



CHORUS

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

High accuracy binary black hole simulations with an extended wave zone

Denis Pollney, Christian Reisswig, Erik Schnetter, Nils Dorband, and Peter Diener

Phys. Rev. D **83**, 044045 — Published 23 February 2011

DOI: [10.1103/PhysRevD.83.044045](https://doi.org/10.1103/PhysRevD.83.044045)

High accuracy binary black hole simulations with an extended wave zone

Denis Pollney,¹ Christian Reisswig,² Erik Schnetter,^{3,4} Nils Dorband,² and Peter Diener^{4,3}

¹*Max-Planck-Institut für Gravitationsphysik, Albert-Einstein-Institut, Potsdam-Golm, Germany*

²*Max-Planck-Institut für Gravitationsphysik, Albert-Einstein-Institut, Potsdam-Golm, Germany*

³*Center for Computation & Technology, Louisiana State University, Baton Rouge, LA, USA*

⁴*Department of Physics & Astronomy, Louisiana State University, Baton Rouge, LA, USA*

We present results from a new code for binary black hole evolutions using the moving-puncture approach, implementing finite differences in generalised coordinates, and allowing the spacetime to be covered with multiple communicating non-singular coordinate patches. Here we consider a regular Cartesian near zone, with adapted spherical grids covering the wave zone. The efficiencies resulting from the use of adapted coordinates allow us to maintain sufficient grid resolution to an artificial outer boundary location which is causally disconnected from the measurement. For the well-studied test-case of the inspiral of an equal-mass non-spinning binary (evolved for more than 8 orbits before merger), we determine the phase and amplitude to numerical accuracies better than 0.010% and 0.090% during inspiral, respectively, and 0.003% and 0.153% during merger. The waveforms, including the resolved higher harmonics, are convergent and can be consistently extrapolated to $r \rightarrow \infty$ throughout the simulation, including the merger and ringdown. Ringdown frequencies for these modes (to $(\ell, m) = (6, 6)$) match perturbative calculations to within 0.01%, providing a strong confirmation that the remnant settles to a Kerr black hole with irreducible mass $M_{\text{irr}} = 0.884355 \pm 20 \times 10^{-6}$ and spin $S_f/M_f^2 = 0.686923 \pm 10 \times 10^{-6}$.

PACS numbers: 04.25.dg, 04.30.Db, 04.30.Tv, 04.30.Nk

I. INTRODUCTION

The numerical solution of Einstein’s equations has made great progress in recent years. Orbits and mergers of binary systems of black holes and neutron stars are now routinely published by a number of research groups, using independent codes and methodologies [1–4]. A number of important astrophysical phenomena associated with binary black hole mergers have been studied in some detail. In particular, the recoil of the merger remnant has been studied for a number of different initial data models [5–12], and its final mass and spin has been mapped out for fairly generic merger models involving spinning and unequal mass black holes [13–18]. Since these quantities are determined by the last few quasi-circular orbits before merger, they can be calculated to good approximation with fairly short evolutions, and with minimal influence of an artificial outer boundary.

Of particular topical relevance, however, is the construction of long waveforms which can be used for gravitational-wave analysis of the binary [19], and also to construct a family of templates [20–23], so to inform and improve gravitational wave detection algorithms. Here the requirements are particularly challenging for numerical simulations, requiring waveforms which are accurate in phase and amplitude over multiple cycles to allow for an unambiguous matching to post-Newtonian waveforms at large separation. Some recent studies have shown very promising results in this direction for particular binary black hole models [24–32]. However, they have also highlighted the problems associated with producing long waveforms of sufficient accuracy.

First of all, for binaries with a larger separation, systematic errors associated with gravitational waveform extraction at a finite radius become more pronounced. Typically a number of extraction radii are used, and the results extrapolated to infinite radius (assuming such a consistent extrapolation ex-

ists given potential issues of gauge). In order to have some confidence in the results, the outermost “extraction sphere” needs to be at a large radius, say on the order of $150 - 200M$ (where M is the mass of the system and sets the fiducial length scale). Even at this radius, the amplitude of the extrapolated waveform differs significantly from the measured waveform. Unfortunately, extracting at larger radii comes at a computational expense. One of the standard methods in use today is finite differencing in conjunction with “mesh refinement”, in which the numerical resolution is chosen based on the length scale of the problem. A minimum number of discrete data points are required to resolve a feature of a given size accurately, which sets a limit on the minimum resolution which should be applied in a region. Thus, even with mesh refinement there is a limit on the coarseness of the grid which can be allowed in the wave-zone. For a Cartesian grid, the number of computational points scales as r^3 , so that requiring a sufficient resolution to $200M$ already comes at significant expense, and increasing this distance further becomes impractical.

An additional difficulty arises from the requirement that the outer boundary have minimal influence on the interior evolution, since it is (in all current binary black hole codes) an artificial boundary. This places an additional requirement on the size of the computational grids, so that even outside the wave-zone region where the physics is accurately resolved, it is conventional to place several even coarser grids. This is done in the knowledge that the physical variables can not be resolved in these regions, but the grids are helpful as a numerical buffer between the outer boundary and interior domain. Again, adding these outer zones comes at a computational expense. The boundaries with under-resolved regions also lead to unphysical reflections which can contaminate the solution. The problem of increasing the grid size can be significantly reduced if, rather than a Cartesian coordinate system, one uses a discretisation which has a radial coordinate. Then, for a fixed

angular resolution, the number of points on the discrete grid increases simply as a linear function of r , rather than the r^3 of the Cartesian case. This has two advantages. The gravitational wave-zone has spherical topology and therefore, a numerical approximation would be most efficiently represented by employing a spherical grid. A further computational motivation comes from the fact that non-synchronous mesh-refinement (such as the Berger-Oliger algorithm) can greatly complicate the parallelisation of an evolution scheme, and thus having many levels of refinement clearly has an impact on the efficiency of large scale simulations. This will become particularly relevant for the coming generations of peta-scale machines which strongly emphasise parallel execution (possibly over several thousand cores) over single processor performance.

The use of spherical-polar coordinates has largely been avoided in 3-dimensional general relativity due to potential problems associated with the coordinate singularity at the poles. Additionally, even if regularisation were simple, the inhomogeneous areal distribution of latitude-longitude grid points over the sphere make spherical-polar coordinates sub-optimal. A number of alternative coordinate systems have been proposed and implemented for studies of black holes in 3D. The Pittsburgh null code avoids the problem of regularisation at the poles by implementing a 2D stereographic patch system [33]. Cornell/Caltech have developed a multipatch system which has been used for long binary black hole evolutions [4, 34]¹. This code, using spectral spatial differentiation, uses an intricate patch layout in order to reduce the overall discretisation error. The boundary treatment between patches is based on the transfer of characteristic variables. A similar approach was implemented by the LSU group, for the case of finite differences with penalty boundary conditions [38], and used to successfully evolve single perturbed black holes with a fixed background [39] and have recently been attempted for binary black hole systems [40].

In this paper we describe a binary black hole evolution code based on adapted radial coordinates in the wave zone, for generic evolution systems. In particular, we demonstrate stable and accurate binary black hole evolutions using BSS-NOK in conjunction with this coordinate system. The grids in the wave zone follow a prescription which was first used by Thornburg [41], in which six regular patches cover the sphere, and data at the boundaries of the patches are filled by interpolation. Such a patch system has also been successfully applied to characteristic evolutions [42, 43]. In this work, the six patch wave zone is coupled to an interior Cartesian code, which covers the domain in which the bodies move, and optionally allows for mesh refinement around each of the individual bodies. The resulting code has the advantages of making use of established techniques for moving puncture evolutions on Cartesian grids, while having excellent efficiency (and consequently accuracy) in the wave zone due to the use of adapted radially-oriented grids.

In the following sections we detail the coordinate structures which we use. We then describe our Einstein evolution code, which is largely based on conventional techniques common to Cartesian puncture evolutions. Finally we perform evolutions of a binary black hole system in order to validate the code against known results, as well as demonstrate the ability to extract accurate waves at a large radius with comparatively low computational cost.

II. SPACETIME DISCRETISATION

This section describes the implementation of a generic code infrastructure for evolving spacetimes which are covered by multiple overlapping grid patches. A key feature of our approach is its flexibility. It is not restricted to any particular formulation of the Einstein equations; the mechanism for passing data between patches (interpolation) is also formulation independent (though characteristic [44] or penalty-patch boundaries [40, 45, 46] are also an option); the size, placement and local coordinates of individual patches are completely adaptable, provided that there is sufficient overlap between neighbours to transfer boundary data. Further, each patch is a locally Cartesian grid with the ability to perform mesh-refinement to better resolve localised steep gradients, if necessary. The particular application demonstrated in this paper is to provide a more efficient covering of the wave-zone of an isolated binary black hole inspiral.

At the same time, we would like to take advantage of the fact that black hole evolutions via the “moving puncture” approach are well established as a simple and effective method for stably evolving black hole spacetimes [2, 3]. By this method, gauge conditions are applied to prevent the spacetime from reaching the curvature singularity, so that an artificial boundary is not required within the horizons [47]. The usual approach is to discretise using Cartesian grids which cover the black holes with an appropriate resolution, without special treatment or boundary conditions for the black hole interiors, relying rather on the causal structure of the evolution system to prevent error modes from emerging [48]. The Cartesian grids are extended to cover the wave zone (at reduced resolution for the sake of efficiency), extending to a cubical grid outer boundary where an artificial condition is applied.

A principal difficulty faced by this method is that the discretisation is not well suited to model radial waves at large radii. In order to resolve the wave profile, a certain minimum radial resolution is required and must be maintained as the wave propagates to large radii. The angular resolution, however, can remain fixed – if a wave is resolved at a certain angular resolution as small radii, then it is unlikely to develop significant angular features as it propagates to large distances from the isolated source. Cartesian grids with fixed spacing, however, resolve spheres with an angular resolution which scales according to r^2 . Thus, to maintain a given required radial resolution, the angular directions become extremely over-resolved at large radii, and this comes at a large computational cost. Namely, for a Cartesian grid to extend in size or increase its resolution by a factor n , the cost in memory and number of

¹ Multi-domain spectral methods have previously been applied to the problem of generating initial data for binaries in [35–37].

computations per timestep increase by n^3 , while for a radial grid with fixed angular resolution, the increase is linear, n^2 .

For the near-zone, in the neighbourhood of the orbits of the individual bodies, the geometrical situation is not as straightforward, since there is no clearly defined radial propagation direction between the bodies. If the local geometry is reasonably well known (for instance, the location of horizon surfaces), adapted coordinates can also be considered in this regime. The technical requirements of such coordinate systems can, however, be high. Since the bodies are moving, the grids must move with them, or dynamical gauges chosen such that the bodies remain in place relative to the numerical coordinates. Potential problems arise from the coordinate singularity if the grids are extended to $r = 0$, as is the case with the standard puncture approach. Thus, in the near-zone, Cartesian coordinates can provide significant simplification to the overall infrastructure requirements, while the previously mentioned drawbacks of Cartesian coordinates are less prevalent, as it is useful to have homogeneous resolution in each direction in situations where there is no obvious symmetry.

The evolution code which we have constructed for the purpose of modelling waveforms from an isolated system is based on a hybrid approach, involving a Cartesian mesh-refined region covering the near zone in which the bodies orbit, and a set of adapted radial grids which efficiently cover the wave zone. The overall structure is illustrated in Fig. 1 (top), which shows an equatorial slice of the numerical grid. Computations on each local patch are carried out in a globally Cartesian coordinate system. In the particular implementation considered here, the grids overlap by some distance so that data at the boundaries between each local coordinate patch can be communicated by interpolation from neighbouring patches. The resulting code is both efficient, but also simple in structure and able to take advantage of well established methods for evolving moving puncture black holes. If suitable interpolation methods are used, then such a system can also be used for solutions with discontinuities and shocks as are present in hydrodynamics.

The code has been implemented within the `Cactus` framework [49, 50] via extensions to the `Carpet` driver [51–53], which handles parallelisation via domain decomposition of grids over processors, as well as providing the required interpolation operators for boundary communication and analysis tools.

A. Coordinate systems

The configuration displayed in Fig. 1 consists of seven local coordinate patches: an interior Cartesian grid, and six outer patches corresponding to the faces of the interior cube. Each patch consists of a uniformly spaced (in local coordinates) grid which can be refined (though in practise we will only use

this feature for the interior grid). The outer patches have a local coordinate system (ρ, σ, R) corresponding to the “inflated cube” coordinates which have previously been used in relativity for single black hole evolutions [41], and are displayed in the lower plot of Fig. 1. The local angular coordinates (ρ, σ) range over $(-\pi/4, +\pi/4) \times (-\pi/4, +\pi/4)$ and can be related to global angular coordinates (μ, ν, ϕ) which are given by

$$\mu \equiv \text{rotation angle about the x-axis} = \arctan(y/z), \quad (1a)$$

$$\nu \equiv \text{rotation angle about the y-axis} = \arctan(x/z), \quad (1b)$$

$$\phi \equiv \text{rotation angle about the z-axis} = \arctan(y/x). \quad (1c)$$

For example, on the $+z$ patch, the mapping between the local (ρ, σ, R) and Cartesian (x, y, z) coordinates is given by:

$$\rho \equiv \nu = \arctan(x/z), \quad (2a)$$

$$\sigma \equiv \mu = \arctan(y/z), \quad (2b)$$

$$R = f(r), \quad (2c)$$

with appropriate rotations for each of the other cube faces, and where $r = \sqrt{x^2 + y^2 + z^2}$. As emphasised by Thornburg [41], in addition to avoiding pathologies associated with the axis of standard spherical polar coordinates, this choice of local coordinates has a number of advantages. In particular, the angular coordinates on neighbouring patches align so that interpolation is only 1-dimensional, in a line parallel to the face of the patch. This potentially improves the efficiency of the interpolation operation as well as the accuracy. The coordinates also cover the sphere more uniformly than, say, a stereographic 2-patch system.

The local radial coordinate, R , is determined as a function of the global coordinate radius, r . We can use this degree of coordinate freedom to increase the physical (global) extent of the wave-zone grids, at the cost of some spatial resolution. In practise, we use a function of the form

$$f(r) = A(r - r_0) + B\sqrt{1 + (r - r_0)^2/\epsilon}, \quad (3a)$$

with

$$R = f(r) - f(0). \quad (3b)$$

in order to transition between two almost constant resolutions (determined by the parameters A and B) over a region whose width is determined by ϵ , centred at r_0 .

The effect of the radial transformation is illustrated in Fig. 2. The coordinate R is a nearly constant rescaling of r at small and large radii. The change in the scale factor is largely confined to a transition region. Note that since we apply the same global derivative operators (described below) to analysis tools as are used for the the evolution, it is possible to do analysis (e.g., measure waveforms, horizon finding) within regions where the radial coordinate is non-uniform. The regions of near-constant resolution are, however, useful in order to make comparisons of measurements at different radii without the additional complication of varying numerical error due to the underlying grid spacing.

Data on each patch are evaluated at grid-points which are placed at uniformly spaced points of a Cartesian grid. Thus,

² Note that the Courant limit introduces an additional factor of n in each case due to the requirement of a reduced timestep with increasing resolution.




FIG. 1. A schematic view of the $z = 0$ slice of the grid setup that is used. The upper plot shows the central Cartesian grid surrounded by six “inflated-cube” patches (the four equatorial patches are shown, shaded). The boundaries of the nominal grids owned by each patch are indicated by thick lines. The lower plot shows an $r = \text{constant}$ surface of the exterior patches, indicating their local coordinate lines.

local derivatives can be calculated via standard finite difference techniques. These are then transformed to a common underlying Cartesian coordinate system by applying an appropriate Jacobian which relates the local to global coordinates. That is, the global (Cartesian) coordinates, x_i , are related to the local coordinates, a_i , by

$$x_i = x_i(a_j), \quad i, j = 0, 1, 2. \quad (4)$$

and derivatives, $\partial/\partial a_i$, of fields are determined using finite differences in the regularly spaced a_i coordinates, which are




FIG. 2. The local radial coordinate, R (solid line), can be stretched as a function of the global coordinate, r , in order to increase the effective size of the grid. The function specified by Eqs. (3) transitions between two almost constant radial resolutions over a distance ϵ centred at r_0 .

then transformed using

$$\frac{\partial}{\partial x_i} = \left(\frac{\partial a_j}{\partial x_i} \right) \frac{\partial}{\partial a_j}, \quad (5a)$$

$$\frac{\partial^2}{\partial x_i \partial x_j} = \left(\frac{\partial^2 a_k}{\partial x_i \partial x_j} \right) \frac{\partial^2}{\partial a_k^2} + \left(\frac{\partial a_k}{\partial x_i} \frac{\partial a_l}{\partial x_j} \right) \frac{\partial^2}{\partial a_k \partial a_l}, \quad (5b)$$

in order to determine their values in the global frame. We store and evaluate tensor components and their evolution equations in the common global frame, so that there is no need to apply transformations when relating data across patch boundaries. In addition to the obvious simplification of the inter-patch boundary treatment, this has a number of other advantages, particularly when it comes to analysis tools (surface finding, gravitational wave measurements, visualisation) which may reference data on multiple patches. Since the data is represented in the common global basis, these tools do not need to know anything about the local grid structures or coordinates.

B. Inter-patch interpolation

Data is communicated between patches by interpolating onto overlapping points. Each patch, p , is responsible for determining the numerical solution for a region of the spacetime. The boundaries of these patches can overlap neighbouring patches, q , (and in fact must do so for the case of the interpolating boundaries considered here), creating regions of the spacetime which are covered by multiple patches. We define three sets of points on a patch. The *nominal* regions, \mathcal{N}_p , contain the points where the numerical solution is to be determined. The nominal regions of the patches do not overlap, $\bigcap_p \mathcal{N}_p = \emptyset$, so that if data is required at any point in the spacetime, it can be obtained without ambiguity by referencing the single patch in whose nominal region it resides. A patch, p , is bounded by a layer of *ghost* points, \mathcal{G}_p , which overlap the nominal zones of neighbouring patches, q , $\mathcal{G}_p \cap \bigcup_q \mathcal{N}_q = \mathcal{G}_p$, and are filled by interpolation. (These points are conceptually similar to the inter-processor ghost-zones used by domain decomposition parallelisation algorithms in order to divide grids over processors.) The size of these regions is determined by the width of the finite difference stencil to be used in finite




FIG. 3. Schematic of interpolating patch boundaries in 1-dimension, assuming 4-point finite difference and interpolation stencils. Points in the nominal zones, $\mathcal{N}_{p,q}$, are indicated by filled circles, points in ghost zones, $\mathcal{G}_{p,q}$, by open squares, and points in overlap zones, $\mathcal{O}_{p,q}$, by closed squares. The vertical dotted line demarcates the boundary between nominal zones on each patch. Ghost points on patch p are evaluated by centred interpolation operations from points in \mathcal{S}_q on the overlapping patch (arrows) and *vice versa*.

differencing the evolution equations on the nominal grid. Finally, an additional layer of *overlap* points, \mathcal{O}_q , is required: i) to ensure that the set of stencil points, $\mathcal{S}_q \subset \mathcal{O}_q \cup \mathcal{N}_q$, used to interpolated to the ghost zone does not itself originate from the ghost zone of the interpolating patch, $\mathcal{S}_q \cap \mathcal{G}_q = \emptyset$, $\mathcal{O}_q \cap \bigcup_p \mathcal{N}_p = \mathcal{O}_q$; and ii) to compensate for any difference in the grid spacing between the local coordinates on the two patches. An illustration of the scheme in 1-dimension the scheme is provided in Fig. 3.

Note that points in $\bigcup_q \mathcal{O}_q \subset \bigcup_p \mathcal{N}_p$ are not interpolated, but rather are evolved in the same way as nominal grid points within $\bigcup_p \mathcal{N}_p$. That is, in these regions points on each grid are evolved independently, and is in principle multi-valued. However, since the union of set of nominal points on each patch $\bigcup_p \mathcal{N}_p$ uniquely and unambiguously covers the entire simulation domain, *i.e.* $\bigcap_p \mathcal{N}_p = \emptyset$, and since the overlap regions are a subset of the nominal grid, if data is required at a point within these overlap zones, there is exactly one patch owing this point on its nominal grid, and it will be returned uniquely from this patch. The differences between evolved field values evaluated in these overlap points converge away with the finite difference order of the evolution scheme.

The use of additional overlap points makes this inter-patch interpolation algorithm somewhat simpler than the one implemented by Thornburg in [41]. That algorithm required inter-patch boundary conditions to be applied in a specific order to ensure that all interpolation stencils are evaluated without using undefined grid points, and requires off-centring interpolation stencils under certain circumstances, which is not necessary when overlap points are added. It also relies on the particular property of the inflated-cube coordinates which ensured that the ghost-zones could be filled using 1-dimensional interpolation in a direction orthogonal to the boundary. This property would be non-trivial (and often impossible) to generalise to match arbitrary patch boundaries, such as that between the Cartesian and radially oriented grids of Fig. 1.

Another significant difference between Thornburg's approach and the approach presented here is that former stores tensor components in the patch-local frame, while we store them in the global coordinate frame. Evaluating components in the patch-local frame requires a basis transformation while interpolating. This is further complicated in the case of non-tensorial quantities (such as the $\bar{\Gamma}^i$ of the BSSNOK formula-

tion) which have quite complicated basis transformation rules involving spatial derivatives. Instead, we store tensor components in the global coordinate frame, which requires no basis transformation during inter-patch interpolations.

The number of ghost points in \mathcal{G}_p can be reduced using finite difference stencils which become lop-sided towards the boundaries of the patch, and may provide an important optimisation since interpolation between grids can be expensive, particularly if processor communication is involved. However, this tends to be at the cost of increased numerical error in the finite difference operations towards the grid boundaries. We have generally found it preferable to use centred stencils throughout the nominal, \mathcal{N}_p , and overlap, \mathcal{O}_p , zones and have applied certain optimisations to the interpolation operators as described below. Another optimisation can be achieved by using lower order interpolation so that it is possible to reduce the number of overlapping points in \mathcal{O}_p .

The interpolation scheme for evaluating data in ghost zones is based on Lagrange polynomials using data from the overlapping patch. In 1-dimension, the Lagrange interpolation polynomial can be written as

$$\mathcal{L}_x[\phi](x) = \sum_i^N b_i(x) \phi_i, \quad (6a)$$

where the coefficients are

$$b_i(x) = \prod_{k \neq i} \frac{(x - x_k)}{(x_i - x_k)}. \quad (6b)$$

In these expressions, $x \in \mathcal{G}_p$ is the coordinate of the interpolation point and $\phi_i \in \mathcal{S}_q \subset \mathcal{N}_q \cup \mathcal{O}_q$ are the values at grid-points in the interpolation molecule surrounding x . The number of grid-points in the interpolation molecule, N , determines the interpolation order, and interpolation of n -th order accuracy is given by $N = n + 1$ stencil points in the molecule.

For interpolation in d -dimensions, the interpolation polynomial can be constructed as a tensor product of 1-dimensional Lagrange interpolation polynomials along coordinate directions, $\mathbf{x} = (x^1, \dots, x^d)$:

$$\begin{aligned} \mathcal{L}[\phi](\mathbf{x}) &= \mathcal{L}_{x^1}[\phi](x^1) \otimes \dots \otimes \mathcal{L}_{x^d}[\phi](x^d) \\ &= \left(\sum_i^N b_i(x^1) \phi_i \right) \cdots \left(\sum_j^N c_j(x^d) \phi_j \right). \quad (7) \end{aligned}$$

Therefore, for d -dimensional interpolation of order n , one has to determine N^d neighbouring stencil points and associated interpolation coefficients, Eq. (6b), for *each* point in the ghost-zone of a given patch. Most generally, full 3-dimensional interpolation is required, though in particular cases coordinates between two patches can be constructed such that they align locally so that only 1-dimensional interpolation is needed. This is, for instance, the case for the overlap region between the inflated-cube spherical patches used here (see Fig. 1). We have optimised the current code to automatically take advantage of this.

In order to interpolate to a point for which the coordinates a_i^p given in the basis of patch p are given, we need to know

the patch owning the nominal region containing this point. For this we first convert a_i^p to the global coordinate basis x_i , then determine which patch q owns the corresponding nominal region \mathcal{N}_q , and then convert x_i to the local coordinate bases this patch a_i^q . By construction, patch q has sufficient overlap points to evaluate the interpolation stencil there:

$$x_i := \text{local-to-global}_p(a_i^p), \quad (8a)$$

$$q := \text{owning-patch}(x^i), \quad (8b)$$

$$a_i^q := \text{global-to-local}_q(x^i). \quad (8c)$$

The three operations “local-to-global”, “owning-patch”, and “global-to-local” depend on the patch system and their local coordinate systems.

We can then find the points of patch q that are closest to the interpolation point a_i^q in the local coordinates this patch. In order to find these points, we exploit the uniformity of the grid in local coordinates. The grid indices of the stencil points in a given direction are determined via

$$i \in \left\{ \text{floor}(j+k) \mid j = \frac{x-x_0}{\Delta x}, k = -\frac{n}{2}, \dots, \frac{n}{2} \right\}, \quad (9)$$

where x_0 is the origin of the local grid, n is the interpolation order, and “floor” denotes rounding downwards to the nearest integer.

There remains to be determined the refinement level which contains the region surrounding the interpolation point, as well as the processor that owns this part of the grid. For this purpose, an efficient tree-search algorithm has been implemented. In this algorithm, the individual patches and refinement levels are defined as “super-regions”, i.e., bounding boxes that delineate the global grid extent before processor decomposition. Each of these super-regions is recursively split into smaller regions. The leaves of the resulting tree structure represent the individual local components of the processor decomposition. Locating a grid point in this tree structure requires $O(\log n)$ operations on n processors, whereas a linear search (that would be necessary without a tree structure) would require $O(n)$ operations.

While the corresponding tree structure is generic, the actual algorithm used in `Carpet` splits the domain into a kd tree of depth d in $d = 3$ dimensions. That is, the domain is first split into k sub-domains in the x direction, each of these sub-domains is then independently split into several in the y direction, and each of these is then split in the z direction. This leads to cuboid sub-domains for each processor, where the sub-domains do not overlap, and where each sub-domain can have a different shape. `Carpet` balances the load so that each processor receives approximately the same number of grid points, while keeping the sub-domains’ shapes as close to a cube as possible.

Our implementation pre-calculates and stores most of the above information when the grid structure is set up, saving a significant amount of time during interpolation. In particular, the following are stored:

- For each ghost-point, the source patch (where the interpolation is performed), and the local coordinates on this patch;

- For each ghost-point, the interpolation stencil coefficients (6b);
- For each processor, the communication schedule specifying which interpolation points need to be sent to what other processor.

When the grid structure changes, for example, when a mesh-refinement grid is moved or resized, these quantities have to be recalculated.

Altogether, the inter-patch interpolation therefore consists of applying processor-local interpolation stencils, sending the results to other processors, receiving results from other processors, and storing these results in the local ghost-points. These are all operations requiring no look-up in complex data structures, and which consequently execute very efficiently on modern hardware.

C. Finite differencing

Spatial derivatives are computed using standard finite difference stencils, which have currently been implemented up to 8th-order [46]. The stencils are centred, except for the terms corresponding to an advection by the shift vector, of the form $\beta^i \partial_i u$ (see Sec. III, below). These derivatives are calculated using an “upwind” stencil which is shifted by one point in the direction of the shift, and of the same order. We find that these upwind stencils provide a significant increase in the numerical accuracy of the puncture motion at a given resolution (see Appendix A). The particular stencils which we use can be generated via a recursion algorithm, as outlined in [54].

On each patch we allow the local grids to be refined in order to increase the accuracy of computations in localised regions. For the application of the evolution of an isolated binary considered here, we only refine the central Cartesian grid in the neighbourhood the bodies. This is done using standard 2 : 1 Berger-Oliger mesh refinement techniques via the `Carpet` infrastructure [51–53]. The time step for the outer patches is taken to be the same as the coarse grid step of the interior patch, so that no time-interpolation is required at inter-patch boundaries.

Time integration is carried out using standard method-of-lines integrators. We find that for the time resolution we are using, a 4th-order Runge-Kutta (RK4) method provides a good compromise between sufficient accuracy and a low memory footprint. We set the time resolution of the outer grids according to the timestep of the coarsest Cartesian grid, which is limited by the Courant condition at the specified spatial resolution. By placing the Cartesian-spherical boundary at a small radius (and thus extending to finer Cartesian grids) we attain a high time resolution in the wave zone, potentially important for determining higher harmonics.

D. Surface integration and harmonic decomposition

A number of quantities of physical interest are measured by projecting them onto closed surfaces surrounding the source.

In particular, gravitational wave measurements rely on computations on constant coordinate spheres S^2 , parameterized by local spherical-polar coordinates (θ, ϕ) with constant coordinate radius r . In principle, it would be possible to construct coordinates on these 2-dimensional spheres which correspond to the underlying grid coordinates of the evolution, for instance as portrayed in the lower figure of Fig. 1. In this case, data can be mapped directly onto the spheres. More generally, however, interpolation can be used to obtain data at points on the measurement spheres, according to the procedure outlined in Sec. II B, above.

For the purpose of analysis, it is often convenient to decompose the data on S^2 in terms of (spin-weighted) spherical harmonic modes,

$$A_{\ell m} = \int d\Omega \sqrt{-g} A(\Omega)_s \bar{Y}_{\ell m}(\Omega), \quad (10)$$

where g is the determinant of the surface metric and Ω angular coordinates. In standard spherical-polar coordinates (θ, ϕ) ,

$$\sqrt{-g} = \sin^2 \theta. \quad (11)$$

The integral, Eq. (10), is solved numerically as follows. In the spherical polar case, we can take advantage of an highly accurate Gauss quadrature scheme which is exact for polynomials of order up to $N_\theta/2 - 1$, where N_θ is the number of gridpoints along the θ -direction. On a staggered grid, *i.e.* $\theta_j = (j + 1/2)\pi/N_\theta$, $j = 0, \dots, N_\theta - 1$, the scheme can be written as

$$\int d\Omega f(\Omega) = \sum_i^{N_\theta} \sum_j^{N_\phi} f_{ij} w_j + \mathcal{O}(N_\theta), \quad (12)$$

where N_θ and N_ϕ are the number of angular gridpoints and w_j are weight functions [55, 56],

$$w_j = \frac{2\pi}{N_\phi} \frac{4}{N_\theta} \sum_{l=0}^{N_\theta/2-1} \frac{1}{2l+1} \sin([2l+1]\theta_j), \quad (13)$$

N_θ even.

In our implementation, the weight functions are pre-calculated for fast surface integration.

III. EVOLUTION SYSTEM

We evolve the spacetime using a variant of the ‘‘BSSNOK’’ evolution system [57–60] and a specific set of gauges [61, 62], which have been shown to be effective at treating the coordinate singularities of Bowen-York initial data. We examine the well-posedness and characteristic propagation speeds of our system in [63, 64], where we prove (and also demonstrate numerically) that constraint violations introduced by this singularity treatment cannot escape the horizon.

The 4-geometry of a spacelike slice Σ (with timelike normal, n^α) is determined by its intrinsic 3-metric, γ_{ab} and extrinsic curvature, K_{ab} , as well as a scalar lapse function, α ,

and shift vector, β^a which determine the coordinate propagation. The standard BSSNOK system defines modified variables by performing a conformal transformation on the 3-metric,

$$\phi := \frac{1}{12} \ln \det \gamma_{ab}, \quad \tilde{\gamma}_{ab} := e^{-4\phi} \gamma_{ab}, \quad (14)$$

subject to the constraint

$$\det \tilde{\gamma}_{ab} = 1, \quad (15)$$

and by removing the trace of K_{ab} ,

$$K := \text{tr} K_{ij} = g^{ij} K_{ij}, \quad (16)$$

$$\tilde{A}_{ij} := e^{-4\phi} (K_{ij} - \frac{1}{3} \gamma_{ij} K), \quad (17)$$

with the constraint

$$\tilde{A} := \tilde{\gamma}^{ij} \tilde{A}_{ij} = 0. \quad (18)$$

Additionally, three new variables are introduced, defined in terms of the Christoffel symbols of $\tilde{\gamma}_{ab}$ by

$$\tilde{\Gamma}^a := \tilde{\gamma}^{ij} \tilde{\Gamma}_{ij}^a. \quad (19)$$

In principle the $\tilde{\Gamma}^a$ can be determined from the $\tilde{\gamma}_{ab}$, on a slice however their introduction is key to establishing a strongly hyperbolic (and thus stable) evolution system. In practise, only the constraint Eq. (18) is enforced during evolution [65], while Eq. (15) and Eq. (19) are simply monitored as indicators of numerical error. Thus, the traditional BSSNOK prescription evolves the variables

$$\phi, \quad \tilde{\gamma}_{ab}, \quad K, \quad \tilde{A}_{ab}, \quad \tilde{\Gamma}^a, \quad (20)$$

according to evolution equations which have been written down a number of times (see [66, 67] reviews).

In the context of puncture evolutions, it has been noted that alternative scalings of the conformal factor may exhibit better numerical behaviour in the neighbourhood of the puncture as compared with ϕ . In particular, a variable of the form

$$\hat{\phi}_\kappa := (\det \gamma_{ab})^{-1/\kappa}, \quad (21)$$

has been suggested [3, 68]. In [3], it is noted that certain singular terms in the evolution equations for Bowen-York initial data can be corrected using $\chi := \hat{\phi}_3$. Alternatively, [68] notes that $W := \hat{\phi}_6$ has the additional benefit of ensuring γ remains positive, a property which needs to be explicitly enforced with χ .

In terms of $\hat{\phi}_\kappa$, the BSSNOK evolution equations become:

$$\partial_t \hat{\phi}_\kappa = \frac{2}{\kappa} \hat{\phi}_\kappa \alpha K + \beta^i \partial_i \hat{\phi}_\kappa - \frac{2}{\kappa} \hat{\phi}_\kappa \partial_i \beta^i, \quad (22a)$$

$$\partial_t \tilde{\gamma}_{ab} = -2\alpha \tilde{A}_{ab} + \beta^i \partial_i \tilde{\gamma}_{ab} + 2\tilde{\gamma}_{i(a} \partial_{b)} \beta^i - \frac{2}{3} \tilde{\gamma}_{ab} \partial_i \beta^i, \quad (22b)$$

$$\partial_t K = -D_i D^i \alpha + \alpha (A_{ij} A^{ij} + \frac{1}{3} K^2) + \beta^i \partial_i K, \quad (22c)$$

$$\partial_t \tilde{A}_{ab} = (\hat{\phi}_\kappa)^{\kappa/3} (-D_a D_b \alpha + \alpha R_{ab})^{\text{TF}} + \beta^i \partial_i \tilde{A}_{ab} + 2\tilde{A}_{i(a} \partial_{b)} \beta^i - \frac{2}{3} A_{ab} \partial_i \beta^i, \quad (22d)$$

$$\begin{aligned} \partial_t \tilde{\Gamma}^a = & \tilde{\gamma}^{ij} \partial_i \beta_j \beta^a + \frac{1}{3} \tilde{\gamma}^{ai} \partial_i \partial_j \beta^j - \tilde{\Gamma}^i \partial_i \beta^a \\ & + \frac{2}{3} \tilde{\Gamma}^a \partial_i \beta^i - 2\tilde{A}^{ai} \partial_i \alpha \\ & + 2\alpha (\tilde{\Gamma}_{ij}^a \tilde{A}^{ij} - \frac{\kappa}{2} \tilde{A}^{ai} \frac{\partial_i \hat{\phi}_\kappa}{\hat{\phi}_\kappa} - \frac{2}{3} \tilde{\gamma}^{ai} \partial_i K), \end{aligned} \quad (22e)$$

where D_a is the covariant derivative determined by $\tilde{\gamma}_{ab}$, and ‘‘TF’’ indicates that the trace-free part of the bracketed term is used.

We have implemented the traditional ϕ form of the BSSNOK evolution equations, as well as the χ and W variants implicit in the evolution system, Eqs. (22). All three evolution systems produce stable evolutions of binary black holes, though the χ variant requires some special treatment if, due to numerical error particularly in the neighbourhood of the punctures, $\hat{\phi}_3 \leq 0$ [69]. Generally we find that the advection of the puncture (and thus the phase accuracy of the simulation) exhibits lower numerical error when using the χ and W variants (see Appendix C). Convergence properties of physical variables (such as measured gravitational waves, or constraints measured outside of the horizons), however, are not affected by the choice of conformal variable.

The Einstein equations are completed by a set of four constraints which must be satisfied on each spacelike slice:

$$\mathcal{H} \equiv R^{(3)} + K^2 - K_{ij} K^{ij} = 0, \quad (23a)$$

$$\mathcal{M}^a \equiv D_i (K^{ai} - \gamma^{ai} K) = 0. \quad (23b)$$

Again, we do not actively enforce these equations, but rather monitor their magnitude in order to determine whether our numerical solution is deviating from a solution to the Einstein equations.

The gauge quantities, α and β^a , are evolved using the prescriptions that have been commonly applied to BSSNOK black hole, and particularly puncture, evolutions. For the lapse, we evolve according to the ‘‘1 + log’’ condition [70],

$$\partial_t \alpha - \beta^i \partial_i \alpha = -2\alpha K, \quad (24)$$

while the shift is evolved using the hyperbolic ‘‘ $\tilde{\Gamma}$ -driver’’ equation [61],

$$\partial_t \beta^a - \beta^i \partial_i \beta^a = \frac{3}{4} \alpha B^a, \quad (25a)$$

$$\partial_t B^a - \beta^i \partial_i B^a = \partial_t \tilde{\Gamma}^a - \beta^i \partial_i \tilde{\Gamma}^a - \eta B^a, \quad (25b)$$

where η is a parameter which acts as a (mass dependent) damping coefficient, and is typically set to values on the order of unity for the simulations carried out here. The advective terms in these equations were not present in the original definitions of [61], where co-moving coordinates were used, but have been added following the experience of more recent studies using moving punctures [2, 62].

A. Wave extraction

The Newman-Penrose formalism [71] provides a convenient representation for a number of radiation related quantities as spin-weighted scalars. In particular, the curvature component

$$\psi_4 \equiv -C_{\alpha\beta\gamma\delta} n^\alpha \bar{m}^\beta n^\gamma \bar{m}^\delta, \quad (26)$$

is defined as a particular component of the Weyl curvature, $C_{\alpha\beta\gamma\delta}$, projected onto a given null frame, $\{\mathbf{l}, \mathbf{n}, \mathbf{m}, \bar{\mathbf{m}}\}$.

The identification of the Weyl scalar ψ_4 with the gravitational radiation content of the spacetime is a result of the peeling theorem [71–73], which states that in an appropriate frame and for sufficiently smooth and asymptotically flat initial data near spatial infinity, the ψ_4 component of the curvature has the slowest fall-off with radius, $\mathcal{O}(1/r)$.

The most straight-forward way of evaluating ψ_4 in numerical (Cauchy) simulations is to define an orthonormal basis in the three space $(\hat{r}, \hat{\theta}, \hat{\phi})$, centered on the Cartesian grid center and oriented with poles along \hat{z} . The normal to the slice defines a time-like vector \hat{t} , from which we construct the null frame

$$\mathbf{l} = \frac{1}{\sqrt{2}} (\hat{t} - \hat{r}), \quad \mathbf{n} = \frac{1}{\sqrt{2}} (\hat{t} + \hat{r}), \quad \mathbf{m} = \frac{1}{\sqrt{2}} (\hat{\theta} - i\hat{\phi}). \quad (27)$$

Note that in order to make the vectors $\{\mathbf{l}, \mathbf{n}, \mathbf{m}, \bar{\mathbf{m}}\}$ null, $(\hat{r}, \hat{\theta}, \hat{\phi})$ have to be orthonormal relative to the spacetime metric. In practice, we fix \hat{r} and then apply a Gram-Schmidt orthonormalization procedure to determine $\hat{\theta}$ and $\hat{\phi}$ ³. It is then possible to calculate ψ_4 via a reformulation of (26) in terms of the geometrical variables on the slice [75] via the electric and magnetic parts of the Weyl tensor,

$$\psi_4 = C_{ij} \bar{m}^i \bar{m}^j, \quad (28)$$

where

$$C_{ij} \equiv E_{ij} - iB_{ij} = R_{ij} - KK_{ij} + K_i^k K_{kj} - i\epsilon_i^{kl} \nabla_l K_{jk}. \quad (29)$$

³ Alternative frame constructions have also been used, such as orthonormalising relative to one of the angular basis vectors [74], or omitting the orthonormalisation altogether [4]. We have generally found these modifications do not lead to significantly different measurements

The remaining Weyl scalars can be similarly obtained and read

$$\psi_3 = \frac{1}{\sqrt{2}} C_{ij} \bar{m}^i e_r^j, \quad (30a)$$

$$\psi_2 = \frac{1}{2} C_{ij} e_r^i e_r^j, \quad (30b)$$

$$\psi_1 = -\frac{1}{\sqrt{2}} C_{ij} m^i e_r^j, \quad (30c)$$

$$\psi_0 = C_{ij} m^i m^j, \quad (30d)$$

where $(e_r^j) \equiv \hat{r}$ is the unit radial vector.

In relating ψ_4 to the gravitational radiation, one is limited by the fact that the measurements have been taken at a finite radius from the source. Local coordinate and frame effects can complicate the interpretation of ψ_4 . These problems can largely be alleviated by taking measurements at several radii and performing polynomial extrapolations to $r \rightarrow \infty$. Procedures for doing so have been studied in [76, 77]. In [77] we have shown that if a sufficiently large outermost extrapolation radius is used, the variation in this procedure is of the order $\Delta A = 0.03\%$ and $\Delta\phi = 0.003$ rad in amplitude and phase respectively, and is consistent throughout the evolution, including inspiral, merger and ringdown. The extrapolation error is larger than the numerical error determined in Sec. IV C 2, below, even if it is performed using data at $r = 1000M$ distant from the source, highlighting the need for measurements at large radii. For the ‘‘extrapolated’’ data plotted in this paper, we have performed polynomial extrapolations as detailed in [77], using the six outermost measurements at $r = \{280M, 300M, 400M, 500M, 600M, 1000M\}$.

In a companion paper [78], we use the same dataset to calculate ψ_4 directly at \mathcal{I}^+ using characteristic extraction [79, 80]. Here the traditional approach (which is gauge dependent and has a finite-radius cut-off error) presented here is replaced by a characteristic formulation of the Einstein equations in order to determine the fields out to future null infinity. In this paper, we restrict ourselves to a discussion of the numerical error inherent in the evolution procedure via the multi-patch code, and will report in more detail on systematic measurement errors due to finite radius effects and the characteristic extraction procedure elsewhere [78, 81].

IV. CODE VERIFICATION

A. Initial data

To demonstrate the efficacy of the infrastructure described in the previous sections, we have carried out an evolution of the by now well-studied case of the late-inspiral and merger of a pair of non-spinning equal-mass black holes (see, for example, [82] for an extensive discussion of numerical results involving this model). The particular numerical evolution which we have carried out starts from an initial separation $d/M = 11.0$ and goes through approximately 8 orbits (a physical time of around $1360M$), merger and ring-down. The masses of the punctures are set to $m = 0.4872$ and are initially placed on

the x -axis with momenta $p = (\pm 0.0903, \mp 0.000728, 0)$, giving the initial slice an ADM mass $M_{\text{ADM}} = 0.99051968 \pm 2 \times 10^{-8}$. These initial data parameters were determined using a post-Newtonian evolution from large initial separation, following the procedure outlined in [83], with the conservative part of the Hamiltonian accurate to 3PN, and radiation-reaction to 3.5PN, and determines orbits with a measured eccentricity of $e = 0.004 \pm 0.0005$.

B. Grid setup

The binary black hole evolution was carried out on a 7-patch grid structure, as described in Sec. II, incorporating a Cartesian mesh-refined region which covers the near-zone, and six radially oriented patches covering the wave-zone.

The inner boundary of the radial grids was placed at $r_t = 35.2M$ relative to the centre of the Cartesian grid. As a general rule, this boundary should be made as small as possible to improve efficiency in terms of memory usage. However other factors may make it preferable to move it further out. In particular, since we do not perform time interpolation at grid boundaries, the time step dt of the coarsest Cartesian grid determines the timestep of the radial grids, and thus the wave zone. Updates of the radial grids tend to be expensive, as they are large, so that if dt is too small, computation time may be spent over-resolving (in time) the wave zone. Particularly if the principle interest is in the lower order wave modes, it may be optimal to add an additional Cartesian mesh-refinement grid with a coarser time-step, and thus move r_t outwards.

The outer boundary for the spherical grids was chosen based on the expected time duration of the measurement and radius of the furthest detector, in order to remove any influence of the artificial outer boundary condition. In particular, given that the evolution takes a time T_m for the entire inspiral, merger and ringdown, and gravitational wave measurements taken at a finite radius r_d , we would like to ensure that a disturbance travelling at the speed of light from the outer boundary does not reach the measurement radius (see Fig. 4). For the BSSN evolution system, the physical modes travel at the speed of light, $c = 1$, in normal coordinates [63, 64], which are well-approximated in the wave zone where $\alpha \simeq 1$ and $\beta^r \simeq 0$. The $1 + \log$ slicing condition, Eq. (24) propagates at a speed of $\sqrt{2\alpha}$, however this is a gauge mode and empirically we find it to have negligible effect on measurements. The speed of propagation of numerical error from the boundary may also be larger than that of the physical modes, but does not affect the propagation of physically meaningful (i.e. continuum) quantities if the numerical scheme is consistent and convergent. Thus, taking note of the $v \simeq c = 1$ propagation speed of physical error modes from the boundary, we place the outer boundary at

$$r_b > T_m + 2r_d. \quad (31)$$

For the particular evolution considered here, $T_m \simeq 1350M$, and our outermost measurements are taken at $r_d = 1000M$. We have placed the outer boundary of the evolution domain at $r_b = 3600M$.




FIG. 4. Schematic of the causal propagation of information during the evolution. The gravitational wave source is located in the vicinity of $r = 0$, with waves propagating outward at the speed of light $c = 1$, and are measured at radius r_d for a time of interest which would include the inspiral, merger and ringdown of a binary system. The unphysical outer boundary of the grid is located at r_b , which is chosen to be sufficiently far removed that the future Cauchy horizon of the domain of dependence of the initial slice does not reach r_b until the measurement is complete.

The near-zone grids incorporate 5 levels of 2:1 mesh refinement, covering regions centred around each of the black holes. For the highest resolution we have considered here, the finest grid (covering the black hole horizon) has a grid spacing of $dx = 0.02M$. The wave-zone grids have an inner radial resolution which is commensurate with the coarse Cartesian grid resolution, $dr = 0.64M$ in this case. This resolution is maintained essentially constant to the outermost measurement radius ($r = 1000M$), at which point we apply a gradual decrease in resolution (as described in Sec. II A) over a distance of $r = 500M$. From $r = 1500M$ to the outer boundary, we maintain a resolution of $dx = 2.56M$, sufficient to resolve the inspiral frequencies of the dominant $(\ell, m) = (2, 2)$ mode of the gravitational wave signal. The transition between the resolutions is performed over a distance of $500M$ between $r = 1000M$ and $r = 1500M$. The angular coordinates have 31 points (30 cells) in ν and ϕ on each of the 6 patches. The time-step of the wave-zone grids is $dt = 0.144$, and we take wave measurements at each iteration.

We have carried out evolutions at three resolutions in order to estimate the convergence of our numerical methods. The grid described above is labelled $h_{0.64}$. The lower resolutions, labelled $h_{0.80}$ and $h_{0.96}$ have each of the specified grid spacings scaled by $0.80/0.64$ and $0.96/0.64$, respectively.

C. Results

The binary black hole initial data described in Sec. IV A evolves for about 8 orbits ($\simeq 1350M$) before merger. Various (ℓ, m) modes of ψ_4 are plotted in Fig. 5. We find that for the grids we have used, the modes to $(\ell, m) = (4, 4)$ mode are

quite well resolved throughout the evolution. The $(6, 6)$ mode is also measurable, and shows a clear signal, particularly during ringdown. The $(8, 8)$ mode is dominated by noise for most of the inspiral, though during the merger and ringdown phase, a clear signal is present and the amplitude and frequency can be estimated. Tests with an analytical solution confirm that the angular resolutions which we have used are at best marginal for resolving this mode.

In the following sections, we report results regarding the convergence and accuracy of these measurements, as well as determine the parameters of the merger remnant. By analysing the ring-down behaviour of the waves we conclude that the remnant is indeed a Kerr black hole (see Sec. IV C 4, below).

1. Numerical convergence

We can establish the consistency of our discretisation by showing that it does indeed converge to a unique solution in the continuum limit. Ideally, an exact solution can be used to test this. However, since there are no exact solutions which adequately model the physical scenario which we wish to consider (inspiralling black hole binaries), an alternative is to evaluate numerical solutions at several (at least three) different resolutions and establish that the differences decrease as resolution is increased. For an implementation in which all of the discrete operations are carried out with the same order of accuracy, the convergence test should yield a clear exponent corresponding to that order.

The evolution code incorporates a number of discrete operations, which for various practical reasons, are carried out to different orders of accuracy. These are listed in Table I. The primary operation which is carried out over the bulk of the grid is the computation of finite difference derivative operations in order to evaluate the right-hand side of the evolution equations (22a)–(22e). For the tests carried out in this paper, 8th-order stencils are used for this operation, including the upwinded advection terms. It is common to apply a small amount of artificial dissipation in order to smooth high-frequency effects. This is done at one higher order (9th) than the interior finite differencing, with stencils constructed from 10th derivatives. This maintains the correct continuum limit.⁴ Various boundary operations (inter-patch boundary communication, mesh-refinement boundaries) are carried out at lower order. This is done largely for efficiency reasons, as the communication involved in boundary interpolation can be time-consuming if the stencil widths are large. Intuitively, the numerical error associated with these operations may have reduced influence in any case, as they are applied only at 2D interfaces. In practise this does seem to be the case, for instance, as experiments with 4th and 5th order interpolation operators between patches show similar accuracy in the final

⁴ In our experiments, we have noted that dissipation at this high order has a negligible impact on the solution, and could effectively be omitted. However, we continue to include it in the work presented here.




FIG. 5. The dominant spherical harmonic modes of ψ_4 for $\ell = 2, 4, 6, 8$, measured at $r = 200M$ from the coordinate centre. The plots on the right show amplitude and frequency evolution during the late inspiral, merger and ringdown..

solution. Similarly, operations involving different time-levels are at lower order, again for efficiency reasons. The time resolution of our evolutions tends to be high enough that one might expect a small error coefficient of the RK4 integrator. The lowest order operation which we use is the 2nd-order time interpolation at mesh-refinement boundaries. Applying higher order here would require keeping more time levels in memory (currently we store three). Our results are consistent with previous studies using mesh-refinement for black hole evolution which suggest that the influence of the low order time-interpolation boundary conditions is negligible for the time resolutions which we apply (see, for example, [69]).

For test cases involving a single non-spinning black hole, in fact we find 8th-order convergence in the Hamiltonian constraints. This is likely due to the relatively constant values (except for some gauge evolution) maintained by the evolution variables during the evolution, which minimises error due to time-integration or propagation across boundaries.

A more relevant situation is that of a binary black hole inspiral, which we have tested using the parameters described above in Sec. IV A. For this model, we have measured the gravitational waveform, ψ_4 , integrated over spheres at radii from $r = 100M$ to $r = 1000M$, at the three resolutions $h_{0.96}$, $h_{0.80}$ and $h_{0.64}$. Results for the $(\ell, m) = (2, 2)$ mode are shown in Fig. 6. The evolution lasts for about $1350M$

Numerical method	Order
Grid interior finite differencing	8
Inter-patch interpolation	5
Kreiss-Oliger dissipation	9
Time integration (RK4)	4
Mesh-refinement:	
Spatial prolongation	5
Spatial restriction	n/a
Time interpolation	2
Analysis tools:	
Interpolation	4
Finite differencing	8
Surface integration	$N_\theta/2 - 1$

TABLE I. Table of convergence order of various numerical aspects of the evolution code. Spatial restriction is carried out by a direct copy. The surface integration is exact for polynomials up to degree $N_\theta/2 - 1$, where N_θ is the number of grid-points along the θ -direction on the sphere.

before merger, and the plots encompass the inspiral, merger (at $t = 0M$ on this time axis), and ringdown. The figure plots the error in phase $\Delta\phi$ and relative amplitude ΔA for the $(2, 2)$ mode extracted at $r = 100M$ and $r = 1000M$, respec-

tively, between medium $h_{0.80}$ and low $h_{0.96}$ resolutions and high $h_{0.64}$ and medium $h_{0.80}$ resolutions in the wave-zone. The latter error is scaled such that the curves will overlap in the case of a 4th-order convergent solution. At both radii, we find that during the inspiral phase, the rescaled error of the higher resolutions lies below that of the lower resolution, suggesting better than 4th-order convergence (in fact, closer to 8th-order over significant portions of the plot). At later times, around the peak of the waveform, the curves are more closely aligned, indicating 4th-order convergence. The plot suggests that during the very dynamical late stages of the inspiral, the lower order boundary conditions and/or the time integration, play a more important role relative to the early inspiral phase of the evolution, where the convergence order is closer to that of the interior finite differencing. The results are, however, convergent over the entire evolution (including merger and ringdown). As we will see in the next section, the accuracy is excellent for these resolutions so that the rate of convergence is not a particular issue.

We have verified convergence for a number of different modes of the ψ_4 waveform at different radii. For instance, Fig. 7 shows similar results for the $(\ell, m) = (6, 6)$ mode, which is some two orders of magnitude smaller in peak amplitude than the $(\ell, m) = (2, 2)$ mode (see Fig. 5). During the early inspiral, it is difficult to evaluate a convergence order due to high frequency noise which is large relative to the waveform amplitude. However, a measurable signal is clear in the last orbit, merger and ringdown phase, and converges at a clear 3rd order.

2. Accuracy

We estimate the numerical phase and amplitude error by means of a Richardson expansion at a given resolution Δ ,

$$u_{\Delta}(t, x) = u(t, x) + \Delta e_1(t, x) + \Delta^2 e_2(t, x) + \dots, \quad (32)$$

where $u(t, x)$ is the solution of the original differential equation, and the $e_i(t, x)$ are error terms at different orders in Δ . Assuming convergence at a fixed order, n , we can expect some of these error functions to vanish. Using solutions, u , obtained at two resolutions, Δ_1 and Δ_2 , the Richardson expansion implies

$$\begin{aligned} u_{\Delta_1} - u_{\Delta_2} &= e_n(\Delta_1^n - \Delta_2^n) + \mathcal{O}(\Delta^{n+1}) \\ &= e_n \Delta_2^n (C^n - 1) + \mathcal{O}(\Delta^{n+1}) \\ &\sim \epsilon_{\Delta_2} (C^n - 1), \end{aligned} \quad (33)$$

where ϵ_{Δ_2} is the estimated solution error on the higher resolution grid, and where

$$C^n := \left(\frac{\Delta_1}{\Delta_2} \right)^n. \quad (34)$$

We thus obtain an estimate for the solution error that is at least accurate to order $n + 1$,

$$\epsilon_{\Delta_2} \sim \frac{1}{C^n - 1} (u_{\Delta_1} - u_{\Delta_2}), \quad (35)$$

which we use as an estimate of the numerical error in our solutions.

During the inspiral phase (which for this purpose we regard as being the period $t \leq -100M$), we have found roughly 8th-order convergence in the amplitude and phase, as described above. The remaining relative error for the $(\ell, m) = (2, 2)$ mode can be estimated as

$$\max_{t \in [-1350, -100]} \text{err}(A)_{\text{inspiral}} = 0.090\%, \quad (36a)$$

$$\max_{t \in [-1350, -100]} \text{err}(\phi)_{\text{inspiral}} = 0.010\%. \quad (36b)$$

where $\text{err}(A) := \Delta A/A$ and $\text{err}(\phi) := \Delta \phi/\phi$, i.e., the rate of loss of phase with ϕ . During merger and ring-down ($t > -100M$), we observe 4th-order convergence in the amplitude, while maintaining 8th-order convergence in the phase. This results in the estimate

$$\max_{t \in (-100, 150]} \text{err}(A)_{\text{merger}} = 0.153\%, \quad (37a)$$

$$\max_{t \in (-100, 150]} \text{err}(\phi)_{\text{merger}} = 0.003\%. \quad (37b)$$

The time evolution of the numerical error in phase and amplitude is shown in Fig. 8.

We note that these errors are of comparable order to the errors inherent in the extrapolation [77]. Moreover, as is pointed out in [78], the error between extrapolated waveforms and those determined at future null infinity, \mathcal{J}^+ , by characteristic extraction, is an order of magnitude larger than the numerical error determined here. This highlights the importance of reducing systematic errors inherent in finite radius measurements of ψ_4 .

3. Properties of the merger remnant

The merger remnant can be measured with high accuracy, using either the isolated horizon formalism [84, 85], or geometrical measures of the apparent horizon [86, 87]. Some results are reported in Table II, along with estimated numerical errors. The results agree well with previous high-accuracy measurements, such as those obtained by spectral evolution [4, 82], with the spin and irreducible mass agreeing within three decimal and four decimal places, respectively. While this is larger than the reported errors, we note that we have evolved a different initial data set than [4]. As reported in Sec. IV A our evolution has somewhat more eccentricity, and the level of agreement can be used to judge the influence of small amounts of eccentricity on the result.

By comparing the properties of the merger remnant with the integrated radiated energy, E_{rad} , and angular momentum, J_{rad} , determined from the gravitational waveforms, we find the residuals

$$|M_f + E_{\text{rad}} - M_{\text{ADM}}| = 4.1 \times 10^{-5}, \quad (38a)$$

$$|S_f + J_{\text{rad}} - J_{\text{ADM}}| = 1.0 \times 10^{-3}. \quad (38b)$$

Here we have used the extrapolations of the gravitational waveforms to $r \rightarrow \infty$ based on the 6 outermost measurement

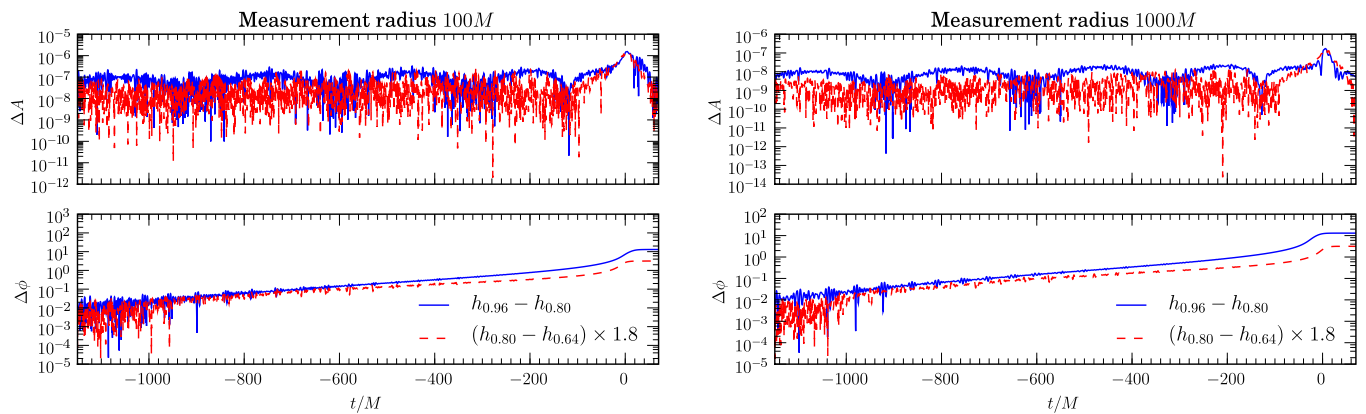


FIG. 6. Convergence in amplitude (top) and phase (bottom) of the $(\ell, m) = (2, 2)$ mode of ψ_4 for detectors at $r = 100M$ and $r = 1000M$. The higher resolution difference, $h_{0.80} - h_{0.64}$, is scaled for 4th-order convergence.




FIG. 7. Convergence in amplitude (top) and phase (bottom) of the $(\ell, m) = (6, 6)$ mode of ψ_4 for detector at $r = 100M$ during the late through merger. The higher resolution difference, $h_{0.80} - h_{0.64}$, is scaled for 3rd-order convergence.




FIG. 8. Absolute numerical error in the amplitude (top) and phase (bottom) accumulated over the course of the evolution for the highest resolution run, determined according to Eq. (35) for the pointwise differences in amplitude and phase between medium and high resolution runs. For the phase we assume the measured 8th-order convergence over the entire evolution, while for the amplitude we use 8th-order before $t \leq -100$, and 4th-order thereafter (see text).

Total ADM mass, M_{ADM}	$0.99051968 \pm 20 \times 10^{-9}$
Total ADM angular momentum, J_{ADM}	$0.99330000 \pm 10 \times 10^{-17}$
Irreducible mass, M_{irr}	$0.884355 \pm 20 \times 10^{-6}$
Spin, S_f/M_f^2	$0.686923 \pm 10 \times 10^{-6}$
Christodoulou mass, M_f	$0.951764 \pm 20 \times 10^{-6}$
Angular momentum, S_f	$0.622252 \pm 10 \times 10^{-6}$
Radiated energy, E_{rad}	$0.038715 \pm 2 \times 10^{-6}$
Radiated angular momentum, J_{rad}	$0.370007 \pm 68 \times 10^{-6}$

TABLE II. Properties of the merger remnant as measured on the apparent horizon (M_{irr} , S_f/M_f^2) and from the gravitational radiation (E_{rad} , J_{rad}). Ranges indicate the estimated numerical error. For the error in J_{ADM} , we have simply quoted machine precision (it is an analytical expression of the input momenta on the conformally flat initial slice).

radii. A more detailed discussion of this procedure is given in [77]. The results can be further improved through taking measurements at \mathcal{J}^+ , as outlined in [78, 81].

4. Quasi-normal modes of the merger remnant

In Fig. 5, we have shown the late-time behaviour of the amplitude and frequency for the dominant spherical harmonic modes of ψ_4 , to $(\ell, m) = (8, 8)$. We note that during ring-down, the frequencies settle to a constant value. If the final black hole is a Kerr black hole, these frequencies are given by the quasi-normal modes of a Kerr black hole with given spin a .

As reported in the previous section, our evolution leads to a merger remnant with $a = 0.686923 \pm 1 \times 10^{-5}$ (see Table II), as measured on the horizon. The real part of the prograde quasi-normal mode (QNM) frequencies for modes up to $(\ell, m) = (7, 7)$, can be found tabulated in [88]. For example, $M\omega_{22} = 0.526891$ for the $(\ell, m) = (2, 2)$ mode, given a final black hole of the measured mass M_f and spin S_f .

At this point it is worth noting that the QNM determined from perturbations of a Kerr black hole are most naturally

expressed in terms of a basis of spin-weighted *spheroidal* harmonics. By contrast, our waveforms have been decomposed relative to a basis of spin-weighted *spherical* harmonics, which are easily calculated via Legendre functions. In order to make an appropriate comparison between these modes with the perturbative results we need to apply a transformation to the wave-modes. We have

$$\hat{\psi}_4^{\ell' m'} = \sum_{\ell, m} \psi_4^{\ell, m} \langle \ell, m | \ell', m' \rangle, \quad (39)$$

where a dash denotes labelling of the spheroidal harmonic modes, and $\langle \ell, m | \ell', m' \rangle$ is the overlap defined by

$$\langle \ell, m | \ell', m' \rangle = \int_{\Omega} d\Omega_{-2} \bar{S}_{\ell' m'}(c_{\ell' m'})_{-2} Y_{\ell m}. \quad (40)$$

The spheroidal harmonics parameter $c_{\ell' m'} = a\omega_{\ell' m'}$ depends on the spin a of the black hole and the corresponding prograde or retrograde QNM frequency $\omega_{\ell' m'}$ of the $(\ell' m')$ spheroidal harmonic mode⁵. If $c = 0$ (as is the case for non-spinning black holes), the spheroidal harmonics reduce to the spherical harmonics. The spin-weighted spheroidal harmonics used here have been implemented following Leaver [89] and are reviewed in [90].

The frequencies measured during the ringdown are plotted in Fig. 9 for the modes $(\ell, m) = (2, 2), (4, 4)$ and $(6, 6)$. We have plotted data for the $r = 1000M$ measurement, as well as the value obtained by extrapolating the waveforms extracted at the outermost 6 measurement spheres to $r \rightarrow \infty$, and find that in fact the extrapolation has little effect on the frequency of the lower order modes at these distances from the source. We note that there is a modulation of the ringdown frequency, particularly apparent in the $(2, 2)$ mode. This is a result of mode mixing, which stems from the use of the spherical harmonic basis for the ψ_4 measurements. By transforming the $r = 1000M$ result to spheroidal harmonics, this modulation visible in the $t < 40M$ signal is largely removed (dashed line).

As the amplitude of the wave declines exponentially to the level of numerical error, the frequencies become difficult to measure accurately. We estimate the ringdown frequency for each mode by performing a least-squares fit of a horizontal line through the measured spheroidal harmonic frequency over the range $t \in [40, 80]M$ (dotted line) with the standard deviation of the fit as a gauge of the error (grey region). These constant lines represent the estimated frequency of the associated QNM modes, and are tabulated as ω^{NR} in Table III. They agree to high precision with the prograde QNM frequencies, $\omega^{\text{lit.}}$, determined Kerr black holes by perturbative methods [88]. We conclude that the merger remnant is compatible with a Kerr black hole within the given error estimates.

V. DISCUSSION

The results of this paper provide a demonstration of the usefulness of adapted coordinates in numerical relativity simula-

(ℓ, m)	$M_f \omega^{\text{lit.}}$	$M_f \omega^{\text{NR}}$	$ M_f \omega^{\text{NR}} - M_f \omega^{\text{lit.}} $
(2, 2)	0.526891	0.5267 ± 0.0011	1.9×10^{-4}
(4, 4)	1.131263	1.1312 ± 0.0028	6.3×10^{-5}
(6, 6)	1.707630	1.7074 ± 0.0062	2.3×10^{-4}

TABLE III. Prograde $N = 0$ QNM frequencies for different modes and spin $a = 0.6869$ as determined by perturbative methods [88], $\omega^{\text{lit.}}$, and as measured during ringdown in the numerical relativity simulation, ω^{NR} .

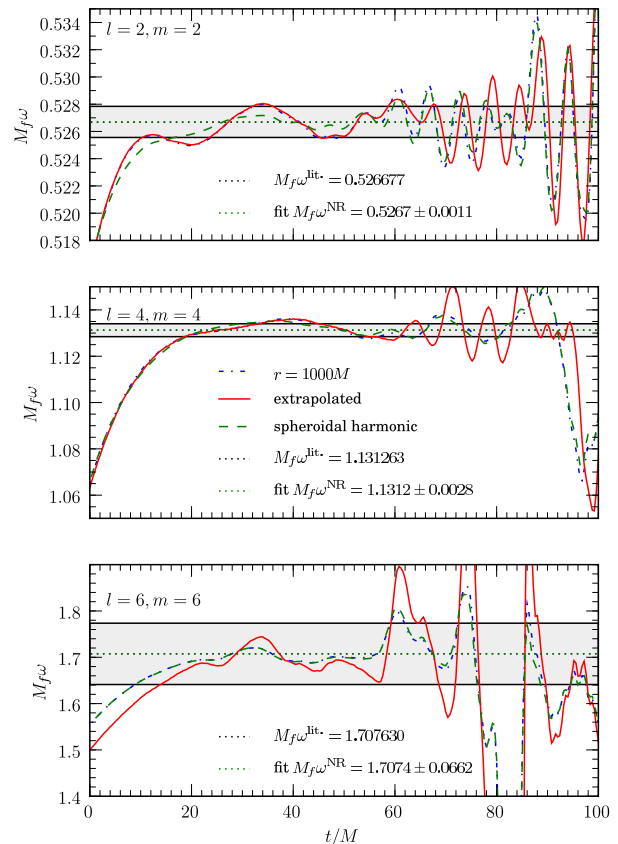


FIG. 9. The ringdown frequencies for the dominant ψ_4 modes to $\ell = 6$ of the merger remnant. From top to bottom, the plots show the frequencies of the $(\ell, m) = (2, 2), (4, 4)$ and $(6, 6)$ modes respectively, over a timescale from the $(2, 2)$ waveform peak to $100M$ later, at which point the waveform amplitude is too small to measure an accurate frequency. The ψ_4 data measured at $r = 1000M$ is plotted, in addition to the value extrapolated to $r \rightarrow \infty$, and the transformation to spheroidal harmonics. The expected quasi-normal mode frequency is plotted as a dotted line, as well as a fit to the spheroidal harmonic data over the range $t \in [40M, 80M]$, with error-bars determined by the standard deviation of the fit.

tions. The precision of the calculations have allowed us to obtain convergent modes to $\ell = 6$, through merger and ringdown, with accurate predictions of the quasi-normal ringdown frequencies of the remnant.

Our implementation of non-singular radially adapted coordinates for the wave zone is based on the use of multiple grid patches with interpolating boundaries, coupled to a BSSNOK

⁵ We restrict attention to the $N = 0$ harmonic only.

evolution code. Thornburg [41] first demonstrated that such a setup could lead to stable evolutions in the case of a spinning black hole in Kerr-Schild coordinates. We have demonstrated that the approach is also effective and robust for dynamical puncture evolutions, and in particular the problem of binary black holes.

The implementation described here has a number of advantages, principle among them being its flexibility. While we have presented results for a particular grid structure adapted to radially propagating waves, there are no principle problems with restructuring the grids to cover any required domain, for instance adapted to excision boundaries or toroidal fields. Since data is stored in the underlying Cartesian basis, and passed by interpolation across boundaries, the coordinates used on each patch are largely independent of the others, and there is no need for numerical grid generating schemes. While we have used the BSSNOK formalism to evolve the Einstein equations, in principle any stable strongly hyperbolic system can be substituted. The BSSNOK system has, however, proven particularly useful for evolving black holes via the puncture approach, which itself has proven to be a very flexible methodology. We have demonstrated results for the most well-studied test case, non-spinning equal-mass black holes, the same techniques can be applied to different mass ratios and spinning black holes, simply by changing the physical input parameters. (The appendices include some examples of spinning black hole evolutions.)

Finally, we emphasise again the accuracies which can be attained by this approach. Our finite difference results show numerical error estimates which are on par with those achieved using spectral spatial discretisation [4]. The adapted radial coordinate allows us to take measurements at radii much larger than have been used before, as well as obtain accurate measurements of higher ℓ modes during merger, which have an amplitude more than two orders of magnitude smaller than the dominant $(\ell, m) = (2, 2)$ mode. One of the aspects which makes this possible is the fact that we are able to extend our grids to a distance such that the measurements are included in the future domain of dependence of the initial data (causally disconnected outer boundaries), and the waves are reasonably well resolved over this entire domain so that internal reflections are minimised. Further, we note that our results are consistent with other puncture-method calculation in that the results are convergent and can be consistently extrapolated to $r \rightarrow \infty$ throughout the entire evolution, including late inspiral and ringdown [77], where other approaches have had difficulties.

The absence of artificial boundaries, as well as dissipative regions in the wave zone, removes an important source of potential error in solving the Einstein equations as an initial-boundary value problem. The remaining errors can be categorised in three forms. First, numerical error due to the discretisation. This can be reduced through the use of higher order methods for the operations performed in various parts of the code, and fortunately is also easy to quantify by performing tests at multiple resolutions. We note that for finite differences, the largest improvement in accuracy occurs in going from 2nd to 4th-order for the interior computations, and be-

yond that there are diminishing returns [91]. While it does not yet seem to be a limiting factor, except possibly during the merger, the RK4 time-stepping will at some level of resolution be a determining factor in the accuracy regardless of the spatial order (and this is also the case for current implementations of spectral methods). The second source of error is a physical error, inherent in the choice of initial data parameters for the binary evolution. At the separations which are practical for numerical relativity (say $d < 20M$), the physical model is expected to have shed all of its eccentricity. We have used post-Newtonian orbital parameters to attempt to place our black holes in low eccentricity trajectories, and this is quite effective. Alternative approaches, involving iteratively correcting the initial data parameters until a tolerable eccentricity has been reached, are able to reduce the eccentricity still further [92]. This technique can in principle also be adapted to the moving puncture approach. The final source of error arises in the measurement of ψ_4 , which is done at a finite radius, and then extrapolated to $r \rightarrow \infty$ by some procedure. We have attempted to minimise this error by placing detectors at large radii, well into the region where the perturbations are linear, and have shown that the extrapolations are consistent with measurements at larger radii, as well as with each other in the $r \rightarrow \infty$ limit [77]. However, there remain ambiguities particularly in gauge-dependent quantities such as the choice of surface on which measurements are taken, and the definition of time and radial distance to be used in the extrapolation. In a companion paper [78], we have demonstrated that these ambiguities can be removed entirely by the procedure of *characteristic extraction*, whereby evolution data on a world-tube is used as an inner boundary condition for a fully relativistic characteristic evolution, extending to null infinity, \mathcal{J}^+ . The results suggest that systematic errors inherent in finite radius measurements of ψ_4 are more than an order of magnitude larger than the numerical errors reported here.

ACKNOWLEDGMENTS

We dedicate this paper to the memory of Thomas Radke, who has made invaluable contributions to the development and optimisation of Cactus, Carpet and the code described here. The authors are pleased to thank: Ian Hinder, Sascha Husa, Badri Krishnan, Philipp Moesta, Christian D. Ott, Luciano Rezzolla, Jennifer Seiler, Jonathan Thornburg, and Burkhard Zink for their helpful input; the developers of Cactus [49, 50] and Carpet [51–53] for providing an open and optimised computational infrastructure on which we have based our code; Nico Budewitz for optimisation work with our local compute cluster, *damiana*; support from the DFG SFB/Transregio 7, the VESF, and by the NSF awards no. 0701566 *XiRel* and no. 0721915 *Alpaca*. Computations were performed at the AEI, at LSU, on LONI (numrel03), on the TeraGrid (TG-MCA02N014), and the Leibniz Rechenzentrum München (h0152).




FIG. 10. Trajectories of the two inspiralling punctures for a spinning configuration $a_1 = -a_2 = 0.8$, with upwinded advection terms (solid lines) and without (dashed lines). In the case where no upwinding has been used, the black holes do not inspiral, due to the accumulation of numerical error.

Appendix A: The influence of upwinded advection stencils

It has long been recognised that for BSSNOK evolutions employing a shift vector, β^a , the overall accuracy can be improved by “upwinding” the finite difference stencils for advective terms of the form $\beta^i \partial_i u$ [61]. The upwind derivatives employ stencils which are off-centred by some number of grid points in the direction of β^a . The drawback of the method is that in order to maintain the same order of accuracy in the derivatives, the stencil must have the same width as a centred stencil, but since it is offset in either a positive or negative direction, it effectively requires an additional number of points to be available to the derivative operator equal to the size of the offset. For parallel codes which physically decompose the grid over processors and communicate ghost-zone boundaries, this means that a larger number of points must be communicated and can impact the overall efficiency. Further, a larger number of points must be translated at inter-patch and refinement level boundaries.

The original observation that upwinding is helpful was made with a code that used 2nd-order spatial finite differences. In that case, the centred stencils are small (three points) and the upwind derivatives correspond to sideways derivatives in the direction of the shift, i.e., no “downwind” information is used. For higher order schemes, the importance of upwinding may be less significant, since the stencils are large relative to the size of the shift vector. In practise, some implementations have empirically determined that upwinding by 1 point at 6th-order is helpful [83]. However, this is not done universally, particularly in conjunction with 8th-order centred differencing [12, 93].

We have found upwinding to be important in reducing numerical error in the black hole motion for every order of accuracy we have tried. The effect is demonstrated in Fig. 10, which plots the motion of the black hole punctures for a data set involving a pair of equal-mass binaries with spins $a_1 = -a_2 = 0.8$ evolved at a relatively low resolution with 8th-order spatial finite differencing. The results of two evolutions are plotted, one using fully centred stencils, and the other upwinding the advection terms with a one-point offset. Whereas the latter evolution displays the expected inspiral behaviour, at this resolution the binary evolved with centred advection actually flies apart. This is purely a result of accumulated numerical error, and at higher resolutions both tracks can be made to inspiral and merge. Our observation, however, is that for a given fixed resolution, the one-point offset advection has a significantly reduced numerical error in the phase as compared to the fully centred derivatives.

Based on some limited experimentation with larger offsets, we have the general impression that the one point offset provides the optimal accuracy for each of the finite difference orders we have tried (4th, 6th, 8th). We do not exclude the possibility that there may be situations in which the fully centred stencils perform as well as upwinded advection, however we have not come across a situation where the latter method performs worse.

As an alternative, we have also tested lower order upwinded derivatives as a potential scheme which would allow us to maintain a smaller stencil width. We generally find that the resultant numerical errors are of the same magnitude or larger than if we had not done the upwind at all.

We note parenthetically the fact that the off-centering is most important in the immediate neighbourhood of the black holes, where the shift has a non-trivial amplitude. It is possible that a scheme where the stencils are off-centred only on grids where the shift is larger than some threshold would also be effective, and not suffer the drawbacks mentioned above over the bulk of the grid. We have not experimented with such a scheme, however.

Appendix B: High order finite differencing

A recent trend in the implementation of finite difference codes for relativity has been the push towards higher order spatial derivatives. It is now common to use 6th or 8th-order stencils. The benefit of higher order stencils is that the convergence rate can be dramatically increased, so that a small increase in resolution leads to a large gain in accuracy. And while not guaranteed, it is often the case that for a given fixed resolution, a higher order derivative will be more accurate, requiring fewer points to accurately represent a wavelength [91].

In moving to high order stencils, there is a trade-off between the possible accuracy improvements, and the extra computational cost. High order stencils generally involve two extra floating point operations per order. Since they require a larger stencil width, they also incur a cost in communication of larger ghost zones, as well as requiring wider overlap zones at grid boundaries. In practice, we find that higher order




FIG. 11. Phase evolution of the $(\ell, m) = (2, 2)$ mode ψ_4 for the aligned-spin model with $a_1 = -a_2 = 0.8 h = 0.64M$. The 6th-order case at $h_{0.64}$ has a trajectory between the low resolution ($h_{0.80}$) and high resolution ($h_{0.64}$) 8th-order evolution.

stencils can also have a more strict Courant limit, requiring a smaller timestep (and thus more computation to reach a given physical time). While it is possible to demonstrate a large gain in accuracy in switching from 2nd to 4th-order operators, there are diminishing returns in the transition to 6th and higher order [91].

We have experimented with 4th, 6th and 8th-order finite differencing for the evolution equations. Generally we find that the 8th-order operators can indeed provide a notable benefit, particularly in the phase accuracy, at low resolution. In Fig. 11, we plot the phase evolution for an equal mass model with spins $a_1 = -a_2 = 0.8$. The evolution covers the last three orbits and ringdown. We find that for this high-spin case, even over this short duration, a significant dephasing takes place. Assuming 8th-order convergence, the 6th-order evolution at the $h_{0.64}$ resolution would be comparable to the 8th-order at approximately $h_{0.77}$ resolution. We can get some idea of the relative amount of work required for each calculation by noting there would be $N = (0.64/0.77)^3$ fewer grid points in the $h_{0.77}$ evolution, but the 8th-order derivatives require $9/7$ times as many floating point computations for a derivative in one coordinate direction, and requires a Courant factor which is 0.9 times that of the 6th-order run. Taken together, this suggests an 8th-order run at $h_{0.77}$ would require a factor 0.68 of the amount of work of the 6th-order case to achieve comparable accuracy. Note that this computation does not take into account potential additional communication overhead associated with the wider 8th-order stencils. But assuming this is not dominant, the conclusion seems to be that for this level of accuracy, the 6th-order evolution is somewhat less efficient than the 8th-order version would be.

For a given situation, it may be that these factors change significantly. Implementation, and even hardware, details can shift the balance of costs between various operations. Further, the test case considered here involves a fairly high spin. Lower spin models (such as that considered in the main body of the paper), are accurate at modest resolutions, and in such cases the 6th-order evolutions may in fact prove to be relatively more efficient if the accuracy is already sufficient for a given purpose. On the other hand, if grid sizes and memory consumption are limiting factors, the 8th-order operators do




FIG. 12. Amplitude and phase evolution of the $(\ell, m) = (2, 2)$ mode of ψ_4 for the equal-mass aligned-spin model, comparing 8th-order spatial finite differencing with a scheme in which 8th-order is used only on the fine meshes surrounding the bodies, and 4th-order on the wave-zone grids.

give a consistent accuracy benefit for a fixed grid size. Our expectation, however, is that implementing yet higher order stencils (for example, 10th-order) may not be justified on the basis of efficiency.

As a final point, we note that the required high-order accuracy appears to be largely a consequence of the field gradients in the near-zone, immediately surrounding the black holes. An alternative scheme, then, could be to apply high-order finite differencing in this region, while using a lower order (and thus more efficient) scheme in the wave zone. Results from such a test are displayed in Fig. 12, where we have used 8th-order only on the finest refinement level, *i.e.*, the mesh surrounding the black holes, but 4th-order on all coarser Cartesian and radial wave-zone grids. This, in turn, allows for a slightly less restrictive Courant limit, so that it becomes possible to run with a slightly larger time-stepping. The phase evolution of ψ_4 is almost identical to that of the fully 8th-order case, but we found that the speed of the run was increased by more than 25% (similar to that of the fully 6th-order evolution). Further optimisations, such as decreasing ghost-zone sizes of the 4th-order grids and consequently the communication overhead, might improve this further. While the errors and convergence order of this scheme have not been tested in detail, we suggest it as a potentially quite effective scheme for the impatient.




FIG. 13. Differences in phase of a spinning configuration with resolution $h = 0.80M$ and conformal variables ϕ and W against a simulation with $h = 0.64M$ and conformal variable W . The dephasing is significant as we are on the coarse limit of resolution for this particular configuration.

Appendix C: Choice of conformal variable

In Sec. III, we have described our implementation of the BSSNOK evolution system, and note that currently three variations are in use, based on the use of different variables to represent the conformal scalar. The original formulation is based on the use of $\phi := \log \gamma/12$. An issue with this variable in the context of puncture evolutions is that it has an $O(\ln r)$ singularity which can lead to large numerical error in finite differences calculated in the neighbourhood of the puncture. More recently, the use of alternative variables $\chi = \gamma^{-1/3}$ [3]

and $W = \gamma^{-1/6}$ [68] have been proposed as a means of improving this situation by replacing ϕ with variables that are regular everywhere on the initial data slice. In terms of the evolution system outlined in Eqs. (22), the χ and W options correspond to the choices $\kappa = 3$ and $\kappa = 6$, respectively.

The influence of this change of variable can be seen in improved phase accuracy of binary evolutions carried out with either χ or W . In Fig. 13, we show results from an evolution of the equal-mass aligned-spin ($a_1 = -a_2 = 0.8$) test case presented in the previous appendices, using ϕ and W as evolution variables. Plotted are the phase errors, $\Delta\phi$, between runs at low resolution, $h_{0.80}$, using both ϕ and W with a higher resolution, $h_{0.64}$, evolution using W . The numerical error associated with the low resolution ϕ evolution is significantly larger than that of the corresponding W evolution.

The reason for this may be related to that of the benefit seen from upwind advective differences in Appendix A. The phase accuracy of the waveforms is crucially dependent on correctly modelling the motion of the bodies, and this requires accurate advective derivatives in the neighbourhood of the punctures. The reduced numerical error associated with the regular χ and W variables is important.

Note that even in the ϕ case, numerical error generated at the puncture seems to be confined to within the horizon. Quantities such as constraints measured outside the horizon, or the horizon properties itself, are not significantly affected. However, it seems that a clear reduction in phase error can be attained through the use of either the χ or W variants of BSSNOK, and we have used the latter for the tests carried out in this paper.

-
- [1] F. Pretorius, Phys. Rev. Lett. **95**, 121101 (2005).
 - [2] J. G. Baker, J. Centrella, D.-I. Choi, M. Koppitz, and J. van Meter, Phys. Rev. Lett. **96**, 111102 (2006).
 - [3] M. Campanelli, C. O. Lousto, P. Marronetti, and Y. Zlochower, Phys. Rev. Lett. **96**, 111101 (2006).
 - [4] M. A. Scheel et al., Phys. Rev. **D79**, 024003 (2009).
 - [5] J. A. Gonzalez, U. Sperhake, B. Bruegmann, M. Hannam, and S. Husa, Phys. Rev. Lett. **98**, 091101 (2007).
 - [6] J. A. Gonzalez, M. D. Hannam, U. Sperhake, B. Bruegmann, and S. Husa, Phys. Rev. Lett. **98**, 231101 (2007).
 - [7] M. Campanelli, C. O. Lousto, Y. Zlochower, and D. Merritt, Astrophys. J. **659**, L5 (2007).
 - [8] M. Campanelli, C. O. Lousto, Y. Zlochower, and D. Merritt, Phys. Rev. Lett. **98**, 231102 (2007).
 - [9] F. Herrmann, I. Hinder, D. Shoemaker, P. Laguna, and R. A. Matzner, Astrophys. J. **661**, 430 (2007).
 - [10] M. Koppitz et al., Phys. Rev. Lett. **99**, 041102 (2007).
 - [11] D. Pollney, C. Reisswig, L. Rezzolla, B. Szilágyi, M. Ansorg, B. Deris, P. Diener, E. N. Dorband, M. Koppitz, A. Nagar, et al., Phys. Rev. **D76**, 124002 (2007).
 - [12] C. O. Lousto and Y. Zlochower, arXiv:0805.0159 (2008).
 - [13] L. Rezzolla et al., Astrophys. J. **679**, 1422 (2008).
 - [14] L. Rezzolla, E. Barausse, E. N. Dorband, D. Pollney, C. Reisswig, J. Seiler, and S. Husa, Astrophys. J. **674**, L29 (2008).
 - [15] L. Rezzolla et al., Phys. Rev. **D78**, 044002 (2008).
 - [16] W. Tichy and P. Marronetti, Physical Review D **78**, 081501 (2008).
 - [17] C. O. Lousto and Y. Zlochower, Phys. Rev. **D76**, 041502 (2007).
 - [18] E. Barausse and L. Rezzolla, Astrophys. J. Lett. **704**, L40 (2009).
 - [19] C. Reisswig, S. Husa, L. Rezzolla, E. N. Dorband, D. Pollney, and J. Seiler, Phys. Rev. D **80**, 124026 (2009).
 - [20] P. Ajith et al., Class. Quant. Grav. **24**, S689 (2007).
 - [21] P. Ajith et al., Phys. Rev. D **77**, 104017 (2008).
 - [22] P. Ajith, Class. Quant. Grav. **25**, 114033 (2008).
 - [23] P. Ajith et al., arXiv:0909.2867 (2009).
 - [24] J. G. Baker, J. R. van Meter, S. T. McWilliams, J. Centrella, and B. J. Kelly (2006).
 - [25] A. Buonanno, G. B. Cook, and F. Pretorius (2006).
 - [26] M. Hannam, S. Husa, U. Sperhake, B. Bruegmann, and J. A. Gonzalez, Phys. Rev. **D77**, 044020 (2008).
 - [27] M. Hannam, S. Husa, B. Bruegmann, and A. Gopakumar, Phys. Rev. **D78**, 104007 (2008).
 - [28] T. Damour, A. Nagar, E. N. Dorband, D. Pollney, and L. Rezzolla, Phys. Rev. **D77**, 084017 (2008).
 - [29] A. Buonanno et al., Phys. Rev. **D76**, 104049 (2007).
 - [30] T. Damour, A. Nagar, M. Hannam, S. Husa, and B. Bruegmann, Phys. Rev. **D78**, 044039 (2008).
 - [31] A. Buonanno et al., Phys. Rev. **D79**, 124028 (2009).
 - [32] M. Boyle et al., Phys. Rev. **D76**, 124038 (2007).
 - [33] N. T. Bishop, R. Gómez, L. Lehner, M. Maharaj, and J. Winicour, Phys. Rev. D **56**, 6298 (1997).
 - [34] M. A. Scheel et al., Phys. Rev. **D74**, 104006 (2006).

- [35] E.ourgoulhon, P. Grandclement, K. Taniguchi, J.-A. Marck, and S. Bonazzola, *Phys. Rev.* **D63**, 064029 (2001).
- [36] E.ourgoulhon, P. Grandclement, and S. Bonazzola, *Phys. Rev. D* **65**, 044020 (2002).
- [37] P. Grandclement, E.ourgoulhon, and S. Bonazzola, *Phys. Rev.* **D65**, 044021 (2002).
- [38] E. Schnetter, P. Diener, E. N. Dorband, and M. Tiglio, *Class. Quant. Grav.* **23**, S553 (2006).
- [39] E. N. Dorband, E. Berti, P. Diener, E. Schnetter, and M. Tiglio, *Phys. Rev.* **D74**, 084028 (2006).
- [40] E. Pazos, M. Tiglio, M. D. Duez, L. E. Kidder, and S. A. Teukolsky, *Phys. Rev.* **D80**, 024027 (2009).
- [41] J. Thornburg, *Class. Quant. Grav.* **21**, 3665 (2004).
- [42] C. Reisswig, N. T. Bishop, C. W. Lai, J. Thornburg, and B. Szilágyi, *Class. Quantum Grav.* **24**, S327 (2007).
- [43] R. Gomez, W. Barreto, and S. Frittelli, *Phys. Rev.* **D76**, 124029 (2007).
- [44] H. P. Pfeiffer, L. E. Kidder, M. A. Scheel, and S. A. Teukolsky, *Comput. Phys. Commun.* **152**, 253 (2003).
- [45] M. Carpenter, D. Gottlieb, and S. Abarbanel, *J. Comput. Phys.* **111**, 220 (1994).
- [46] P. Diener, E. N. Dorband, E. Schnetter, and M. Tiglio, *J. Sci. Comput.* **32**, 109 (2007).
- [47] M. Hannam, S. Husa, D. Pollney, B. Brugmann, and N. O'Murchadha, *Phys. Rev. Lett.* **99**, 241102 (2007).
- [48] J. D. Brown (2009).
- [49] T. Goodale, G. Allen, G. Lanfermann, J. Massó, T. Radke, E. Seidel, and J. Shalf, in *Vector and Parallel Processing – VECPAR'2002, 5th International Conference, Lecture Notes in Computer Science* (Springer, Berlin, 2003).
- [50] Cactus Computational Toolkit home page, URL <http://www.cactuscode.org/>.
- [51] E. Schnetter, S. H. Hawley, and I. Hawke, *Class. Quantum Grav.* **21**, 1465 (2004).
- [52] E. Schnetter, P. Diener, N. Dorband, and M. Tiglio, *Class. Quantum Grav.* **23**, S553 (2006).
- [53] Mesh Refinement with Carpet, URL <http://www.carpetcode.org/>.
- [54] B. Fornberg, *Mathematics of Computation* **51**, 699 (1988).
- [55] J. R. Driscoll and D. M. Healy, Jr., *Adv. Appl. Math.* **15**, 202 (1994), ISSN 0196-8858.
- [56] H. Bateman, *Higher transcendental functions* (1955).
- [57] T. Nakamura, K. Oohara, and Y. Kojima, *Prog. Theor. Phys. Suppl.* **90**, 1 (1987).
- [58] M. Shibata and T. Nakamura, *Phys. Rev. D* **52**, 5428 (1995).
- [59] T. W. Baumgarte and S. L. Shapiro, *Phys. Rev. D* **59**, 024007 (1998).
- [60] M. Alcubierre, B. Brügmann, T. Dramlitsch, J. A. Font, P. Papadopoulos, E. Seidel, N. Stergioulas, and R. Takahashi, *Phys. Rev. D* **62**, 044034 (2000).
- [61] M. Alcubierre, B. Brügmann, P. Diener, M. Koppitz, D. Pollney, E. Seidel, and R. Takahashi, *Phys. Rev. D* **67**, 084023 (2003).
- [62] J. R. van Meter, J. G. Baker, M. Koppitz, and D.-I. Choi, *Phys. Rev. D* **73**, 124011 (2006).
- [63] D. Brown, O. Sarbach, E. Schnetter, M. Tiglio, P. Diener, I. Hawke, and D. Pollney, *Phys. Rev. D* **76**, 081503(R) (2007).
- [64] D. Brown, P. Diener, O. Sarbach, E. Schnetter, and M. Tiglio, *Phys. Rev. D* **79**, 044023 (2009).
- [65] M. Alcubierre, G. Allen, B. Brügmann, E. Seidel, and W.-M. Suen, *Phys. Rev. D* **62**, 124011 (2000).
- [66] T. W. Baumgarte and S. L. Shapiro, *Physics Reports* **376**, 41 (2003).
- [67] M. Alcubierre, *Introduction to 3 + 1 Numerical Relativity* (Oxford University Press, Oxford, UK, 2008).
- [68] P. Marronetti, W. Tichy, B. Bruegmann, J. Gonzalez, and U. Sperhake, *Phys. Rev.* **D77**, 064010 (2008).
- [69] B. Bruegmann et al., *Phys. Rev.* **D77**, 024027 (2008).
- [70] C. Bona, J. Massó, E. Seidel, and J. Stela, *Phys. Rev. Lett.* **75**, 600 (1995).
- [71] E. T. Newman and R. Penrose, *J. Math. Phys.* **3**, 566 (1962), erratum in *J. Math. Phys.* **4**, 998 (1963).
- [72] R. Sachs, *Proc. Roy. Soc. London* **A264**, 309 (1961).
- [73] R. Penrose, *Phys. Rev. Lett.* **10**, 66 (1963).
- [74] J. Baker, M. Campanelli, and C. O. Lousto, *Phys. Rev. D* **65**, 044001 (2002).
- [75] L. Gunnarsen, H. Shinkai, and K. Maeda, *Class. Quantum Grav.* **12**, 133 (1995).
- [76] M. Boyle and A. H. Mroue (2009).
- [77] D. Pollney, C. Reisswig, N. Dorband, E. Schnetter, and P. Diener, *Phys. Rev.* **D80**, 121502 (2009).
- [78] C. Reisswig, N. T. Bishop, D. Pollney, and B. Szilágyi, *Phys. Rev. Lett.* **103**, 221101 (2009).
- [79] N. Bishop, R. Isaacson, R. Gómez, L. Lehner, B. Szilágyi, and J. Winicour, in *Black Holes, Gravitational Radiation and the Universe*, edited by B. Iyer and B. Bhawal (Kluwer, Dordrecht, The Netherlands, 1999), p. 393.
- [80] M. C. Babiuc, N. T. Bishop, B. Szilágyi, and J. Winicour, *Phys. Rev.* **D79**, 084011 (2009).
- [81] C. Reisswig, N. T. Bishop, D. Pollney, and B. Szilágyi, *Class. Quant. Grav.* **27**, 075014 (2010).
- [82] M. Hannam et al., *Phys. Rev.* **D79**, 084025 (2009).
- [83] S. Husa, M. Hannam, J. A. Gonzalez, U. Sperhake, and B. Bruegmann, *Phys. Rev.* **D77**, 044037 (2008).
- [84] O. Dreyer, B. Krishnan, D. Shoemaker, and E. Schnetter, *Phys. Rev. D* **67**, 024018 (2003).
- [85] A. Ashtekar and B. Krishnan, *Living Rev. Relativ.* **7**, 10 (2004).
- [86] S. Brandt and E. Seidel, *Phys. Rev. D* **52**, 870 (1995).
- [87] M. Alcubierre, B. Brügmann, P. Diener, F. S. Guzmán, I. Hawke, S. Hawley, F. Herrmann, M. Koppitz, D. Pollney, E. Seidel, et al., *Phys. Rev. D* **72**, 044004 (2005).
- [88] E. Berti, V. Cardoso, and A. O. Starinets, *Class. Quant. Grav.* **26**, 163001 (2009).
- [89] E. Leaver, *Proc. R. Soc. London, Ser. A* **402**, 285 (1985).
- [90] E. Berti, V. Cardoso, and M. Casals, *Phys. Rev. D* **73**, 024013 (2006).
- [91] B. Gustafsson, H.-O. Kreiss, and J. Olinger, *Time dependent problems and difference methods* (Wiley, New York, 1995).
- [92] H. P. Pfeiffer et al., *Class. Quant. Grav.* **24**, S59 (2007).
- [93] M. Campanelli, C. O. Lousto, H. Nakano, and Y. Zlochower, *Phys. Rev.* **D79**, 084010 (2009).