual nor algorithmic searches [56] have identified any better formulae, which leads us to believe that for the best predictive results at high spins, a non-parametric approach may be preferred. Such an approach could be implemented using, for example, a correlated Gaussian process [57], which would provide a way to predict final masses and spins without the use of a parametric fitting formula.

The analysis comparing the subset to the full dataset does not change in any appreciable way if we use a near-

extremal spin, e.g. $\chi_i = 0.97$, instead of the most extreme case, $\chi_i = 1$.

V. CONCLUSIONS

In this paper, we present and analyze a family of numerical relativity simulations performed using SpEC in order to construct improved fitting formulae for the final spin and radiated energy as a function of initial spin. We consider a physically motivated, one-dimensional subset of the binary black hole parameter space, in which the black holes have equal masses and equal spins aligned or anti-aligned with the orbital angular momentum. The improvement in these fitting formulae is most dramatic in the regime where the initial spin is above the ‘‘Bowen-York limit,’’ since for the first time data from simulations above this limit have been included in the fits.

For the final spin, we improve on the second-order polynomial fitting formula by using a fourth-order polynomial to capture the statistically significant cubic and quartic features. For the radiated energy, we find that a 3-parameter hyperbolic fitting formula is greatly preferred to a second-order polynomial. The qualitatively different behavior at large, aligned spins in the new fit to E_{rad} implies that there is somewhat more power in gravitational waves from nearly extremal sources than previously thought, perhaps because of higher-order effects that become relevant at very high spins.

We have shown that performing more nearly extremal simulations is the most effective way to reduce the uncertainty in the fitting formula parameters. However, we have also observed that the systematic uncertainty in E_{rad} may prohibit a simple fitting formula from providing any further significant improvement to the prediction uncertainty of E_{rad} for high, aligned spins.

Analytic models, such as the aligned-spin EOB model, are needed to generate templates for gravitational-wave detectors (e.g. LIGO), because of the prohibitive expense of generating sufficient numerical relativity data to adequately cover the parameter space. The fitting formulae we define in this paper can be used to better calibrate these models, and therefore improve future template waveforms.

Appendix A: Method for constructing our fitting formulae

Our goal is to find a convenient but reasonably accurate function that predicts the final black hole spin, χ_f , or the fractional radiated energy, E_{rad} , as a function of the initial spin, χ_i . We will specify one or more simple parametric candidate functions, find the best parameter values, quantify uncertainties in the parameters and predictions, and compare rival candidate functions. To treat both the χ_f and E_{rad} problems in generic notation, we let ξ denote the *predictor* (i.e., χ_i), and η denote

the *response* we seek to predict (i.e., χ_f or E_{rad}). We have one or more parametric models for the relationship, $\eta \approx f(\xi; \theta)$, with parameters θ (we sometimes suppress the parameter dependence below to simplify notation).

The data for the analysis are from post-processing outputs from deterministic numerical calculations of the binary black hole merger. A complex computation produces initial data (ID) targeting a specified value of ξ , but the actual value of ξ that the generated ID corresponds to necessarily differs from the target value. A processing algorithm estimates the actual value to be x . Evolution of the ID produces high-dimensional outputs that are processed to produce the computed response, y , that estimates the result, η , that would be obtained by solving the PDEs exactly. A set of (x, y) pairs constitutes the basic data we must use to find $f(\xi; \theta)$.

A variety of parameters govern the accuracy of the ID, evolution, and processing algorithms. These are summarized via a resolution level k (defined in Sec. II) assigned to each (x, y) pair, with the x and y values likely to be closer to the ξ and η values when k is larger. For every choice of ID, we have results for multiple values of k , comprising repeated measurements of (ξ, η) of varying accuracy.

We have developed a Bayesian nonlinear measurement error model for the analysis.⁵ Letting the index n label the choice of ID, the model specification is:

$$x_{nk} = \xi_n + \epsilon_{nk}, \quad (\text{A1})$$

$$y_{nk} = \eta_n + \delta_{nk} \quad (\text{A2})$$

$$= f(\xi_n; \theta) + \Delta_n + \delta_{nk}, \quad (\text{A3})$$

for N ID cases ($n = 1$ to N), and $k \in L_n$ for ID case n , where L_n denotes a set of levels run for case n (for most cases, $L_n = \{2, 3, 4\}$, but for runs targeting $\chi_i = \pm 0.44$, $L_n = \{1, 2, 3\}$). Here ϵ_{nk} and δ_{nk} denote *level error* terms reflecting the difference between numerical results at finite resolution and the actual solution to the differential equations we are studying. For Eq. (A3) we set $\eta_n = f(\xi_n; \theta) + \Delta_n$, where Δ_n is a *discrepancy* term representing the difference between the true response and the prediction based on the fitting function.

To complete the model we must assign (prior) probability density functions (PDFs), i.e. *priors*, to a number of random variables: the level errors, ϵ_{nk} and δ_{nk} ; the latent predictor variables, ξ_n ; and the latent discrepancy variables, Δ_n .

We assign independent, zero-mean normal distributions to the level error terms, ϵ_{nk} and δ_{nk} , with standard deviations σ_x/α_k and σ_y/α_k (respectively). We assign α_k scale factors to capture the notion that we expect the errors to be smaller (on average) for higher levels. For the

⁵ For introductions aimed at physicists, see Refs. [58–60] for Bayesian inference and Ref. [61] for multilevel Bayesian modeling. Multilevel measurement error models inspiring our approach here are covered in Refs. [62, 63].

calculations here, we took $\alpha_k = (1/2)^{4-k}$, so the standard deviations for the highest-resolution $k = 4$ results are σ_x and σ_y , and the error scales double for each decreasing level. We did not explore this assignment except to verify that this choice has a much higher likelihood than taking $\alpha_k = 1$, i.e., the data themselves show clear evidence for convergence as k grows. Although in principle we could let the error scale be different for each ID case, for simplicity we assign a common error scale across ID cases; the modest amount of data we have does not indicate a strong variation of error scale with ID. We adopt normal distributions, partly for convenience, but also because we are modeling relationships between scalar quantities calculated from high-dimensional computational outputs with complicated algorithms. Presuming the final errors result from numerous additive contributions whose uncertainties have finite variance, the central limit theorem motivates the normal choice.

We assign informative but relatively broad priors for the ξ_n values, reflecting the ability to produce ID corresponding to a ξ_n value close to a desired target value, μ_n . The priors are normal with means μ_n (equal to the target value for ID case n) and common standard deviation $w = 0.002$, reflecting the typical change in mass and spin as a result of the initial relaxation (as seen in Fig. 3). These values do not strongly impact the results.

We also assign independent, zero-mean normal distribution PDFs to the discrepancy terms, with common standard deviation σ_Δ . The quantity σ_Δ represents the typical scale of systematic error magnitude in the model. A more flexible and realistic choice would be to assign a correlated Gaussian process prior over the space of discrepancy *functions*, $\Delta(\xi_n)$, and to identify $\Delta_n = \Delta(\xi_n)$. This would resemble the practice in the literature on Bayesian emulation of input/output response surfaces, the prevailing approach in the literature on the statistical analysis of the results of deterministic numerical simulations (see, e.g., [64, 65]). But the goal of that literature is not to find simple and tractable fitting functions; it instead builds nonparametric emulators that, while simpler than the simulators being emulated, are still computationally nontrivial. Moreover, the vast majority of existing work on emulation addresses cases with precisely known inputs, which is not the case here; uncertainty in the predictor significantly complicates implementation of Gaussian process regression [66, 67]. The independent normal PDF for Δ_n will enable us to invoke a simple approximation leading to analytical results.

Finally, we adopt flat priors for the fitting function model parameters, θ .

The conditional dependency structure of such a multi-level model can be represented by a directed acyclic graph (DAG). A graphical model of this type can be readily coded in a DAG-oriented statistical modeling language (e.g., WinBUGS or JAGS) to enable Bayesian computation via Markov chain Monte Carlo (MCMC) posterior sampling. Here the focused goal (finding a simple fitting function) and the small uncertainties in the level error

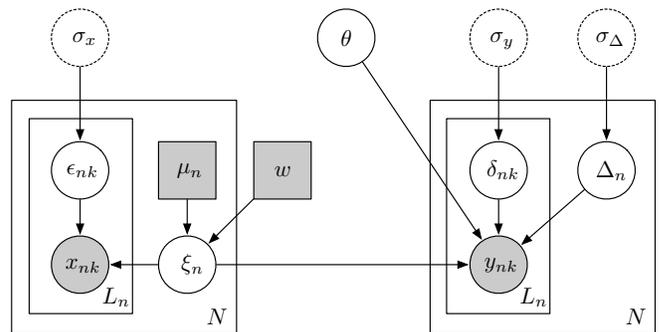


FIG. 10: Directed acyclic graph displaying the conditional dependence structure of the Bayesian nonlinear measurement error model adopted for the fitting function analysis. See text for a detailed description.

terms (well below 1% for nonzero spins) motivated an analytical approach based on linearization of $f(\xi)$. This lets us avoid the complexity of MCMC, producing a fast algorithm that is relatively simple to use.

Figure 10 shows the DAG for our model. Circles denote random variables (RVs, uncertain quantities with assigned or computed PDFs). Shaded circles are the data (x, y), and shaded squares are fixed constants that help define the model. We marginalize over the error RVs (ϵ_{nk} , δ_{nk} , and Δ_n) and the uncertain input variables (ξ_n), and then we solve for the remaining non-shaded variables simultaneously. The plates (enclosing boxes) denote parts of the graph that are replicated as indicated by the quantity in the lower right corner of each plate.

Dashed circles indicate RVs that play the role of *hyperparameters*, i.e., parameters defining prior PDFs for lower-level RVs. Formally, we could account for uncertainty in the hyperparameters by assigning them priors of their own and marginalizing over them (the *hierarchical Bayes* approach). As a simpler approximation, we optimized these hyperparameters (the *empirical Bayes* approach), using constant prior PDFs for them.

Directed edges (arrows) in a Bayesian network DAG are used to indicate the dependency structure. The top-level RVs have no dependencies (incoming arrows); their PDFs would be specified a priori for a full hierarchical Bayesian analysis. The conditional PDFs of lower-level RVs depend only on the values of their dependencies. The full joint PDF for all RVs is the product of the prior and conditional PDFs. Therefore, Fig. 10 indicates that the joint PDF for the RVs comprising our model may be written

$$\begin{aligned}
 p(\theta, \sigma_x, \sigma_y, \sigma_\Delta, \xi, \Delta, \epsilon, \delta, x, y) &= p(\theta)p(\sigma_x)p(\sigma_y)p(\sigma_\Delta) \\
 &\times \prod_{n=1}^N \left[p(\xi_n | \mu_n, w) p(\Delta_n | \sigma_\Delta) \right. \\
 &\times \left. \prod_{k \in L_n} p(\epsilon_{nk} | \sigma_x) p(x_{nk} | \xi_n, \epsilon_{nk}) \right]
 \end{aligned}$$

$$\times p(\delta_{nk}|\sigma_y)p(y_{nk}|\theta, \xi_n, \Delta_n, \delta_{nk}) \Big], \quad (\text{A4})$$

where $(\xi, \Delta, \epsilon, \delta, x, y)$ is shorthand notation for the indexed collections of those variables. Since we are adopting a constant prior PDF for θ , and an empirical Bayes treatment of the hyperparameters $\psi = (\sigma_x, \sigma_y, \sigma_\Delta)$, the quantity of interest is the conditional PDF for the data, (x, y) , and the latent parameters, $(\xi, \Delta, \epsilon, \delta)$, given the fitting function parameters and the hyperparameters,

$$p(\xi, \Delta, \epsilon, \delta, x, y|\theta, \psi) = \prod_{n=1}^N \left[p(\xi_n|\mu_n, w) p(\Delta_n|\sigma_\Delta) \right. \\ \times \prod_{k \in L_n} p(\epsilon_{nk}|\sigma_x) p(x_{nk}|\xi_n, \epsilon_{nk}) \\ \left. \times p(\delta_{nk}|\sigma_y) p(y_{nk}|\theta, \xi_n, \Delta_n, \delta_{nk}) \right], \quad (\text{A5})$$

The model Eqs. (A1) and (A3) imply that the conditional PDFs for x_{nk} and y_{nk} in these equations are δ -functions. This lets us trivially marginalize over ϵ and δ , giving a marginal PDF for the remaining variables,

$$p(\xi, \Delta, x, y|\theta, \psi) = \prod_{n=1}^N \left[p(\xi_n|\mu_n, w) p(\Delta_n|\sigma_\Delta) \right. \\ \times \prod_{k \in L_n} p(\epsilon_{nk} = x_{nk} - \xi_n|\sigma_x) \\ \left. \times p(\delta_{nk} = y_{nk} - f(\xi_n; \theta) - \Delta_n|\sigma_y) \right]. \quad (\text{A6})$$

Marginalizing over ξ and Δ gives the *marginal likelihood function* (the probability for the data, conditioned on parameter values) for the fitting parameters and hyperparameters,

$$\mathcal{L}_M(\theta, \psi) = \prod_{n=1}^N \int d\xi_n \int d\Delta_n \left[\mathcal{N}(\xi_n|\mu_n, w) \mathcal{N}(\Delta_n|\sigma_\Delta) \right. \\ \left. \times \prod_{k \in L_n} \mathcal{N}(x_{nk} - \xi_n|0, \sigma_x) \mathcal{N}(y_{nk} - f(\xi_n; \theta) - \Delta_n|0, \sigma_y) \right], \quad (\text{A7})$$

where $\mathcal{N}(z|\mu, \sigma)$ denotes the normal distribution PDF for z with mean μ and standard deviation σ ,

$$\mathcal{N}(z|\mu, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} e^{-(z-\mu)^2/2\sigma^2}. \quad (\text{A8})$$

When $f(\xi)$ is a nonlinear function of ξ , the ξ integral in Eq. (A7) is in general intractable. However, the x and y errors are small, so we expect a local linear approximation of $f(\xi)$ to be very accurate over regions of ξ_n that have significant probability density. So we use

$$f(\xi_n; \theta) \approx f(\tilde{\xi}_n; \theta) + (\xi_n - \tilde{\xi}_n) f'(\tilde{\xi}_n; \theta) \quad (\text{A9})$$

in Eq. (A7), where $f'(\xi; \theta)$ denotes the derivative of the fitting function with respect to ξ , and $\tilde{\xi}_n$ is a fixed reference value of ξ_n based on the x_{nk} values for a particular n (we use a weighted mean of the x_{nk}). With this linearization, the integrals in the marginal likelihood function can be performed analytically.

We estimate the parameters for a candidate fitting function by maximizing the marginal likelihood function over both θ and ψ :

$$(\hat{\theta}, \hat{\psi}) = \arg \max \mathcal{L}_M(\theta, \psi). \quad (\text{A10})$$

For the fitting functions studied here, the θ dependence of the marginal likelihood function is approximately multivariate Gaussian. To quantify the θ uncertainties, we calculate the observed Fisher information matrix (with ψ fixed at $\hat{\psi}$),

$$I_{\alpha\beta} = \frac{\partial^2}{\partial\theta_\alpha \partial\theta_\beta} \mathcal{L}_M(\theta, \psi) \Big|_{\hat{\theta}, \hat{\psi}}, \quad (\text{A11})$$

where θ_α denotes the α th parameter of the fitting function. The posterior PDF for θ (conditional on $\hat{\psi}$) is then approximately a multivariate normal PDF with mean $\hat{\theta}$ and covariance matrix $\Sigma = I^{-1}$.

To predict the value of the response at a specified value of ξ , we calculate an approximate predictive distribution (also conditioned on $\hat{\psi}$) using the multivariate normal PDF and the delta method (propagation of errors). The model assumes the response is given by the sum of the fitting function and a discrepancy term with zero mean. The most probable value of the response is simply $f(\xi; \hat{\theta})$. There are two components to the uncertainty in the prediction. One comes from propagating the θ uncertainty (accounting for correlations between the parameters, which can be large). The resulting standard deviation in the fitting function evaluated at ξ is $\sigma_f(\xi)$, satisfying

$$\sigma_f^2(\xi) = \sum_{\alpha\beta} \frac{\partial f(\xi; \theta)}{\partial\theta_\alpha} \Sigma_{\alpha\beta} \frac{\partial f(\xi; \theta)}{\partial\theta_\beta}. \quad (\text{A12})$$

The full uncertainty in the predicted response must also account for the uncertainty in the discrepancy term, which is given by the hyperparameter σ_Δ that we estimate from the data. The full uncertainty in the prediction is

$$\sigma_{\text{tot}}(\xi) = \sqrt{\sigma_f^2(\xi) + \hat{\sigma}_\Delta^2}. \quad (\text{A13})$$

This calculation ignores the uncertainty in the value of σ_Δ , but that uncertainty is relatively small in our calculations.

To compare rival parametric fitting functions, a formal model comparison could be implemented, e.g., using Bayes factors (which would require assigning normalized priors to θ for each candidate fitting function, and integrating the product of the prior and the marginal

likelihood function over θ), or an information criterion such as the Bayesian information criterion (BIC) or the Akaike information criterion (AIC). The BIC and AIC rank models according to their maximum likelihoods, penalized by a term depending on the number of parameters in each model (and the sample size in the case of the BIC). These criteria were developed for comparing simple parametric models, not multilevel models with many latent parameters. We adopt a less formal approach here. We simply calculate the logarithm of ratios of the maximum marginal likelihood function. For the models we consider, the log-ratio for the best model vs. the next-best competitor is large (well over 10), far larger than the typical penalty terms in information criteria, so the choice of best model is unambiguous.

Acknowledgments

We thank Matt Giesler, Tony Chu, and Bryant Garcia for providing data from their simulations. We grate-

fully acknowledge support from the Sherman Fairchild Foundation; from NSF grants PHY-0969111 and PHY-1005426 at Cornell; and from NSF grants PHY-1068881, PHY-1005655, and DMS-1065438 at Caltech. Loredó's work for this project was supported by NSF grant AST-0908439 and NASA grant NNX09AK60G. Simulations used in this work were computed with SpEC [32]. Computations were performed on SHC at Caltech, which is supported by the Sherman Fairchild Foundation; on the Zwicky cluster at Caltech, which is supported by the Sherman Fairchild Foundation and by NSF award PHY-0960291; on the NSF XSEDE network under grant TG-PHY990007N; and on the GPC supercomputer at the SciNet HPC Consortium [68]. SciNet is funded by: the Canada Foundation for Innovation under the auspices of Compute Canada; the Government of Ontario; Ontario Research Fund–Research Excellence; and the University of Toronto.

-
- [1] B. Abbott et al. (LIGO Scientific Collaboration), Rept. Prog. Phys. **72**, 076901 (2009), 0711.3041.
- [2] H. Grote (LIGO Scientific Collaboration), Class.Quant.Grav. **25**, 114043 (2008).
- [3] F. Acernese, M. Alshourbagy, P. Amico, F. Antonucci, S. Aoudia, et al., Class.Quant.Grav. **25**, 184001 (2008).
- [4] J. Abadie et al. (The LIGO Scientific Collaboration and the Virgo Collaboration), Phys. Rev. D **83**, 122005 (2011), 1102.3781.
- [5] The LIGO Scientific Collaboration, the Virgo Collaboration, J. Aasi, J. Abadie, B. P. Abbott, R. Abbott, T. D. Abbott, M. Abernathy, T. Accadia, F. Acernese, et al., ArXiv e-prints (2012), 1209.6533.
- [6] K. S. Thorne, Astrophys. J. **191**, 507 (1974).
- [7] C. F. Gammie, S. L. Shapiro, and J. C. McKinney, Astrophys. J. **602**, 312 (2004).
- [8] M. Kesden, G. Lockhart, and E. S. Phinney, Phys. Rev. D **82**, 124045 (2010), arXiv:1005.0627.
- [9] J. E. McClintock, R. Shafee, R. Narayan, R. A. Remillard, S. W. Davis, and L.-X. Li, Astrophys. J. **652**, 518 (2006).
- [10] J.-M. Wang, Y.-M. Chen, L. C. Ho, and R. J. McLure, Astrophys. J. **642**, L111 (2006).
- [11] L. Brenneman, C. Reynolds, M. Nowak, R. Reis, M. Trippe, et al., Astrophys.J. **736**, 103 (2011), 1104.1172.
- [12] P. Ajith, M. Hannam, S. Husa, Y. Chen, B. Brügmann, N. Dorband, D. Müller, F. Ohme, D. Pollney, C. Reisswig, et al., Phys. Rev. Lett. **106**, 241101 (2011), URL <http://link.aps.org/doi/10.1103/PhysRevLett.106.241101>.
- [13] C. Van Den Broeck, D. A. Brown, T. Cokelaer, I. Harry, G. Jones, et al., Phys.Rev. **D80**, 024009 (2009), 0904.1715.
- [14] J. Aasi et al. (the LIGO Scientific Collaboration, the Virgo Collaboration) (2013), 1304.1775.
- [15] A. Taracchini, Y. Pan, A. Buonanno, E. Barausse, M. Boyle, et al., Phys. Rev. D **86**, 024011 (2012), 1202.0790.
- [16] J. F. Steiner and J. E. McClintock, Astrophys.J. **745**, 136 (2012), 1110.6849.
- [17] T. Bogdanovic, C. S. Reynolds, and M. C. Miller, Astrophys.J.Lett. (2007), astro-ph/0703054.
- [18] T. Fragos, M. Tremmel, E. Rantsiou, and K. Belczynski (2010), 1001.1107.
- [19] G. Lovelace, M. Boyle, M. A. Scheel, and B. Szilágyi, Class. Quant. Grav. **29**, 045003 (2012), arXiv:1110.2229.
- [20] G. Lovelace, M. A. Scheel, and B. Szilágyi, Phys. Rev. D **83**, 024010 (2011), 1010.2777.
- [21] B. Garcia, G. Lovelace, L. E. Kidder, M. Boyle, S. A. Teukolsky, M. A. Scheel, and B. Szilágyi, Phys. Rev. D **86**, 084054 (2012), 1206.2943.
- [22] T. Chu, H. P. Pfeiffer, and M. A. Scheel, Phys. Rev. D **80**, 124051 (2009), 0909.1313.
- [23] T. Chu, Ph.D. thesis, California Institute of Technology (2012).
- [24] G. B. Cook and J. W. York, Jr., Phys. Rev. D **41**, 1077 (1990).
- [25] S. Dain, C. O. Lousto, and R. Takahashi, Phys. Rev. D **65**, 104038 (2002).
- [26] M. Hannam, S. Husa, and N. O. Murchadha, Phys. Rev. D **80**, 124007 (2009).
- [27] L. Rezzolla et al., Astrophys. J. **679**, 1422 (2008), 0708.3999.
- [28] E. Barausse and L. Rezzolla, Astrophys.J. **704**, L40 (2009), 0904.2577.
- [29] W. Tichy and P. Marronetti, Phys. Rev. D **78**, 081501 (2008), 0807.2985.
- [30] C. Reisswig, S. Husa, L. Rezzolla, E. N. Dorband, D. Pollney, et al., Phys. Rev. D **80**, 124026 (2009), 0907.0462.
- [31] E. Barausse, V. Morozova, and L. Rezzolla, Astrophys.J.

- 758**, 63 (2012), 1206.3803.
- [32] <http://www.black-holes.org/SpEC.html>.
- [33] J. W. York, Phys. Rev. Lett. **82**, 1350 (1999).
- [34] G. B. Cook, Phys. Rev. D **65**, 084003 (2002).
- [35] G. B. Cook and H. P. Pfeiffer, Phys. Rev. D **70**, 104016 (2004).
- [36] M. Caudill, G. B. Cook, J. D. Grigsby, and H. P. Pfeiffer, Phys. Rev. D **74**, 064011 (2006), gr-qc/0605053.
- [37] E.ourgoulhon, P. Grandclément, and S. Bonazzola, Phys. Rev. D **65**, 044020 (2002).
- [38] P. Grandclément, E.ourgoulhon, and S. Bonazzola, Phys. Rev. D **65**, 044021 (2002).
- [39] G. Lovelace, R. Owen, H. P. Pfeiffer, and T. Chu, Phys. Rev. D **78**, 084017 (2008).
- [40] G. Lovelace, Class. Quantum Grav. **26**, 114002 (2009).
- [41] H. P. Pfeiffer, L. E. Kidder, M. A. Scheel, and S. A. Teukolsky, Comput. Phys. Commun. **152**, 253 (2003).
- [42] L. T. Buchman, H. P. Pfeiffer, M. A. Scheel, and B. Szilágyi, Phys. Rev. D **86**, 084033 (2012), 1206.3015.
- [43] L. Lindblom and B. Szilágyi, Phys. Rev. D **80**, 084019 (2009), arXiv:0904.4873.
- [44] M. W. Choptuik and F. Pretorius, Phys. Rev. Lett. **104**, 111101 (2010), 0908.1780.
- [45] B. Szilágyi, L. Lindblom, and M. A. Scheel, Phys. Rev. D **80**, 124010 (2009), 0909.3557.
- [46] A. H. Mroué, H. P. Pfeiffer, L. E. Kidder, and S. A. Teukolsky, Phys. Rev. D **82**, 124016 (2010), arXiv:1004.4697 [gr-qc].
- [47] D. A. Hemberger, M. A. Scheel, L. E. Kidder, B. Szilágyi, G. Lovelace, N. W. Taylor, and S. A. Teukolsky, Class. Quantum Grav. **30**, 115001 (2013), 1211.6079, URL <http://stacks.iop.org/0264-9381/30/i=11/a=115001>.
- [48] C. Gundlach, Phys. Rev. D **57**, 863 (1998), URL <http://link.aps.org/abstract/PRD/v57/p863>.
- [49] I. Booth and S. Fairhurst, Phys. Rev. D **77**, 084005 (2008), arXiv:0708.2209v3.
- [50] L. Lindblom, M. A. Scheel, L. E. Kidder, R. Owen, and O. Rinne, Class. Quantum Grav. **23**, S447 (2006).
- [51] K. Alvi, Phys. Rev. D **64**, 104020 (2001), URL <http://link.aps.org/doi/10.1103/PhysRevD.64.104020>.
- [52] M. Campanelli, C. O. Lousto, and Y. Zlochower, Phys. Rev. D **74**, 084023 (2006), astro-ph/0608275.
- [53] C. O. Lousto, M. Campanelli, Y. Zlochower, and H. Nakano, Class. Quant. Grav. **27**, 114006 (2010), arXiv:0904.3541 [gr-qc].
- [54] L. A. Gergely and P. L. Biermann (2012), 1208.5251.
- [55] S. S. Wilks, The Annals of Mathematical Statistics **9**, 60 (1938), ISSN 00034851, URL <http://dx.doi.org/10.1214/aoms/1177732360>.
- [56] M. Schmidt and H. Lipson, Science **324**, 81 (2009), URL <http://www.sciencemag.org/content/324/5923/81.abstract>.
- [57] C. E. Rasmussen and C. K. I. Williams, *Gaussian processes for machine learning*, Adaptive Computation and Machine Learning (MIT Press, Cambridge, MA, 2006), ISBN 978-0-262-18253-9.
- [58] E. T. Jaynes and G. L. Bretthorst, *Probability Theory* (2003).
- [59] D. S. Sivia, *Data analysis*, Oxford Science Publications (Oxford University Press, Oxford, 2006), 2nd ed., ISBN 978-0-19-856832-2, a Bayesian tutorial, With J. Skilling.
- [60] P. Gregory, *Bayesian Logical Data Analysis for the Physical Sciences* (2010).
- [61] T. Loredo, in *Astrostatistical Challenges for the New Astronomy*, edited by J. M. Hilbe (Springer New York, 2013), vol. 1 of *Springer Series in Astrostatistics*, pp. 15–40, ISBN 978-1-4614-3507-5, URL http://dx.doi.org/10.1007/978-1-4614-3508-2_2.
- [62] R. J. Carroll, D. Ruppert, L. A. Stefanski, and C. M. Crainiceanu, *Measurement error in nonlinear models*, vol. 105 of *Monographs on Statistics and Applied Probability* (Chapman & Hall/CRC, Boca Raton, FL, 2006), 2nd ed., ISBN 978-1-58488-633-4; 1-58488-633-1, a modern perspective, URL <http://dx.doi.org/10.1201/9781420010138>.
- [63] D. Giltinan, *Nonlinear models for repeated measurement data*, Monographs on Statistics and Applied Probability (Chapman and Hall, 1995), ISBN 9780412983412, URL <http://books.google.com/books?id=0eSIBPAL4qsC>.
- [64] M. C. Kennedy and A. O'Hagan, J. R. Stat. Soc. Ser. B Stat. Methodol. **63**, 425 (2001), ISSN 1369-7412, URL <http://dx.doi.org/10.1111/1467-9868.00294>.
- [65] D. Higdon, M. Kennedy, J. C. Cavendish, J. A. Cafoe, and R. D. Ryne, SIAM J. Sci. Comput. **26**, 448 (2004), ISSN 1064-8275, URL <http://dx.doi.org/10.1137/S1064827503426693>.
- [66] A. McHutchon and C. E. Rasmussen, in *Advances in Neural Information Processing Systems 24*, edited by J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, and K. Weinberger (2011), pp. 1341–1349, URL <http://books.nips.cc/nips24.html#P0780>.
- [67] P. Dallaire, C. Besse, and B. Chaib-draa, Neurocomputing **74**, 1945 (2011), ISSN 0925-2312, URL <http://www.sciencedirect.com/science/article/pii/S0925231211000440>.
- [68] C. Loken, D. Gruner, L. Groer, R. Peltier, N. Bunn, M. Craig, T. Henriques, J. Dempsey, C.-H. Yu, J. Chen, et al., J. Phys.: Conf. Ser. **256**, 012026 (2010).