

Comparing gravitational waveform extrapolation to Cauchy-characteristic extraction in binary black hole simulations

Nicholas W. Taylor,¹ Michael Boyle,² Christian Reisswig,¹ Mark A. Scheel,¹ Tony Chu,³
Lawrence E. Kidder,² and Béla Szilágyi¹

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

²*Center for Radiophysics and Space Research, Cornell University, Ithaca, New York 14853, USA*

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

(Received 13 September 2013; published 3 December 2013)

We extract gravitational waveforms from numerical simulations of black hole binaries computed using the Spectral Einstein Code. We compare two extraction methods: direct construction of the Newman-Penrose (NP) scalar Ψ_4 at a finite distance from the source and Cauchy-characteristic extraction (CCE). The direct NP approach is simpler than CCE, but NP waveforms can be contaminated by near-zone effects—unless the waves are extracted at several distances from the source and extrapolated to infinity. Even then, the resulting waveforms can in principle be contaminated by gauge effects. In contrast, CCE directly provides, by construction, gauge-invariant waveforms at future null infinity. We verify the gauge invariance of CCE by running the same physical simulation using two different gauge conditions. We find that these two gauge conditions produce the same CCE waveforms but show differences in extrapolated- Ψ_4 waveforms. We examine data from several different binary configurations and measure the dominant sources of error in the extrapolated- Ψ_4 and CCE waveforms. In some cases, we find that NP waveforms extrapolated to infinity agree with the corresponding CCE waveforms to within the estimated error bars. However, we find that in other cases extrapolated and CCE waveforms disagree, most notably for $m = 0$ “memory” modes.

DOI: [10.1103/PhysRevD.88.124010](https://doi.org/10.1103/PhysRevD.88.124010)

PACS numbers: 04.25.D-, 02.70.Hm, 04.25.dg, 04.30.-w

I. INTRODUCTION

In the next few years, the second generation of ground-based gravitational-wave interferometers is expected to make the first direct detection of gravitational waves (GWs) from the inspiral and coalescence of compact binaries, marking the beginning of the era of gravitational wave astronomy [1–4]. Because of the very low compact binary coalescence rate [5], observable GW events are expected to originate from sources at the edge of the detectable range, with signal to noise ratios of order unity. Detecting these exceptionally weak GW signals requires the use of matched filtering, in which the noisy data are compared with a template bank of expected waveforms (see, e.g., Ref. [6] and references therein). For black hole binaries, these expected waveforms can be accurately computed only by using full numerical solutions of Einstein’s equations. However, because these simulations are computationally expensive, analytical or phenomenological models of GW emission are required in order to densely cover the parameter space. Because these models must be calibrated using results from numerical simulations [7–12], it is essential that accurate waveforms from numerical simulations are available. Moreover, it is crucial that the uncertainties in these numerical waveforms are well understood.

There are several sources of uncertainty in numerical waveforms. Perhaps the most straightforward to understand and measure is the numerical truncation error in

the binary black hole simulation itself, which we refer to as the “Cauchy error.” Numerical relativity codes for black hole binaries solve the full nonlinear Einstein equations. These are formulated as an initial value (Cauchy) problem, in which initial data (satisfying the Einstein constraints) are provided on some spacelike surface labeled by coordinate time t . The Einstein evolution equations are then used to determine data at subsequent times. The Cauchy error is the error made in solving these evolution equations numerically. It depends on the truncation error of the employed numerical scheme and the coarseness of the computational grid.

Another source of uncertainty in numerical simulations is the error associated with the use of a finite outer boundary. In principle the solution of Einstein’s equations is needed for the entire spacetime, but most simulations solve the equations only on a finite spatial domain. For example, simulations performed using the Spectral Einstein Code (SpEC) typically have outer boundaries located at about $500M$ (where M is the total mass of the system), while the total simulation time may be thousands of M [13–16]. The effects of a finite outer boundary can be mitigated by choosing constraint-preserving boundary conditions (see, e.g., Ref. [17]). However, such boundary conditions are not exact and cannot account for physical effects such as the backscatter of GWs off the spacetime curvature from regions outside the boundary. Previous studies have shown that this outer boundary error is typically comparable to or smaller than the Cauchy error [15,18].

Yet another source of uncertainty is the error associated with waveform extraction from finite-radius numerical data to future null infinity (I^+). A waveform at I^+ represents what would be measured by an Earth-based GW observatory that detects an astrophysical source. The simplest approach to waveform extraction is to compute the Newman-Penrose scalar Ψ_4 (see Sec. II A for details) at a large but finite distance from the source [19], and to use this as an approximation to the waveform at I^+ . This can be inaccurate, because the quantity Ψ_4 represents measurable outgoing gravitational radiation only in the limit of infinite distance from the source (see, e.g., Refs. [20–22]) and in the Bondi gauge [20] (rather than the gauge of the simulation).

A better approximation is a popular refinement of this single-extraction-radius method: Ψ_4 is extracted as before, but at several different radii instead of at a single radius, and this information is used to extrapolate Ψ_4 to I^+ (see Sec. II A 3 for details). This extrapolation procedure can remove near-zone effects and some gauge effects from the resulting waveform. However, as we show below, extrapolation does not always succeed in a convergent way, and even when it does, it is possible for some near-zone and gauge effects to remain. Estimating the magnitude of these remaining effects is difficult; it currently requires either repeating simulations using multiple gauge conditions or comparing with an independent wave-extraction method. Most of the currently available numerical-relativity waveforms, either published or used by groups working on calibration of analytical methods, employ (low-order) extrapolation of Ψ_4 or simply Ψ_4 extracted at a finite radius [23–26].

A more robust method of waveform extraction is Cauchy-characteristic extraction (CCE). This procedure uses a characteristic evolution code to solve Einstein's equations on a foliation of outgoing null hypersurfaces rather than on spacelike hypersurfaces [27–30]. Radial compactification enables the use of null hypersurfaces that extend all the way to future null infinity, so a waveform at I^+ can be directly computed. Furthermore, the waveform at I^+ can be computed in a gauge-invariant way [27]. In practice, the strong-field region near the source is evolved using a Cauchy code, while the asymptotic region is evolved with a characteristic code. The Cauchy evolution supplies data on a timelike, finite-radius world tube, which serves as the inner boundary for the characteristic evolution (see Fig. 1). This technique has been used in Refs. [31–33] for simulations of binary black hole mergers and in Refs. [34–37] for simulations of stellar collapse, binary neutron star mergers, and black hole formation. The primary disadvantages of CCE are its computational expense and its complexity (because it requires two separate methods of solving the Einstein equations). Our binary black hole simulations typically require weeks of wall time, and performing CCE can add several additional days of computation time. By comparison, the extrapolation procedure requires only minutes.




FIG. 1. Spacetime diagram illustrating CCE, with two spatial dimensions suppressed. The Cauchy evolution code advances its solution of Einstein's equations on successive spatial hypersurfaces Σ bounded by the outer boundary R_B . The wavy, dashed line on the left represents the small-radius part of the Cauchy simulation, whose details are not important here. The characteristic code advances its solution of Einstein's equations on successive null hypersurfaces (labeled $u = \text{const.}$). It uses data from the Cauchy code on the inner boundary (the world tube labeled by R_Γ) to produce a solution that is valid all the way to I^+ , where gravitational radiation is well defined. The characteristic code requires initial data on the null surface u_0 .

Other methods of waveform extraction have been considered in the literature. In addition to the methods discussed above, the most widely used is the Regge-Wheeler-Zerilli-Moncrief method [38–41], in which the far-field solution is treated as a perturbation about a fixed background (typically Schwarzschild or Minkowski), and the perturbed solution is constructed by reading off gauge-invariant perturbation coefficients from the numerical solution on a finite extraction sphere. See Ref. [42] for a review. Related methods for finding the asymptotic form of the waves from the finite-radius behavior were considered by Abrahams and Evans [43,44] and Lousto *et al.* [45], and have recently been generalized by Benedict *et al.* [46]. However, these analyses rely on certain assumptions about gauge that we do not make. In Ref. [34], a comparison between CCE, Ψ_4 , and Regge-Wheeler-Zerilli-Moncrief extraction was performed in the context of stellar collapse. In this paper we consider only two extraction methods: Ψ_4 extrapolation and CCE.

The goal of this paper is to compare extrapolated- Ψ_4 and CCE waveforms for binary black hole simulations performed using SpEC. We estimate the uncertainties in the waveforms associated with each extraction method, and we examine the differences between the waveforms relative to these estimated errors. In particular, by comparing extrapolated- Ψ_4 and CCE waveforms, we can estimate the unknown gauge error that may be present in the former. One important question we wish to address is whether it

suffices to use the (simpler and less computationally expensive) extrapolation method, or whether the gauge invariance of CCE is necessary, given the current accuracy of our simulations.

Some previous comparisons of CCE and extrapolation have been done using binary black hole simulations performed with the finite-difference code LLAMA [47]. In Refs. [31,48], it was found that differences between extrapolated and CCE quantities were on the order of the discretization error of the Cauchy simulation. Additionally, the differences were found to be nonconvergent, suggesting that the waveform extraction error could become dominant for high-accuracy simulations. These previous studies focused on short simulations of equal-mass and spin-aligned binaries. Here, we also consider longer unequal mass and generic precessing configurations, we compare multiple $Y_{\ell m}$ modes, and we use more sophisticated extrapolation and waveform-alignment methods.

This paper is organized as follows. In Sec. II we review different methods of waveform extraction. We discuss direct construction of Ψ_4 on finite-radius extraction spheres, extrapolation of Ψ_4 to infinity, and waveform extraction using CCE. In Sec. III we describe the black hole binary simulations that we use, briefly discussing the initial data, gauge conditions, and evolution algorithms. In Sec. IV we discuss how to estimate the various sources of error in the gravitational waveforms, including errors in the Cauchy evolution as well as in the waveform extraction methods. In Sec. V we verify that CCE is indeed gauge invariant by comparing waveforms from two simulations with identical physics but with different gauge conditions. We compare the relative magnitudes of the various errors, and we show that the error associated with the location of the CCE inner boundary (which we attribute to a mismatch of characteristic and Cauchy initial data) is typically greater than the numerical error in the characteristic evolution. We also show that, except for modes with $m = 0$, extrapolated- Ψ_4 and CCE waveforms agree to within the estimated error bars. We summarize in Sec. VI. Note that we will refer to extrapolated- Ψ_4 waveforms simply as “extrapolated waveforms,” and we will use the terms uncertainty and error interchangeably when discussing error estimates.

II. GRAVITATIONAL WAVE EXTRACTION

In this section, we review some of the mathematical preliminaries as well as the GW extraction methodology. We discuss how gravitational radiation content is extracted from the finite-radius numerical simulation, and we review the extrapolation and CCE methods.

A. Direct extraction of Newman-Penrose Ψ_4

1. The Newman-Penrose scalar Ψ_4

The GW content of a spacetime can be defined in terms of a particular component of the Weyl tensor using the

Newman-Penrose (NP) formalism [19]. This formalism is based on a complex tetrad of null vectors $\{l^\mu, n^\mu, m^\mu, \bar{m}^\mu\}$ that satisfy $l^\mu n_\mu = -m^\mu \bar{m}_\mu = 1$. Here, a bar denotes complex conjugation. The Weyl tensor $C_{\alpha\beta\gamma\delta}$ can be uniquely represented via five complex scalars by contracting with elements of the null tetrad:

$$\Psi_0 := l^\alpha m^\beta l^\gamma m^\delta C_{\alpha\beta\gamma\delta}, \quad (1a)$$

$$\Psi_1 := l^\alpha n^\beta l^\gamma m^\delta C_{\alpha\beta\gamma\delta}, \quad (1b)$$

$$\Psi_2 := l^\alpha m^\beta \bar{m}^\gamma n^\delta C_{\alpha\beta\gamma\delta}, \quad (1c)$$

$$\Psi_3 := l^\alpha n^\beta \bar{m}^\gamma n^\delta C_{\alpha\beta\gamma\delta}, \quad (1d)$$

$$\Psi_4 := n^\alpha \bar{m}^\beta n^\gamma \bar{m}^\delta C_{\alpha\beta\gamma\delta}. \quad (1e)$$

In asymptotically flat spacetimes, by virtue of the peeling theorem, the Weyl tensor obeys

$$C_{\alpha\beta\gamma\delta} \sim \frac{[N]}{r} + \frac{[III]}{r^2} + \frac{[II]}{r^3} + \frac{[I]}{r^4} + \mathcal{O}(r^{-5}), \quad (2)$$

where the letters in brackets denote Petrov types (see, e.g., Refs. [20–22]). As the distance from the source tends toward infinity, the spacetime approaches type N . Petrov type N spacetimes are outgoing plane-wave solutions, with Ψ_4 the only nonzero component of the Weyl tensor for a suitable choice of null tetrad. Consequently, in the limit of infinite distance from the source, Ψ_4 is identified as containing purely outgoing gravitational radiation. Assuming Bondi gauge [20], Ψ_4 can be directly related to the measurable plus and cross polarization modes of the strain h via two time integrals,

$$h_+ - ih_\times = \lim_{r \rightarrow \infty} \int_{-\infty}^t dt' \int_{-\infty}^{t'} dt'' \Psi_4|_{S^2}, \quad (3)$$

on a spherical surface S^2 at I^+ .

2. Ψ_4 extraction at finite distance

To extract Ψ_4 from a numerical simulation, one chooses a tetrad $\{l^\mu, n^\mu, m^\mu, \bar{m}^\mu\}$, computes the Weyl tensor by differentiating the metric, and then constructs Ψ_4 via Eq. (1e). Since the computational domain is of finite size, it is not possible to compute Ψ_4 at an infinite distance from the source. Instead, we typically compute Ψ_4 on finite-radius coordinate spheres. On each of these spheres, we expand Ψ_4 in spin-weighted spherical harmonics,

$$\Psi_4(t, r, \vartheta, \varphi) = \sum_{\ell, m} \Psi_4^{\ell, m}(t, r) {}_{-2}Y_{\ell, m}(\vartheta, \varphi), \quad (4)$$

where (ϑ, φ) are the usual polar coordinates on the sphere, in the coordinate system used by the simulation. In SpEC, we choose a coordinate tetrad that is only asymptotically null and orthonormal, in anticipation of extrapolation to infinity (see Sec. II A 3). Details of the Ψ_4 extraction method used by SpEC are described in Refs. [18,49,50].

This procedure has three drawbacks. First, it computes Ψ_4 at a finite radius where the spacetime is not necessarily of Petrov type N . This means that even if in the proper gauge, Ψ_4 may not be the only nonzero component in Eq. (2), and furthermore Ψ_4 does not necessarily correspond only to purely outgoing gravitational radiation. Second, we choose a coordinate-based tetrad $\{l^\mu, n^\mu, m^\mu, \bar{m}^\mu\}$, which only asymptotically has the properties that lead to the peeling theorem, Eq. (2). Third, we do not impose the Bondi gauge, but instead we use whatever gauge is used by the code that evolves Einstein's equations. This may lead to mixing of the Ψ_n , and hence it can invalidate Eq. (3), which relates Ψ_4 to the GW strain h even in the limit of infinite distance from the source.

The first two of these drawbacks can be reduced by extracting Ψ_4 on multiple coordinate spheres with different radii, and then extrapolating these results to $r \rightarrow \infty$, as described in Sec. II A 3 below. This extrapolation procedure not only handles the problem of the finite extraction radius, but it also corrects error terms introduced by the choice of a coordinate-based tetrad, since these error terms scale like higher powers of $1/r$. Extrapolation can also correct some gauge errors, provided that they fall off faster than $1/r$. However, it is possible that some gauge choices may produce effects that persist even after extrapolation, and some gauge choices may prevent accurate extrapolation altogether. This could occur, for example, if the gauge-induced leading-order falloff of the extracted Ψ_4 were slower than $1/r$. We will show an example of the latter case in Sec. V C.

3. Extrapolation

To extrapolate Ψ_4 to infinite radius, using data extracted on a series of finite spheres of different radii, we follow the procedure of Ref. [51]. In this section we summarize the technique, including certain minor improvements.

We measure the coefficients $\Psi_4^{\ell,m}$ of Eq. (4) at a set of coordinate times $\{t_i\}$ on a set of coordinate spheres of radii $\{R_j\}$, using the procedure described in Sec. II A 2. At each time, we also compute the areal radius r_{ar} of each sphere by integrating over the sphere using the full spatial metric, and we compute the average value of the metric component g^{tt} over each sphere. From the initial data we compute the Arnowitt-Deser-Misner (ADM) mass [52] M_{ADM} of the spacetime.

We then construct a retarded time that slightly generalizes the usual Schwarzschild definition to account for simple time dependence of the lapse and the radial coordinate. We define the retarded time as

$$t_{\text{ret}} := t_{\text{corr}} - r_*, \quad (5a)$$

where

$$r_* := r_{\text{ar}} + 2M_{\text{ADM}} \ln\left(\frac{r_{\text{ar}}}{2M_{\text{ADM}}} - 1\right), \quad (5b)$$

and

$$t_{\text{corr}} := \int_0^t \sqrt{\frac{-1/g^{tt}}{1 - 2M_{\text{ADM}}/r_{\text{ar}}}} dt'. \quad (5c)$$

Here, r_* is the standard tortoise coordinate of the Schwarzschild metric, with the Schwarzschild radial coordinate replaced by the areal radius r_{ar} , and the Schwarzschild mass parameter replaced by the ADM mass M_{ADM} of the initial data for simplicity. The corrected time t_{corr} is constructed so that if the metric in the given coordinates has the standard Schwarzschild form except for the lapse, then t_{ret} will be precisely a null coordinate. This does not account for other departures of the metric from Schwarzschild.

The quantities t_{ret} and r_{ar} defined above may not be the most optimal choices of coordinates; for instance, there may be other choices that make t_{ret} more nearly a null coordinate. The final result of extrapolation, however, will not be affected by imperfect choices of t_{ret} and r_{ar} as long as two conditions are satisfied: (1) our choices differ from the optimal choices by factors of at most $1 + \mathcal{O}(1/r_{\text{ar}})$, and (2) the extrapolated quantities can be expanded in a convergent power series in $1/r_{\text{ar}}$. In most cases it appears that these conditions are satisfied. However, in Sec. V C we show an example where at least one of them fails.

Having measured $\Psi_4^{\ell,m}(t, R)$, which is a function of coordinate time and coordinate radius, we can instead express Ψ_4 as a function of retarded time and areal radius: $\Psi_4^{\ell,m}(t_{\text{ret}}, r_{\text{ar}})$. The straightforward way to extrapolate to infinity is to fit $\Psi_4^{\ell,m}(t_{\text{ret}}, r_{\text{ar}})$ to a polynomial in $1/r_{\text{ar}}$ at a fixed value of t_{ret} , and then to evaluate the polynomial in the limit $1/r_{\text{ar}} \rightarrow 0$, thus obtaining $\Psi_4^{\ell,m}(t_{\text{ret}}, \infty)$.

Because $\Psi_4^{\ell,m}(t_{\text{ret}}, r_{\text{ar}})$ may be rapidly oscillating in t_{ret} , however, errors made in computing t_{ret} can lead to large errors in $\Psi_4^{\ell,m}(t_{\text{ret}}, r_{\text{ar}})$ and subsequently in the extrapolated value $\Psi_4^{\ell,m}(t_{\text{ret}}, \infty)$. For this reason, it would be better to extrapolate a function that is slowly varying in t_{ret} . For most modes of nonprecessing systems, a slowly varying representation is obtained by decomposing the complex quantity into amplitude and phase as

$$r_{\text{ar}} M \Psi_4^{\ell,m}(t_{\text{ret}}, r_{\text{ar}}) := A^{\ell,m}(t_{\text{ret}}, r_{\text{ar}}) e^{i\phi^{\ell,m}(t_{\text{ret}}, r_{\text{ar}})}, \quad (6)$$

where M is the sum of the initial Christodoulou masses of the two holes. We include the factor of M to make the amplitude dimensionless. (The use of Christodoulou mass is simply a conventional choice; we could have also used the ADM mass here.) For these modes, we extrapolate $A^{\ell,m}$ and $\phi^{\ell,m}$ rather than the real and imaginary components of $\Psi_4^{\ell,m}$, and then we reconstruct the extrapolated $\Psi_4^{\ell,m}(t_{\text{ret}}, \infty)$ from $A^{\ell,m}(t_{\text{ret}}, \infty)$ and $\phi^{\ell,m}(t_{\text{ret}}, \infty)$. For other modes, such as modes in which $\Psi_4^{\ell,m}$ is purely real ($m = 0$ modes in certain cases), or modes in which the amplitude $A^{\ell,m}$ passes through zero, the phase $\phi^{\ell,m}$ can be

discontinuous, ill-defined, or numerically difficult to determine. In these cases, the real and imaginary parts of $\Psi_4^{\ell,m}$ are extrapolated directly.

Similarly, it is possible to decompose the modes in a corotating frame [53], so that the modes show very little time dependence—and in particular, essentially no oscillations. Because they are slowly varying, the real and imaginary parts are extrapolated directly for all modes. This is the preferred method for precessing systems (although it can be applied to nonprecessing systems as well).

To find the form of the extrapolating functions, we consider standard expressions for the general formal radiative solution of the Einstein vacuum equations [54,55]. It turns out [56] that we can expect solutions to have finite-radius behavior in the form of expansions in λ/r_{ar} , where $\lambda = \lambda/2\pi$ is the typical (reduced) wavelength of a given mode. Because λ may be several hundred times the mass of the system, fitting to polynomials in $1/r$ would be numerically problematic—the fit coefficients for high-order terms would quickly become very large. Therefore we fit to polynomials in λ/r_{ar} , measuring λ from the frequency of the $(\ell, m) = (2, 2)$ mode. Note that the purpose of this correction is only to improve numerical behavior; fitting to $1/r_{\text{ar}}$ should produce the same answer modulo numerical issues.

To reiterate, our extrapolation of nonprecessing systems involves the following steps. First, the extracted waveform $\Psi_4^{\ell,m}(t, R)$ is reexpressed as $\Psi_4^{\ell,m}(t_{\text{ret}}, r_{\text{ar}})$. A set of retarded times $\{t_{\text{ret},i}\}$ is then constructed—the times at which we want the final extrapolated waveform. Next, for each time $t_{\text{ret},i}$, the waveforms at each radius are interpolated in retarded time to produce $\Psi_4^{\ell,m}(t_{\text{ret}}, r_{\text{ar}})$. At each time $t_{\text{ret},i}$, the reduced wavelength of the $(2, 2)$ mode is read off as $\lambda^{2,2} = 1/\dot{\phi}^{2,2}$, as measured on the outermost extraction sphere. The set of finite-radius waveforms is then fit to a polynomial in r_{ar} using

$$A^{\ell,m}(t_{\text{ret},i}, r_{\text{ar}}) \approx \sum_{k=0}^N A_{(k)}^{\ell,m}(t_{\text{ret},i}) \left(\frac{2\lambda^{2,2}}{mr_{\text{ar}}} \right)^k, \quad (7a)$$

$$\phi^{\ell,m}(t_{\text{ret},i}, r_{\text{ar}}) \approx \sum_{k=0}^N \phi_{(k)}^{\ell,m}(t_{\text{ret},i}) \left(\frac{2\lambda^{2,2}}{mr_{\text{ar}}} \right)^k, \quad (7b)$$

for oscillatory modes ($m \neq 0$), or

$$r_{\text{ar}} M \Psi_4^{\ell,0}(t_{\text{ret},i}, r_{\text{ar}}) \approx \sum_{k=0}^N \psi_{(k)}^{\ell,0}(t_{\text{ret},i}) \left(\frac{1}{r_{\text{ar}}} \right)^k, \quad (8)$$

for nonoscillatory modes ($m = 0$), where the ψ are complex fitting coefficients. The time-dependent $k = 0$ coefficients are then used as the amplitude and phase (or for $m = 0$, the real and imaginary parts) of the extrapolated waveform.

Our extrapolation of the precessing system follows the same steps, except that the finite-radius data are

transformed to the corotating frame [53] of the outermost extracted data, and modes with $m \neq 0$ are fit to polynomials of the form

$$r_{\text{ar}} M \Psi_4^{\ell,m}(t_{\text{ret},i}, r_{\text{ar}}) \approx \sum_{k=0}^N \psi_{(k)}^{\ell,m}(t_{\text{ret},i}) \left(\frac{1}{m\Omega r_{\text{ar}}} \right)^k, \quad (9)$$

where Ω is the angular velocity of the waveform [53] as measured on the outermost extraction sphere. Modes with $m = 0$ are again extrapolated using Eq. (8). The final result is then transformed back to the inertial frame.

In all cases, the choice of order of the extrapolating polynomial N is somewhat arbitrary. Early in the simulation, during the slow inspiral, λ/r_{ar} is typically relatively large, so higher-order terms may still be important. This means that the polynomial approximation will converge slowly, suggesting that higher N may be necessary. On the other hand, during the merger and ringdown, λ/r_{ar} will typically be quite small. In this case, we generally find that small N is sufficient; using large N simply overfits the data. In practice, the extrapolation procedure is never strictly convergent, because we have data at a finite number of extraction radii (typically about 20), and because these data inevitably contain some amount of truncation-level noise. This leads to extrapolating polynomials that converge for the first few orders, but eventually begin to diverge because of overfitting. The effect of the choice of N is discussed further in Sec. IV C.

When the wavelength of a given mode is comparable to or larger than the extraction radii, it is possible that the convergence of extrapolation will be adversely affected. In particular, the convergence for nonoscillatory modes ($m = 0$) tends to be slow because of their large wavelength (except possibly during the merger and ringdown). Even though we expand such modes in powers of $1/r_{\text{ar}}$ in Eq. (8) (rather than in powers of λ/r_{ar}), the coefficients in the expansion will accordingly be large. As previously mentioned, this can be numerically problematic and can limit the accuracy of the extrapolating fits. Indeed, we will see below that the quality of extrapolation is poor when $m = 0$.

B. Cauchy-characteristic GW extraction

Cauchy-characteristic extraction (CCE) is a method of computing gravitational radiation unambiguously and gauge invariantly at future null infinity [27,28,57]. This method is by construction immune to uncertainties associated with finite-radius and gauge effects. The essential idea is to couple a Cauchy evolution used to evolve the strong field region containing the black holes to a characteristic evolution evolving the far gravitational field (see Fig. 1). As opposed to the spatial hypersurface foliation in Cauchy evolutions, characteristic evolutions are based on null hypersurface foliations of spacetime. Without loss of accuracy, this allows one to apply a compactification of the radial coordinate to include infinity on the computational grid. Note that in CCE, the interface between Cauchy and

characteristic foliation is only a one-way boundary. Metric data are propagated from the Cauchy domain onto the characteristic domain, but not vice versa. The full two-way coupling is achieved by Cauchy-characteristic *matching* [28,57], which has been implemented in the linearized limit in Ref. [58].

1. Characteristic evolutions

We use the PITTNull characteristic code [27,28] to evolve the gravitational far field out to future null infinity. This code uses the framework established by Bondi and Sachs [27,28,59,60]. In this framework, the metric is written in the form

$$ds^2 = -(e^{2\beta}(1 + r\hat{W}) - r^2 h_{AB} U^A U^B) du^2 - 2e^{2\beta} du dr - 2r^2 h_{AB} U^B du dy^A + r^2 h_{AB} dy^A dy^B, \quad (10)$$

where $u = r - t$ is a retarded time coordinate, r is an areal radial coordinate, and y^A with $A = 2, 3$ are angular coordinates. The variables β , \hat{W} , U^A , and h_{AB} are free metric coefficients that must satisfy the Einstein equations. In addition, h_{AB} satisfies $h^{AB} h_{BC} = \delta_C^A$ and $\det(h_{AB}) = \det(q_{AB})$, where q_{AB} is the unit sphere metric. In the PITTNull code, angular components A are represented by means of complex spin-weighted scalars:

$$J \equiv q^A q^B h_{AB}, \quad U \equiv q_A U^A, \quad (11)$$

where q^A is a complex dyad satisfying $q^A = q^{AB} q_B$, $q^A q_B = 0$, and $q^A \bar{q}_A = 2$.

Recasting the Einstein equations in terms of the line element above results in a set of hypersurface equations, evolution equations, and constraint equations. The hypersurface and evolution equations are solved to determine the metric variables β , U , \hat{W} , and J between a world tube Γ at a radius R_Γ and future null infinity I^+ . To place future null infinity on the computational grid, a compactified radial coordinate $x(r) = r/(R_\Gamma + r)$ is introduced.

On the world tube Γ , inner boundary data in the form of the metric coefficients β , U , \hat{W} , and J must be supplied. Following the prescription of Ref. [57], these quantities are obtained via a transformation of metric data produced by the Cauchy evolution (see below). In addition, the metric variable J is required on the initial null hypersurface. Currently, there exists no solution for binary black hole initial data for the characteristic system. Instead, we impose a reasonable approximation: we use the value of J obtained from metric data on the initial Cauchy hypersurface at the world tube, and we smoothly blend J to zero on the initial null hypersurface so that $J|_{I^+} = 0$ [32]. Note that for conformally flat Cauchy initial data, this corresponds to $J = 0$ everywhere on the initial null hypersurface (see Ref. [61] for a discussion). We have also tried setting $J = 0$ on the initial null hypersurface for a case with nonconformally flat Cauchy initial data (Case 4 in Table I, described below). We find that, at least in this case,

it makes no significant difference to any of the results whether $J = 0$ or J is smoothly blended to $J|_{I^+} = 0$.

This choice of characteristic initial data will in general be inconsistent with the Cauchy initial data: the time evolution of Cauchy initial data in the region $R > R_\Gamma$ yields a solution on the outgoing initial null hypersurface (see Fig. 1), and this solution does not generally agree with the supplied characteristic initial data there. In the error analysis in Sec. IV E, we refer to the associated waveform uncertainty as the ‘‘CCE initial-data error.’’

The characteristic equations are solved on a finite difference grid consisting of N_x radial points that discretize the compactified radial direction. For each radial point, the angular discretization of S^2 consists of two overlapping stereographic patches. Each patch contains N_{ang} points per angular direction. The two angular patches use circular boundaries to eliminate noise from patch corners [62].

As detailed in Ref. [32], the radial and time directions are evolved using second-order finite difference derivatives together with a second-order null-parallelogram integration algorithm (see Ref. [30] for a new full fourth-order algorithm with spectral angular derivatives). The angular derivative operators are discretized using fourth-order finite differences. Interpatch boundary data are obtained via fourth-order interpolation.

2. World tube boundary data

We obtain boundary data from 3 + 1 Cauchy metric data as described analytically in Ref. [57] and as implemented in Refs. [32,48]. We define the world tube Γ as a time succession of spheres of constant coordinate radius $R_\Gamma = \sqrt{x^2 + y^2 + z^2}$, with surface normal s^α (see Fig. 1). On Γ , we construct outgoing null rays ℓ^α that induce the null foliation. As detailed in Refs. [32,48,57], the transformation from Cauchy to characteristic metric data requires two steps. The first step involves transformation of the 4-metric from a Cartesian to an affine null coordinate system. The second step involves transformation of the affine 4-metric to the characteristic Bondi coordinate system (u, r, y^A) . The intermediate transformation step to the affine coordinate system is necessary since the areal radius of the Bondi coordinates can only be computed once angular metric components are known.

The characteristic code requires Cauchy metric data in the form of the spherical harmonic modes of the 3-metric g_{ij} , lapse α , shift β^i , and their radial and time derivatives. In the evolutions we have performed using SpEC (see Sec. III below), we decompose the required quantities into modes up to $\ell = 16$.

3. Wave extraction in Bondi gauge

We extract waveforms at I^+ using the methods described in Refs. [27,62]. The original wave extraction method of Ref. [27] computes the gravitational news

function \mathcal{N} , which in Bondi gauge is related to the metric component J by $\mathcal{N} = J_{,ur}$. An alternative and independent method computes the Weyl scalar Ψ_4 [62], which is related to the news function by $\Psi_4 = \mathcal{N}_{,u}$. Note that this last relation is not used in the characteristic code; the two quantities Ψ_4 and \mathcal{N} are computed independently. It is also possible to directly extract the strain h at I^+ as well, and this could potentially remove the problems associated with the time integration of Ψ_4 or \mathcal{N} [63]. An algorithm to accomplish this has recently been implemented [64], but is not used here.

During a simulation, the gauge at I^+ is induced by the boundary data at the world tube, and the assumption of Bondi gauge generally does not hold. As detailed in Refs. [27,62], it is necessary to apply a transformation from the induced gauge, denoted by coordinates (u, r, y^A) , to Bondi gauge, denoted by coordinates (u_B, r_B, y_B^A) . The code presented in Refs. [27,62] applies the relevant transformation to spatial Bondi gauge (r_B, y_B^A) and computes the Bondi time $u_B(u, y_B^A)$ as a function of coordinate time u and angular coordinates y_B^A . In a final step, it is necessary to make the transformation $u \rightarrow u_B$ to constant Bondi time $u_B(y_B^A) = \text{const}$ by means of time interpolation at each point on the sphere at I^+ [48].

4. Convergence order

The characteristic evolution algorithm of the PITTNull code is expected to exhibit at least second-order convergence (see, e.g., Ref. [65] for tests with linearized solutions). In combination with the algorithm for obtaining world tube boundary data from a Cauchy evolution (Sec. II B 2), however, we observe first-order convergence in certain quantities [32,48]. This may be due to a term at the world tube that is only known to first order. In addition, the numerical algorithm for evaluating Ψ_4 at I^+ (Sec. II B 3) involves a large number of terms, some of them including one-sided finite difference derivatives. As noted in Ref. [32], the convergence order may be negatively affected by this, particularly for quantities measured at I^+ .

III. BINARY BLACK HOLE SIMULATIONS

In this section we describe the binary black hole (BBH) simulations that we use for comparing wave extraction

techniques. All simulations were performed using SpEC [13] described in Refs. [14–16,66,67] and references therein. This code evolves a first-order representation [68] of the generalized harmonic system [69–71] with constraint damping [68,71,72]. Outgoing-wave boundary conditions [17,68,73] designed to preserve the constraints [74–80] are imposed at the outer boundary. Interdomain boundary conditions are enforced with a penalty method [81,82].

We consider four simulations, which are listed in Table I. The first two are equal-mass nonspinning binary simulations that have identical initial data but different gauge conditions; these are used to test the gauge dependence of the two wave-extraction methods in Sec. V C. Case 1 is described in Ref. [18], and Case 2 is the $q = 1$ run discussed in Refs. [11,15]. Case 3 is a BBH with no spin but with a mass ratio of 6, and is the $q = 6$ run discussed in Refs. [11,15]. Case 4 is a generic, precessing BBH with a mass ratio of 3, and spins on both holes in generic directions; this simulation is new and has not been presented elsewhere.

In the generalized harmonic system, the gauge is chosen by freely specifying four gauge source functions H_a . The simulations in Table I utilize several different gauge choices. For Case 1, H_a is fixed (F) in the corotating frame during inspiral and smoothly transitions to a solution of an auxiliary wave equation (W) of the form $\nabla^c \nabla_c H_a = \dots$ during the plunge and ringdown. The gauge used in Case 1 is described in detail in Ref. [18]. Case 2 begins with the same fixed gauge as Case 1, but transitions smoothly to the harmonic (H) gauge $H_a = 0$ very quickly (after about $t \sim 40M$) and remains in the harmonic gauge throughout the inspiral. It then transitions to the damped harmonic (DH) gauge [14,83,84] of Ref. [14] before the merger, and maintains the DH gauge through the merger and ringdown. Case 3 uses the same fixed gauge as Case 1 during the inspiral, and transitions to the damped harmonic gauge about 1.5 orbits before the merger. Case 4 uses the fixed gauge for only the first $t \sim 40M$ of the inspiral, and transitions directly to the damped harmonic gauge for the remainder of the simulation.

The simulations in Table I employ two different methods of constructing the initial data. For the nonspinning cases, we use conformally flat (CF) data, as described in Ref. [50]. For the spinning, precessing case we use superposed Kerr-Schild (SKS) data [85]. Both of these methods

TABLE I. Parameters of BBH runs. Columns indicate mass ratio q , dimensionless spins χ_1, χ_2 , type of initial data, gauge conditions, number of orbits before the merger, initial orbital eccentricity (ECC), and the initial gravitational-wave frequency $M\omega_{\text{ini}}$ of the (2, 2) mode.

Case	q	χ_1	χ_2	ID	Gauge	Orbits	ECC	$M\omega_{\text{ini}}$
1	1	0	0	CF	F \rightarrow W	16	5×10^{-5}	0.034
2	1	0	0	CF	F \rightarrow H \rightarrow DH	16	5×10^{-5}	0.034
3	6	0	0	CF	F \rightarrow DH	22	4×10^{-5}	0.038
4	3	$(0.7, 0, 0.7)/\sqrt{2}$	$(-0.3, 0.3, 0)/\sqrt{2}$	SKS	F \rightarrow DH	26	1×10^{-3}	0.032

can produce astrophysically relevant initial data, but the superposed Kerr-Schild method is more flexible and (for example) allows the construction of initial data with higher spins [66,85,86].

For all cases in Table I, the initial orbital parameters are adjusted via the iterative method of Refs. [49,87] so as to reduce the orbital eccentricity of the binary. In addition, all of the simulations were done at multiple numerical resolutions in order to provide a means of estimating Cauchy error.

IV. ESTIMATING ERRORS IN WAVEFORMS

A main goal of this paper is to estimate the gauge-related error in extrapolated waveforms by comparing to CCE waveforms, which are gauge invariant. For this comparison to be meaningful, we must first estimate the other sources of error in the numerical waveforms.

We first consider the numerical truncation error of the Cauchy simulation (“Cauchy error”); this contributes to both extrapolated and CCE waveforms. For waveforms extrapolated to infinity, we also estimate the uncertainty introduced by the extrapolation procedure. For CCE waveforms, we estimate two sources of error in addition to Cauchy error: the numerical truncation error of the characteristic evolution and the error associated with the location of the CCE extraction world tube. The latter error is due to incompatibility of the Cauchy solution and the data chosen on the initial null hypersurface of the characteristic code.

We do not estimate the error associated with imperfect outer-boundary conditions in the Cauchy simulation. This error has previously been estimated [15,18] by comparing otherwise-identical Cauchy simulations with the outer boundary placed at different locations; this outer-boundary error was found to be comparable to or smaller than the Cauchy error.

For most of this section, we concentrate on errors in the amplitude and phase of the waveform, as these are the errors most often quoted by the numerical relativity community. However, in some cases the phase of a waveform can become ill-defined. Therefore, in Sec. IV F we consider alternative error measures.

A. Waveform alignment

Our error estimates for a given (complex) waveform $\psi(t)$ are obtained by computing the difference between two versions of that waveform, $\psi_A(t)$ and $\psi_B(t)$, that are generated by slightly different methods (for instance, extrapolation versus CCE, or two different numerical resolutions). In matched filtering, the procedure for comparing a signal waveform against a template waveform includes a global time and phase shift of the template in order to best match the signal. These shifts effectively account for the arrival time of the signal and the orbital phase at that time. Therefore, when computing the

difference between two waveforms $\psi_A(t)$ and $\psi_B(t)$ that might ultimately be used as templates, it is appropriate to likewise introduce a global time and phase shift between $\psi_A(t)$ and $\psi_B(t)$, which are chosen to minimize some measure of the difference between the waveforms. This procedure is referred to as waveform alignment.

Waveform alignment in matched filtering is done implicitly by Fourier transforming and working in the frequency domain. The measure of comparison is typically an overlap integral that includes the noise spectrum of the detector [88,89]. The integral and alignment may be done simultaneously by inverse Fourier transforming the integrand, taking the absolute value, and finding the maximum value as a function of time. In this paper we choose instead to work in the time domain, and we do not include noise from a specific detector.

For nonprecessing systems, we use an alignment procedure described in Ref. [90], in which $\psi_A(t)$ is given a time shift Δt and a phase shift $\Delta\Phi$ that are chosen to minimize

$$\Xi(\Delta t, \Delta\Phi) := \int_{t_1}^{t_2} (\phi_A(t) - \phi_B(t + \Delta t) - 2\Delta\Phi)^2 dt, \quad (12)$$

where the waveform phases are defined according to Eq. (6). We choose the range $[t_1, t_2]$ to be early in the inspiral, but late enough to avoid the junk radiation, and wide enough to average over numerical noise ($t_2 - t_1 > 700M$, where M is the sum of the Christodoulou masses of the two holes). We determine the phase and time offsets $\Delta\Phi$ and Δt by matching only the $(\ell, m) = (2, 2)$ spin-weighted harmonic modes of $\psi_A(t)$ and $\psi_B(t)$; we then use the same Δt and $\Delta\Phi$ (the latter scaled by m for each mode) to shift all other spin-weighted harmonic modes (ℓ, m) .

This method is a special case of the more general one needed for precessing systems. For precessing systems, the alignment must apply an arbitrary rotation rather than the simple one shown above [53,91–95]. The waveform modes transform just as ordinary spherical harmonics do under rotations, by application of the Wigner \mathcal{D} matrices.¹ Reference [53] describes the method we use for achieving this alignment in the precessing case. Essentially, the corotating frame of each waveform is found. Because these frames are physically and geometrically meaningful measures of the waveform, it is meaningful to compare them. We can define a phase difference between the two frames using the logarithms of their orientations, which are represented as unit quaternions \mathbf{R}_A and \mathbf{R}_B . This phase difference is inserted into an expression that is the appropriate

¹In the case of nonprecessing systems, the symmetry allows us to pick out a preferred direction: the axis of rotation, which we choose to coincide with the z axis. For a rotation about the z axis through an angle γ , the Wigner matrices simplify to $\mathcal{D}_{m',m}^{(\ell)} = \exp[im\gamma]\delta_{m',m}$, which is why we simply multiply the modes by $\exp[im\Delta\phi/2]$ in the nonprecessing case.

generalization of Eq. (12) to full three-dimensional rotations:

$$\Xi(\Delta t, \mathbf{R}_\Delta) := \int_{t_1}^{t_2} 4|\log[\mathbf{R}_A(t)\bar{\mathbf{R}}_B(t+\Delta t)\bar{\mathbf{R}}_\Delta]|^2 dt, \quad (13)$$

which is then minimized over Δt and all three degrees of freedom in the unit quaternion \mathbf{R}_Δ . Once the optimum rotation is found, it is applied to the waveform $\psi_A(t)$.

For some purposes, alignment need not be done at all when estimating errors. For example, when estimating extrapolation error by subtracting waveforms of different extrapolation orders, alignment is not strictly necessary because all finite-radius waveforms have already been shifted by r_* when expressing them as functions of retarded time. However, our goal is to compare extrapolated and CCE waveforms, and these cannot be compared directly without alignment. This is because the extrapolated waveforms are shifted (in retarded time) by some r_* , whereas the CCE waveforms are shifted by a different offset that depends on the radius of the CCE world tube. Therefore, for consistency we estimate every source of error using the same waveform alignment procedure that is used to compare CCE with extrapolated waveforms.

Because small time shifts can lead to large accumulated phase differences, especially for the relatively long waveforms that we consider, it is important that the alignment procedure be robust. For example, we find that aligning waveforms at peak amplitude is sensitive to small amounts of noise in the waveforms. For the procedure we use, we have verified that small changes in the alignment window $[t_1, t_2]$ do not affect the results. Furthermore, we have repeated all of the analyses in this paper with an alignment window near the merger, $[t_{\text{merger}} - 450M, t_{\text{merger}} + 50M]$, instead of in the early inspiral. We find that although this changes the shapes of error-versus-time plots, the main results of this paper (namely, the relative magnitudes of different sources of error) are not affected.

B. Cauchy error

To estimate the waveform uncertainties associated with the numerical truncation error in the Cauchy simulations, we use waveforms computed at different numerical resolutions. Each case in Table I was evolved at three different resolutions (not necessarily the same in different cases), which we refer to as *low*, *medium*, and *high* resolution.

For a sufficiently fast convergence rate (we expect exponential convergence for spectral simulations of smooth problems), the difference between the waveforms at low and medium resolution is a good estimate for the low-resolution Cauchy error, while the difference between the medium- and high-resolution waveforms is a good estimate for the medium-resolution Cauchy error. We prefer to err on the side of caution, so we use the difference between the medium- and high-resolution waveforms as an estimate for the high-resolution Cauchy error.




FIG. 2 (color online). Phase differences in the extrapolated $\Psi_4^{2,2}$ mode between successive Cauchy resolutions for simulation 2 of Table I, using extrapolation order $N = 5$. The waveforms at different resolutions have been aligned over the interval $[1000M, 2000M]$. The maximum amplitude occurs at $t_{\text{ret}} \approx 3952M$, shown here as the dotted vertical line.

Figure 2 shows phase differences between the $\Psi_4^{2,2}$ modes from Case 2 of Table I at different Cauchy resolutions. For each resolution, $\Psi_4^{2,2}$ has been extrapolated to infinity using $N = 5$ in Eq. (7b). The waveforms for different resolutions are aligned early in the inspiral before taking differences. These phase differences represent the estimated Cauchy error in the medium- and high-resolution extrapolated (2, 2) modes. Relative amplitude differences between Cauchy resolutions show similar convergence. We compute the Cauchy error for each (ℓ, m) mode in the extrapolated and CCE waveforms in an analogous way.

C. Extrapolation fit error

The extrapolation fit error is the uncertainty in the extrapolated waveform $\Psi_4^{\ell,m}(t_{\text{ret}})$ computed by the procedure of Sec. II A 3, given $\Psi_4^{\ell,m}(t, R)$ on extraction spheres of several radii R . Recall that this procedure involves fitting the modulus and argument of $r_{\text{ar}} M \Psi_4^{\ell,m}(t_{\text{ret}}, r_{\text{ar}})$ to N th-order polynomials in λ/r_{ar} (where λ is the reduced wavelength), and that the extrapolated result is the coefficient of the constant term in the polynomial.

There are several ways one might seek to estimate this error, although we are not aware of a method that can provide a rigorous estimate. One approach is to compare waveforms extrapolated using different polynomial orders N . Phase differences from such a comparison are shown in Fig. 3, where

$$\Delta\phi^{2,2} := \phi_N^{2,2} - \phi_{N-1}^{2,2} \quad (14)$$

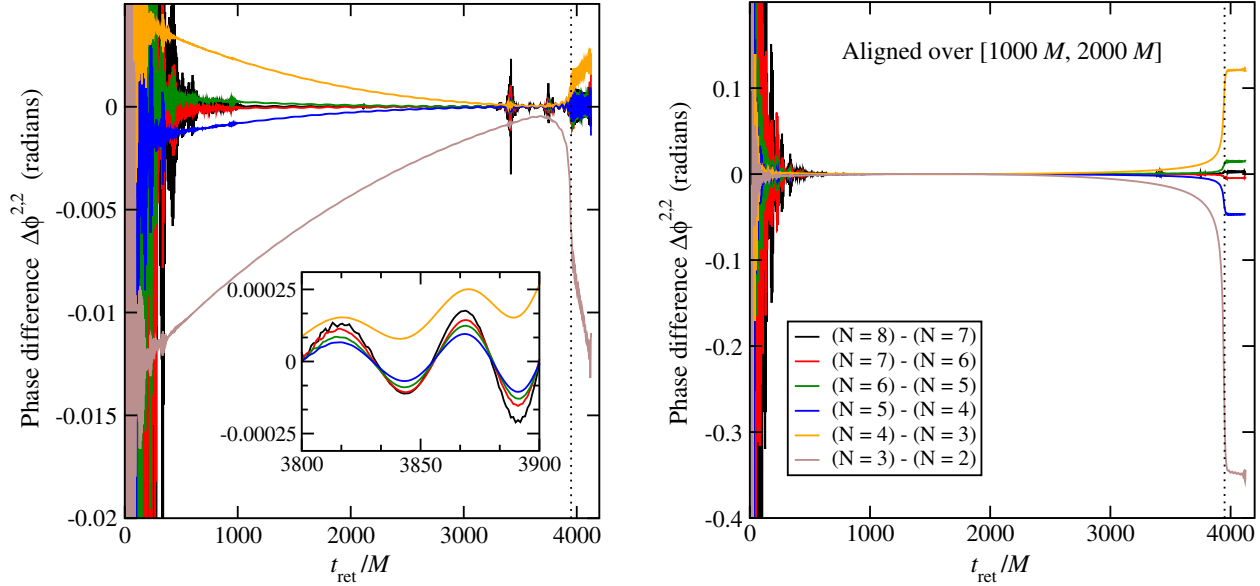


FIG. 3 (color online). Phase differences in the high-resolution, extrapolated $\Psi_4^{2,2}$ waveform between different extrapolation orders N , for simulation 2 of Table I. In the left panel, no alignment has been done; in the right, each pair of waveforms has been aligned over $[1000M, 2000M]$. Each curve is the phase difference between a waveform with order $N + 1$ and an otherwise identical waveform with order N . The maximum amplitude occurs at $t_{\text{ret}} \approx 3952M$, shown in each plot as a vertical dotted line. Note the difference in vertical scales. Because the frequency is greatest near the merger, small time shifts in the alignment window produce large phase differences in the plot on the right. The two noisy regions in the left panel between $t_{\text{ret}} \approx 3400$ and 3700 correspond to gauge or grid changes in the Cauchy simulation.

is the error estimate for extrapolation of order N . The left panel of this figure shows the comparisons without alignment; the right panel shows the comparisons after aligning the waveforms in the early inspiral. Although the phase errors generally decrease with N , the amount of noise in the extrapolated waveforms increases with N . The noise is largest in the first few hundred M , during the junk-radiation phase, and near the merger at times corresponding to grid or gauge changes in the Cauchy simulation.

In the nonaligned case (left panel), the phase differences are well described as constant multiples of λ^N during the inspiral. This is presumably due to near-field effects [56] and is the reason for our choice of λ/r as the extrapolation variable. There is very little work for extrapolation to do near the merger, when λ becomes comparatively small. In fact, as shown in the inset, the differences grow slightly with an increasing order of extrapolation. Presumably, this is because the higher-order polynomial coefficients are fitting to noise in the data when there are no significant physical features present.

Note the very different vertical scales in the two panels. The large phase differences near the merger in the right panel are a result of aligning waveforms in the early inspiral. Alignment introduces time offsets between waveforms, which are necessary to make the phase and frequency agree as much as possible in the alignment window. But even small time offsets in the early inspiral can result in large phase differences near the merger, because the frequency is large there.

It may seem that alignment unfairly inflates the estimate of the extrapolation fit error, but the relevant error for many applications is the one that includes alignment. For example, if we were to attach the numerical waveform to an analytic waveform for hybridization, we would have to do so at a time when both waveforms are valid—presumably during the early inspiral. The relevant uncertainty in the numerical waveform for this situation is the one computed with alignment in the hybridization region.

As an alternative measure of the uncertainty, one might consider the variance σ_N of extrapolation at order N , as inferred from the least-squares fit to the data. In a classical model, with unbiased and uncorrelated errors, the variance gives the standard error in the fit coefficients. But here we have no reason to assume that the errors are unbiased and uncorrelated. If we simply assign equal, arbitrary errors to the input waveforms, then even in the best case, this leaves the overall scale of σ_N arbitrary (although it would at least provide some relative measure of goodness of fit).

Yet another approach is to obtain an error estimate by Richardson extrapolation instead of simply comparing neighboring values of N . The idea is to first estimate the waveform that one would obtain with $N \rightarrow \infty$, and then construct the error for order N by subtracting the order- N waveform from the order- ∞ waveform. This approach, and to a lesser extent the approach used in Fig. 3, assumes that the extrapolated waveform converges as $N \rightarrow \infty$. However, the extrapolation series usually begins to diverge at some

order (which is time dependent), as shown in the inset of the left panel of Fig. 3. We can take the difference between two orders as some kind of measure, but we cannot justify a rigorous error bound because of the lack of convergence with extrapolation order N .

The above considerations indicate a need for further investigation into the complicated issue of extrapolation fit error. For now, we defer these issues to a future work, and we henceforth choose the simplest approach of estimating extrapolation fit errors: taking the difference between two waveforms of successive extrapolation orders.

Finally, there is the question of how to choose the value of N when constructing the nominal extrapolated waveform. One must balance the desire for small error in smooth regions (such as in the left panel of Fig. 3) with the desire for low noise. For concreteness in this paper we choose $N = 5$, and we estimate the error as the difference between the $N = 5$ and $N = 4$ waveforms. An alternative method would be to vary the extrapolation order N as the simulation progresses, choosing a large value of N during the smooth inspiral, and a smaller value of N to reduce the noise in the merger and ringdown. We do not consider this refinement here.

D. CCE truncation error

The waveform uncertainty associated with the numerical truncation error on the characteristic grid can be estimated by considering a sequence of CCE resolutions, which we label $r0$, $r1$, $r2$. The actual time step size Δu , the number of radial points N_x , and the number of angular points N_{ang} for each of these resolutions are listed in Table II.

Let $\phi_k(t_{\text{ret}})$ denote the phase of a CCE waveform computed with resolution k , and let $\Delta\phi_{k,k+1} := |\phi_k - \phi_{k+1}|$ be the phase difference between the waveforms from different resolutions. If we measure the convergence rate of $\Delta\phi_{k,k+1}$ with increasing k and find a consistent convergence order, then we can use Richardson extrapolation to estimate the error in the highest resolution (see, e.g., Ref. [96]).

The top panel of Fig. 4 shows the phase differences $\Delta\phi_{k,k+1}$ for $\Psi_4^{2,2}$ CCE waveforms from simulation 2 in Table I. To estimate the convergence order, we assume that the phase obeys

$$\phi(h) = \phi(0) + \mathcal{O}(h^n), \quad (15)$$

TABLE II. Resolution of the characteristic grid.

Resolution	Δu [M]	N_x	N_{ang}
$r0$	0.37500	101	41
$r1$	0.25000	151	61
$r2$	0.18750	201	81




FIG. 4 (color online). Top: Phase differences between $\Psi_4^{2,2}$ near the merger, computed using different CCE resolutions (as labeled in Table II), for simulation 2 in Table I. Bottom: Ratio of the phase differences from the top panel. We find roughly first-order convergence, i.e. a ratio of about 0.5. All waveforms use the same high-resolution Cauchy data and a world tube radius of $R = 385M$. Waveforms are aligned in the interval $[1000M, 2000M]$. The maximum amplitude occurs at $t_{\text{ret}} \approx 3952M$, denoted here by the vertical dotted line.

where h represents the grid spacing and n the convergence order. Note that because time, radial, and angular resolutions are all refined by the same factor between successive resolutions, we can use a single measure h here. We then compute

$$\frac{\Delta\phi_{12}}{\Delta\phi_{01}} = \frac{h_1^n - h_2^n}{h_0^n - h_1^n} + \mathcal{O}(h_0^{n+1}), \quad (16)$$

where h_0 , h_1 , and h_2 represent the grid spacings in resolutions $r0$, $r1$, and $r2$, respectively. For the values shown in Table II, we expect $\Delta\phi_{12}/\Delta\phi_{01} = 0.5$ for first-order convergence, and $\Delta\phi_{12}/\Delta\phi_{01} = 0.35$ for second-order convergence.

As shown in the bottom panel of Fig. 4, the ratio of phase differences is roughly consistent with first-order convergence. We note here that without any alignment of the waveforms, the phase convergence of the CCE data is very cleanly first order. Aligning early in the inspiral renders the phase convergence somewhat less uniform. Doing the same for the amplitude error, we find good second-order convergence (independent of alignment). These measured convergence orders are consistent with the theoretically expected convergence discussed in Sec. II B 4.

Assuming first-order convergence in phase, we estimate the phase error in the $r2$ waveform using Richardson extrapolation to be $3 \cdot \Delta\phi_{12}$. Similarly, assuming second-order

convergence, we estimate the (relative) amplitude error to be $9/7 \cdot \Delta A_{12}$. Henceforth, we use resolution $r2$ as the nominal CCE waveform.

E. CCE initial-data error

Waveforms evolved using CCE may depend on the location of the characteristic world tube (the surface labeled R_F in Fig. 1) for two reasons. Most significantly, the characteristic evolution requires data on an initial null hypersurface (the surface labeled u_0 in Fig. 1). In the simulations we consider, these data are chosen to be blended to conformally flat as described in Sec. II B. However, this does not necessarily agree with the Cauchy evolution, which may contain physical backscattered radiation, junk radiation, and ingoing radiation from imperfect outer boundary conditions. This incompatibility is a source of uncertainty in the CCE waveforms. As the radius of the world tube is increased, this mismatch and the resulting error should decrease.

Another reason one might expect a CCE waveform to depend on the world tube location is that the length scale of dynamical features in the spacetime decreases as the world tube is moved closer to the source. Unless there is a corresponding increase in the resolution of the characteristic code, one would therefore expect a smaller world tube radius to result in larger truncation errors. However, we find this contribution to the overall error to be insignificant; the estimated error is essentially independent of the characteristic-code resolution.

Because of these observations, we refer to this error as the ‘‘CCE initial-data error,’’ even though we measure it by varying the finite-radius world tube location. One method for estimating this error is simply to take the difference between waveforms computed using two different world tube radii. This approach is inadequate because it depends too heavily on which radii are chosen. If the two radii are very near to each other, then this would result in an arbitrarily small estimate. On the other hand, if the two radii were very far apart, this method might yield an incorrectly large estimate of the error.

For the high-resolution run of simulation 2 in Table I, we have computed CCE waveforms from 28 different world tube radii (ranging from $R = 77.5M$ to $R = 385M$). We calculate the phase difference between the waveform from each radius and the waveform from the outermost radius, where the two waveforms are aligned over $[1000M, 2000M]$. Figure 5 shows these phase differences at a particular time ($t_{\text{ret}} \approx 2600M$), plotted against the inverse world tube radius $1/R$. Note that the outermost world tube radius has a phase difference of zero in this plot by definition. It is immediately evident that the phase differences decrease predominantly like $1/R^2$ as R increases. We find this same feature at other times and for relative amplitude differences as well as phase differences.




FIG. 5 (color online). Phase differences as a function of inverse world tube radius R . Each of the 28 circles (both open and closed) is the phase difference, evaluated at $t_{\text{ret}} \approx 2600M$, between the CCE waveform $\Psi_4^{2,2}$ computed with world tube radius $385M$ and the same waveform computed with world tube radius R . Solid curves are polynomial fits of different orders in powers of $1/R$. The four closed circles represent the typical world tube radii used in most of the simulations, and the dashed curve shows the second-order fit to just these four points. All waveforms are from the high-resolution run of simulation 2 in Table I, using CCE resolution $r2$. Waveform alignment is done using the interval $[1000M, 2000M]$.

We can estimate the CCE initial-data error at each time by fitting such phase differences to a polynomial in $1/R$ and then extrapolating $1/R \rightarrow 0$. The solid curves in Fig. 5 show these fits for polynomials of different orders. We see that this extrapolation diverges as the polynomial order is increased. This is presumably the same issue (overfitting to noisy data) that arises in waveform extrapolation, as discussed in Sec. IV C. Since we are interested here only in an estimate of the CCE finite-radius error and not in extrapolating the CCE waveforms to infinite world tube radius, we simply choose a quadratic fit.

In the above procedure for estimating the CCE initial-data error, we extrapolate CCE phase differences (such as those shown in Fig. 5) to infinity. One may ask why we do not instead extrapolate these phase differences to the outer boundary of the Cauchy simulation. After all, placing the world tube at the outer boundary would seem to eliminate any mismatch between characteristic and Cauchy initial data. But imagine a perfect Cauchy code with infinite resolution, and with perfect outer boundary conditions so that even with a finite outer boundary, it exactly reproduces the true solution of Einstein’s equations including effects such as backscatter and tails. If the CCE world tube were

placed at the outer boundary of this perfect Cauchy domain, then there would still be a mismatch between the (blended to conformally flat) characteristic initial data and the true solution. Extrapolating phase differences to infinity estimates the error induced by this mismatch.

It is important to verify that the procedure for estimating CCE initial-data error works well when using fewer world tube radii, because most of the runs we consider have CCE data from only four radii. The dashed black line in Fig. 5 shows the second-order fit using only the four radii $R = 75, 100, 244, 385M$ (the four solid black dots in the figure). As can be seen in the figure, this fit is quite consistent with the fit using all 28 radii. We find this to be the case at other times (not only at the time shown in the figure) and for (relative) amplitude differences as well. We therefore use this quadratic extrapolation procedure to estimate the CCE initial-data error, and we use the waveform computed from the outermost world tube as the nominal CCE waveform.

F. Alternatives to measuring phase error

In the previous sections we estimated errors by computing phase and amplitude differences between otherwise-identical waveforms computed using different resolutions, world tube radii, or extrapolation orders. However, phase errors are not always well-defined. In this section we illustrate some of the ways in which phase errors can become difficult to measure, and in Sec. IV F 2 we describe another error measure that obviates this difficulty.

1. Problems with phase differences

The phase of a $\Psi_4^{\ell,m}$ mode may become ill-defined because the amplitude momentarily vanishes, or it may simply vary rapidly as the amplitude passes near zero [93]. The imaginary part of the waveform may be zero analytically, but at truncation level numerically. This can cause the phase to change randomly and discontinuously between ϵ and $2\pi - \epsilon$, depending on the numerical errors. Issues like these can cause trouble even when the waveform at I^+ has a well-defined phase, because the waveforms computed from (some of) the world tube or extraction radii may exhibit such problems.

For example, in the precessing case (simulation 4 in Table I), we find that phase differences between otherwise-identical CCE waveforms computed from different world tube radii sometimes jump by 2π (similar examples can be found in most cases). Such jumps enter into the CCE initial-data error estimate, as described in Sec. IV E, where they can lead to estimated phase uncertainties of $\mathcal{O}(\pi)$. This renders the error estimate much less meaningful (although it is at least consistent).

It is interesting to examine the real and imaginary parts of the waveform from different world tube radii near the time of such a jump. Figure 6 shows an example of this. We plot $\Psi_4^{3,2}$ in the complex plane for times corresponding to an observed 2π jump in phase differences. Evidently, the




FIG. 6 (color online). The CCE $rM\Psi_4^{3,2}$ modes computed from four different world tube radii, plotted in the complex plane for times in the approximate interval $[5495M, 5565M]$, for Case 4 in Table I. The trajectory of each waveform is traversed in a clockwise direction with time, entering on the right and leaving on the bottom. The maximum amplitude [of the $(2, 2)$ mode] occurs at $t_{\text{ret}} \approx 6722M$. The waveform from each radius has been aligned over $[1000M, 2000M]$ with the waveform from the largest radius ($455M$). A jump of 2π in the phase differences between the waveforms occurs because only the waveform from $R = 100M$ encircles the origin (shown as + in the figure).

jump corresponds to the trajectories of the waveforms from different radii encircling the origin a different number of times. This occurs even though the trajectories shown in Fig. 6 are clearly converging to a nonvanishing amplitude as the world tube radius is increased.

Similar problems can occur for extrapolation, but they can be even worse because the nearly discontinuous jumps in phase end up not only in the error estimate (as in the case of CCE) but in the extrapolated waveform itself. Even in mildly precessing cases, one may thus encounter nonextrapolatable waveforms—at least with the naive extrapolation algorithm, which extrapolates phase and amplitude separately according to Eqs. (7a) and (7b). We solve this problem for the precessing system by first transforming the data at all radii to a common corotating frame (the corotating frame of the outermost extraction radius) before extrapolation [53], as described in Sec. II A 3.

The corotating frame method also gives rise to another way to measure phase error, because the phase information is almost entirely recorded in the orientation of the corotating frame. The phase difference between the two frames is described completely² by the logarithm of the ratio of the

²See Eq. (19) and surrounding discussion in Ref. [53].

two orientations, as in the integrand of Eq. (13). This difference is not subject to the sudden phase jumps seen above, and it is invariant under overall rotations of the physical system or the coordinate system. This provides a robust and uniform method that can be used in nonprecessing and precessing systems alike. However, this definition of phase error applies to an entire waveform including all (ℓ, m) modes. We prefer to use an error quantity that can be defined separately for each mode, as described below.

2. Error measure in the complex plane

Motivated by the difficulty of defining phase errors in some generic BBH simulations, we employ an alternative error measure that is an L^2 norm of the difference between two (complex) waveforms, integrated over all positions on the sky:

$$\|\Psi_4^A - \Psi_4^B\|^2 = \int_{S^2} |\Psi_4^A - \Psi_4^B|^2 d\Omega. \quad (17)$$

Expanding each waveform in spin-weighted spherical harmonics using Eq (4), and using orthonormality relations, one obtains

$$\|\Psi_4^A - \Psi_4^B\|^2 = \sum_{\ell, m} |\Psi_4^{\ell, mA} - \Psi_4^{\ell, mB}|^2. \quad (18)$$

This quantity could be normalized by a norm of the individual waveforms (computed using the same measure), such as $\|\Psi_4^A\| + \|\Psi_4^B\|$. However, if one is interested in comparing errors in a particular spin-weighted harmonic mode, then the normalization (which is the same for each mode) can be neglected. In this case, amplitude and phase errors are combined into a single measure,

$$\begin{aligned} \Delta_{\ell m}^2 &= |\Psi_4^{\ell, mA} - \Psi_4^{\ell, mB}|^2 \\ &= (\Delta A^{\ell, m})^2 + 2A^{\ell, mA}A^{\ell, mB}(1 - \cos \Delta\phi^{\ell, m}), \end{aligned} \quad (19)$$

which has the advantage of being immune to ill-defined phase errors, as well as properly ignoring phase differences when amplitudes are small. It also provides the option of combining all errors for a mode-independent measure. The sum in Eq. (18) is invariant under overall rotations of both waveforms, making this a particularly useful measure in precessing systems (this is true even when considering a single value of ℓ). Using this measure, we can estimate the various sources of error in the same way as we did above for phase and amplitude errors.

G. Combination of errors

In the preceding discussion, we concentrated on computing various error quantities: Cauchy error, extrapolation fit error, CCE truncation error, and CCE initial-data error. In this section we discuss how to combine these quantities into a single error bar. Here we still consider each (ℓ, m) mode separately.

In addition to constructing a combined total error bar for a waveform, we would also like to compare the relative magnitudes of the different sources of error. The above error measures are all time dependent, so either we must compare them at each value of t or we must construct a time-averaged error measure. We choose the latter and average the absolute value of each error over an interval $[t_1, t_2]$, where t_1 represents a time after junk radiation (usually about $500M$), and t_2 represents the time after the merger when the amplitude of the waveform has decayed to the truncation level. The early-time and late-time cutoffs avoid portions of the waveform where the phase is ill-conditioned and difficult to measure, or where the waveform is unphysical. The relative magnitudes of these time-averaged errors then allow us to see at a glance how the different sources of error compare.

1. Error bar for an individual waveform

To determine the uncertainty in an individual waveform, we combine the various sources of error using an L_1 norm—i.e., we add the absolute values of each source of error. For independent, normally distributed random errors it would be more appropriate to sum the errors in quadrature (see, e.g., Ref. [97]). In the present case, however, we do not know how the errors are distributed, and we have no reason to expect them to be independent or normally distributed. So, we assume the worst case and combine errors by adding magnitudes.

For the uncertainty in a CCE waveform, we combine the Cauchy error (measured using CCE waveforms), the CCE truncation error, and the CCE initial-data error. Similarly, for the error in an extrapolated waveform, we combine the Cauchy error (measured using extrapolated waveforms) and extrapolation fit error. This error bar is incomplete for extrapolated waveforms, as it does not include any contribution from gauge error; we estimate the magnitude of the gauge error in Sec. V below by comparing extrapolated waveforms with CCE.

2. Error bar for difference between CCE and extrapolated waveforms

We wish to determine whether a CCE waveform and an extrapolated waveform agree to within some error bar. If they do, then we can regard the gauge error in extrapolated waveforms as small, and we can use the extrapolation procedure instead of the more complicated and computationally expensive CCE procedure to obtain waveforms at J^+ . The estimated error bar for the difference between CCE and extrapolated waveforms is constructed as the L_1 norm of the CCE truncation error, the CCE initial-data error, the extrapolation fit error, and the Cauchy error.

Because the CCE and extrapolated waveforms each have their own Cauchy error, it is not immediately clear which Cauchy error should enter into the error bar. Let C^C and C^E denote the Cauchy error determined using CCE and

extrapolated waveforms, respectively. Both C^C and C^E arise from the same source: the truncation error in the Cauchy simulation. To define the Cauchy error for the difference between a CCE and an extrapolated waveform, we take the average of C^C and C^E .

It is not obvious that averaging C^C and C^E is the correct procedure: the issue is whether they are correlated. To pursue this further, note that there are two contributions to both C^C and C^E . The first contribution corresponds to the error made in determining the motion of the black holes; this affects C^C and C^E in an identical way. The second contribution corresponds to the error made in propagating waves through the grid; this affects C^C and C^E differently, because the extraction radii and the quantities read from the Cauchy code are different for extrapolated waveforms than for CCE waveforms. If the first contribution is dominant, then C^C and C^E are correlated, and so it would be appropriate to use their average. But if the second contribution is dominant, then C^C and C^E are uncorrelated, and so it would be appropriate to use their sum.

We can determine which part of the Cauchy error is dominant by plotting the difference between CCE and extrapolated waveforms taken from a low-resolution simulation, and comparing with the difference between CCE and extrapolated waveforms taken from a high-resolution simulation. Such a plot is shown in Fig. 7. We find that the

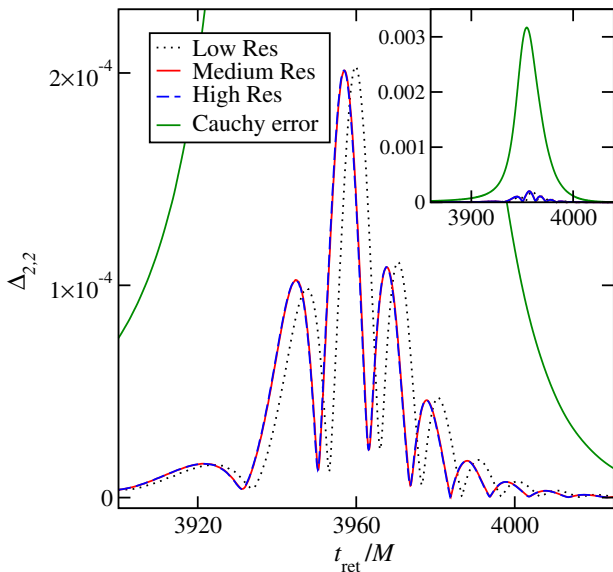


FIG. 7 (color online). Cauchy-resolution dependence of $\Delta_{\ell,m}$ [cf. Eq. (19)] between CCE and extrapolated $\Psi_4^{2,2}$, shown near peak amplitude at $\approx 3952M$ for simulation 2 in Table I. At each lower Cauchy resolution, the extrapolated waveform is aligned with the high-resolution extrapolated waveform. Then the CCE waveform for each Cauchy resolution is aligned with the corresponding extrapolated waveform and $\Delta_{\ell,m}$ is computed. The differences are nearly independent of resolution. Also shown (labeled “Cauchy error”) is the difference between the extrapolated $\Psi_4^{2,2}$ waveforms from the high and medium resolutions.

difference between CCE and extrapolated waveforms is largely independent of resolution, indicating that the dominant effect of Cauchy error is to change the trajectories of the black holes, and that C^C and C^E are highly correlated rather than independent. Therefore, we are justified in computing the combined Cauchy error as the average of C^C and C^E , rather than their sum.

Additionally, we find that the difference between CCE and extrapolated waveforms is significantly smaller than the estimated Cauchy error, as shown in the figure—at least for the (2, 2) mode. Accordingly, the measures C^C and C^E are not merely correlated, but are also nearly identical to each other. This continues to hold even for subdominant modes, for which the Cauchy error can be comparable to the difference between CCE and extrapolated quantities (cf. Fig. 11 and the discussion below).

V. RESULTS

In this section we compare the relative magnitudes of the various error quantities for both extrapolated and CCE waveforms. We verify the gauge dependence of extrapolated waveforms and the gauge invariance of CCE by examining waveforms from two physically identical simulations performed using different gauge conditions. For each BBH configuration in Table I, we evaluate the quality of extrapolated waveforms by comparing with CCE waveforms. This allows us to determine whether the gauge error in extrapolated waveforms is smaller than the other sources of error, and hence whether we can justify using the extrapolation method instead of CCE.

A. Is waveform extraction to I^+ necessary?

We first address the question of whether waveform extraction to I^+ is even necessary, or whether finite-radius waveforms are sufficient, given the accuracy of our simulations. Consider the finite-radius $\Psi_4^{2,2}$ mode, for Case 2 in Table I, computed from the outermost extraction radius ($R = 385M$). Figure 8 shows the phase difference between this finite-radius waveform and the corresponding CCE waveform. Also shown are the difference between the CCE and extrapolated $\Psi_4^{2,2}$ waveforms and the estimated error bar for the phase of the CCE waveform. The phase of the finite-radius waveform falls far outside of the estimated error bar, while in this case the extrapolated and CCE waveforms agree very well. This indicates that the finite-radius waveform is a poor proxy for the waveform at I^+ , and that some form of waveform extraction (either extrapolation of CCE) is required.

B. Comparing different sources of uncertainty

Here we examine the average magnitudes of errors from different sources for both extrapolated and CCE waveforms. To illustrate the typical sizes of these errors, Fig. 9 shows the estimated phase errors in $\Psi_4^{2,2}$ for the




FIG. 8 (color online). Magnitude of phase difference between the CCE and the outermost finite-radius ($R = 385M$) $\Psi_4^{2,2}$ waveforms, for Case 2 in Table I. The error bar for the phase of the CCE waveform and the phase difference between extrapolated and CCE waveforms are also shown. The error bar includes the Cauchy error (measured using CCE waveforms), the CCE truncation error, and the CCE initial-data error. Waveforms are aligned over $[1000M, 2000M]$. The maximum amplitude occurs at $t_{\text{ret}} \sim 3952M$, indicated by the dotted vertical line.

equal-mass, nonspinning simulation (Case 2 in Table I). All uncertainties are computed using the procedures described in Sec. IV. The errors shown in Fig. 9 include the Cauchy error measured using extrapolated and CCE waveforms, as well as the extrapolation fit error, the CCE truncation error (on the null grid), and the CCE initial-data error.

We find that the Cauchy error measured from CCE waveforms is essentially indistinguishable from that measured from extrapolated waveforms, on the scale of Fig. 9. This is consistent with the discussion in Sec. IV G 2, i.e. that these Cauchy errors are highly correlated. For extrapolated waveforms, we find that the Cauchy and extrapolation fit errors are about equal. For CCE waveforms, the Cauchy error dominates, followed by the CCE initial-data error, and finally by the very small CCE truncation error.

Figure 10 shows relative amplitude errors for the same simulation as Fig. 9. During the merger and ringdown, the relative contributions of each error source are the same as for the phase error, with the Cauchy error being the largest and the CCE truncation error the smallest. Interestingly, we find that the CCE initial-data error is the dominant source of amplitude error during the inspiral, although in absolute terms it is still a small error at $\mathcal{O}(10^{-3})$. Near the merger and during the ringdown, the relative amplitude errors are small compared with the phase errors shown in Fig. 9. Hence, during this portion of the waveform, the error




FIG. 9 (color online). Phase errors in the $\Psi_4^{2,2}$ waveform at I^+ from various sources, for simulation 2 in Table I. Cauchy errors determined from both extrapolated (C^E) and CCE (C^C) waveforms are shown (see Sec. IV G 2). The outermost extraction radius is $R = 385M$, and all waveforms are aligned over $[1000M, 2000M]$. The maximum amplitude occurs at $t_{\text{ret}} \sim 3952M$, shown here as the dotted vertical line.

measure $\Delta_{\ell,m}$ given by Eq. (19) will be essentially the same as the phase error.

Having investigated the error in $\Psi_4^{2,2}$ for simulation 2 of Table I, we now consider the errors for the other simulations and for other spin-weighted spherical harmonic modes. To condense information from many modes and




FIG. 10 (color online). Same as Fig. 9 except showing relative amplitude errors instead of phase errors.

several simulations into a smaller number of figures, we compute time-averaged errors as described in Sec. IV G, and we use the error measure $\Delta_{\ell,m}$ of Eq. (19) instead of measuring phase and amplitude errors separately. This reduces each error measure for a given (ℓ, m) mode to a single number.

Figure 11 shows the time-averaged Cauchy errors in extrapolated and CCE waveforms for all $\Psi_4^{\ell,m}$ up to $\ell = 4$ and for all simulations in Table I. Although only modes up to $\ell = 4$ have been included in this figure, the qualitative features are unchanged if modes up to $\ell = 8$ (the maximum mode we have computed) are included.

There are a few general features evident in the figure. First, the Cauchy errors in CCE and extrapolated waveforms have similar magnitudes. In addition, the modes with $|m| = \ell$ have the largest errors. This is to be expected, because these are the modes with the greatest amplitudes. Along the same lines, we see that in the $q = 1$ cases, the average error is very small for the modes with odd m , because by symmetry (rotation through π) these modes should have vanishing amplitude.

Figure 12 shows the time-averaged Cauchy error, the extrapolation fit error, the CCE truncation error, and the CCE initial-data error in $\Psi_4^{\ell,m}$ for all (ℓ, m) up to $\ell = 4$. The Cauchy error shown here is the average of those

computed from the CCE and extrapolated waveforms. As was the case for Fig. 11, we only show results up to $\ell \leq 4$, but the qualitative features are the same for all modes we have examined (up to $\ell = 8$).

The truncation error on the CCE null grid is by far the smallest source of error in each case. The largest source of error varies, depending both on the simulation and on the mode. For most cases, the CCE initial-data error and the Cauchy error are the largest, except in the $q = 1$ cases where the extrapolation fit error dominates.

C. Gauge dependence

In principle, extrapolated waveforms may be contaminated by gauge effects, whereas CCE waveforms should be gauge invariant. Here we directly investigate the gauge dependence of both extraction methods by comparing two equal-mass, zero-spin BBH simulations (the first two cases in Table I) with identical initial data but with different gauge conditions. The first simulation is the one described in Ref. [18]. It uses a gauge in which the gauge-source function obeys a wave equation, and the source terms of this wave equation are fine-tuned by hand. We have found previously that this gauge does not work well for black-hole binaries with unequal masses or large spins [14,83,98], so current BBH simulations using SpEC

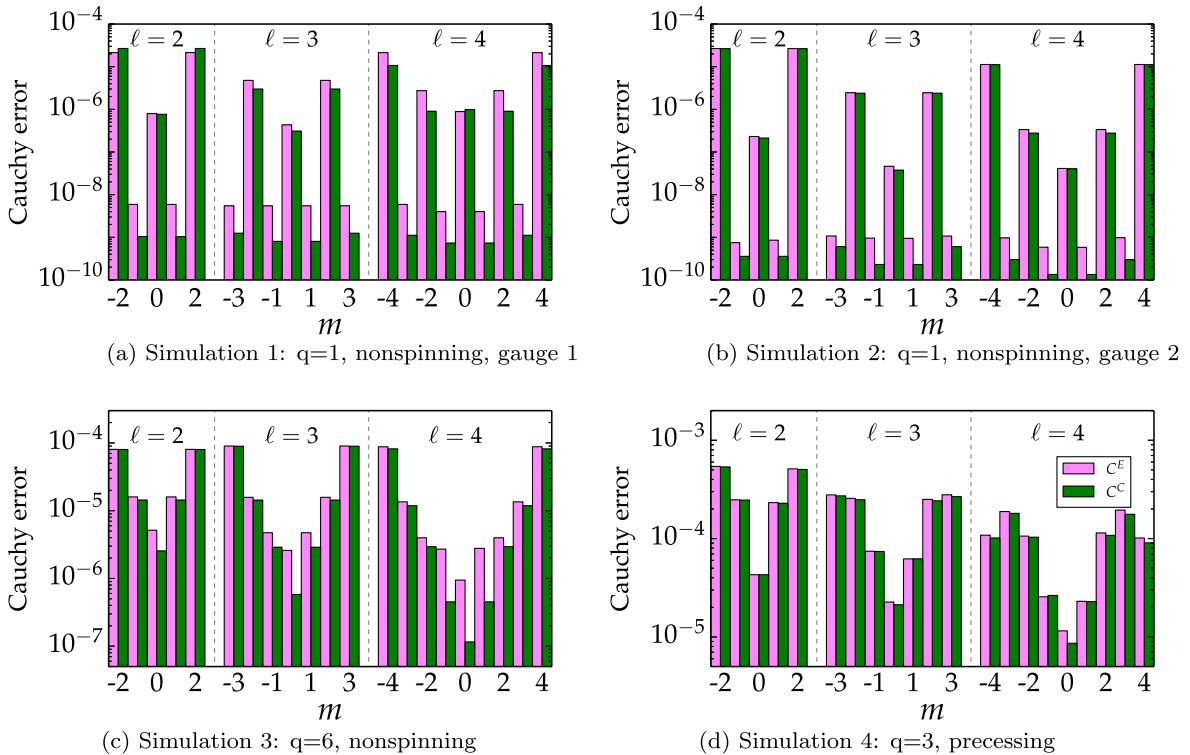


FIG. 11 (color online). Cauchy errors C^C and C^E (see Sec. IV G 2) as a function of (ℓ, m) spherical harmonic mode for different simulations in Table I. The vertical axis is the error measure $\Delta_{\ell,m}$ of Eq. (19), time-averaged so that each source of error is described by a single number for each (ℓ, m) mode. The horizontal axis represents the spherical harmonic m index; vertical dashed lines separate $\ell = 2$, $\ell = 3$, and $\ell = 4$ modes, and for each ℓ , every other value of m is labeled on the horizontal axis. The pink bars represent the Cauchy error C^E in the extrapolated waveforms, and the dark green bars represent the Cauchy error C^C in the CCE waveforms.

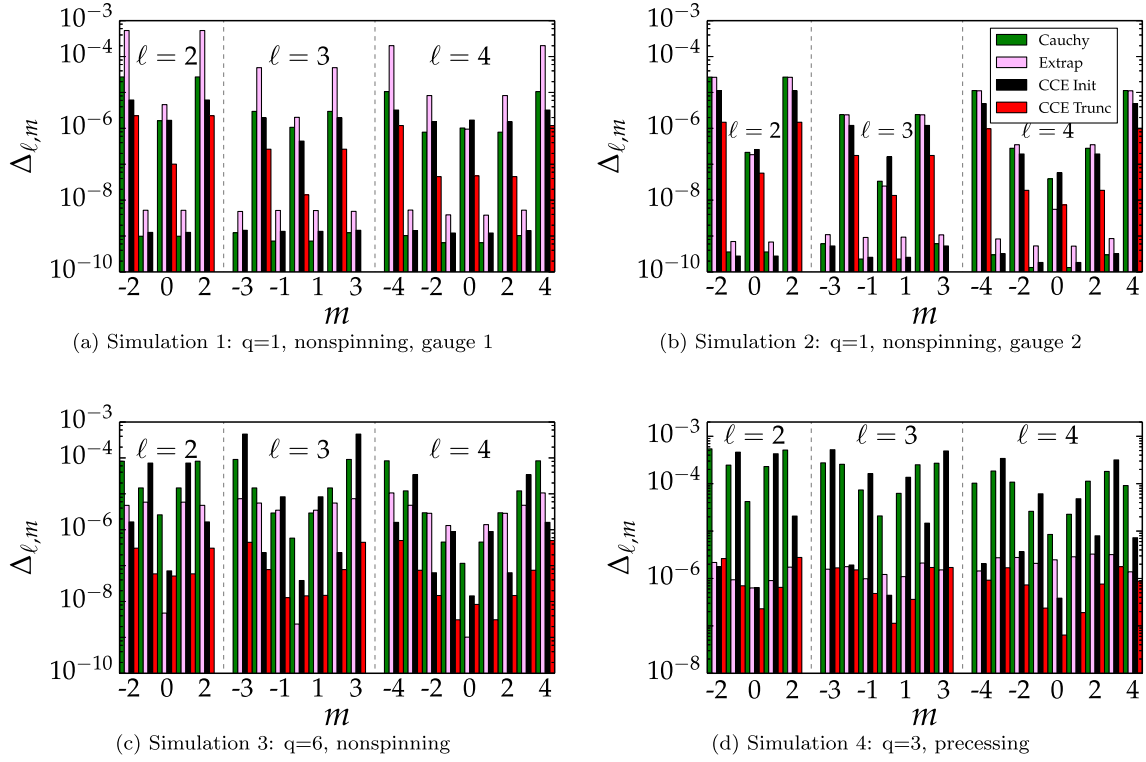


FIG. 12 (color online). Same as Fig. 11, but showing multiple sources of error. The Cauchy error shown here is the average of those computed using CCE and extrapolated waveforms, for each (ℓ, m) .

employ a damped harmonic gauge condition [14,83,84], which is the gauge used in simulation 2 of Table I.

Figure 13 shows the dominant mode $\Psi_4^{2,2}$ as a function of time for both gauge choices, and for both extrapolated and CCE waveforms. All four plots in this figure agree well, suggesting that both CCE and extrapolated waveforms for this dominant mode are independent of gauge, at least on the scale of the figure.

On the other hand, the extrapolated waveform for the subdominant mode $\Psi_4^{2,0}$ differs significantly between simulations 1 and 2. In particular, for simulation 1, the gauge effects appear to be so strong that it is difficult to even define the extrapolated $\Psi_4^{2,0}$ waveform. To understand the difficulty, recall that the extrapolation procedure assumes that $rM\Psi_4^{\ell,m}$ approaches a finite limit as $r \rightarrow \infty$. However, if $rM\Psi_4^{2,0}$ from simulation 1 is plotted at different extraction radii r , we find that it appears to grow without limit as r increases, as illustrated in Fig. 14. The assumption that the finite-radius waveforms $rM\Psi_4^{\ell,m}(t_{\text{ret}}, r)$ can be expanded in a convergent series in $1/r$ is thus violated in this case. Note that this problem occurs only for the gauge used in simulation 1; for the gauge used in the other simulations, $rM\Psi_4^{2,0}$ approaches a finite limit as r increases.

Although extrapolation fails to converge for the $\Psi_4^{2,0}$ waveform in simulation 1, we nevertheless compute the $N = 5$ extrapolant for this mode for comparison purposes. Based on Fig. 14, we do not expect this $N = 5$ extrapolant

to be very accurate. It is worth noting, however, that this extrapolated waveform nevertheless agrees better with CCE than the unextrapolated waveform measured at the outermost extraction radius.




FIG. 13 (color online). The real part of $rM\Psi_4^{2,2}$ for both extrapolated and CCE waveforms, for the first two simulations in Table I. Waveforms are aligned in the interval $[1000M, 2000M]$. The four curves agree very well. Time-averaged differences between these curves are shown in Figs. 16–18 below.




FIG. 14 (color online). The real part of $rM\Psi_4^{2,0}$ extracted at multiple radii, before extrapolation, for the first simulation in Table I. Waveforms are shown only near peak amplitude because they are very small elsewhere. The waveform for each extraction radius r is plotted versus time, rather than t_{ret} , so that waveforms extracted at larger r reach their peak amplitude later. The increase in amplitude with extraction radius r indicates that $\Psi_4^{2,0}$ falls off more slowly than $1/r$. We attribute this slow falloff to the gauge condition used for simulation 1. The other simulations, which use a more robust gauge condition, do not exhibit this behavior.

Our expectations are confirmed by Fig. 15, which shows $\Psi_4^{2,0}$ as a function of time for CCE and extrapolated waveforms, and for simulations 1 and 2. This figure is the same as Fig. 13, except that it shows $\Psi_4^{2,0}$ instead of $\Psi_4^{2,2}$. The extrapolated waveforms in Fig. 15 are very different for the two simulations, whereas the CCE waveforms are almost indistinguishable. This provides a demonstration both of the gauge invariance of CCE and of the gauge dependence of extrapolated waveforms.

To make the above conclusions more precise, the differences between these waveforms should be compared to the various sources of error discussed in Sec. VB. We construct a measure $\epsilon_{\ell m}$ of the fractional difference between the waveforms for each mode, computed as the difference between the extrapolated $\Psi_4^{\ell,m}$ from simulation 1 and the same waveform from simulation 2, divided by a combined error bar for the difference. The combined error bar is defined as the L^1 norm of the various sources of uncertainty that enter into the difference (cf. Sec. IV G). To obtain a single measure of the fractional agreement between the waveforms for each mode, we also perform a time averaging of these fractional differences. In other words, we define

$$\epsilon_{\ell m} = \langle |\Psi_4^{l,mA} - \Psi_4^{l,mB}|/E \rangle, \quad (20)$$




FIG. 15 (color online). The real part of $rM\Psi_4^{2,0}$ for both extrapolated and CCE waveforms, for the first two simulations in Table I. Waveforms are aligned in the interval $[1000M, 2000M]$. We show only times near the merger because the waveform is very small elsewhere. Although the difference between CCE and extrapolated waveforms for Case 2 is far smaller than for Case 1, even in Case 2 this difference is several times the numerical error. Note that the time-averaged difference shown below in Fig. 18 for Case 2 is dominated by the inspiral portion of the waveforms.

where A and B refer to the different simulations, and angle brackets represent a time average. The numerator is the error measure $\Delta_{\ell,m}$ of Eq. (19), and the error bar in the denominator is computed here as

$$E = \frac{1}{2} (|C_A^E| + |C_B^E|) + |F_A| + |F_B|, \quad (21)$$

where C^E represents the Cauchy error computed using the extrapolated waveforms in simulation A or B , and F represents the extrapolation fit error. Note that each of these error measures is computed separately for each (ℓ, m) mode and for each time, and that the division in Eq. (20) is done before the time averaging.

In Fig. 16, we plot these time-averaged fractional differences for all modes. Values less than unity indicate differences that are (on average) within the error bars. Figure 16 shows that for most (ℓ, m) modes, extrapolated waveforms for the two different gauge choices are essentially indistinguishable (i.e. within the error bars). However, for the $m = 0$ modes, extrapolated waveforms are contaminated by significant gauge effects that are larger than other sources of error. As ℓ increases, the average fractional difference between $m = 0$ modes $\epsilon_{\ell 0}$ decreases. This is simply because the amplitude of the modes decreases with increasing ℓ , so eventually the differences fall within the error bars.




FIG. 16 (color online). Fractional differences $\epsilon_{\ell m}$ [cf. Eqs. (20) and (21)] between extrapolated $\Psi_4^{\ell, m}$ from physically equivalent simulations with different gauge conditions (i.e., the first two simulations in Table I), as a function of (ℓ, m) . The (ℓ, m) modes are labeled as in Fig. 11. Waveforms are aligned in the interval $[1000M, 2000M]$.

Figure 17 shows fractional differences between waveforms from the same two simulations as Fig. 16, but for CCE waveforms. Because CCE waveforms have different sources of error than extrapolated waveforms, the denominator of Eq. (20) is computed in this case as

$$E = \frac{1}{2}(|C_A^C| + |C_B^C|) + |T_A| + |T_B| + |I_A| + |I_B|, \quad (22)$$

where C^C represents the Cauchy error computed using the CCE waveforms in simulation A or B , T represents the CCE truncation error, and I represents the CCE initial-data error.

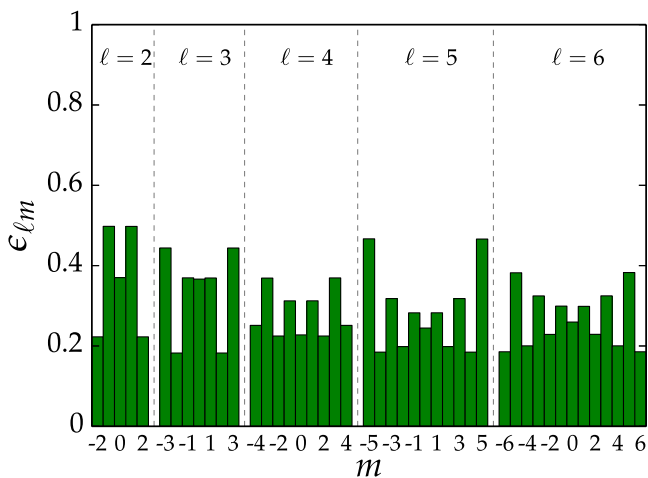


FIG. 17 (color online). Fractional differences $\epsilon_{\ell m}$ [cf. Eqs. (20) and (22)] between CCE waveforms from the same simulations as shown in Fig. 16. Labels are the same as Fig. 16, except here the differences are shown on a linear plot.

The differences shown in Fig. 17 are smaller than unity, verifying that CCE is indeed gauge invariant to the level of our numerical error, even for a gauge (the gauge from simulation 1) that is sufficiently ill-behaved that extrapolation fails to converge (cf. Fig. 14). Moreover, comparing Fig. 16 with Fig. 17 shows that the differences between CCE waveforms from simulations 1 and 2 are on average smaller than the differences between extrapolated waveforms from the same two simulations.

D. When is CCE necessary?

In the previous section, we found an example of extrapolated waveforms being significantly contaminated by gauge effects. In particular, the gauge used in simulation 1 of Table I results in waveforms for which some spherical harmonic modes (namely, those with $m = 0$) cannot be reliably extrapolated because they fall off more slowly than $1/r$.

This example raises the question of how reliable the extrapolation method is in general. It should be possible to find (or construct) other examples in which extrapolation yields the wrong waveform. But will all of these examples exhibit clear erroneous behavior such as the slow falloff shown in Fig. 14, or is it possible for extrapolation to yield the incorrect result without any indication of a problem? In principle, the latter should be possible for a sufficiently pathological gauge. For instance, a gauge pulse traveling outward and falling off exactly like $1/r$ would allow convergent extrapolation, but would still contaminate the extrapolated waveform.

Here we focus on a more specific question: for simulations using the damped harmonic gauge condition [14,83,84] as currently implemented in SpEC, how reliable are extrapolated waveforms? We answer this question for the simulations in Table I, by comparing extrapolated waveforms to gauge-invariant CCE waveforms.

This comparison is shown in Fig. 18, where we plot the average fractional differences between extrapolated and CCE waveforms for all simulations in Table I, and for all modes with $\ell \leq 6$. The quantity plotted is $\epsilon_{\ell m}$ as defined in Eq. (20), with the error bar defined by

$$E = \frac{1}{2}(|C^C| + |C^E|) + |F| + |T| + |I|. \quad (23)$$

Here C^C and C^E are the respective Cauchy errors computed using CCE and extrapolated waveforms, T is the CCE truncation error, F is the extrapolation fit error, and I is the CCE initial-data error.

If the magnitudes of the fractional errors plotted in Fig. 18 are less than unity, then the differences between CCE and extrapolated waveforms are smaller on average than the estimated error bars, and we can conclude that gauge errors in extrapolated waveforms are unimportant.

We find that this is indeed the case for almost all (ℓ, m) modes, including the dominant $(2, 2)$ modes. However, for

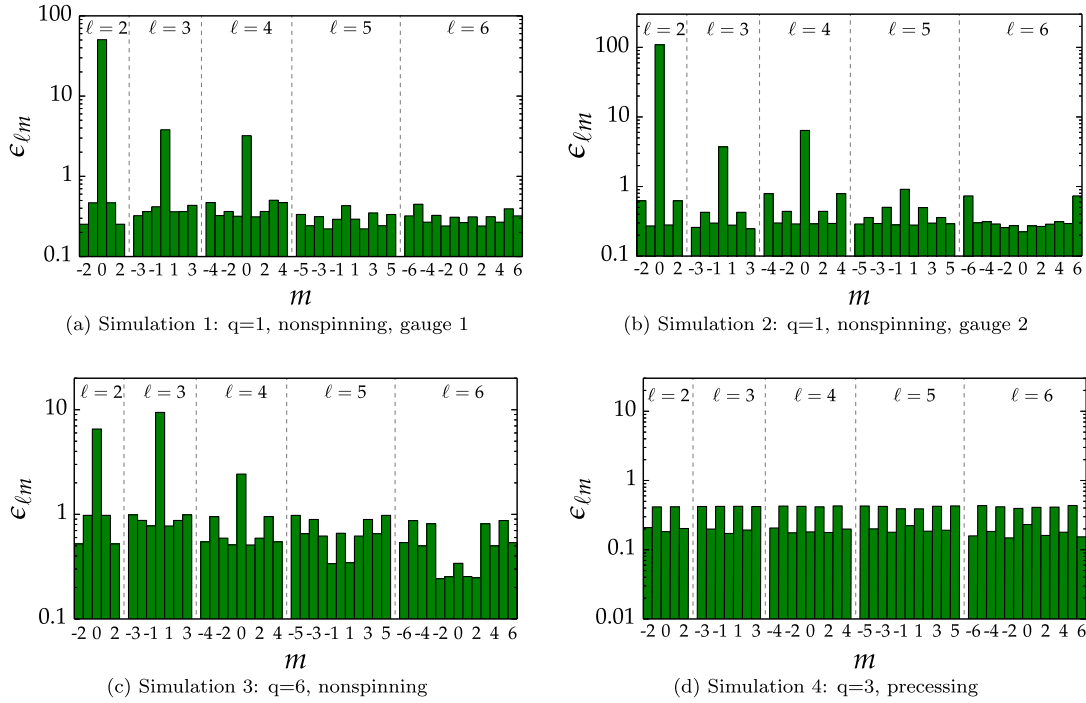


FIG. 18 (color online). Fractional differences between extrapolated and CCE $\Psi_4^{\ell,m}$ for all four cases in Table I, as a function of (ℓ, m) . The (ℓ, m) modes are labeled as in Fig. 11. Waveforms are aligned in the interval $[1000M, 2000M]$.

the first few modes with $m = 0$, we find that the difference between CCE and extrapolated waveforms is larger than the estimated error, suggesting that for these modes the gauge contamination in extrapolated waveforms is significant.

Earlier in Fig. 15 we compared the $(2, 0)$ mode between CCE and extrapolated waveforms, and we found that the agreement was much better for simulation 2 than for simulation 1. But Fig. 18 appears to support the opposite conclusion; the fractional differences between CCE and extrapolated waveforms in this figure are smaller for simulation 1 than simulation 2. This discrepancy can be explained by noting that the quantities in Fig. 18 are normalized by the error bar, defined in Eq. (23), which is much larger for simulation 1 than simulation 2. Figure 12 shows that for most modes, the largest contribution to this error measure in simulation 1 is the extrapolation fit error.

For most modes with $m \neq 0$, the average fractional differences in Fig. 18 are less than unity. For the $q = 6$ simulation (Case 3 in Table I), however, many of these modes have average fractional differences that are very close to unity. Upon further examination, we find that for this case, using lower-order extrapolation seems to improve the agreement with CCE for most modes. In fact, if we use order $N = 2$ extrapolation, the average fractional differences between extrapolated and CCE waveforms fall markedly below unity for every mode, including $m = 0$ modes. This is potentially misleading, however, because the primary reason for the improvement is that the estimated extrapolation fit error is erroneously small when

using higher-order extrapolation. The actual difference (not normalized by the error bar) between CCE and extrapolation is in fact greater for lower-order extrapolation. Evidently, the accuracy of the estimated extrapolation fit error decreases as order is increased.

This behavior could be at least somewhat anticipated by inspecting the convergence of the extrapolated waveforms with the extrapolation order. We find that both the amplitude and phase of many modes exhibit clear divergence as the extrapolation order is increased, particularly for times near the merger. Increasing extrapolation order produces increasing amounts of higher-frequency noise, as shown in Fig. 19. This casts significant doubt on the reliability of any extrapolation error estimate in this case. Note that the extrapolated $(2, 2)$ mode in this simulation actually does converge for the first few extrapolation orders, and it also agrees well with CCE. Note also that for the other BBH cases, there is no clear lack of convergence in the extrapolated waveforms (for $m \neq 0$ modes), and lower-order extrapolation does not improve the agreement with CCE.

So far we have considered different $Y_{\ell m}$ modes separately. Let us now briefly examine the difference between CCE and extrapolated waveforms when summing over all modes, as is done when computing the waveform in a particular sky direction. In particular, we would like to investigate whether the large errors in the extrapolated $m = 0$ modes shown in Fig. 18 correspond to large errors after summing over all modes. Instead of choosing a single direction on the sky, we integrate the difference between CCE and extrapolated waveforms over all sky directions,




FIG. 19 (color online). Merger portion of the real part of the extrapolated $rM\Psi_4^{3,2}$ mode for the $q = 6$ case (simulation 3 in Table I). Divergence of the extrapolated waveform is evident as extrapolation order is increased. Note that the order $N = 2$ extrapolated waveform agrees well with CCE in this case. Maximum amplitude (of $rM\Psi_4^{2,2}$) is at $t_{\text{ret}} \approx 4901M$.

and use Eqs. (17) and (18) to write this integral as a sum over modes. We then normalize by the quadrature sum of the errors in each mode. Thus we compute the expression

$$\epsilon = \frac{\|\Psi_4^A - \Psi_4^B\|}{\sum_{\text{sources}} (\sum_{\ell,m} E_{\ell,m}^2)^{1/2}}, \quad (24)$$

where A and B in the numerator represent CCE and extrapolated waveforms, and where the numerator is evaluated using Eq. (18). The sum in the denominator is over all sources of error, with the individual mode contributions summed in quadrature, for each source of error. The sources of error that enter into this calculation include the Cauchy error, the extrapolation fit error, the CCE initial-data error, and the CCE truncation error. Figure 20 shows the quantity ϵ for each of the four numerical simulations we consider. In the figure, curves have been shifted in time so that the merger occurs at $t \approx 0$ for each case.

To estimate the importance of $m = 0$ modes in the sum over all modes, we compute the sums in Eq. (24) twice—once with all modes included (up to $L = 8$), and again with $m = 0$ modes omitted. As shown in Fig. 20, including $m = 0$ modes substantially changes the waveform agreement for the equal-mass, nonspinning configurations (Cases 1 and 2 in Table I): in both cases $\epsilon < 1$ when omitting the $m = 0$ modes, and $\epsilon > 1$ when including them. For the $q = 6$ simulation (Case 3), the difference between CCE and extrapolated waveforms is the same size as the combined error bar. Including the $m = 0$ modes makes no noticeable difference in this case, even though $m = 0$ modes were in




FIG. 20 (color online). Differences between CCE and extrapolated waveforms integrated over the sky and normalized by error bars, computed according to Eq. (24) for the four cases of Table I. Curves have been shifted so that the merger in each case is at $t \approx 0$. For each case, solid lines are computed using all modes up to $L = 8$, while dotted lines are the same but with $m = 0$ modes omitted. For Cases 3 and 4, the dotted lines are indistinguishable from the corresponding solid lines.

disagreement (albeit not by as much) in Fig. 18. Including $m = 0$ modes makes no noticeable difference in the generic configuration (Case 4) as well, although this is to be expected because of the good agreement between CCE and extrapolated waveforms for all modes in this case.

It may be somewhat surprising that the curves for Cases 3 and 4 are largely constant in time. This is because for many modes both the difference between CCE and extrapolated waveforms in the numerator of Eq. (24) and the estimated error bars in the denominator are dominated by the CCE initial-data error, as shown in Fig. 12. This error manifests as a largely constant in time amplitude offset, as illustrated in Fig. 10. This accounts both for the flatness of the Case 3 and 4 curves in Fig. 20 as well as the negligible impact of $m = 0$ modes for these cases.

The above considerations indicate that the question of whether CCE is necessary to achieve accurate waveforms depends not only on the various sources of error, but also on which (ℓ, m) modes are of interest. For general applications in which one is interested in all (ℓ, m) modes, we find that without CCE, (presumed) gauge errors can dominate the errors in our waveforms.

VI. DISCUSSION

Comparisons between different methods of waveform extraction are meaningful only when considering the various sources of uncertainty that affect the final waveform. We have estimated the key error contributions for a handful

of simulations. In all of the cases we considered, the CCE null-grid truncation error was by far the smallest source of uncertainty. This suggests that the relative expense of CCE could be reduced by running at lower CCE resolution, without significant impact on the results. The extrapolation fit error was the most significant source of error in the first equal-mass simulation (Case 1 in Table I), presumably because of the gauge condition used (described in Sec. III). In the other cases, which used harmonic or damped harmonic gauge, the extrapolation fit, the CCE initial data, and the Cauchy errors were generally comparable.

A potential improvement to extrapolation would be to use a time-varying extrapolation order, with higher order in the early inspiral and lower order near the merger, so that the order decreases with decreasing wavelength. This could be achieved smoothly by combining extrapolants of different orders, each weighted by the (inverse) variances of the polynomial fit, and suitably normalized. Such a procedure would not only provide for more accurate extrapolation but would also reduce the magnitude of the estimated extrapolation fit error.

We were somewhat surprised to find that the CCE initial-data error was often quite significant, sometimes dominating the other source of error. Reducing the magnitude of this error could be achieved by using a larger Cauchy computational domain (so that the world tube radius could be larger), which would increase the computational cost of the simulations. The extra cost would be modest for codes (like SpEC) that use spherical outer domains rather than Cartesian-aligned grids, except for the extra evolution time necessary for the gravitational waves to reach the more distant world tube. The CCE initial-data error could also be reduced, in principle, by using improved initial data in the characteristic code [61].

By explicitly comparing two simulations with identical physical parameters, differing only in the gauge condition used for the Cauchy simulation, we showed in Sec. VC that CCE waveforms are gauge independent to within uncertainties. We found that extrapolated waveforms, on the other hand, had significant gauge dependence for $m = 0$ modes. It was clear from Fig. 14 that extrapolation would fail for $m = 0$ for the simulation with the gauge condition of Case 1 from Table I, and that therefore another method such as CCE was required. In the $q = 6$ simulation, the poor convergence of extrapolation made it clear that an alternate extraction method was required. However, for $m = 0$ modes in Case 2, there was no *a priori* indication that extrapolated waveforms would be inaccurate.

We find that large-amplitude modes (such as Ψ_4^{22}) generally agree well between CCE and extrapolated waveforms. However, the $m = 0$ “memory” modes disagree significantly in almost every case. This disagreement is not necessarily a result of gauge effects alone. The long wavelength of the $m = 0$ modes may lead to inherent difficulties

in the polynomial fit, resulting in poor extrapolation, as discussed at the end of Sec. IIA3. Indeed, we find that most of the difference in the $(2, 0)$ mode in the upper right panel of Fig. 18, for example, comes from the inspiral, where the wavelength is longer. The fractional difference between the extrapolated and CCE waveforms is greater than unity during the merger in this case as well, but it is orders of magnitude less there than it is during the inspiral.

Unlike in the other cases, extrapolated and CCE waveforms were found to agree quite well for all (including $m = 0$) modes in the precessing configuration (Case 4 in Table I). One reason for this is that the uncertainties are larger in this case than in the others, as shown in Fig. 12. Even with larger error bars, however, it is somewhat surprising that $m = 0$ modes do not stand out in Fig. 18, as they do for the other cases. We do not know the reason for this, but we note that this is the only simulation we consider that utilized a damped harmonic gauge condition for the majority of the inspiral, as described in Sec. III.

Because of the potential disagreement in $m = 0$ modes, we recommend using CCE in applications for which all modes are important. Additionally, Even though we found above in the $q = 6$ simulation that extrapolated waveforms agreed with CCE for $m \neq 0$ modes, we do not consider this a confirmation of the reliability of extrapolated waveforms. When no convergence at any order is evident in the extrapolation procedure, the waveforms and error estimates simply cannot be trusted. For this reason, we also recommend CCE in cases where extrapolation fails to show convergence for at least the first few orders. We caution that each mode of interest must be individually checked for convergence. For instance, as discussed above for the $q = 6$ simulation, the extrapolated $(2, 2)$ mode was convergent, while the other modes were not.

When extrapolation does show reasonable convergence, however, the uncertainties in the two waveform extraction methods are comparable. In this case, because of the simplicity and reduced computational expense, extrapolated waveforms are preferred for $m \neq 0$ modes. Nevertheless, even if the extrapolation is convergent, we recommend double checking with CCE waveforms for simulations that use new gauge conditions or for new regions of parameter space.

ACKNOWLEDGMENTS

We thank Nigel Bishop, Ian Hinder, Lee Lindblom, Harald Pfeiffer, and Jeffrey Winicour for helpful discussions. We thank Christian Ott for help in initiating and completing this project. We gratefully acknowledge support from the Sherman Fairchild Foundation; from NSF Grants No. PHY-1068881, No. PHY-1005655, and No. DMS-1065438 at Caltech; and from NSF Grants No. PHY-0969111 and No. PHY-1005426, and NASA Grant No. NNX09AF96G at Cornell. C.R. acknowledges support by NASA through Einstein Postdoctoral

Fellowship Grant No. PF2-130099 awarded by the Chandra X-ray Center, which is operated by the Smithsonian Astrophysical Observatory for NASA under Contract No. NAS8-03060. Simulations used in this work were computed with SpEC [13]. Computations were performed on the Zwicky cluster at Caltech, which is supported by the Sherman Fairchild Foundation and by NSF

Award PHY-0960291; on the NSF XSEDE network under Grant No. TG-PHY990007N; and on the GPC supercomputer at the SciNet HPC Consortium [99]. SciNet is funded by the Canada Foundation for Innovation under the auspices of Compute Canada, the Government of Ontario, Ontario Research Fund—Research Excellence, and the University of Toronto.

-
- [1] G.M. Harry for the LIGO Scientific Collaboration, *Classical Quantum Gravity* **27**, 084006 (2010).
- [2] K. Somiya and the KAGRA Collaboration, *Classical Quantum Gravity* **29**, 124007 (2012).
- [3] The Virgo Collaboration, *Advanced Virgo Baseline Design* (2009), <https://tds.ego-gw.it/ql/?c=6589>.
- [4] H. Grote (LIGO Scientific Collaboration), *Classical Quantum Gravity* **27**, 084003 (2010).
- [5] J. Abadie *et al.* (LIGO Scientific), *Classical Quantum Gravity* **27**, 173001 (2010).
- [6] B.J. Owen and B.S. Sathyaprakash, *Phys. Rev. D* **60**, 022002 (1999).
- [7] A. Buonanno and T. Damour, *Phys. Rev. D* **59**, 084006 (1999).
- [8] P. Ajith *et al.*, *Phys. Rev. D* **77**, 104017 (2008).
- [9] P. Ajith, *Phys. Rev. D* **84**, 084037 (2011).
- [10] A. Buonanno, Y. Pan, H.P. Pfeiffer, M.A. Scheel, L.T. Buchman, and L.E. Kidder, *Phys. Rev. D* **79**, 124028 (2009).
- [11] Y. Pan, A. Buonanno, M. Boyle, L.T. Buchman, L.E. Kidder, H.P. Pfeiffer, and M.A. Scheel, *Phys. Rev. D* **84**, 124052 (2011).
- [12] A. Taracchini, Y. Pan, A. Buonanno, E. Barausse, M. Boyle, T. Chu, G. Lovelace, H.P. Pfeiffer, and M.A. Scheel, *Phys. Rev. D* **86**, 024011 (2012).
- [13] <http://www.black-holes.org/SpEC.html>.
- [14] B. Szilágyi, L. Lindblom, and M.A. Scheel, *Phys. Rev. D* **80**, 124010 (2009).
- [15] L.T. Buchman, H.P. Pfeiffer, M.A. Scheel, and B. Szilágyi, *Phys. Rev. D* **86**, 084033 (2012).
- [16] D.A. Hemberger, M.A. Scheel, L.E. Kidder, B. Szilágyi, G. Lovelace, N.W. Taylor, and S.A. Teukolsky, *Classical Quantum Gravity* **30**, 115001 (2013).
- [17] O. Rinne, L. Lindblom, and M.A. Scheel, *Classical Quantum Gravity* **24**, 4053 (2007).
- [18] M.A. Scheel, M. Boyle, T. Chu, L.E. Kidder, K.D. Matthews, and H.P. Pfeiffer, *Phys. Rev. D* **79**, 024003 (2009).
- [19] E. Newman and R. Penrose, *J. Math. Phys. (N.Y.)* **3**, 566 (1962).
- [20] J.M. Stewart, *Advanced General Relativity* (Cambridge University Press, Cambridge, 1990).
- [21] R. Penrose and W. Rindler, *Spinors and Space-time* (Cambridge University Press, Cambridge, 1986), Vol. 2.
- [22] R. Penrose, *Proc. R. Soc. A* **284**, 159 (1965).
- [23] I. Hinder *et al.* (The NRAR Collaboration, Perimeter Institute for Theoretical Physics), [arXiv:1307.5307](https://arxiv.org/abs/1307.5307).
- [24] P. Ajith *et al.*, *Classical Quantum Gravity* **29**, 124001 (2012).
- [25] P. Ajith *et al.*, *Phys. Rev. Lett.* **106**, 241101 (2011).
- [26] J. Healy, P. Laguna, L. Pekowsky, and D. Shoemaker, *Phys. Rev. D* **88**, 024034 (2013).
- [27] N.T. Bishop, R. Gómez, L. Lehner, M. Maharaj, and J. Winicour, *Phys. Rev. D* **56**, 6298 (1997).
- [28] J. Winicour, *Living Rev. Relativity* **8**, 10 (2005), <http://www.livingreviews.org/lrr-2005-10>.
- [29] R. Gomez, W. Barreto, and S. Frittelli, *Phys. Rev. D* **76**, 124029 (2007).
- [30] C. Reisswig, N.T. Bishop, and D. Pollney, *Gen. Relativ. Gravit.* **45**, 1069 (2013).
- [31] C. Reisswig, N.T. Bishop, D. Pollney, and B. Szilágyi, *Phys. Rev. Lett.* **103**, 221101 (2009).
- [32] M.C. Babiuc, B. Szilágyi, J. Winicour, and Y. Zlochower, *Phys. Rev. D* **84**, 044057 (2011).
- [33] D. Pollney and C. Reisswig, *Astrophys. J. Lett.* **732**, L13 (2011).
- [34] C. Reisswig, C. Ott, U. Sperhake, and E. Schnetter, *Phys. Rev. D* **83**, 064008 (2011).
- [35] C. Reisswig, C. Ott, E. Abdikamalov, R. Haas, P. Moesta, and E. Schnetter, *Phys. Rev. Lett.* **111**, 151101 (2013).
- [36] C. Reisswig *et al.*, *Phys. Rev. D* (to be published).
- [37] C.D. Ott, C. Reisswig, E. Schnetter, E. O’Connor, U. Sperhake, F. Löffler, P. Diener, E. Abdikamalov, I. Hawke, and A. Burrows, *Phys. Rev. Lett.* **106**, 161103 (2011).
- [38] T. Regge and J.A. Wheeler, *Phys. Rev.* **108**, 1063 (1957).
- [39] F.J. Zerilli, *Phys. Rev. Lett.* **24**, 737 (1970).
- [40] F.J. Zerilli, *Phys. Rev. D* **2**, 2141 (1970).
- [41] V. Moncrief, *Ann. Phys. (N.Y.)* **88**, 323 (1974).
- [42] A. Nagar and Rezzolla, *Classical Quantum Gravity* **22**, R167 (2005).
- [43] A.M. Abrahams and C.R. Evans, *Phys. Rev. D* **37**, 318 (1988).
- [44] A.M. Abrahams and C.R. Evans, *Phys. Rev. D* **42**, 2585 (1990).
- [45] C.O. Lousto, H. Nakano, Y. Zlochower, and M. Campanelli, *Phys. Rev. D* **82**, 104057 (2010).
- [46] A.G. Benedict, S.E. Field, and S.R. Lau, *Classical Quantum Gravity* **30**, 055015 (2013).
- [47] D. Pollney, C. Reisswig, E. Schnetter, N. Dorband, and P. Diener, *Phys. Rev. D* **83**, 044045 (2011).
- [48] C. Reisswig, N.T. Bishop, D. Pollney, and B. Szilágyi, *Classical Quantum Gravity* **27**, 075014 (2010).

- [49] H. P. Pfeiffer, D. A. Brown, L. E. Kidder, L. Lindblom, G. Lovelace, and M. A. Scheel, *Classical Quantum Gravity* **24**, S59 (2007).
- [50] M. Boyle, D. A. Brown, L. E. Kidder, A. H. Mroué, H. P. Pfeiffer, M. A. Scheel, G. B. Cook, and S. A. Teukolsky, *Phys. Rev. D* **76**, 124038 (2007).
- [51] M. Boyle and A. H. Mroué, *Phys. Rev. D* **80**, 124045 (2009).
- [52] R. Arnowitt, S. Deser, and C. Misner, *Gen. Relativ. Gravit.* **40**, 1997 (2008).
- [53] M. Boyle, *Phys. Rev. D* **87**, 104006 (2013).
- [54] K. S. Thorne, *Rev. Mod. Phys.* **52**, 299 (1980).
- [55] L. Blanchet and T. Damour, *Proc. R. Soc. A* **320**, 379 (1986).
- [56] M. Boyle, Ph.D. thesis, California Institute of Technology, 2008.
- [57] N. T. Bishop, R. Gómez, R. A. Isaacson, L. Lehner, B. Szilágyi, and J. Winicour, in *Black Holes, Gravitational Radiation and the Universe*, edited by B. R. Iyer and B. Bhawal (Kluwer, Dordrecht, 1998), Chap. 24.
- [58] B. Szilágyi and J. Winicour, *Phys. Rev. D* **68**, 041501 (2003).
- [59] H. Bondi, M. G. J. van der Burg, and A. W. K. Metzner, *Proc. R. Soc. A* **269**, 21 (1962).
- [60] R. K. Sachs, *Proc. R. Soc. A* **270**, 103 (1962).
- [61] N. Bishop, D. Pollney, and C. Reisswig, *Classical Quantum Gravity* **28**, 155019 (2011).
- [62] M. C. Babiuc, N. T. Bishop, B. Szilágyi, and J. Winicour, *Phys. Rev. D* **79**, 084011 (2009).
- [63] C. Reisswig and D. Pollney, *Classical Quantum Gravity* **28**, 195015 (2011).
- [64] N. T. Bishop and C. Reisswig, [arXiv:1308.1521](https://arxiv.org/abs/1308.1521).
- [65] C. Reisswig, N. T. Bishop, C. W. Lai, J. Thornburg, and B. Szilágyi, *Classical Quantum Gravity* **24**, S327 (2007).
- [66] G. Lovelace, M. A. Scheel, and B. Szilágyi, *Phys. Rev. D* **83**, 024010 (2011).
- [67] S. Ossokine, L. E. Kidder, and H. P. Pfeiffer, *Phys. Rev. D* **88**, 084031 (2013).
- [68] L. Lindblom, M. A. Scheel, L. E. Kidder, R. Owen, and O. Rinne, *Classical Quantum Gravity* **23**, S447 (2006).
- [69] H. Friedrich, *Commun. Math. Phys.* **100**, 525 (1985).
- [70] D. Garfinkle, *Phys. Rev. D* **65**, 044029 (2002).
- [71] F. Pretorius, *Classical Quantum Gravity* **22**, 425 (2005).
- [72] C. Gundlach, J. M. Martin-Garcia, G. Calabrese, and I. Hinder, *Classical Quantum Gravity* **22**, 3767 (2005).
- [73] O. Rinne, *Classical Quantum Gravity* **23**, 6275 (2006).
- [74] J. M. Stewart, *Classical Quantum Gravity* **15**, 2865 (1998).
- [75] H. Friedrich and G. Nagy, *Commun. Math. Phys.* **201**, 619 (1999).
- [76] J. M. Bardeen and L. T. Buchman, *Phys. Rev. D* **65**, 064037 (2002).
- [77] B. Szilágyi, B. Schmidt, and J. Winicour, *Phys. Rev. D* **65**, 064015 (2002).
- [78] G. Calabrese, J. Pullin, O. Reula, O. Sarbach, and M. Tiglio, *Commun. Math. Phys.* **240**, 377 (2003).
- [79] B. Szilágyi and J. Winicour, *Phys. Rev. D* **68**, 041501(R) (2003).
- [80] L. E. Kidder, L. Lindblom, M. A. Scheel, L. T. Buchman, and H. P. Pfeiffer, *Phys. Rev. D* **71**, 064020 (2005).
- [81] D. Gottlieb and J. S. Hesthaven, *J. Comput. Appl. Math.* **128**, 83 (2001).
- [82] J. S. Hesthaven, *Applied Numerical Mathematics* **33**, 23 (2000).
- [83] L. Lindblom and B. Szilágyi, *Phys. Rev. D* **80**, 084019 (2009).
- [84] M. W. Choptuik and F. Pretorius, *Phys. Rev. Lett.* **104**, 111101 (2010).
- [85] G. Lovelace, R. Owen, H. P. Pfeiffer, and T. Chu, *Phys. Rev. D* **78**, 084017 (2008).
- [86] G. Lovelace, M. Boyle, M. A. Scheel, and B. Szilágyi, *Classical Quantum Gravity* **29**, 045003 (2012).
- [87] A. Buonanno, L. E. Kidder, A. H. Mroué, H. P. Pfeiffer, and A. Taracchini, *Phys. Rev. D* **83**, 104034 (2011).
- [88] T. Damour, B. R. Iyer, and B. S. Sathyaprakash, *Phys. Rev. D* **57**, 885 (1998).
- [89] E. E. Flanagan and S. A. Hughes, *Phys. Rev. D* **57**, 4535 (1998).
- [90] 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).
- [91] P. Schmidt, M. Hannam, S. Husa, and P. Ajith, *Phys. Rev. D* **84**, 024046 (2011).
- [92] R. O'Shaughnessy, B. Vaishnav, J. Healy, Z. Meeks, and D. Shoemaker, *Phys. Rev. D* **84**, 124002 (2011).
- [93] M. Boyle, R. Owen, and H. P. Pfeiffer, *Phys. Rev. D* **84**, 124011 (2011).
- [94] E. Ochsner and R. O'Shaughnessy, *Phys. Rev. D* **86**, 104037 (2012).
- [95] P. Schmidt, M. Hannam, and S. Husa, *Phys. Rev. D* **86**, 104063 (2012).
- [96] M. Alcubierre, *Introduction to 3 + 1 Numerical Relativity* (Oxford University Press, New York, 2008).
- [97] J. R. Taylor, *An Introduction to Error Analysis: The Study of Uncertainties in Physical Measurements* (University Science Books, Sausalito, CA, 1997).
- [98] L. Lindblom, K. D. Matthews, O. Rinne, and M. A. Scheel, *Phys. Rev. D* **77**, 084001 (2008).
- [99] C. Loken *et al.*, *J. Phys. Conf. Ser.* **256**, 012026 (2010).