

Numerical relativity waveform surrogate model for generically precessing binary black hole mergers

Jonathan Blackman,¹ Scott E. Field,^{2,3} Mark A. Scheel,¹ Chad R. Galley,¹ Christian D. Ott,^{1,4} Michael Boyle,³ Lawrence E. Kidder,³ Harald P. Pfeiffer,⁵ and Béla Szilágyi^{1,6}

¹*Theoretical Astrophysics 350-17, California Institute of Technology, Pasadena, California 91125, USA*

²*Mathematics Department, University of Massachusetts Dartmouth, Dartmouth, Massachusetts 02747, USA*

³*Cornell Center for Astrophysics and Planetary Science, Cornell University, Ithaca, New York 14853, USA*

⁴*Center for Gravitational Physics and International Research Unit of Advanced Future Studies, Yukawa Institute for Theoretical Physics, Kyoto University, Kyoto, Japan*

⁵*Canadian Institute for Theoretical Astrophysics, 60 St. George Street, University of Toronto, Toronto, Ontario M5S 3H8, Canada*

⁶*Caltech JPL, Pasadena, California 91109, USA*

(Received 22 May 2017; published 31 July 2017)

A generic, noneccentric binary black hole (BBH) system emits gravitational waves (GWs) that are completely described by seven intrinsic parameters: the black hole spin vectors and the ratio of their masses. Simulating a BBH coalescence by solving Einstein's equations numerically is computationally expensive, requiring days to months of computing resources for a single set of parameter values. Since theoretical predictions of the GWs are often needed for many different source parameters, a fast and accurate model is essential. We present the first surrogate model for GWs from the coalescence of BBHs including all seven dimensions of the intrinsic noneccentric parameter space. The surrogate model, which we call NRSur7dq2, is built from the results of 744 numerical relativity simulations. NRSur7dq2 covers spin magnitudes up to 0.8 and mass ratios up to 2, includes all $\ell \leq 4$ modes, begins about 20 orbits before merger, and can be evaluated in ~ 50 ms. We find the largest NRSur7dq2 errors to be comparable to the largest errors in the numerical relativity simulations, and more than an order of magnitude smaller than the errors of other waveform models. Our model, and more broadly the methods developed here, will enable studies that were not previously possible when using highly accurate waveforms, such as parameter inference and tests of general relativity with GW observations.

DOI: [10.1103/PhysRevD.96.024058](https://doi.org/10.1103/PhysRevD.96.024058)

I. INTRODUCTION

With LIGO's two confident detections of gravitational waves (GWs) from binary black hole (BBH) systems [1,2], we have entered the exciting new era of GW astronomy. The masses and spins of the source black holes (BHs) can be determined by comparing the signal to waveforms predicted by general relativity (GR) [3,4], and new strong-field tests of GR can be performed [5]. These measurements and tests require GW models that are both accurate and fast to evaluate. The total mass of the system M can be scaled out of the problem, leaving a seven-dimensional noneccentric intrinsic parameter space over which the waveform must be modeled, consisting of the mass ratio and two BH spin vectors.

Numerical relativity (NR) simulations of BBH mergers [6–12] solve the full Einstein equations and produce the most accurate waveforms. These simulations are computationally expensive, requiring weeks to months on dozens of CPU cores for a waveform beginning ~ 20 orbits before the merger. Analytic and semianalytic waveform models [13–21] are quick to evaluate, but they make approximations that can introduce differences with respect to the true

waveform predicted by GR. These differences could lead to parameter biases or inaccurate tests of GR for some high signal-to-noise ratio detections that could be made in the near future [22,23].

A surrogate waveform model [16,24–27] is a model that takes a set of precomputed waveforms that were generated by some other model (e.g., NR or a semianalytic model), and interpolates in parameter space between these waveforms to quickly produce a waveform for any desired parameter values. A surrogate waveform can be evaluated much more quickly than the underlying model, and can be made as accurate as the underlying model given a sufficiently large set of precomputed waveforms that cover the parameter space. Previous surrogate models based on NR waveforms were built for nonspinning BBH systems [25] and for a four-dimensional parameter subspace containing precession [24]. Here, we present the first NR surrogate model including all seven dimensions of the parameter space. The model, which we call NRSur7dq2, produces waveforms nearly as accurate as those from NR simulations, but can be evaluated in ~ 50 ms on a single CPU core for a speedup of more than 8 orders of magnitude compared

to NR. Our method enables performing high accuracy GW data analysis, including parameter inference for astrophysics and tests of GR.

II. NUMERICAL RELATIVITY DATA

The NR simulations used to build the surrogate model are performed using the Spectral Einstein Code (SpEC) [8,28–33]. The simulations begin at a coordinate time $\tau = 0$, where we specify the BH mass ratio $q = m_1/m_2 \geq 1$ and initial dimensionless spin vectors

$$\vec{\chi}_i(\tau = 0) = \vec{S}_i(\tau = 0)/m_i^2, i \in \{1, 2\}. \quad (1)$$

The system is evolved through merger and ringdown, and the GWs are extracted at multiple finite radii from the source. These are extrapolated to future null infinity [34] using quadratic polynomials in $1/r$, where r is a radial coordinate. The effects of any drifts in the center of mass that are linear in time are removed from the waveform [35–38]. The waveforms at future null infinity use a time coordinate $\tilde{\tau}$, which is different from the simulation time τ , and begins approximately at $\tilde{\tau} = 0$. The spins $\vec{\chi}_i(\tau)$ are also measured at each simulation time. To compare spin and waveform features, we identify τ with $\tilde{\tau}$. While this identification is not gauge independent, the spin directions are already gauge dependent. We note that the spin and orbital angular momentum vectors in the damped harmonic gauge used by SpEC agree quite well with the corresponding vectors in post-Newtonian (PN) theory [39].

Once we have the spins $\vec{\chi}_i(\tau)$ and spin-weighted spherical harmonic modes of the waveform $h^{\ell,m}(\tau)$, we perform the same alignment discussed in Sec. III.D of Ref. [24]. Briefly, for each simulation, we first determine the time τ_{peak} which maximizes the total amplitude of the waveform

$$A_{\text{tot}}(\tau) = \sqrt{\sum_{\ell,m} |h^{\ell,m}(\tau)|^2}. \quad (2)$$

We determine τ_{peak} by fitting a quadratic function to five adjacent samples of $A_{\text{tot}}(\tau)$, consisting of the largest sample and two neighbors on either side. We choose a new time coordinate

$$t = \tau - \tau_{\text{peak}}, \quad (3)$$

which maximizes A_{tot} at $t = 0$. We then rotate the waveform modes such that at our reference time of $t = t_0 = -4500M$, \hat{z} is the principal eigenvector of the angular momentum operator [40] and the phases of $h^{2,2}(t_0)$ and $h^{2,-2}(t_0)$ are equal. We sample the waveform and spins in steps of $\delta t = 0.1M$, from t_0 to $t_f = 100M$, by interpolating the real and imaginary parts of each waveform mode, as well as the spin components, using cubic splines. The initial separations and velocities of the BBH systems

were chosen such that, after aligning the peak amplitude to $t = 0$ as above, the waveforms begin at $t \approx -5000M$. Choosing $t_0 = -4500M$ discards the first $\sim 500M$ which is contaminated by junk radiation [41].

We first include all 276 NR simulations used in the NRSur4d2s surrogate model and the nine additional simulations used in Sec. IV.D and Table V of Ref. [24]. We perform 459 additional NR simulations. The first 361 of these are chosen based on sparse grids [42,43] and include combinations of extremal parameter values (such as $q \in \{1, 2\}$) and intermediate values as detailed in Appendix A. The parameters for the remaining 98 simulations are chosen as follows. We randomly sample 1000 points in parameter space uniformly in mass ratio, spin magnitude, and spin direction on the sphere. We compute the distance between points a and b using

$$ds^2 = (0.3(q_a - q_b))^2 + \sum_{i \in \{1,2\}} \|\vec{\chi}_{ia} - \vec{\chi}_{ib}\|^2. \quad (4)$$

The coefficients multiplying each term in this expression have been chosen somewhat arbitrarily, although our expectation is that any choice of order unity should provide a reasonable criteria for point selection. For each sampled parameter, we compute the minimum distance to all previously chosen parameters. We then choose the sampled parameter maximizing this minimum distance. We then resample the 1000 parameters for the next of the 98 iterations. This results in a total of 744 NR simulations. For simulations with equal masses and unequal spins, we use the results twice by reversing the labeling of the BHs and rotating the waveform accordingly. There are 142 such simulations, leading to 886 NR waveforms.

III. WAVEFORM DECOMPOSITION

The goal of a surrogate model is to take a precomputed set of waveform modes $\{h_i^{\ell,m}(t)\}$ at a fixed set of points in parameter space $\{\vec{\lambda}_i\}$, and to produce waveform modes $\{h^{\ell,m}(t)\}$ at new desired parameter values. Because $h^{\ell,m}(t)$ is highly oscillatory and changes in a complicated way as one varies the masses and spins, it is not feasible to directly interpolate $\{h_i^{\ell,m}(t)\}$ in parameter space with only $886^{1/7} \approx 2.64$ available points per dimension. Instead, we decompose each waveform $h(t)$ into many waveform data pieces. Each waveform data piece is a simpler function that varies slowly over parameters. Once we have interpolated each waveform data piece to a desired point in parameter space, we recombine them to form $h(t)$. Our decomposition is similar to but improves upon the one used in Ref. [24].

We first determine the unit quaternions $\hat{q}(t)$ that define the coprocessing frame [40,44,45], and we determine the waveform modes $\{h_C^{\ell,m}(t)\}$ in this frame. This is done using the transformation T_C given by Eq. (25) of Ref. [24]. The orbital phase

$$\varphi(t) = \frac{1}{4}(\arg[h_C^{2,-2}(t)] - \arg[h_C^{2,2}(t)]) \quad (5)$$

is computed from the coprecessing waveform modes. This is expected to be superior to computing the orbital phase from the BH trajectories because unlike the coordinate-dependent trajectories, the waveform can be made gauge invariant up to Bondi-Metzner-Sachs transformations [35].

We filter the spins $\vec{\chi}_i$ in the inertial frame using the same Gaussian filter that was used to filter $\hat{q}(t)$ in Ref. [24], and note that here we do not filter $\hat{q}(t)$. The filtered spins are given by

$$\vec{\chi}_i^{\text{filt}}(t) = \frac{\int_{\varphi_-}^{\varphi_+} \vec{\chi}_i(t'(\varphi))G(\varphi)d\varphi}{\int_{\varphi_-}^{\varphi_+} G(\varphi)d\varphi}, \quad (6)$$

where $t'(\varphi)$ is the inverse function of Eq. (5), and φ_{\pm} and $G(\varphi)$ are given by Eqs. (39)–(43) of Ref. [24]. Note that φ_{\pm} and $G(\varphi)$ implicitly depend on t . The spins are then transformed to the coprecessing frame using

$$\vec{\chi}_i^{\text{copr}}(t) = \hat{q}^{-1}(t)\vec{\chi}_i^{\text{filt}}(t)\hat{q}(t). \quad (7)$$

Note that quaternion multiplication is used here, and vectors are treated as quaternions with zero scalar component. We find that filtering the spins leads to a more accurate surrogate model by suppressing the orbital time scale oscillations in the spin components [39] and therefore making the spin time derivatives easier to model.

We then transform the spins and waveform modes to a co-orbital frame, in which the BHs are nearly on the x axis. The co-orbital frame is just the coprecessing frame rotated by $\varphi(t)$ about the z axis. Specifically, we have

$$\hat{q}_r(t) = \cos\left(\frac{\varphi(t)}{2}\right) + \hat{z} \sin\left(\frac{\varphi(t)}{2}\right), \quad (8)$$

$$\vec{\chi}_i^{\text{coorb}}(t) = \hat{q}_r^{-1}(t)\vec{\chi}_i^{\text{copr}}(t)\hat{q}_r(t), \quad (9)$$

$$h_{\text{coorb}}^{\ell,m}(t) = h_C^{\ell,m}(t)e^{im\varphi(t)}, \quad (10)$$

where $\hat{q}_r(t)$ is a unit quaternion representing a rotation about the \hat{z} axis by φ . Finally, using fourth-order finite differences, we compute the orbital frequency

$$\omega(t) = \frac{d}{dt}\varphi(t) \quad (11)$$

and the spin time derivatives in the coprecessing frame, which we then transform to the co-orbital frame

$$\dot{\vec{\chi}}_i^{\text{coorb}}(t) = \hat{q}_r^{-1}(t)\dot{\vec{\chi}}_i^{\text{copr}}(t)\hat{q}_r(t), \quad (12)$$

where a dot means d/dt . For the precession dynamics, we compute the angular velocity of the coprecessing frame

$$\frac{1}{2}\vec{\Omega}^{\text{copr}}(t) = \lim_{dt \rightarrow 0} \frac{1}{dt}(\hat{q}^{-1}(t)\hat{q}(t+dt) - 1) \quad (13)$$

$$= s(t)\dot{\vec{v}}(t) - \dot{s}(t)\vec{v}(t) - \vec{v}(t) \times \dot{\vec{v}}(t), \quad (14)$$

where $s(t)$ and $\vec{v}(t)$ are the scalar and vector components of $\hat{q}(t)$. Fourth-order finite difference stencils are used to compute the time derivatives appearing in Eq. (14). We also transform $\vec{\Omega}^{\text{copr}}(t)$ to the co-orbital frame to obtain $\vec{\Omega}^{\text{coorb}}(t)$ as in Eq. (12). The minimal rotation condition of the coprecessing frame ensures

$$\Omega_z^{\text{coorb}}(t) = \Omega_z^{\text{copr}}(t) = 0 \quad (15)$$

up to finite difference errors.

Given a waveform data piece $X(t)$ evaluated at a set of parameters, one would be tempted to parametrize $X(t)$ at any fixed time t_i by the mass ratio and the initial spins, and then construct a fit to $X(t_i)$ as a function of these parameters. However, we find much better fits if we instead parametrize $X(t_i)$ by the spins at time t_i and the mass ratio. While this is easy to do during the inspiral where we still have two BHs with individual spins, we seek a way to extend this parametrization through the merger and ringdown, where individual BH spins are no longer available. We extend, unphysically, the spin evolution through the merger and ringdown using the PN expressions

$$\frac{d}{dt}\vec{\chi}_i = \vec{\Omega}_i^{\text{Spin}} \times \vec{\chi}_i, \quad (16)$$

where $\vec{\chi}_i$ is the spin in the inertial frame, and $\vec{\Omega}_i^{\text{Spin}}$ is a PN expression given by Eq. (A32) of Ref. [39]. $\vec{\Omega}_i^{\text{Spin}}$ is a




FIG. 1. The x component of a spin extended through merger and ringdown with PN expressions. These spins are not physically meaningful, but provide a parametrization of the system leading to accurate fits. The thick solid blue curve and dashed orange curve show the unfiltered and filtered spins from a NR simulation, and are not measured past $t = -6M$ due to the merger of the BHs. Each thin line is identical to the filtered NR curve before some time t_{PN} indicated with a dot, after which the spins are evolved using Eq. (16). The spins during the ringdown are affected somewhat by the choice of t_{PN} , but the overall phasing is quite similar.

function of the orbital angular momentum vector $\hat{l}(t)$, a vector pointing from one BH to the other $\hat{n}(t)$, and the PN parameter $x(t)$. Evaluating $\vec{\Omega}_i^{\text{Spin}}$ requires several quantities that are typically computed from BH trajectories in PN theory. Since the trajectories are also not available after the merger, we compute them from the waveform. We take \hat{l} and \hat{n} to be the \hat{z} and \hat{x} axes of the co-orbital frame, and we take the PN parameter x to be $\omega^{2/3}$, where ω is defined in Eq. (11) [see Eq. (230) of Ref. [46]]. We choose $t_{\text{PN}} = -100M$ and begin the PN integrations from the spins at t_{PN} . The extended spins are somewhat robust to the choice of t_{PN} as seen in Fig. 1. We stress that these extended spins are not physically meaningful for $t > t_{\text{PN}}$, but provide a convenient parametrization of the system that leads to accurate parametric fits.

IV. BUILDING THE MODEL

In this section, we describe the quantities that are computed from the waveform data pieces and stored when building the NRSur7dq2 surrogate model. The subsequent section then describes how the NRSur7dq2 surrogate model uses these stored quantities to generate waveforms.

We first construct surrogate models for the waveform modes in the co-orbital frame $h_{\text{coorb}}^{\ell,m}(t)$. For $m = 0$ modes, we directly model the real and imaginary components without any additional decompositions. For $m > 0$, we compute

$$h_{\pm}^{\ell,m} = \frac{1}{2}(h_{\text{coorb}}^{\ell,m} \pm h_{\text{coorb}}^{\ell,-m*}) \quad (17)$$

and model the real and imaginary parts of $h_{\pm}^{\ell,m}$. Each of these modeled components is considered a waveform data piece. We proceed according to Sec. V of Ref. [24]: For each waveform data piece, we construct a compact linear basis using singular value decomposition with a RMS tolerance of 3×10^{-4} . We then construct an empirical interpolant and determine one empirical node time T_j for each basis vector. The times T_j are chosen differently for each waveform data piece. Finally, for each T_j , we construct a *parametric fit* for the waveform data piece evaluated at T_j , which is described below. The fits are functions of the mass ratio and the co-orbital spin components $\vec{\chi}_i^{\text{coorb}}(t)$ evaluated at T_j . Note that the x component of a vector in the co-orbital frame is roughly the component in the direction of a vector pointing from one BH to the other, the z component is along the axis of orbital angular momentum, and the y component is the remaining orthogonal direction. In addition to the resulting fit data, the empirical interpolation matrix [see Eq. (B7) of Ref. [26]] for each of these waveform data pieces is stored in the NRSur7dq2 surrogate model.

These parametric fits use the forward-stepwise greedy fitting method described in Appendix A of Ref. [24]. One

benefit of this fitting method is that it automatically selects for higher order fits whenever the data are high quality (e.g., in the inspiral), and lower order fits whenever the data are more noisy (e.g., during the ringdown). We choose the basis functions to be a tensor product of one-dimensional monomials in the spin components and

$$x = \frac{q - 1.5}{0.51}, \quad (18)$$

which is an affine mapping from $[0.99, 2.01]$ to the standard interval $[-1, 1]$. We consider up to cubic functions in x and up to quadratic functions in the spin components. We perform 20 trials using 50 validation points each. The fit coefficients and the basis functions selected during the fitting procedure are stored in the NRSur7dq2 surrogate model.

We also construct parametric fits for $\omega(t)$, $\Omega_{\{x,y\}}^{\text{coorb}}(t)$, and $\dot{\chi}_{j\{x,y,z\}}^{\text{coorb}}(t)$ at selected time nodes t_i . These quantities describe the dynamics of the binary and the spins, so we call these t_i the dynamics time nodes. We attempt to choose the time nodes t_i to be approximately uniformly spaced in $\varphi(t)$ with ten nodes per orbit. Because $\varphi(t)$ is different for different simulations, and we choose the same time nodes for all simulations, in practice our choice of 238 time nodes gives us between 8 and 15 nodes per orbit. We find that this is sufficient—including additional nodes per orbit does not improve the accuracy of the surrogate model. Our time nodes are labeled $t_0 < t_1 < \dots < t_{234} = 100M$ plus three additional nodes $t_{\frac{1}{2}}$, $t_{\frac{3}{2}}$, and $t_{\frac{5}{2}}$, which are the midpoints of their adjacent integer time nodes. The reason for including the fractional time nodes is for Runge-Kutta time integration at the beginning of the time series, which is made clear in the next section. In Appendix B, we describe in detail the algorithm for choosing t_i , but any choice that is roughly uniformly spaced in $\varphi(t)$ and sufficiently dense should yield a surrogate with comparable accuracy.

V. EVALUATING THE MODEL

To evaluate the NRSur7dq2 surrogate model, we provide the mass ratio q and initial spins $\vec{\chi}_j(t_0)$ as inputs. The evaluation consists of three steps: we first integrate a coupled ODE system for the spins, the orbital phase, and the coprecessing frame, then we evaluate the co-orbital waveform modes, and finally we transform the waveform back to the inertial frame. We describe each of these steps below.

We initialize the ODE system with

$$\varphi(t_0) = 0, \quad \hat{q}(t_0) = 1, \quad \vec{\chi}_j^{\text{copr}}(t_0) = \vec{\chi}_j(t_0).$$

To integrate this system forward in time using a numerical ODE solver (described below), we need to evaluate the time derivatives of φ , \hat{q} , and $\vec{\chi}_j^{\text{copr}}$ at a time node t_i , given the values of those variables at t_i . To do this, we first determine

$\tilde{\chi}_j^{\text{coorb}}(t_i)$ by rotating the x and y components of $\tilde{\chi}_j^{\text{copr}}(t_i)$ by an angle $\varphi(t_i)$ as in Eq. (9). We then evaluate the fits for $\omega(t_i)$, $\Omega_{\{x,y\}}^{\text{coorb}}(t_i)$, and $\dot{\chi}_{j\{x,y,z\}}^{\text{coorb}}(t_i)$ using the mass ratio q and the current co-orbital spins $\tilde{\chi}_j^{\text{coorb}}(t_i)$. We set $\Omega_z^{\text{coorb}}(t_i) = 0$, and obtain $\tilde{\chi}_j^{\text{copr}}(t_i)$ and $\tilde{\Omega}^{\text{copr}}(t_i)$ by rotating the x and y components of the corresponding co-orbital quantities by an angle of $-\varphi(t_i)$. We evolve the coprecessing vectors instead of the co-orbital vectors because the former evolve on the longer precession time scale, allowing us to take large time steps. Finally, after computing

$$\left. \frac{d}{dt} \hat{q}(t) \right|_{t_i} = \frac{1}{2} \hat{q}(t_i) \tilde{\Omega}^{\text{copr}}(t_i), \quad (19)$$

we obtain the time derivatives of φ , \hat{q} , and $\tilde{\chi}_j^{\text{copr}}$ at $t = t_i$.

These time derivatives are then used to integrate φ , \hat{q} , and $\tilde{\chi}_j^{\text{copr}}$ using an ODE solver. We desire an ODE integration method that uses few evaluations of the time derivatives to keep the computational cost of evaluating the model low. We use a fourth-order Adams-Bashforth method [47,48] detailed in Appendix C, which determines the solutions at the next node based on the time derivatives at the current and three previous nodes. This allows us to reuse fit evaluations from the previous nodes, and requires only one additional evaluation of the fits per node compared to four evaluations for a fourth-order Runge-Kutta scheme. The Adams-Bashforth integration is initialized by performing the first integration steps with the fourth-order Runge-Kutta scheme. This is why we include the three additional time nodes $t_{\frac{1}{2}}$, $t_{\frac{3}{2}}$ and $t_{\frac{5}{2}}$; they enable evaluating the midpoint increments of the initial Runge-Kutta scheme. Once we have evaluated the solutions at the time nodes t_i , we use cubic spline interpolation to determine the solutions at all times.

Now that we have φ , \hat{q} , and $\tilde{\chi}_j^{\text{copr}}$ for all t , we evaluate each co-orbital waveform data piece. This is done by first evaluating the fits at the empirical nodes T_i using the mass ratio q and the co-orbital spins at the empirical nodes $\tilde{\chi}_j^{\text{coorb}}(T_i)$, and then evaluating the empirical interpolant to obtain the waveform data piece at all times. Finally, we transform the co-orbital frame waveform modes back to the coprecessing frame using $\varphi(t)$ and then to the inertial frame using $\hat{q}(t)$. The NRSur7dq2 surrogate data and Python evaluation code can be found at [49].

To reduce the computational cost of transforming the coprecessing waveform modes to the inertial frame using $\hat{q}(t)$, which takes ~ 1 s using all $\ell \leq 4$ modes sampled with $\delta t = 0.1M$, we reduce the number of time samples of the co-orbital waveform data pieces by using nonuniform time steps. We choose 2000 time samples that are roughly uniformly spaced in the orbital phase, using the same method used to choose the dynamics time nodes described in Appendix B. This is sufficiently many time samples to

yield negligible errors when interpolating back to the dense uniformly spaced time array using cubic splines on the real and imaginary parts of the waveform modes.

Integrating the ODE system takes ~ 3 ms, where the numerical computations are performed by a Python extension written in C. Interpolating the results of the ODE integration to the 2000 time samples described above takes ~ 2 ms using cubic splines. Evaluating the co-orbital waveform surrogate takes ~ 4 ms, and transforming the modes to the inertial frame takes ~ 16 ms, for a total of ~ 25 ms. Variations in the evaluation time can increase this up to ~ 30 ms. Restricting to only $\ell = 2$ modes can reduce this time to ~ 10 ms. If we wish to sample the surrogate waveform at the same time nodes as the original numerical relativity simulations, which is a uniformly spaced time array with $\delta t = 0.1M$, the modes are interpolated to these points using cubic splines. This requires ~ 6 ms per mode, for a total time of ~ 150 ms when all $\ell \leq 4$ modes are interpolated in this way. We note, however, that the original NR simulations are oversampled for typical GW data analysis purposes. For example, a sampling rate of 4096 Hz for a $M = 60 M_\odot$ binary has $\delta t \approx 0.83M$, leading to an evaluation time of ~ 50 ms. All timings were done on Intel Xeon E5-2680v3 cores running at 2.5 GHz.

VI. SURROGATE ERRORS

We use two error measures to quantify the accuracy of the surrogate model. Given two sets of waveform modes h_1 and h_2 , we first compute

$$\mathcal{E}[h_1, h_2] = \frac{1}{2} \frac{\sum_{\ell,m} \int_{t_0}^{t_f} |h_1^{\ell,m}(t) - h_2^{\ell,m}(t)|^2 dt}{\sum_{\ell,m} \int_{t_0}^{t_f} |h_1^{\ell,m}(t)|^2 dt}, \quad (20)$$

which is introduced in Eq. (21) of [24]. Since we have aligned all the NR waveforms at $t = t_0$ and the surrogate model reproduces this alignment, we do not perform any time or phase shifts when computing \mathcal{E} .

For these comparisons, we use modes $\ell \leq 5$; if a mode is not included in a particular waveform model, we assume this mode is 0 for that model. Since the NRSur7dq2 model does not contain $\ell = 5$ modes, this ensures that the errors discussed below include the effect of neglecting $\ell = 5$ and higher modes.

Histograms of \mathcal{E} for all 886 NR waveforms are given in Fig. 2. For all curves in the figure, h_1 is the highest available resolution NR waveform. For the thick solid black curve, h_2 is the same NR waveform as h_1 , except computed at a lower numerical resolution, so this curve represents an estimate of the numerical truncation error in the NR waveforms used to build the surrogate model. For the solid blue curve, h_2 is the NRSur7dq2 surrogate waveform evaluated with the same mass ratio and initial spins of h_1 . Note that since the surrogate was trained using all NR waveforms, this is an *in-sample* error.




FIG. 2. Error histograms for \mathcal{E} defined in Eq. (20), normalized such that the area under each curve is 1 when integrated over $\log_{10}(\mathcal{E})$. The largest surrogate errors are comparable to the largest NR resolution errors, which compare high and medium resolution NR simulations to estimate the error in the NR waveforms. The error in the orbital phase φ is the dominant error contribution to the surrogate.

The remaining curves in Fig. 2 indicate the in-sample error contribution from each of the three main waveform data pieces in the surrogate waveform: the orbital phase φ (dash-dotted green curve), the quaternions \hat{q} representing the precession (dashed orange curve), and the waveform modes in a co-orbital frame h_{coorb} (thin solid red curve). For these curves, h_2 is computed by using the surrogate evaluation for one waveform data piece and the NR evaluation of the other pieces. The orbital phase errors give rise to the largest surrogate errors, indicating that efforts to improve the surrogate model should be focused on improving the orbital phasing.

We then compute mismatches [50]

$$1 - \frac{\langle h_1, h_2 \rangle}{\sqrt{\langle h_1, h_1 \rangle \langle h_2, h_2 \rangle}}, \quad (21)$$

where $\langle \cdot, \cdot \rangle$ is a noise-weighted inner product computed in the frequency domain, as in Sec. VI.B of Ref. [24]. We use a flat power spectral density to avoid a dependence on the total mass of the system. The mismatches are minimized over time shifts, polarization angle shifts, and shifts in the azimuthal angle of the direction of GW propagation, where the system's orbital angular momentum is initially aligned with the \hat{z} axis. We randomly sample 30 directions of gravitational wave propagation on the sphere, and use a pair of detectors with idealized orientations such that one detector measures h_+ and the other detector measures h_\times . Histograms of the mismatches are given in Fig. 3 and are comparable to the top panel of Fig. 17 in Ref. [24]. To estimate the *out-of-sample* errors of the surrogate model, we perform a 20-fold cross-validation test. This is done by first randomly dividing the 886 NR waveforms into 20 sets of 44 or 45 waveforms. For each set, we build a trial surrogate using the waveforms from the other 19 sets. The trial surrogate is then evaluated at the parameters corresponding to the waveforms in the chosen validation set, and the results are compared to the NR waveform. These




FIG. 3. Mismatch histograms computed in the frequency domain with a flat power spectral density. The NR resolution mismatches compare waveforms from high and medium resolution NR simulations. This can be an overestimate of the error in the high resolution NR waveform, leading to some NR resolution mismatches being larger than the surrogate mismatches. We note that the IMRPhenomPv2 model does not contain all spin components.

cross-validation mismatches are given by the dashed purple curve. They are quite similar to the in-sample errors given by the solid blue curve, indicating that we are not overfitting the data. We also compute mismatches for a fully precessing effective one-body model (SEOBNRv3 [17]), and for a phenomenological waveform model that includes some, but not all, effects of precession (IMRPhenomPv2 [19]). These models have mismatches more than an order of magnitude larger than our NRSur7dq2 surrogate model. Both IMRPhenomPv2 and SEOBNRv3 depend on a parameter f_{ref} , which is a reference frequency at which the spin directions are specified. For SEOBNRv3, which is a time-domain model, we choose f_{ref} so that the waveform begins at $t = t_0$. For IMRPhenomPv2, which is a frequency-domain model, we minimize the mismatches over f_{ref} , using an initial guess of twice the orbital frequency of the NR waveform at $t = t_0$. While all of the mismatches can be decreased by minimizing over additional parameters such as BH masses and spins, this would result in biased parameters when measuring the source parameters of a detected GW signal.

We then compute mismatches using the advanced LIGO design sensitivity noise curve [51,52] using various total masses M . For each mass M , we obtain histograms as in Fig. 3, and we show the median and 95th percentile mismatches from these histograms in Fig. 4. We note that for $M \lesssim 114 M_\odot$ some or all waveforms begin above 10 Hz and do not cover the full design sensitivity frequency band. We find that the 95th percentile mismatches of our surrogate model are similar to the corresponding NR mismatches, except for total masses above $160 M_\odot$ where the NR mismatches are slightly smaller. The NRSur7dq2 surrogate yields mismatches at least an order of magnitude smaller than the other waveform models for all total masses investigated.

Figure 5 shows the real part of $h^{2,2}(t)$ for the cases leading to the largest mismatches in Fig. 3. The top panel shows the case leading to the largest surrogate




FIG. 4. Median (circles) and 95th percentile (triangles) mismatches of all 866 cases computed with the advanced LIGO design sensitivity curve. The surrogate mismatches are computed using trial surrogates, as in the cross-validation curve of Fig. 3.

cross-validation mismatch, and the bottom panel shows the case leading to the largest SEOBNRv3 mismatch. The surrogate waveforms shown are evaluated using the appropriate trial surrogate, so that they were not trained on the NR waveforms they are compared with. All waveforms are aligned to have their peak amplitude at $t = 0$ and are rotated to have their orbital angular momentum aligned with the z axis at $t = t_0 = -4500M$. In the top panel, we see that both the SEOBNRv3 and surrogate waveforms have a similar phasing error around $t = -50M$. The phasing error of the surrogate does not grow significantly larger through merger and ringdown, so most of this error can be removed with a time and phase shift. For the SEOBNRv3 waveforms in both the upper and lower panels, the phasing error changes significantly during the merger; therefore, this error does not decrease significantly even after performing a time and phase shift. In the top panel of Fig. 5, the IMRPhenomPv2 waveform does as well as the surrogate; in the bottom panel, the IMRPhenomPv2 waveform has large errors in both phase and amplitude.




FIG. 5. The real part of time-domain waveforms for the case leading to the largest surrogate mismatch (top) and the largest SEOBNRv3 mismatch (bottom). The surrogate waveforms are evaluated using trial surrogates which were not trained with the NR waveform shown. The top panel uses SXS:BBH:0922 with $q \approx 2$, $\vec{\chi}_1(t_0) \approx 0.8\hat{z}$, and $\vec{\chi}_2(t_0) \approx -0.8\hat{y}$. The lower panel uses SXS:BBH:0900 with $q \approx 2$, $\vec{\chi}_1(t_0) \approx (0.29, -0.74, 0.02)$ and $\vec{\chi}_2(t_0) \approx (0.43, -0.34, 0.58)$.

VII. DISCUSSION AND CONCLUSIONS

Within its range of validity, our NRSur7dq2 surrogate model is nearly as accurate as performing new NR simulations. The surrogate model takes only ~ 50 ms to evaluate on a single CPU core, making it sufficiently fast for current GW data analysis applications such as parameter estimation. This evaluation time can be compared to $\mathcal{O}(\text{weeks})$ on dozens of CPU cores to perform a new NR simulation, decreasing the cost in CPU hours by $\mathcal{O}(10^8)$. The NRSur7dq2 surrogate model data along with Python evaluation code is publicly available for download at [49].

Our surrogate model is limited to mass ratios $q \leq 2$ and spin magnitudes $|\vec{\chi}_{1,2}| \leq 0.8$. While in principle the parametric fits can be extrapolated to more extreme mass ratios and spin magnitudes, we do not expect extrapolation to yield accurate waveforms. However, these limits can be extended in future versions of our surrogate model by performing NR simulations with larger mass ratios and spins.

Additionally, the waveforms produced by NRSur7dq2 are limited in duration to $4500M$ before the peak amplitude. This covers frequencies $f \geq 20$ Hz for all systems with $M \gtrsim 57 M_\odot$. For systems with lower total masses, or for systems with $M \lesssim 114 M_\odot$ when including frequencies down to 10 Hz, longer waveforms are needed. In future work, we plan to overcome this limitation by hybridizing with either PN or SEOBNRv3 [53–56], either by hybridizing the NR waveforms before building the surrogate or by hybridizing the surrogate waveforms. Longer NR waveforms would then be needed to test the accuracy of the hybridization step.

ACKNOWLEDGMENTS

We thank Matt Giesler for helping to carry out the new SpEC simulations used in this work. We thank Saul Teukolsky, Patricia Schmidt, Rory Smith, and Vijay Varma for helpful discussions. This work was supported in part by the Sherman Fairchild Foundation and by NSF Grants No. CAREER PHY-1151197, No. PHY-1404569, No. AST-1333129, and No. PHY-1606654. Computations were performed on NSF/NCSA Blue Waters under allocation PRAC ACI-1440083; on the NSF XSEDE network under Grant No. TG-PHY100033; and on the Zwicky cluster at Caltech, which is supported by the Sherman Fairchild Foundation and by NSF Grant No. PHY-0960291. This paper has been assigned YITP report number YITP-17-44.

APPENDIX A: SPARSE GRID PARAMETERS

We take the polar and azimuthal spin angles of the inertial frame spins $\vec{\chi}_i$ to be θ_i and ϕ_i , respectively, for $i \in \{1, 2\}$. We can then parametrize our seven-dimensional parameter space by

- (i) $q \in [1, 2]$,
- (ii) $|\vec{\chi}_i| \in [0, 0.8]$,

- (iii) $\theta_i \in [0, \pi]$,
- (iv) $\phi_i \in [0, 2\pi]$.

The range of each of these variables is some closed interval $[a, b]$. For a variable x with range $[a, b]$, we define a grid of N uniformly spaced points

$$g_x^N = \left\{ a + \frac{n}{N-1}(b-a) : n = 0, \dots, N-1 \right\}, \quad (\text{A1})$$

where $N \geq 2$. We then define a sequence of grids

$$G_x \equiv G_x^0, G_x^1, \dots, \quad (\text{A2})$$

where

$$G_x^n = g_x^{f_x(n)} \quad (\text{A3})$$

for some monotonically increasing function $f_x(n)$. We call G_x^n the level n grid for x . We take

$$f_q(n) = f_{|\vec{\chi}_1|}(n) = 1 + 2^n, \quad (\text{A4})$$

$$f_{\theta_i}(n) = 1 + 2^{n+1}, \quad (\text{A5})$$

$$f_{\phi_i}(n) = 1 + 3 \cdot 2^n. \quad (\text{A6})$$

These choices ensure that $G_x^n \subset G_x^{n+1}$, and that the level 0 grids already give a description of the parameter space that does not leave out any phenomenology; the level 0 grids for θ_i contain the midpoint $\pi/2$ leading to precession, and the level 0 grids for ϕ_i contain three unique points (since $\phi_i = 0$ and $\phi_i = 2\pi$ lead to the same physical spin) in order to get at least some resolution of features that behave like $\sin(\phi_i + \phi + 0)$.

We have already seen that $\phi_i = 0$ and $\phi_i = 2\pi$ correspond to the same physical spin, but we have many other scenarios where two combinations of variables lead to the same physical configuration. For example, if $|\vec{\chi}_1| = 0$, all combinations of θ_1 and ϕ_1 lead to the same physical configuration. We ignore these degenerate combinations for now, and remove them later on.

Dense grids in parameter space could be constructed as

$$G_{\text{dense}}^n = G_q^n \times G_{|\vec{\chi}_1|}^n \times G_{\theta_1}^n \times G_{\phi_1}^n \times G_{|\vec{\chi}_2|}^n \times G_{\theta_2}^n \times G_{\phi_2}^n,$$

where \times denotes the Cartesian product. While the one-dimensional grids grow in size as $\mathcal{O}(2^n)$, these dense grids grow in size as $\mathcal{O}(2^{7n})$ or as the seventh power of the size of the one-dimensional grids. This is known as the curse of dimensionality; the amount of data needed often grows exponentially with the dimensionality. Sparse grids [42,43] overcome the curse of dimensionality by using a sparse product such that the grids grow in size as $\mathcal{O}(2^n(\log 2^n)^6)$. If G_x and G_y are two sequences of grids, we define the sparse product of G_x with G_y to be $G_{x,y} = G_x \bullet G_y$, where

$$G_{x,y}^n = \bigcup_{k=0}^n G_x^k \times G_y^{n-k}. \quad (\text{A7})$$

We now define the sparse grids for our parameter space from the sequence of grids

$$G = G_q \bullet G_{|\vec{\chi}_1|} \bullet G_{\theta_1} \bullet G_{\phi_1} \bullet G_{|\vec{\chi}_2|} \bullet G_{\theta_2} \bullet G_{\phi_2} \quad (\text{A8})$$

such that

$$G^n = \bigcup_{\sum_{i=1}^7 k_i = n} G_q^{k_1} \times G_{|\vec{\chi}_1|}^{k_2} \times G_{\theta_1}^{k_3} \times G_{\phi_1}^{k_4} \times G_{|\vec{\chi}_2|}^{k_5} \times G_{\theta_2}^{k_6} \times G_{\phi_2}^{k_7}.$$

Starting with the parameters in G^1 we removed physically identical configurations. We also removed configurations with $|\vec{\chi}_2| \propto \hat{z}$, which are within the parameter space of the NRSur4d2s surrogate model, which was already covered by the 276 NRSur4d2s NR simulations. We performed 361 new NR simulations based on the remaining set of parameter values.

APPENDIX B: TIME SAMPLING

We wish to choose time nodes $t_0 < t_1 < \dots < t_f$ that are roughly uniformly spaced in the orbital phase $\varphi(t)$ for all cases. Given some number N , we choose time nodes yielding roughly N nodes per orbit. Since different NR waveforms have different orbital frequencies, they have a different number of time nodes per orbit. Our scheme for choosing the time nodes given N is based on the leading order PN expression for the orbital angular frequency $\omega(t)$ during the inspiral, smoothly transitioning to a maximum value of $\omega = 2\pi/(20M)$ during the ringdown. We do this by computing a bounded time

$$\tilde{t}(t) = -1.7 + \frac{1}{2} \left((t+5) - \sqrt{(t+5)^2 + 25} \right), \quad (\text{B1})$$

and then choosing

$$\omega_{\text{ref}}(t) = \omega_{0\text{PN}}(\tilde{t}(t)) = \left(\frac{64}{5} (-\tilde{t}(t)) \right)^{-\frac{3}{8}}. \quad (\text{B2})$$

We then use spacings between nodes $t_{j+1} - t_j = \omega_{\text{ref}}(t_j)$.

APPENDIX C: FOURTH-ORDER ADAMS-BASHFORTH METHOD

We integrate the system of ordinary differential equations on a nonuniformly spaced grid of time nodes $t_0 < t_1 < \dots < t_f$ using a fourth-order Adams-Bashforth scheme [47,48]. We denote the solution $\vec{y}(t)$, and at each time node t_i , we can evaluate fits to determine

$$\frac{d\vec{y}}{dt} = \vec{f}(t; \vec{y}). \quad (\text{C1})$$

We first integrate up to t_3 using a Runge-Kutta fourth-order scheme.

Once we have integrated up to t_i for $i \geq 3$, we have previously evaluated

$$\vec{k}_j = \vec{f}(t_j; \vec{y}(t_j)) \quad (\text{C2})$$

for $0 \leq j < i$, and we now evaluate \vec{k}_i . We approximate $\vec{g}(t) = \vec{f}(t; \vec{y}(t))$ by a cubic function

$$\vec{g}(t) \approx \vec{g}_3(t) = \vec{A} + \vec{B}(t - t_i) + \vec{C}(t - t_i)^2 + \vec{D}(t - t_i)^3. \quad (\text{C3})$$

The coefficients are chosen such that $\vec{g}_3(t_j) = \vec{g}(t_j) = \vec{k}_j$ for $i - 3 \leq j \leq i$, giving $\vec{A} = \vec{k}_i$, and

$$\begin{bmatrix} \vec{B} \\ \vec{C} \\ \vec{D} \end{bmatrix} = \begin{bmatrix} \frac{\delta_{-1,0}\delta_{-2,0}}{\Delta_1} & \frac{\delta_{-1,0}\delta_{-3,0}}{\Delta_2} & \frac{\delta_{-2,0}\delta_{-3,0}}{\Delta_3} \\ \frac{\delta_{-2,0} + \delta_{-1,0}}{\Delta_1} & \frac{\delta_{-1,0} + \delta_{-3,0}}{\Delta_2} & \frac{\delta_{-2,0} + \delta_{-3,0}}{\Delta_3} \\ \frac{1}{\Delta_1} & \frac{1}{\Delta_2} & \frac{1}{\Delta_3} \end{bmatrix} \begin{bmatrix} \vec{k}_i - \vec{k}_{i-3} \\ \vec{k}_i - \vec{k}_{i-2} \\ \vec{k}_i - \vec{k}_{i-1} \end{bmatrix}.$$

Here, $\delta_{n,m} = t_{i+m} - t_{i+n}$ and

$$\Delta_1 = \delta_{-3,-2}\delta_{-3,-1}\delta_{-3,0}, \quad (\text{C4})$$

$$\Delta_2 = \delta_{-3,-2}\delta_{-2,-1}\delta_{-2,0}, \quad (\text{C5})$$

$$\Delta_3 = \delta_{-2,-1}\delta_{-3,-1}\delta_{-1,0}. \quad (\text{C6})$$

Finally, we approximate

$$\begin{aligned} \vec{y}(t_{i+1}) &= \vec{y}(t_i) + \int_{t_i}^{t_{i+1}} g(t) dt \\ &\approx \vec{y}(t_i) + \int_{t_i}^{t_{i+1}} g_3(t) dt \\ &= \vec{y}(t_i) + \delta_{0,1}\vec{A} + \frac{1}{2}\delta_{0,1}^2\vec{B} + \frac{1}{3}\delta_{0,1}^3\vec{C} + \frac{1}{4}\delta_{0,1}^4\vec{D}. \end{aligned}$$

-
- [1] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), *Phys. Rev. Lett.* **116**, 061102 (2016).
- [2] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), *Phys. Rev. Lett.* **116**, 241103 (2016).
- [3] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), *Phys. Rev. Lett.* **116**, 241102 (2016).
- [4] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), *Phys. Rev. D* **94**, 064035 (2016).
- [5] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), *Phys. Rev. Lett.* **116**, 221101 (2016).
- [6] F. Pretorius, *Phys. Rev. Lett.* **95**, 121101 (2005).
- [7] Y. Zlochower, J. Baker, M. Campanelli, and C. Lousto, *Phys. Rev. D* **72**, 024021 (2005).
- [8] Spectral Einstein Code, <http://www.black-holes.org/SpEC.html>.
- [9] Einstein Toolkit home page, <http://einsteintoolkit.org>.
- [10] S. Husa, J. A. González, M. Hannam, B. Brügmann, and U. Sperhake, *Classical Quantum Gravity* **25**, 105006 (2008).
- [11] B. Brügmann, J. A. González, M. Hannam, S. Husa, U. Sperhake, and W. Tichy, *Phys. Rev. D* **77**, 024027 (2008).
- [12] F. Herrmann, I. Hinder, D. Shoemaker, and P. Laguna, *Classical Quantum Gravity* **24**, S33 (2007).
- [13] T. Damour, P. Jaranowski, and G. Schaefer, *Phys. Rev. D* **78**, 024009 (2008).
- [14] T. Damour and A. Nagar, *Phys. Rev. D* **79**, 081503 (2009).
- [15] A. Taracchini, A. Buonanno, Y. Pan, T. Hinderer, M. Boyle, D. A. Hemberger, L. E. Kidder, G. Lovelace, A. H. Mroue, H. P. Pfeiffer, M. A. Scheel, B. Szilágyi, N. W. Taylor, and A. Zenginoglu, *Phys. Rev. D* **89**, 061502 (2014).
- [16] M. Pürrer, *Phys. Rev. D* **93**, 064041 (2016).
- [17] Y. Pan, A. Buonanno, A. Taracchini, L. E. Kidder, A. H. Mroué, H. P. Pfeiffer, M. A. Scheel, and B. Szilágyi, *Phys. Rev. D* **89**, 084006 (2014).
- [18] A. Bohé, L. Shao, A. Taracchini, A. Buonanno, S. Babak, I. W. Harry, I. Hinder, S. Ossokine, M. Pürrer, V. Raymond, T. Chu, H. Fong, P. Kumar, H. P. Pfeiffer, M. Boyle, D. A. Hemberger, L. E. Kidder, G. Lovelace, M. A. Scheel, and B. Szilágyi, *Phys. Rev. D* **95**, 044028 (2017).
- [19] M. Hannam, P. Schmidt, A. Bohé, L. Haegel, S. Husa, F. Ohme, G. Pratten, and M. Pürrer, *Phys. Rev. Lett.* **113**, 151101 (2014).
- [20] S. Khan, S. Husa, M. Hannam, F. Ohme, M. Prer, X. Jiménez Forteza, and A. Bohé, *Phys. Rev. D* **93**, 044007 (2016).
- [21] S. Husa, S. Khan, M. Hannam, M. Prer, F. Ohme, X. Jiménez Forteza, and A. Bohé, *Phys. Rev. D* **93**, 044006 (2016).
- [22] L. Lindblom, B. J. Owen, and D. A. Brown, *Phys. Rev. D* **78**, 124020 (2008).
- [23] B. P. Abbott *et al.* (Virgo, LIGO Scientific), *Classical Quantum Gravity* **34**, 104002 (2017).
- [24] J. Blackman, S. E. Field, M. A. Scheel, C. R. Galley, D. A. Hemberger, P. Schmidt, and R. Smith, *Phys. Rev. D* **95**, 104023 (2017).

- [25] J. Blackman, S. E. Field, C. R. Galley, B. Szilágyi, M. A. Scheel, M. Tiglio, and D. A. Hemberger, *Phys. Rev. Lett.* **115**, 121102 (2015).
- [26] S. E. Field, C. R. Galley, J. S. Hesthaven, J. Kaye, and M. Tiglio, *Phys. Rev. X* **4**, 031006 (2014).
- [27] M. Pürrer, *Classical Quantum Gravity* **31**, 195010 (2014).
- [28] H. P. Pfeiffer, L. E. Kidder, M. A. Scheel, and S. A. Teukolsky, *Comput. Phys. Commun.* **152**, 253 (2003).
- [29] G. Lovelace, R. Owen, H. P. Pfeiffer, and T. Chu, *Phys. Rev. D* **78**, 084017 (2008).
- [30] L. Lindblom, M. A. Scheel, L. E. Kidder, R. Owen, and O. Rinne, *Classical Quantum Gravity* **23**, S447 (2006).
- [31] B. Szilágyi, L. Lindblom, and M. A. Scheel, *Phys. Rev. D* **80**, 124010 (2009).
- [32] M. A. Scheel, M. Boyle, T. Chu, L. E. Kidder, K. D. Matthews, and H. P. Pfeiffer, *Phys. Rev. D* **79**, 024003 (2009).
- [33] B. Szilágyi, *Int. J. Mod. Phys. D* **23**, 1430014 (2014).
- [34] M. Boyle and A. H. Mroué, *Phys. Rev. D* **80**, 124045 (2009).
- [35] M. Boyle, *Phys. Rev. D* **93**, 084031 (2016).
- [36] M. Boyle, *Phys. Rev. D* **87**, 104006 (2013).
- [37] M. Boyle, L. E. Kidder, S. Ossokine, and H. P. Pfeiffer, arXiv:1409.4431.
- [38] Scri, <https://github.com/moble/scri>.
- [39] S. Ossokine, M. Boyle, L. E. Kidder, H. P. Pfeiffer, M. A. Scheel, and B. Szilágyi, *Phys. Rev. D* **92**, 104028 (2015).
- [40] M. Boyle, R. Owen, and H. P. Pfeiffer, *Phys. Rev. D* **84**, 124011 (2011).
- [41] B. Aylott, J. G. Baker, W. D. Boggs, M. Boyle, P. R. Brady *et al.*, *Classical Quantum Gravity* **26**, 165008 (2009).
- [42] S. A. Smolyak, in *Doklady Akademii Nauk SSSR* (1963), Vol. 4, pp. 240–243.
- [43] H.-J. Bungartz and M. Griebel, *Acta Numer.* **13**, 147 (2004).
- [44] P. Schmidt, M. Hannam, S. Husa, and P. Ajith, *Phys. Rev. D* **84**, 024046 (2011).
- [45] R. O’Shaughnessy, B. Vaishnav, J. Healy, Z. Meeks, and D. Shoemaker, *Phys. Rev. D* **84**, 124002 (2011).
- [46] L. Blanchet, *Living Rev. Relativ.* **17**, 2 (2014).
- [47] J. C. Butcher, *Numerical Methods for Ordinary Differential Equations* (Wiley, New York, 2003).
- [48] F. Bashforth and J. C. Adams, *An Attempt to Test the Theories of Capillary Action by Comparing the Theoretical and Measured Forms of Drops of Fluid* (Cambridge University Press, Cambridge, 1883).
- [49] SXS Surrogates Data Server, <http://www.black-holes.org/surrogates/>.
- [50] S. Babak, A. Taracchini, and A. Buonanno, *Phys. Rev. D* **95**, 024010 (2017).
- [51] D. Shoemaker (LIGO Collaboration), LIGO Report No. T0900288-v3, 2010.
- [52] J. Aasi *et al.* (LIGO Scientific Collaboration), *Classical Quantum Gravity* **32**, 115012 (2015).
- [53] J. C. Bustillo, A. Boh, S. Husa, A. M. Sintes, M. Hannam *et al.*, arXiv:1501.00918.
- [54] M. Boyle, *Phys. Rev. D* **84**, 064013 (2011).
- [55] F. Ohme, M. Hannam, and S. Husa, *Phys. Rev. D* **84**, 064029 (2011).
- [56] I. MacDonald, S. Nissanke, and H. P. Pfeiffer, *Classical Quantum Gravity* **28**, 134002 (2011).