



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Uncertainty in hybrid gravitational waveforms: Optimizing initial orbital frequencies for binary black-hole simulations

Michael Boyle

Phys. Rev. D **84**, 064013 — Published 12 September 2011

DOI: [10.1103/PhysRevD.84.064013](https://doi.org/10.1103/PhysRevD.84.064013)

The uncertainty in hybrid gravitational waveforms: Optimizing initial orbital frequencies for binary black-hole simulations

Michael Boyle

Center for Radiophysics and Space Research, Cornell University, Ithaca, New York 14853, USA

A general method is presented for estimating the uncertainty in hybrid models of gravitational waveforms from binary black-hole systems with arbitrary physical parameters, and thence the highest allowable initial orbital frequency for a numerical-relativity simulation such that the combined analytical and numerical waveform meets some minimum desired accuracy. The key strength of this estimate is that no prior numerical simulation in the relevant region of parameter space is needed, which means that these techniques can be used to direct future work. The method is demonstrated for a selection of extreme physical parameters. It is shown that optimal initial orbital frequencies depend roughly linearly on the mass of the binary, and therefore useful accuracy criteria must depend explicitly on the mass. The results indicate that accurate estimation of the parameters of stellar-mass black-hole binaries in Advanced LIGO data or calibration of waveforms for detection will require much longer numerical simulations than are currently available or more accurate post-Newtonian approximations—or both—especially for comparable-mass systems with high spin.

Numerical relativists face a thorny dilemma when creating initial data for simulations of binary black holes. Two competing motivations vie for control of one key parameter: the initial orbital frequency of the binary, Ω_0 . On one hand, the larger the initial frequency is, the more quickly the simulation will run. In fact, at lowest order, doubling Ω_0 will shorten a simulation by a factor of six. On the other hand, the smaller the initial frequency is, the clearer the correspondence will be between the numerical simulation and the system found in nature. Post-Newtonian approximations will be more accurate; the velocity and spin of each black hole will be more well defined and easily measured; even the junk radiation will be smaller [1]. In practice, numerical relativists choose Ω_0 largely by intuition, with primary considerations being the available computer time and the roundness of the number. This lack of precision can lead to simulations that are too short and should ideally be redone, or are longer than necessary—a waste of resources in either case. More objective choices are possible, and will be needed to improve the effectiveness of numerical relativity. This paper demonstrates a technique¹ to estimate the optimal value of Ω_0 , even for systems in unexplored regions of parameter space.

The field of numerical relativity (NR) exists because of the failure of post-Newtonian models (PN); at some point NR must take over from PN. Of course, PN approximations don't simply break down at one catastrophic instant, having been perfectly accurate before. Rather, the approximations

gradually deteriorate as they approach merger. The question of exactly where NR needs to replace PN is thus a question of how accurate the model needs to be. In the context of designing model waveforms for detection, we are given precise objectives. This allows us to *quantitatively* resolve the conflict between decreasing the length of a simulation and improving the quality of the final modeled waveform. The quality of the final waveform is impacted by the accuracy of both PN and NR data. Because we already have the PN data that will be used in the final waveform, we can test how much of it can be used if we are to achieve a target accuracy. Where that ends is where the NR simulation must begin.

Estimating the impact of PN errors depends on understanding how the data will be used in the finished product. In this context, that means understanding waveforms used in data analysis for gravitational-wave detectors. Advanced detectors of the near future will be more sensitive over broader ranges of frequencies than current detectors, which tightens the requirements for accurate modeling of physical waveforms [7, 8]. In particular, analysis of detector data will require waveforms that are not only more precisely coherent, but also coherent over a greater range of frequencies. PN waveforms alone are expected to be sufficient *for detection* up to a total system mass somewhere² around $12 M_\odot$ [9], and are expected to be entirely inadequate for parameter estimation. Exclusively numerical waveforms, on the other hand, are limited in their usefulness to high-mass systems; current numerical simulations can only cover the Advanced LIGO band down to 10 Hz for masses greater

¹ The technique expands on one introduced in Ref. [2]—which was partially implemented in Ref. [3]—to apply to the time-domain methods currently in use by most numerical relativists, to use the most accurate inspiral models available, and to include more general accuracy requirements. The results obtained here broadly agree with the results of Refs. [4–6], which test whether completed numerical simulations are long enough.

² The exact mass delimiting the range of validity depends, of course, on the parameters of the particular system and the detector in question. Moreover, this says nothing about parameter estimation. The methods of this paper are primarily applicable to parameter estimation, but do have some relevance to detection.

than about $100 M_{\odot}$ [10]. This leaves a large gap in an astrophysically interesting range [11]. To close the gap using NR alone would require simulations roughly 250 times longer than are currently available—or more for parameter estimation. Notwithstanding possible dramatic improvements to the efficiency of long simulations [12–14], this will be impractical for some time to come, especially for large surveys of parameter space [15].

Instead of running such dramatically long simulations, we “hybridize”—synthesizing a single coherent waveform by combining the long PN inspiral with the short NR merger and ringdown. In the interest of simplicity of presentation, let us assume for the moment that a hybrid waveform is constructed by aligning the PN and NR waveforms at a single point in time—using PN data before that point and NR data after—and that the NR portion is essentially perfect. Now, as the alignment point moves closer to merger, the hybrid waveform will become less accurate since it includes more of the deteriorating PN waveform, as demonstrated in Fig. 1. In fact, because the hybridization procedure uses only the worst data when aligning at very late times, errors in the PN waveform just prior to the alignment point will unavoidably taint the data, meaning that the long inspiral and the merger will be dephased relative to the physically correct waveform. The error in the complete hybrid, then, depends crucially on the error in the PN waveform and particularly on the growth in the error at late times.

We cannot know the *error* in a PN waveform—being the difference from the unknown correct waveform it attempts to model—without unreasonably long numerical simulations. We can, however, estimate our *uncertainty* by creating a range of equally plausible PN waveforms. We can then attach these to some ersatz NR waveform,³ such as effective one-body (EOB) or phenomenological waveforms [16–22], to form a range of plausible hybrids. We can compare that range quantitatively to the given error budget. By repeating this process using each of many possible hybridization frequencies, we can discover which hybridization frequencies produce waveforms that satisfy the error budget. The highest such frequency minimizes the length of the simulation, and is thus the optimal value.

To summarize, the work to be done in this scheme consists of three basic steps:

1. Choose an appropriate ersatz NR waveform and construct a range of plausible inspiral waveforms (Sec. I);
2. Hybridize the inspiral and ersatz NR waveforms, evaluate the mismatches among them, and repeat for various hybridization frequencies (Sec. II);
3. Choose Ω_0 to agree with the highest hybridization frequency that satisfies the error requirements (Sec. III).

³ As shown later in this paper, the final results will not depend strongly on the particular choice of ersatz NR waveform.




FIG. 1. **Error caused by aligning at late times.** This plot shows the phase error in hybrid waveforms created by aligning an ersatz NR waveform (EOB) to a PN model at various times, relative to the hybrid when aligned very far in the past. As the alignment point is moved closer to merger ($t/M = 0$), the total phase error increases because the hybrid waveform incorporates more and more of the inaccurate PN data. Note that, in each case, the error grows most rapidly near merger. If the merger occurs at a frequency to which the detector is sensitive, this phase error will negatively impact the match found by data analysts (see Sec. I A). Evaluating that impact using artificial NR data is the essence of the method presented in this paper.

This, of course, assumes that there is some freedom in choosing the initial orbital frequency, as will be the case for NR groups setting out to run a simulation; there is always a choice between running many short simulations and fewer long simulations—or requesting a larger allocation of computer time. This method is designed to help in making that choice. On the other hand, given a completed waveform, steps 1 and 2 above can be used to evaluate the uncertainty in the resulting hybrids. Finally, understanding the results can aid in the design of reasonable and effective accuracy goals, before any simulation is undertaken.

The method will be demonstrated for a few interesting cases, probing the “corners” of a simple parameter space: equal-mass nonspinning, equal-mass high-spin, large mass-ratio nonspinning, and large mass-ratio high-spin systems. The key result of this paper is the plot of the uncertainty of the hybrids for those systems, Fig. 4, discussed in Sec. III C. Finally, Sec. IV summarizes the conclusions and outlines possible applications and extensions to this method.

The general method presented here can be applied quite broadly. However, to demonstrate the method, we need to

make several specific choices. Details of the implementation are given in Appendix B. In particular, this paper uses a simplified EOB model, described in Appendix B 2, to supply the ersatz NR waveforms. In Appendix A, the results for the equal-mass nonspinning case are redone using actual numerical data, to check—to the extent possible—that the final results are not sensitive to this choice for the ersatz NR waveform. The uncertainty will be measured by a range of plausible waveforms formed by hybridizing the EOB merger and ringdown with the inspiral portion of EOB and TaylorT1–T4 approximants [23–25] using all known information (full PN orders), as recently recalculated and summarized by members of the NINJA-2 collaboration [26]. Specifically, the amplitude includes terms up to 3.0 PN order, and the phase includes terms up to 3.5 PN order. See Appendix B 1 for more detail. The hybridization will be done in the time domain by aligning at particular frequencies [25, 27], then blending the PN and NR waveforms as described in Ref. [10]. The resulting hybrids will be compared along the positive z axis, computing the match [see Eq. (3)] using the Advanced LIGO high-power noise curve with no detuning [8], scaling the total system mass between $5 M_\odot$ and $50 M_\odot$.

Throughout this paper, the uppercase Greek letters Φ and Ω refer to the orbit of a binary, contrasting with the lowercase Greek letters ϕ and ω which refer to the phase and frequency of the emitted gravitational waves. Unless otherwise specified, ϕ and ω refer to the $(\ell, m) = (2, 2)$ mode in a spin $s = -2$ spherical harmonic decomposition of the gravitational wave.

I. CREATING THE WAVEFORMS

The first task before us is to construct a large group of model waveforms to be compared to each other. In this section, a review of the standard error measure used in data analysis motivates the use of complete inspiral-merger-ringdown waveforms for our tests, while finding encouraging signs that high accuracy of these waveforms is not essential. We then examine in greater detail the construction of a credibly broad selection of PN waveforms.

A. Motivation

In constructing a range of plausible model waveforms, we need to understand the ultimate form of measurement when designing templates for gravitational-wave detection: the match. This quantity is based on the inner product between two waveforms defined as the integral of the noise-weighted product of the signals in the frequency domain [28]:

$$\langle h_a | h_b \rangle = 2 \Re \int_{-\infty}^{\infty} \frac{\tilde{h}_a(f) \tilde{h}_b^*(f)}{S_n(|f|)} df \quad (1a)$$

$$= 4 \int_0^{\infty} \frac{|\tilde{h}_a(f)| |\tilde{h}_b^*(f)|}{S_n(|f|)} [\cos \delta\phi(f)] df, \quad (1b)$$

where $S_n(f)$ is the power spectral density (PSD) of noise in the detector, and $\delta\phi(f)$ is the phase difference between the two waveforms in the frequency domain. In the following, we will assume $f > 0$. The second form shown here demonstrates a useful way to understand the inner product, by separating the integrand into two factors. We begin with the first factor, the ratio of amplitudes to noise. The amplitude for an example system and the Advanced LIGO noise curve are plotted in Fig. 2. The height of the amplitude above the noise curve shows when the inner product can rapidly accumulate; more height means more rapid contribution to the inner product. Of course, rapid contributions cannot result in a large inner product unless those contributions are also coherent. This requires the second factor of the integrand in Eq. (1b) to be constant. In particular, a large inner product requires $\delta\phi$ to be very close to zero across the full range of frequencies for which the signal amplitude is significantly larger than the noise.

While the amplitude of the waveform from a given system is fixed, $\delta\phi$ has two inherent degrees of freedom we can adjust to improve the inner product. These degrees of freedom are related to the fact that the merger time and orientation of an astrophysical binary are unknowns that must simply be measured. Specifically, we are free to shift the time and phase of either time-domain waveform by ΔT and $\Delta\Phi$. In the frequency domain, this has no effect on the amplitude of the waveform (the curves of Fig. 2 will not be affected), but $\delta\phi$ changes roughly as

$$\delta\phi(f) \rightarrow \delta\phi(f) + 2 \Delta\Phi + 2 \pi f \Delta T, \quad (2)$$

We can use this time- and phase-shift freedom to ensure that the phase difference between two waveforms is smallest at frequencies for which the detector is most sensitive to that particular system, and thus maximize the inner product. The maximum possible (normalized) inner product [28] is called the match, which will be our basic measure of error:

$$\langle h_a | h_b \rangle = \max_{\Delta T, \Delta\Phi} \frac{\langle h_a | h_b \rangle}{\sqrt{\langle h_a | h_a \rangle \langle h_b | h_b \rangle}}. \quad (3)$$

This quantity takes a value between 0 (for completely dissimilar waveforms) and 1 (for identical waveforms). Because many of the matches encountered below will be very close to 1, it is preferable to use another quantity called the mismatch [29–31], which is given by

$$\text{MM}(h_a, h_b) := 1 - \langle h_a | h_b \rangle. \quad (4)$$

Here, values close to 0 indicate the waveforms are similar. The maximum possible signal-to-noise ratio (SNR) at which a given signal s can be detected is given by

$$\rho_s := \sqrt{\langle s | s \rangle}. \quad (5)$$

The mismatch $\text{MM}(s, h)$ between a signal s and a template h is essentially the percentage loss in SNR due to errors in the template [32].




FIG. 2. **Waveforms in the frequency domain.** Amplitudes from an equal-mass nonspinning binary are shown, scaled to total system masses of $10 M_{\odot}$ and $100 M_{\odot}$ at 100 Mpc, and compared to the noise spectral density of Advanced LIGO (both quantities in units of strain \cdot Hz $^{-1/2}$). The factors multiplying the amplitudes are chosen to account for the logarithmic scaling of the horizontal axis, the factor of 2 in Eq. (1), and the fact that only positive frequencies are plotted. The triangles on the waveforms show the approximate initial frequency of the longest numerical simulation currently available. The circles show the frequency of the innermost stable circular orbit (ISCO)—basically the frequency at which PN approximations are expected to be useless. This plot shows that, for the $10 M_{\odot}$ system, nearly half the contribution to the inner product of Eq. (1) comes from the PN data (to the left of the triangle), and the rest from the NR data (to the right). For the $100 M_{\odot}$ system, on the other hand, the inner product is given almost exclusively by NR data.

This understanding of the match teaches us two very important lessons that will guide our approach to our problem. First, the optimal phase shift—and thus the value of the match—will depend on the relative distribution of power at different frequency bands. If, for example, our model waveform simply ends at ISCO, it will fail to model a substantial portion of the physical waveform. In that case, the maximization of Eq. (3) will not need to balance the dephasing between the two portions of the waveform, for example. The requirements for phase coherence will be much looser than they, in fact, need to be. Thus, our model waveforms must have roughly the same distribution of amplitude as the physical waveform, across the entire sensitive frequency band of the detector. If our objective is to evaluate matches before any numerical simulation is done, we will need a suitable approximation to the merger/ringdown waveform. For that purpose, this paper

uses the EOB waveform [16–19], which extends through merger to ringdown. Other complete waveforms could also be used [20–22]. Fortunately, we will see that the final results will not depend strongly on the choice of ersatz NR waveform. For example, after a simulation is done, we can go back and check that the results agree if we use the NR waveform itself. This is done in Appendix A for the equal-mass nonspinning system. Even more extreme, the original proof-of-principle demonstration of this method [2] used stationary-phase approximated (SPA) waveforms terminated at the light-ring frequency. For small mismatches, the results achieved in that test compare well to the results achieved in this paper, even though the SPA waveform mismatches the numerical and EOB waveforms by more than 8% over the relevant mass range [10].

The second lesson gleaned from these considerations of the match is that the *phase* of a model waveform does not come into the match; only the *phase difference* between models matters, as shown explicitly in Eq. (1b). This fine distinction has real importance for us because it implies that the phase error in our ersatz NR waveform relative to the correct physical waveform is not important. Certainly the final model waveform should resemble the physical waveform as closely as possible, but we will assume that errors in the portion of the final waveform covered by NR data will be accounted for separately in the error budget—or are essentially negligible compared to the PN errors.⁴ When comparing two plausible waveforms hybridized with the same ersatz NR data at some frequency f_{hyb} , the phase difference during the NR portion ($f > f_{\text{hyb}}$) will be zero to a very good approximation, at least until the waveforms are shifted in time and phase to maximize the match. But even then, the phase difference will not depend in any way on the phase of the ersatz NR waveform; by Eq. (2), it will be

$$\delta\phi(f) = 2 \Delta\Phi + 2\pi f \Delta T \quad \text{for } f > f_{\text{hyb}}. \quad (6)$$

Thus, the mismatch between plausible hybrids is not directly sensitive to the particular phasing of the ersatz NR waveform. Of course, that phasing will affect the alignment during hybridization, which can affect the relative power in different portions of the waveform or the function $\delta\phi(f)$ for frequencies $f < f_{\text{hyb}}$. However, the results below show that the ersatz waveform is not dominating the uncertainty, which suggests that even this phase error is not important.

These considerations lead us to conclude that the ersatz NR waveform must be reasonably accurate, especially in terms of modeling the relative power in different parts of the waveform. However, we are also given hope that our

⁴ Of course, this method readily applies to quite general error budgets. For example, if the expected uncertainty in the numerical waveform can be estimated, even if that error depends on the initial frequency of the simulation, the error budget can be trivially extended to include that estimate. See Sec. III E for more details.

final result—the predicted value for Ω_0 —does not depend strongly on the accuracy of that ersatz NR data. We can test this expectation and will see in Appendix A that, at least in the case of the equal-mass nonspinning system, it is indeed well founded.

B. The range of plausible PN waveforms

Reliable final results depend on an accurate range of plausible hybrids, correctly depicting the possible error in the analytical models. The range must be neither too broad nor too narrow: too broad, and we will conclude that the error is large, and thus begin the simulation earlier than necessary; too narrow, and we will be overconfident in our models, and thus waste time producing an inaccurate hybrid waveform. Unfortunately, we have no way of knowing the error in our analytical models before the fact. If, however, we assume that our models are not wrong, but are simply incomplete, we should be able to trust the uncertainties in the model to estimate the error.

Analytical relativity has produced multiple methods of calculating waveforms for black-hole binaries. Roughly speaking, these different methods should be equivalent at the level of our knowledge of the true waveform. To the extent that they are different, they are uncertain. In fact, we will use precisely this range of differences as our range of uncertainty, and thus our estimate for the errors in the analytical waveforms. Thus, choosing our range of plausible hybrids comes down to choosing representatives of the various methods for calculating analytical waveforms. The representative methods of calculation we will use are the TaylorT1–T4 [23–25] and EOB [16, 33] models.

An objection might be raised that the EOB waveform is more accurate and, in particular, “breaks down” more elegantly than the basic PN waveforms; the implication being that the TaylorT n waveforms should not be included. It would certainly be possible to use EOB waveforms alone, employing so-called flexibility parameters [34] to delimit a range of plausible waveforms, for example. Unfortunately, while EOB waveforms can be tuned very precisely to resemble the late-time behavior of numerical waveforms after the fact [33, 35], there is no evidence that the *inspiral*—which is of more interest here—will be more accurate after this tuning [36], or that any portion of the EOB waveforms will be more accurate before such tuning can be done. In fact, Blanchet has suggested that EOB appears to converge toward a theory which is different from general relativity [37]. Because the results will need to apply in regions of parameter space where no simulation has yet been done, and because this paper attempts to reflect methods currently in use by the community [26], we will take the more conservative approach of including the TaylorT n approximants.

II. EVALUATING UNCERTAINTY

With a selection of plausible waveforms in hand, we can now evaluate the differences between them in terms of the match. The aim is to treat the artificial data (the ersatz NR waveform) exactly as data from the actual simulation will be treated. First, methods of hybridization will be reviewed. This hybridization will be performed for a variety of hybridization frequencies, ω_{hyb} . Then, we can simply evaluate the match between the various hybrids at each ω_{hyb} , as a function of the total mass of the system.

A. Hybridization techniques

Combining inspiral and merger/ringdown waveforms is a delicate process, beginning with the procedure for aligning the waveforms by matching up the arbitrary time and phase offsets in the data. As described by MacDonald et al. [6], this part of the process has large potential effects on the accuracy of the final result; in their example a misalignment of just 1 M in the time values of the two waveforms resulted in a mismatch of up to $\text{MM} = 0.01$. Clearly, this part of the process must be handled carefully. Many techniques have been devised for doing so, resulting in a variety of choices to be made.

First, alignment of the time and phase offsets may be done in either the time domain or the frequency domain. For the particular case of hybridizing to numerical waveforms, the numerical data is short, beginning at high frequency and—in particular—having large amplitude. Transforming such data into the time domain will either introduce Gibbs phenomena, which will spoil much of the NR data, or require windowing, which will waste much of the NR data. Therefore, we use time-domain alignment for our purposes, as is used throughout most of the current literature.⁵

Second, we must choose a criterion for deciding how well the two waveforms are aligned after offsetting the time and phase. Many possibilities have been suggested for this purpose, including the magnitude of the difference in the complex h data [20]; the gravitational-wave phase and frequency [25]; even the *orbital* phase and frequency [24]. The gravitational-wave phase and frequency are often chosen by numerical-relativity groups to produce hybrids because the post-Newtonian phase is known to higher order than the amplitude, and is thus more likely to result in an accurate alignment. Because of its popularity and simplicity, phase alignment is used in this paper.

Finally, the alignment procedure depends on the width of the region over which the criterion chosen above is evaluated. For example, to align the phase (and implicitly

⁵ In the original proof-of-principle demonstration of the method described in this paper [2], frequency-domain alignment was used. In that case, the ersatz NR and PN waveforms were various versions of the TaylorF2 waveform—which is calculated in the frequency domain, which means Gibbs phenomena are not relevant. See also Ref. [4].

frequency) of two waveforms, a common method [10] is to minimize the squared difference between them:

$$\Xi(\Delta T, \Delta\phi) = \int_{t_1}^{t_2} [\phi_{\text{NR}}(t) - \phi_{\text{PN}}(t + \Delta T) - \Delta\phi]^2 dt. \quad (7)$$

The alignment then depends on t_1 and t_2 . There is a certain trade-off here, between using a short region so that less of the inaccurate PN waveform can be used, and using a long region to smooth out any irregularities in the numerical data, such as junk radiation or residual eccentricity. Additionally, the range $[t_1, t_2]$ must capture some curvature in the graph of $\phi(t)$ for an accurate alignment, which means that the range must become larger at lower frequencies. Reference [6] suggests a simple but robust method of choosing this range, where t_1 and t_2 extend to frequencies 5% above and below some central frequency. This ensures that the range is neither too large at high frequencies, nor too small at low frequencies.

In our case, the ‘‘numerical’’ data is the EOB waveform, which has essentially no eccentricity or noise. Thus, for simplicity of presentation and implementation, we take the limit of this procedure as t_2 approaches t_1 .⁶ To do so stably, we adjust the time offset ΔT so that the frequencies are the same at some time t_{hyb} , then adjust the phase offset $\Delta\phi$ so that the phases are the same at t_{hyb} :

$$\omega_{\text{NR}}(t_{\text{hyb}}) = \omega_{\text{PN}}(t_{\text{hyb}} + \Delta T), \quad (8a)$$

$$\phi_{\text{NR}}(t_{\text{hyb}}) = \phi_{\text{PN}}(t_{\text{hyb}} + \Delta T) + \Delta\phi. \quad (8b)$$

Here, the optimal offsets will depend on the time at which the alignment condition is imposed. Below, this dependence is described using the frequency itself: $\omega_{\text{hyb}} := \omega_{\text{NR}}(t_{\text{hyb}})$.

The alignment just described is typically applied to the $(\ell, m) = (2, 2)$ mode in a spin-weighted spherical harmonic decomposition of the gravitational waves. Other modes must also be aligned. However, we have fixed the only degrees of freedom, which means that the other modes are already determined. In general, the amplitude and phase of any mode of the PN waveform is transformed according to

$$A_{\text{PN}}^{\ell, m}(t) \rightarrow A_{\text{PN}}^{\ell, m}(t - \Delta T), \quad (9a)$$

$$\phi_{\text{PN}}^{\ell, m}(t) \rightarrow \phi_{\text{PN}}^{\ell, m}(t - \Delta T) - m \Delta\phi/2 + 2\pi n, \quad (9b)$$

for some integer n that ensures continuity of the phase.

Now, having aligned the two waveforms, we need to produce a single waveform. Because only $\phi^{2,2}$ is guaranteed to be continuous, discontinuities are possible in other

quantities, so a hybrid is usually formed using a transition function to blend the two waveforms. Here, because the discontinuities are mild, we use a basic linear transition [10] of width $10M$ centered on the alignment point, for amplitudes and phases of all modes.

A special case arises when the PN waveform is the EOB approximant. Rather than actually splitting the EOB waveform into two parts and recombining them, the complete EOB waveform is used as the EOB ‘‘hybrid’’.

B. Mismatches

Now, having formed a series of hybrids using various PN approximants, and a range of hybridization frequencies ω_{hyb} , we can evaluate the difference between them using the same criterion as will be used with the numerical data. In particular, we take the mismatch [Eq. (4)] using the Advanced LIGO zero-detuning, high-power noise curve [8]. The waveforms are projected onto the positive z axis using all available modes of the spin-weighted spherical-harmonic decomposition. Post-Newtonian calculations have been carried out through $\ell = 8$:

$$h(t) = \Re \left[\sum_{\ell=2}^8 \sum_{m=-\ell}^{\ell} A^{\ell, m}(t) e^{i\phi^{\ell, m}(t)} {}_{-2}Y_{\ell, m}(0, 0) \right]. \quad (10)$$

Note, however, that the quasinormal-mode portion of the EOB waveform has only been extended to include $(\ell, m) \in \{(2, \pm 2), (2, \pm 1), (3, \pm 3), (3, \pm 2), (4, \pm 4)\}$ [33]. Other modes are set to zero during ringdown.

As a first example, Fig. 3 shows the mismatches between each pair of hybrids for the equal-mass nonspinning system scaled to a total mass of $20M_{\odot}$. At any particular value of ω_{hyb} there is a range of mismatches, indicating that some pairs of hybrids happen to agree with each other very closely, while some are quite different. There is no reason to suspect that a pair of hybrids in close agreement with each other also agree with the exact waveform. Rather, if these are all plausible waveforms, the uncertainty in our model is given by the maximum mismatch between any pair. For this particular system, that pair happens to be the TaylorT1 and T3 waveforms for most values of ω_{hyb} , though in general TaylorT1 is most dissimilar from the other waveforms.

We can follow the maximum mismatch as a function of frequency, and notice a general trend: increasing hybridization frequency results in increasing uncertainty. This is to be expected for two reasons. First, the PN approximation should be very accurate at low frequencies, but break down at higher frequencies. For example, the time and phase at which ω_{hyb} occurs in the PN waveform will become more uncertain as ω_{hyb} increases. This results in uncertainty in the alignment between the two parts of the hybrid.

Second, as the hybridization frequency increases, the detector will simply be more sensitive to the differences between hybrids. The upper horizontal axis of Fig. 3 shows the physical hybridization frequency f_{hyb} . Comparison

⁶ It has been checked that the results of this paper are essentially identical when using Eq. (7) with values of t_1 and t_2 as prescribed in Ref. [6]. This is a better choice for data from simulations, and is more typical of hybridization as practiced by numerical-relativity groups, but would introduce an unnecessary layer of complexity to the discussion here.




FIG. 3. **Mismatches between hybrids as a function of ω_{hyb} .** This plot shows the mismatch between pairs of hybrids using different approximants, for the equal-mass nonspinning system with $M_{\text{tot}} = 20 M_{\odot}$. At any particular value of ω_{hyb} , the maximum mismatch between each pair of hybrids is the uncertainty in the final waveform hybridized at that frequency. If our target accuracy were, for example, $\text{MM} \leq 10^{-3}$ for this system mass, this plot shows us that the NR waveform would need to contain the GW frequency of $G M_{\text{tot}} \omega_{\text{hyb}} / c^3 = 0.02$, which naturally implies the initial orbital frequency of the simulation, Ω_0 . Achieving a smaller uncertainty requires lower ω_{hyb} , corresponding to times at which the PN approximants agree more closely.

with the noise curve plotted in Fig. 2 demonstrates that the mismatch grows very quickly just as the hybridization point passes the “seismic wall” of the detector ($f_{\text{seismic}} = 10$ Hz) where the sensitivity is improving rapidly with increasing frequency. The mismatch then begins to level out as the detector sensitivity levels out.

In the next section, we discuss how to use plots like the one in Fig. 3, along with a target accuracy, to find the optimal initial orbital frequency of a numerical simulation.

Before moving on, however, let us pause to note an interesting feature of the last plot. The various comparisons separate into distinct groups. The largest mismatches involve TaylorT1 hybrids (solid lines), and *all* hybrids using T1 have large mismatches with other hybrids. Setting aside the T1 waveforms, we see a similar trend develop at high frequencies for T3 hybrids compared to hybrids other than T1 (dashed lines). We could even push this categorization to the T2 hybrids compared to waveforms other than T1 and T3 (dotted lines), though this leaves little room for comparison. We might say that T1 is dominating the uncertainty, in the sense that it is furthest from the consensus

of other hybrids. At higher frequencies, T3 departs from that consensus, followed by T2, leaving only the T4 and EOB agreeing with each other. What is striking about this pattern of *disagreements* is that it is identical to the pattern in *errors* relative to the numerical waveform [25], where T1 is least accurate, followed by T3, with T2 slightly less accurate than T4; the T4 and EOB waveforms agree with the numerical result nearly within numerical errors.

An optimist might suggest that close systematic agreement between two models like T4 and EOB is unlikely—given the size of the function space through which they are free to roam—unless they also agree with the exact waveform. This would imply that we should take the small mismatch between T4 and EOB as the uncertainty in those waveforms. Or, slightly less optimistically, we might discount T1 as being too far from the other waveforms, and thus a mere anomaly. Unfortunately, similar patterns do not develop for the other systems investigated below. For now, we leave this as a mere observation, and take the largest mismatch as an indicator of the uncertainty in any given model.

III. USING MISMATCH TO FIND Ω_0

Given a target accuracy, the uncertainty implied by Fig. 3 suggests a natural starting point for the numerical simulation of that system, simply because the simulation must include—at a minimum—the corresponding ω_{hyb} . In this section, the uncertainty estimate of the previous section is used to produce an optimal initial orbital frequency for the simulation. This is then generalized to apply across a range of masses and to incorporate more complicated target mismatches. In the process, the uncertainties for a small selection of astrophysical systems are shown.

A. Optimizing Ω_0 for a particular mass

Figure 3 establishes a relationship between the uncertainty in plausible hybrids and the frequency at which they are hybridized, ω_{hyb} . If, for example, we wish to model an equal-mass nonspinning system of total mass $M_{\text{tot}}/M_{\odot} = 20$ with a target accuracy of $\text{MM}_{\text{target}} \leq 10^{-3}$, this plot demonstrates that the final hybrid waveform must be formed with $G M_{\text{tot}} \omega_{\text{hyb}} / c^3 \lesssim 0.02$. Naturally, the numerical simulation must include that frequency. There are two simple methods for turning the GW frequency into an initial orbital frequency for the simulation, Ω_0 .

First, we might use the basic approximation $\Omega_0 \approx \omega_{\text{hyb}}/2$. In this case, the result above would suggest an initial frequency of $G M_{\text{tot}} \Omega_0 / c^3 \lesssim 0.01$.⁷ The actual simulation should probably begin somewhat earlier than this, to

⁷ This is roughly half the initial frequency of current long simulations. To lowest order, the length of a simulation goes as $T \propto \Omega_0^{-8/3}$, which means that a simulation held to this standard needs to be roughly $2^{8/3} \approx 6.3$ times longer than current long simulations.

allow junk radiation to leave the system, and to ensure that the alignment region of Eq. (7) does not extend past ω_{hyb} . These considerations depend on the particular formulation of Einstein's equations and numerical methods used in the simulation, and are thus beyond our scope. Ultimately, the numerical relativist will use his or her judgment to produce some time Δt before $\omega_{\text{hyb}}/2$ at which the simulation should begin. In this regard, an additional PN approximation may be useful:

$$\Omega_0 \approx \frac{\omega_{\text{hyb}}}{2} - \Delta t \left(\frac{\omega_{\text{hyb}}}{2} \right)^{11/3} \frac{96\nu}{5} \left(\frac{G M_{\text{tot}}}{c^3} \right)^{5/3}, \quad (11)$$

where $\nu = m_1 m_2 / (m_1 + m_2)^2$ is the symmetric mass ratio of the individual black holes. The extra term is derived from the lowest-order PN approximation for the evolution of the orbital frequency [37]. Because of the approximations, this method may fail in certain extreme cases.

Alternatively, and more robustly, we might simply refer to any of the PN models contributing to our estimate, to find the orbital frequency corresponding to ω_{hyb} . Moreover, if some Δt is prescribed for the simulation, the PN model can be used to find the orbital frequency Ω_0 occurring at a time Δt before the GW frequency ω_{hyb} .

The particular example just discussed applies only when the target accuracy is required for $M_{\text{tot}} = 20 M_{\odot}$. This would be relevant to the situation where, for example, a source has been detected, and its parameters are known to reasonable accuracy, but further simulations are being done for accurate parameter estimation. More generally, however, we should expect to encounter broader accuracy requirements, which might apply across a range of masses [15]. The rest of this section will extend this example to account for various masses; to demonstrate the uncertainty for a selection of interesting systems; then to allow the target accuracy to vary as a function of mass; and finally to allow the target accuracy to vary as a function of both mass and hybridization frequency.

B. Optimizing Ω_0 for a range of masses

The mismatch curves plotted in Fig. 3 depend strongly on our choice of the total system mass. To generalize this to be a function of both ω_{hyb} and M_{tot} , we create the contour plot in the upper left of Fig. 4. Slicing through that plot at $M_{\text{tot}}/M_{\odot} = 20$ gives the uppermost curve in Fig. 3. For comparison, this quantity is also plotted for a selection of systems with different mass ratios or spins, as discussed in greater detail in Sec. III C.

Again, given an accuracy requirement, we can use this plot to derive the optimal initial orbital frequency. If the requirement is a target accuracy of $\text{MM}_{\text{target}} \leq 10^{-3}$, we can follow the 10^{-3} contour in the plot, and see that it is always above $G M_{\text{tot}} \omega_{\text{hyb}} / c^3 \approx 0.0075$ for the range of masses shown. As before, the initial orbital frequency Ω_0 is then deduced from this value by using Eq. (11) or by consulting the PN model, as described above. This stringent accuracy

requirement calls for numerical simulations roughly 87 times longer than the longest current simulation of this system.

C. Comparing the uncertainty in various systems

While the equal-mass nonspinning case nicely illustrates the method of finding optimal initial frequencies, systems at the boundaries of current numerical capabilities also hold a great deal of interest. Fig. 4 illustrates the uncertainty for four systems:

1. Equal-mass, nonspinning;
2. Equal-mass, aligned spins $\chi_1 = \chi_2 = 0.95$;
3. Mass ratio 10:1, nonspinning;
4. Mass ratio 10:1, aligned spins $\chi_1 = \chi_2 = 0.95$.

The quantities χ_1 and χ_2 are the components of the dimensionless spins along the orbital angular-momentum vector. In the lower-left plot, case (3), frequencies above $G M_{\text{tot}} \omega_{\text{hyb}} / c^3 \approx 0.055$ are not possible because the TaylorT3 approximant ends at that frequency for this system. We can treat those higher frequencies as having $\text{MM} = 1$. Note that the smallest known black holes have masses of at least $M \approx 3 M_{\odot}$ [38]. This suggests that the smallest total mass for systems with $q = 10$ would be $M_{\text{tot}} \approx 33 M_{\odot}$. Thus, the low-mass regions of those two plots may not be interesting astrophysically.

In each case, we see the basic trends noted in Sec. II B. Two factors drive the mismatch: how far ω_{hyb} has entered the sensitive band of the detector, and how poorly the post-Newtonian approximation performs up to that frequency. Each plot of Fig. 4 includes a dotted red line denoting f_{seismic} , the lower bound of sensitivity in Advanced LIGO. Below this line, the mismatch must be zero, because the data in the detector's sensitive band is identical for any two waveforms—it is just the ersatz NR data. As we move above this line, a larger fraction of the data in the corresponding hybrids comes from different approximants. Thus, the mismatch increases. We expect it to increase more quickly, as a function of ω_{hyb} for systems that are not well described by PN approximations. Indeed, comparing the plots, we see that any given contour line moves closer to f_{seismic} as either the mass ratio or spin parameter increases. This is simply an explicit confirmation that the physical system is not well modeled in more extreme cases.

Though less important for our purposes, the contour lines for larger mismatches also obey a similar bound. In each plot, the dashed green line shows $f_{20\%}$, the twentieth percentile of power in the match. That is, assuming $\delta\phi = 0$ in Eq. (1b), $f_{20\%}$ is the frequency for which the integral has




FIG. 4. Maximum mismatch between plausible hybrids for a selection of systems. These plots show the maximum mismatch between any pair of hybrids formed by hybridizing TaylorT1–T4 or EOB with an ersatz NR waveform (EOB) at ω_{hyb} , as described in Sec. II A, when scaled to a total system mass of M_{tot} . This quantity describes the uncertainty in the models used to form the hybrid. The plots show data from four cases with mass ratios denoted by q , which is the ratio of the larger to smaller mass, and components of the dimensionless spins χ aligned with the orbital angular-momentum vector. Note that $M_{\text{tot}} \lesssim 33 M_{\odot}$ may not be interesting astrophysically for black-hole binaries when $q = 10$. The region $G M_{\text{tot}} \omega_{\text{hyb}}/c^3 \gtrsim 0.055$ is inaccessible in the $q = 10, \chi = 0$ case, because the TaylorT3 approximant ends at that frequency; this could be interpreted as complete uncertainty. Generally, the uncertainty is larger in systems with more extreme parameters. The dotted red line in each plot shows $f_{\text{seismic}} = 10 \text{ Hz}$, the lower bound of sensitivity in Advanced LIGO. The dashed green line in each plot shows $f_{20\%}$ —the twentieth percentile of the power of the match, defined by Eq. (12). Comparing any plot to an accuracy requirement, which may depend on both M_{tot} and ω_{hyb} , we can extract the maximum sufficient hybridization frequency, which suggests the optimal initial orbital frequency. The method for doing this is described in Sec. III. These plots are discussed in more detail in Sec. III C.

accumulated 20 % of its final value:

$$4 \int_0^{f_{20\%}} \frac{|\tilde{h}_a(f)| |\tilde{h}_b^*(f)|}{S_n(|f|)} df \\ = 0.20 \times 4 \int_0^\infty \frac{|\tilde{h}_a(f)| |\tilde{h}_b^*(f)|}{S_n(|f|)} df. \quad (12)$$

For simplicity, the lines shown in the plots use EOB data for both waveforms: $h_a = h_b = h_{\text{EOB}}$. The line is roughly a lower bound for mismatches of $\text{MM} = 0.20$ —the white region of each plot. The higher the white region is above the $f_{20\%}$ line, the more accurate the PN approximations are. At very high frequencies, the white region must approach this line simply because the PN approximations break down more quickly there.

D. Mass-dependent target accuracy

In each of the examples above, the accuracy requirement calls for a specific mismatch, regardless of the total mass of the system. This is unrealistic for three reasons. First, and most simply, the result depends sensitively on a choice of mass—in the example of Sec. III B, that choice is the lower bound of the mass range used in Fig. 4. If we were to increase that lower bound to $M_{\text{tot}} \geq 50 M_\odot$, rather than calling for simulations 87 times longer, we would accept the longest current simulations. This sensitivity to the mass range shows that it must be considered more carefully than an arbitrary choice of plotting range.

Second, the SNR of an astrophysical signal will depend on the total mass. The precise dependence is complicated, as it involves the shape of the noise curve in Fig. 2. However, for masses at which the merger frequency is much higher than the detector’s low-frequency sensitivity (a good approximation for the mass range discussed here), the stationary-phase approximation shows that the SNR should scale roughly as $M_{\text{tot}}^{5/6}$ [39]. If we expect the typical low-mass system in real data to have a lower SNR than the typical high-mass system, there is no reason to model the two with the same precision. More precisely, for optimal parameter estimation, the error of a model waveform should scale inversely as the square of the SNR of the expected signal [32, 40].

Finally, the merger rates of real binaries will likely depend on the total mass simply because formation mechanisms for such binaries should depend on the total mass [11, 41–43]. Though current understanding of such mass dependence is not great [44], it is an area of active research, and will no doubt be improved in the future. In that case, we may wish to fold the expected event rate into our target accuracy, so that time is not wasted calculating precise waveforms for systems we are unlikely to observe.

For these reasons, useful and efficient accuracy requirements should depend explicitly on M_{tot} . Incorporating mass dependence involves reinterpreting the target mismatch $\text{MM}_{\text{target}}$ as a function of mass, rather than as a

constant. For example, we might hope to model a particular system well enough to ensure detection of any binary expected in the data, and to allow accurate parameter estimation for 90 % of binaries. A very crude function that implements this idea is

$$\text{MM}_{\text{target}}(M_{\text{tot}}) = \begin{cases} 1 & \frac{M_{\text{tot}}}{M_\odot} \leq 7; \\ 10^{-2} & 7 < \frac{M_{\text{tot}}}{M_\odot} \leq 12; \\ 10^{-4} \left(\frac{M_{\text{tot}}}{20 M_\odot}\right)^{-5/3} & 12 < \frac{M_{\text{tot}}}{M_\odot}. \end{cases} \quad (13)$$

The first two cases are inspired by population-synthesis results [41] suggesting that basically all equal-mass black-hole binaries should have⁸ $M_{\text{tot}} \gtrsim 7 M_\odot$, and that roughly 90 % should have $M_{\text{tot}} \gtrsim 12 M_\odot$. This function ignores binaries with $M_{\text{tot}} \leq 7 M_\odot$, and models binaries with $M_{\text{tot}} \leq 12 M_\odot$ just well enough to ensure detection.⁹ For higher masses, the function scales inversely as the SNR, and is normalized to optimally estimate the parameters of a $20 M_\odot$ system having SNR 70. Clearly, more sophisticated treatments could incorporate the objectives of detection and parameter estimation more smoothly, but this will serve to illustrate the idea.

Regardless of the particular form of the target mismatch, we use it to deduce the sufficient value of ω_{hyb} (and hence Ω_0) by plotting the ratio

$$\frac{\max_{a,b} \{\text{MM}(h_a, h_b)\}}{\text{MM}_{\text{target}}(M_{\text{tot}})}. \quad (14)$$

Where this ratio exceeds 1, the hybridization frequency is insufficient. With the target mismatch of Eq. (13), this ratio is plotted in Fig. 5. The red curve denotes values for which the ratio equals 1. The optimal ω_{hyb} is given by the lowest point this curve reaches, $G M_{\text{tot}} \omega_{\text{hyb}}/c^3 \approx 0.011$. Though this frequency is only 47 % higher than the frequency deduced in the previous section without incorporating mass dependence, it corresponds to a simulation that is nearly 3 times shorter—a significant improvement from the perspective of the numerical relativist, and hopefully a more accurate representation of the accuracy truly required for gravitational-wave detector data analysis. This holds great significance for systems with large mass ratios, where astrophysical considerations suggest that the lower bound of M_{tot} will be large, as mentioned in Sec. III C.

⁸ Note that several of the papers cited here discuss individual black-hole mass or the chirp mass of a binary, rather than the total mass.

⁹ A mismatch of 10^{-2} or less ensures detection 97 % of the time for a signal with SNR at least 7 [40]. However, a real search of detector data will use a template bank which does not need to be as accurate as this, in the sense that an inaccurate template with different physical parameters may happen to match the exact waveform. That is, the error of an “effectual” template bank may be far greater than 10^{-2} . The issue of detection is complicated, and is discussed at length in Ref. [9].




FIG. 5. **Ratio of maximum mismatch to target mismatch.** This plot shows the ratio between the maximum mismatch between various hybrids and a target mismatch given by Eq. (13). For values of this ratio greater than 1, the hybridization frequency is too high to achieve the target accuracy. The optimal sufficient value of ω_{hyb} is given by the lowest frequency at which the ratio is 1—roughly $G M_{\text{tot}} \omega_{\text{hyb}} / c^3 = 0.011$ here.

E. Mass- and frequency-dependent target accuracy

Up to this point, we have assumed that the mismatch between our hybrids completely describes the uncertainty in the final result, after a numerical simulation has been done. However, for the stringent accuracy requirements quoted above, the numerical simulation will be very long, making it difficult to achieve high accuracy in the NR portion of the waveform. We may wish to leave a portion of the error budget for the NR data, and the rest for the PN data and hybrid. Given some understanding of how the NR error depends on the length of the simulation, we can incorporate that error into our determination of the optimal value of Ω_0 . The error of a simulation depends, no doubt, on its length, but also on the particular implementation used for that simulation and even the computational resources available, and is thus beyond the scope of this paper. However, the basic idea is a simple extension of the technique discussed in the previous section: generalize the target mismatch to be a function of both M_{tot} and ω_{hyb} , and plot the ratio

$$\frac{\max_{a,b} \{\text{MM}(h_a, h_b)\}}{\text{MM}_{\text{target}}(M_{\text{tot}}, \omega_{\text{hyb}})}. \quad (15)$$

Again, when this is greater than 1, ω_{hyb} is too large. For example, the target mismatch may be constructed (crudely) by setting an astrophysically motivated target $\text{MM}_{\text{t,astro}}$ for the

final waveform, then subtracting the estimated uncertainty due to the NR data. Then the permissible mismatch in the PN data could be defined as

$$\begin{aligned} \text{MM}_{\text{target}}(M_{\text{tot}}, \omega_{\text{hyb}}) \\ := \text{MM}_{\text{t,astro}}(M_{\text{tot}}) - \text{MM}_{\text{NR}}(M_{\text{tot}}, \omega_{\text{hyb}}). \end{aligned} \quad (16)$$

This target is crude because the mismatch is not additive, but it is a conservative estimate.

This extension of the method to include dependence on ω_{hyb} raises an unfortunate—though realistic—possibility: it could be that the ratio of Eq. (15) will never be less than 1 for masses of interest. For example, if $\text{MM}_{\text{NR}}(M_{\text{tot}}, \omega_{\text{hyb}})$ increases too quickly as ω_{hyb} decreases, the quantity in Eq. (16) will be too small, and thus the ratio in expression (15) will be too large. This would indicate that the modeling methods, both analytical and numerical, are simply too crude to compute the waveform with the desired accuracy.

IV. DISCUSSION

Figure 4 presents the main results of this paper, showing the largest mismatch between any pair of plausible hybrids as a function of the frequency of hybridization and the total mass of the system. The hybrids are formed using the EOB waveform to substitute for the NR waveform. As is argued in Sec. IA and shown explicitly for the equal-mass nonspinning system in Appendix A, the particular choice of substitute does not affect the final results in any significant way. The hybrids' inspiral data are supplied by TaylorT1–T4 and EOB waveforms, which are attached at ω_{hyb} . These hybrids are then scaled to various total masses, and mismatches between each pair are calculated. The maximum such mismatch is the estimated uncertainty in the models. The plots in Fig. 4 assume that the error in the numerical portion of the hybrid is negligible, though they can be expanded to account for estimated numerical errors, as in Sec. III E. This uncertainty is a reasonable proxy for the error—the difference between the model and the exact waveform. Given a target uncertainty for the complete model, we can deduce the minimum initial orbital frequency necessary to achieve that target with a simulation, by noting that the relevant value of ω_{hyb} must be present in the simulation data.

The results show several interesting features. First, the uncertainty generally increases as the modeled system becomes more extreme; for a given value of ω_{hyb} and M_{tot} , increasing either the mass ratio or the spin parameter increases the uncertainty. This is not surprising, since the post-Newtonian order of known spin terms is lower than the order for non-spin terms [45]. Similarly, PN methods are expected to break down for larger mass ratios [37], for which more specific methods are necessary [46].

More quantitatively, we can relate these results to basic accuracy standards for gravitational-wave detectors. To calibrate waveforms for detection, accuracies of $\text{MM} \lesssim 0.01$

are generally called for [32, 40]. Meanwhile, the longest current numerical simulations start with $G M_{\text{tot}} \omega_{\text{hyb}}/c^3 \gtrsim 0.035$. The upper-left panel of Fig. 4 shows that hybrids created using such simulations would only be sufficient for $M_{\text{tot}} \gtrsim 26 M_{\odot}$. Of course, real detector data is searched for a range of system parameters; a real $10 M_{\odot} + 10 M_{\odot}$ system might be detected by an inaccurate $6 M_{\odot} + 18 M_{\odot}$ template, for example [10].¹⁰ This dramatically reduces the accuracy requirements on a template bank for detection [9]. However, while template banks may be subject to loose accuracy requirements, numerical relativity will generally be used for other purposes—most likely calibrating template banks to NR waveforms or hybrids incorporating them. The results presented above show that any such calibration is bound to exhibit very large errors for low-mass systems unless the numerical simulation is very long.

More stringent demands are placed on waveforms for parameter estimation, depending on the SNR of the observed signal. For example, modeling the unequal-mass high-spin system to high accuracy would require simulating nearly the entire in-band signal; the simulation would need to begin roughly 40% above the seismic wall to achieve mismatches of $\text{MM} \approx 10^{-4}$. These grim results present discouraging prospects for accurate modeling of precessing systems, PN approximations for which are known to still-lower order.

On the other hand, for small values of the mismatch, the appropriate value of ω_{hyb} varies almost linearly with M_{tot} . The initial orbital frequency Ω_0 required for a simulation will then be nearly proportional to the total mass of the modeled system, so the length of the simulation will vary roughly as $M_{\text{tot}}^{-8/3}$. This strong dependence shows clearly that target accuracies for a simulation to be used across a range of masses should include carefully considered mass dependence. Such dependence is incorporated into the technique for determining Ω_0 in Sec. III D. For the $q = 10$ systems in Fig. 4, this improves the situation dramatically. The smallest black hole observed to date has $M \gtrsim 3 M_{\odot}$ [38], so a binary with $q = 10$ would have $M_{\text{tot}} \gtrsim 33 M_{\odot}$, substantially raising the value of ω_{hyb} required to achieve a given mismatch for astrophysically likely sources. In this sense, systems with large mass ratios are actually easier to model than comparable-mass systems.

There are several possible flaws in these uncertainty estimates. Most basically, we simply assume that the uncertainty in our range of waveforms is a suitable proxy for the error in the waveforms. Of course, these models—both PN and EOB—may simply be wrong. For example, we can imagine that some fundamental error exists in our

understanding of approximations to Einstein’s equations for black-hole binaries. In that case, our models may be perfectly precise but entirely inaccurate; the exact waveform would lie outside the bounds of our uncertainty estimate. Moreover, these estimates depend on the assumption that the range of plausible hybrid waveforms is neither too narrow nor too broad. The choices made above were based largely on coincidences of history; some other reasonable, equally accurate, but not-yet-imagined approximant may exist, lying far from the approximants used here. Conversely, there may be some subtle error in one or more members of this group of approximants, leading to unnecessarily large uncertainty. Unfortunately, the only obvious way to detect such errors is to test the results using very long and accurate numerical simulations.

Taken together, these results indicate that more work will be needed to produce accurate waveforms for stellar-mass black-hole binaries, even for aligned-spin systems. Improvements may come in the form of higher-accuracy PN or EOB waveforms, longer numerical simulations, or both. This paper has not treated precessing systems simply because the production of full waveforms for such systems is still in its infancy. No doubt, however, the uncertainties are greater than in the cases discussed above. While both analytical and numerical relativity have clearly made great progress in the past decade, much remains to be done.

ACKNOWLEDGMENTS

It is my pleasure to thank Ilana MacDonald and Larne Pekowsky for useful discussions and direct comparisons of the results of our various codes; Larry Kidder, Harald Pfeiffer, Béla Szilágyi, and Saul Teukolsky for very useful discussions; Alessandra Buonanno and Yi Pan for kindly sharing their knowledge of EOB models; Larne Pekowsky and Duncan Brown for providing helpful information on noise curves and accuracy requirements; P. Ajith, Sascha Husa, and the rest of the NINJA-2 collaboration for help validating the Taylor PN approximants used in this work. This project was supported in part by a grant from the Sherman Fairchild Foundation; by NSF Grants No. PHY-0969111 and No. PHY-1005426; and by NASA Grant No. NNX09AF96G. The numerical computations presented in this paper were performed primarily on the Zwicky cluster hosted at Caltech by the Center for Advanced Computing Research, which was funded by the Sherman Fairchild Foundation and the NSF MRI-R² program.

Appendix A: The (un)importance of the choice of ersatz NR waveform

Section I A presents arguments that the mismatch between two hybrid waveforms with ersatz NR data should be almost completely insensitive to the phasing of the ersatz waveform above the hybridization frequency, and only weakly dependent on the amplitude. The key point is

¹⁰ In much of the literature, this difference is highlighted by distinguishing between “effectualness” (the match between a given signal and the best fit in a template bank) and “faithfulness” (the match between a given signal and the particular signal in the template bank with the same physical parameters).




FIG. 6. **Using real NR data for merger and ringdown.** This plot reproduces the top-left plot of Fig. 4, using real NR data in place of the EOB ersatz NR waveform. The NR data starts at frequency $\omega \approx 0.035$, and is extended to lower frequencies by hybridizing with a TaylorT4 waveform. The plots are almost identical, indicating that the procedure is not strongly sensitive to details of the ersatz NR waveform.

that the two hybrids are identical for $\omega > \omega_{\text{hyb}}$. It was thus argued that the particular choice of ersatz NR waveform should not strongly affect the results, as long as the power is fairly correctly distributed in the frequency domain.

One simple way to check this claim is to use a different ersatz NR waveform. Here, we will reproduce the crucial result of Fig. 4 in the equal-mass nonspinning case with different ersatz NR data. In this case, we will substitute the EOB waveform with a numerical waveform hybridized with TaylorT4 to extend to lower frequencies. The numerical waveform is the same one introduced in Ref. [47], except that Regge–Wheeler–Zerilli wave extraction is used to produce h . The waveform is hybridized exactly as in Ref. [10]. This hybrid is then substituted for the EOB waveform wherever it is called for.

The results are shown in Fig. 6. Comparing with the upper-left plot of Fig. 4, we see excellent agreement throughout the plot. The plot shown here does exhibit some jagged lines in the range $0.04 < G M_{\text{tot}} \omega_{\text{hyb}}/c^3 < 0.045$. These are evidently due to noise in the waveform itself, which appears to be related to junk radiation. That noise can easily lead to imperfect hybrids, especially using the frequency-alignment scheme of Eq. (8).

At the very least, this demonstrates that the simplistic ringdown-alignment technique used for the EOB waveform in this paper (see Sec. B 2) does not significantly affect the




FIG. 7. **Mismatch between the EOB waveform and the NR hybrid.** This plot shows the mismatch as a function of total mass between the EOB waveform used in the body of this paper and the NR hybrid used in Fig. 6. The similarity between Figs. 4 and 6 despite the significant mismatches shown here lead us to conclude that the uncertainties shown in those figures are indeed robust with respect to the choice of ersatz NR waveform.

final results. On the other hand, we might worry that the NR hybrid used here is practically identical to the EOB waveform used in the main text of this paper, because the EOB waveform aligns quite accurately to the very late stages of the NR data. In fact, that alignment is misleading, because it requires coherence over the relatively short span of the numerical data. Judged in terms of the mismatch, the NR hybrid and the EOB waveform are quite distinct, shown in Fig. 7 as a function of the total mass of the system.

Appendix B: Details of the implementation

The results of this paper depend sensitively on accurate numerical implementation of the technique. The various approximants, their hybrids, and the mismatch must all be calculated to high accuracy to ensure that the plots of Fig. 4 depict uncertainty in the models, rather than errors in the numerical methods. This section outlines the steps necessary to obtain accurate results. In short, every attempt was made to ensure that the model waveforms were as accurate as possible, and each number quoted in this section was tested to ensure that making it more stringent had no significant effect on the final results.

1. Post-Newtonian ingredients

The Taylor Tn approximants used in this paper are based on the results of Ref. [26]. In the appendix of that reference, the most current and complete PN results are collected and expressed in consistent notation. In particular, the orbital energy, tidal heating, and gravitational-wave flux are given to 3.5-PN order in nonspinning terms, and incomplete 2.5-PN order in spinning terms. The spins are assumed to be aligned or anti-aligned with the orbital angular momentum. These are the basic ingredients to construct the phasing of TaylorT1–T4 approximants, as described succinctly in Ref. [25]. The orbital phase thus derived is then used in the waveform amplitudes described by Ref. [26], which include nonspinning terms up to 3-PN order, and spinning terms through 2-PN order.

The main practical concern in constructing these waveforms is producing data on a sufficiently fine grid that accurate derivatives are available for the alignment step, Eq. (8) of the hybridization procedure. For the TaylorT1 and T4 models, this is accomplished by setting a tight tolerance on the numerical integration scheme, as discussed in Sec. B 3. For TaylorT2, the waveform is evaluated on at least 50 000 uniformly distributed values of the velocity parameter v , which is the independent variable of this model, ranging up to $v = 1$. Similarly, the TaylorT3 waveform was evaluated at 50 000 values of the independent time parameter τ , distributed at roughly uniform intervals of v . Lower values of these numbers led to poor hybrids, characterized by noise at low values of ω_{hyb} in the plots of Fig. 4.

2. EOB model

The EOB model used for this paper was designed to incorporate recent improvements to the inspiral portion of the model, including spin terms, while also remaining robust, allowing its application to the somewhat extreme case of $q = 10$, $\chi = 0.95$. The primary compromises made in the interests of robustness were abandonment of the factorized multipolar waveforms of Ref. [48] and coherent attachment of the ringdown portion of the waveform. The former compromise requires the use of the Padé-expanded flux to calculate the phasing of the system, and the standard PN multipolar waveforms. These are both reasonable substitutions: the flux term is used in the EOB code for the LIGO Algorithm Library; the PN multipolar waveform should still be accurate for most of the inspiral [48]. The latter compromise primarily affects the phase of the waveform during its very last stages. As was argued in Sec. IA, this is unlikely to have any significant effect. In any case, the uncertainty of the plausible waveforms is dominated by the Taylor Tn approximants in all cases shown in this paper, and essentially identical results are obtained when using real numerical data for the merger and ringdown, suggesting that any error in the EOB model does not affect the results.

The EOB Hamiltonian used here is roughly the same as the one given by Ref. [19], except that nonspinning

terms in the metric functions $A(r)$ and $D(r)$ are extended with new terms from Ref. [33]. Thus, in the nonspinning case, the Hamiltonian of this EOB model reduces exactly to the Hamiltonian of Ref. [33]; in the spinning case, it reduces nearly to the Hamiltonian of Ref. [19]. The angular momentum flux is described by Eq. (65) of Ref. [49], where the term F_4^4 is given by the Padé expansion of the flux from Ref. [26]. The standard formula for v_{pole} gives very poor results for high spins. For this paper, the following extension of v_{pole} to the spinning case is used:

$$v_{\text{pole}} = \frac{6 + 2\nu}{\sqrt{(3 + \nu)(36 - 35\nu) - \chi_s(8 - 4\nu)}}. \quad (\text{B1})$$

Initial data is set according to Eqs. (4.6) and (4.13) of Ref. [17]. Eccentricity is then iteratively reduced to $e \leq 10^{-14}$ using Eqs. (71) and (73) of Ref. [50]. For the high-spin cases, this method does not work directly. Instead, the spin is increased in stages. The non-eccentric initial data for the given mass ratio is first obtained with $\chi = 0$, then used as initial data for eccentricity reduction with $\chi = 0.1$. This is repeated, incrementing the value of χ , until the desired spin parameter is reached. That non-eccentric initial data is then used to evolve the full inspiral. Reducing eccentricity is not only more faithful to the scenario modeled by the other approximants, but also allows larger time steps to be taken by the numerical integration scheme; significant eccentricity would require at least a few steps to be taken per orbit.

The integration ends when the EOB radial parameter is smaller than 1, or the radial momentum becomes positive. In all cases explored for this paper, the amplitude of the resulting waveform reaches a peak, roughly where merger is expected, and roughly similar in amplitude to the peak expected from numerical simulations. Previously published EOB models align a sum of decaying quasinormal modes to the rising side of this peak [17, 33, 50]. Those techniques do not seem to be sufficiently robust to apply naively to the extreme cases discussed in this paper. Moreover, such techniques seem to be unnecessary; as argued previously, the particular details of the end of the waveform will not strongly affect the final results, especially for the small portion of the waveform represented by ringdown. For these reasons, a simple—though undoubtedly inaccurate—method is used to attach a single quasinormal mode to the inspiral waveform. The descending side of the amplitude peak is used, and the quasinormal mode with the longest decay time is attached at the unique point such that the amplitude and its first derivative are continuous.

3. Numerical integration of ODEs

The TaylorT1, TaylorT4, and EOB waveforms are integrated numerically by the eighth-order Dormand–Prince method implemented in *Numerical Recipes* [51]. In all cases, the absolute tolerance was set to $\text{atol} = 0$ because of the vastly different scales of the dependent variables.

The value of the relative tolerance `rto1` was chosen by looking at the convergence of the phase of each approximant. Tolerances from 10^{-4} to 10^{-11} gave the same results to within small fractions of a radian over the entire $\sim 100\,000$ rad inspiral. For EOB, `rto1` = 10^{-6} was chosen to be conservative, while still allowing the code to run very quickly (less than one second per waveform). For TaylorT1 and T4, `rto1` = 10^{-10} was used as a crude but effective way of ensuring that output was frequent enough to produce smooth derivatives for the alignment procedure, Eq. (8). Additionally, dense output was used to save 50 intermediate points per time step, which further improved the alignment procedure. Practically identical results were obtained with the Bulirsch–Stoer integration scheme, except that this method could not reliably continue into the delicate final few radians of the EOB integration. Integration continues until the dependent variables or their derivatives reach some unphysical value: for TaylorT1 and T4, angular frequency is required to remain positive; for EOB, radial momentum is required to remain negative, and radius greater than 1.

4. Fourier transforms and mismatches

Fourier transforms find two applications in the calculation of the mismatch. First, and most obviously, time-domain waveforms must be converted to the frequency domain for use in the inner product, Eq. (1). Second, the match itself is then evaluated by taking an inverse Fourier transform. Assuming that the waveforms h_a and h_b are normalized so that $\langle h_a | h_a \rangle = \langle h_b | h_b \rangle = 1$, and combining expression (2) with Eq. (3), we see that

$$\langle h_a | h_b \rangle = \max_{\Delta T, \Delta \Phi} \langle h_a | h_b \rangle \quad (\text{B2a})$$

$$= 2 \max_{\Delta T, \Delta \Phi} \Re \int_{-\infty}^{\infty} \frac{\tilde{h}_a \tilde{h}_b^*}{S_n(|f|)} e^{2i \operatorname{sgn}(f) \Delta \Phi + 2\pi i f \Delta T} df \quad (\text{B2b})$$

$$= 4 \max_{\Delta T} \left| \int_0^{\infty} \frac{\tilde{h}_a \tilde{h}_b^*}{S_n(|f|)} e^{2\pi i f \Delta T} df \right| \quad (\text{B2c})$$

Note that the integral in the last expression is simply the inverse Fourier transform of $\tilde{h}_a \tilde{h}_b^*/S_n(|f|)$. The maximization over ΔT involves selecting the largest element (in absolute value) of the discrete set produced by the fast Fourier transform of that quantity.

Two concerns drive the application of these Fourier transforms: aliasing at high frequencies, and Gibbs artifacts at low frequencies. To avoid aliasing, the sampling interval of the time-domain waveforms must be set by the highest frequency of the lowest-mass system of interest. In our case, that system has a mass of $M_{\text{tot}} = 5 M_{\odot}$. An acceptable sampling frequency is half the Advanced LIGO sampling frequency: $f_s = 8192$ Hz. On the other hand, avoiding Gibbs artifacts at low frequencies requires the waveforms to start early enough that the waveform “turns on” outside of the LIGO band, and its amplitude is very small at that point. Tests with the waveforms used in this paper show that an initial frequency of 8 Hz is sufficient to ensure accuracy of the mismatch to $\text{MM} \lesssim 10^{-7}$. For $M_{\text{tot}} = 5 M_{\odot}$, this corresponds to a dimensionless initial orbital frequency of $G M_{\text{tot}} \Omega_0 / c^3 \approx 1.97 \times 10^{-5}$. The waveforms used in this paper were calculated in dimensionless units, starting with that frequency, hybridized as necessary, scaled to the appropriate total mass, projected to the positive z axis, and interpolated to a uniform time grid with spacing $\Delta t = 1/f_s$. In extreme cases, these waveforms can consume hundreds of megabytes each. Given that five such waveforms need to be compared, and that comparison requires significant additional memory, the full memory usage can easily reach several gigabytes. Because of the large memory requirements, the calculations for Fig. 4 were performed on a cluster having ample memory in a single node. To use CPU resources efficiently OpenMP [52] was employed, which allows very simple alterations of source code to incorporate multiple processes—just three additional lines of code enabled multiprocessing which resulted in a speed improvement by a factor of four.

-
- [1] G. Lovelace, R. Owen, H. P. Pfeiffer, and T. Chu, *Phys. Rev. D* **78**, 084017 (2008).
 - [2] M. Boyle, “A note on initial orbital frequencies for BBH simulations,” (2010), presented to members of the NR–AR collaboration during the teleconference of Feb. 10, 2010.
 - [3] T. Damour, A. Nagar, and M. Trias, *Phys. Rev. D* **83**, 024006 (2011).
 - [4] L. Santamaría, F. Ohme, P. Ajith, B. Brügmann, N. Dorband, M. Hannam, S. Husa, P. Mösta, D. Pollney, C. Reisswig, E. L. Robinson, J. Seiler, and B. Krishnan, *Phys. Rev. D* **82**, 064016 (2010).
 - [5] M. Hannam, S. Husa, F. Ohme, and P. Ajith, *Phys. Rev. D* **82**, 124052 (2010).
 - [6] I. MacDonald, S. Nissanke, and H. P. Pfeiffer, *Class. Quant. Grav.* **28**, 134002 (2011).
 - [7] R. Flaminio, A. Freise, A. Gennai, P. Hello, P. L. Penna, G. Losurdo, H. Lueck, N. Man, A. Masserot, B. Mours, M. Punturo, A. Spallicci, and A. Viceré, *Advanced Virgo White Paper VIR-NOT-DIR-1390-304*, Tech. Rep. (Virgo collaboration, 2005).
 - [8] D. Shoemaker, *LIGO-T0900288-v3: Advanced LIGO anticipated sensitivity curves*, Tech. Rep. (LIGO, 2010).
 - [9] A. Buonanno, B. R. Iyer, E. Ochsner, Y. Pan, and B. S. Sathyaprakash, *Phys. Rev. D* **80**, 084043 (2009).
 - [10] M. Boyle, D. A. Brown, and L. Pekowsky, *Class. Quant. Grav.* **26**, 114006 (2009).

- [11] W. M. Farr, N. Sravan, A. Cantrell, L. Kreidberg, C. D. Bailyn, I. Mandel, and V. Kalogera, “The mass distribution of stellar-mass black holes,” (2010), arXiv:1011.1459 [astro-ph].
- [12] S. R. Lau, H. P. Pfeiffer, and J. S. Hesthaven, *Comm. Comp. Phys.* **6**, 1063 (2009).
- [13] J. Hennig and M. Ansorg, *Journal of Hyperbolic Differential Equations* **06**, 161 (2009).
- [14] S. R. Lau, G. Lovelace, and H. P. Pfeiffer, “Implicit-explicit (IMEX) evolution of single black holes,” (2011), arXiv:1105.3922 [astro-ph].
- [15] The NR–AR collaboration, “Results of the NR–AR collaboration,” (2011), not yet published.
- [16] A. Buonanno and T. Damour, *Phys. Rev. D* **59**, 084006 (1999).
- [17] A. Buonanno and T. Damour, *Phys. Rev. D* **62**, 064015 (2000).
- [18] A. Buonanno, Y. Pan, H. P. Pfeiffer, M. A. Scheel, L. T. Buchman, and L. E. Kidder, *Phys. Rev. D* **79**, 124028 (2009).
- [19] Y. Pan, A. Buonanno, L. T. Buchman, T. Chu, L. E. Kidder, H. P. Pfeiffer, and M. A. Scheel, *Phys. Rev. D* **81**, 084041 (2010).
- [20] P. Ajith, S. Babak, Y. Chen, M. Hewitson, B. Krishnan, J. T. Whelan, B. Brügmann, P. Diener, J. Gonzalez, M. Hannam, S. Husa, M. Koppitz, D. Pollney, L. Rezzolla, L. Santamaría, A. M. Sintes, U. Sperhake, and J. Thornburg, *Class. Quant. Grav.* **24**, S689 (2007).
- [21] P. Ajith, M. Hannam, S. Husa, Y. Chen, B. Brügmann, N. Dorband, D. Müller, F. Ohme, D. Pollney, C. Reisswig, L. Santamaría, and J. Seiler, *Phys. Rev. Lett.* **106**, 241101 (2011).
- [22] R. Sturani, S. Fischetti, L. Cadonati, G. M. Guidi, J. Healy, D. Shoemaker, and A. Viceré, *Journal of Physics: Conference Series* **243**, 012007 (2010).
- [23] T. Damour, B. R. Iyer, and B. S. Sathyaprakash, *Phys. Rev. D* **63**, 044023 (2001).
- [24] A. Buonanno, G. B. Cook, and F. Pretorius, *Phys. Rev. D* **75**, 124018 (2007).
- [25] M. Boyle, D. A. Brown, L. E. Kidder, A. H. Mroue, H. P. Pfeiffer, M. A. Scheel, G. B. Cook, and S. A. Teukolsky, *Phys. Rev. D* **76**, 124038 (2007).
- [26] P. Ajith, M. Boyle, D. A. Brown, S. Fairhurst, M. Hannam, I. Hinder, S. Husa, B. Krishnan, R. A. Mercer, F. Ohme, C. D. Ott, J. S. Read, L. Santamaría, and J. T. Whelan, “Data formats for numerical relativity,” (2011), arXiv:0709.0093v3 [gr-qc].
- [27] J. G. Baker, J. R. van Meter, S. T. McWilliams, J. Centrella, and B. J. Kelly, *Phys. Rev. Lett.* **99**, 181101 (2007).
- [28] L. S. Finn, *Phys. Rev. D* **46**, 5236 (1992).
- [29] B. S. Sathyaprakash and S. V. Dhurandhar, *Phys. Rev. D* **44**, 3819 (1991).
- [30] R. Balasubramanian, B. S. Sathyaprakash, and S. V. Dhurandhar, *Phys. Rev. D* **53**, 3033 (1996).
- [31] B. J. Owen, *Phys. Rev. D* **53**, 6749 (1996).
- [32] L. Lindblom, B. J. Owen, and D. A. Brown, *Phys. Rev. D* **78**, 124020 (2008).
- [33] Y. Pan, A. Buonanno, M. Boyle, L. T. Buchman, L. E. Kidder, H. P. Pfeiffer, and M. A. Scheel, “Inspirational-merger-ringdown multipolar waveforms of nonspinning black-hole binaries using the effective-one-body formalism,” (2011), arXiv:1106.1021 [gr-qc].
- [34] T. Damour, B. R. Iyer, P. Jaranowski, and B. S. Sathyaprakash, *Phys. Rev. D* **67**, 064028 (2003).
- [35] A. H. Mroué, L. E. Kidder, and S. A. Teukolsky, *Phys. Rev. D* **78**, 044004 (2008).
- [36] D. Brown and B. S. Sathyaprakash, “Accuracy standards required for NR waveforms for use in gravitational-wave data analysis,” (2010), presented to members of the NR–AR collaboration during the teleconference of Feb. 10, 2010.
- [37] L. Blanchet, *Living Reviews in Relativity* **9** (2006).
- [38] D. M. Gelino and T. E. Harrison, *Astrophys. J.* **599**, 1254 (2003).
- [39] C. Cutler and É. É. Flanagan, *Phys. Rev. D* **49**, 2658 (1994).
- [40] É. É. Flanagan and S. A. Hughes, *Phys. Rev. D* **57**, 4566 (1998).
- [41] T. Bulik and K. Belczyński, *Astrophys. J.* **589**, L37 (2003).
- [42] R. O’Shaughnessy, V. Kalogera, and K. Belczynski, *Astrophys. J.* **716**, 615 (2010).
- [43] F. Özel, D. Psaltis, R. Narayan, and J. E. McClintock, *Astrophys. J.* **725**, 1918 (2010).
- [44] L. S. Collaboration, *Class. Quant. Grav.* **27**, 173001 (2010).
- [45] K. G. Arun, A. Buonanno, G. Faye, and E. Ochsner, *Phys. Rev. D* **79**, 104023 (2009), note that the 1.5 and 2.5pN spin terms in the flux and energy expressions do not account for an erratum from 2010.
- [46] M. Sasaki and H. Tagoshi, *Living Reviews in Relativity* **6** (2003).
- [47] M. A. Scheel, M. Boyle, T. Chu, L. E. Kidder, K. D. Matthews, and H. P. Pfeiffer, *Phys. Rev. D* **79**, 024003 (2009).
- [48] Y. Pan, A. Buonanno, R. Fujita, E. Racine, and H. Tagoshi, *Phys. Rev. D* **83**, 064003 (2011).
- [49] M. Boyle, A. Buonanno, L. E. Kidder, A. H. Mroué, Y. Pan, H. P. Pfeiffer, and M. A. Scheel, *Phys. Rev. D* **78**, 104020 (2008).
- [50] A. Buonanno, L. E. Kidder, A. H. Mroué, H. P. Pfeiffer, and A. Taracchini, *Phys. Rev. D* **83**, 104034 (2011).
- [51] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes: The Art of Scientific Computing*, 3rd ed. (Cambridge University Press, 2007).
- [52] L. Dagum and R. Menon, *Computational Science and Engineering*, *IEEE* **5**, 46 (1998).