

Simulations of binary black hole mergers using spectral methods

Béla Szilágyi, Lee Lindblom, and Mark A. Scheel

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

(Received 22 September 2009; published 8 December 2009)

Several improvements in numerical methods and gauge choice are presented that make it possible now to perform simulations of the merger and ringdown phases of “generic” binary black hole evolutions using the pseudospectral evolution code SpEC. These improvements include the use of a new damped-wave gauge condition, a new grid structure with appropriate filtering that improves stability, and better adaptivity in conforming the grid structures to the shapes and sizes of the black holes. Simulations illustrating the success of these new methods are presented for a variety of binary black hole systems. These include fairly generic systems with unequal masses (up to 2:1 mass ratios), and spins (with magnitudes up to $0.4M^2$) pointing in various directions.

DOI: [10.1103/PhysRevD.80.124010](https://doi.org/10.1103/PhysRevD.80.124010)

PACS numbers: 04.25.D-, 02.70.-c, 02.70.Hm, 04.25.dg

I. INTRODUCTION

Black hole science took a great stride forward in 2005 when Pretorius [1] performed the first successful full non-linear dynamical numerical simulation of the inspiral, merger, and ringdown of an orbiting black hole binary system; this initial success then stimulated other groups to match this achievement within months [2,3]. These developments lead quickly to advances in our understanding of black hole physics: investigations of the orbital mechanics of spinning binaries [4–9], studies of the recoil from the merger of unequal-mass binary systems [10–14], the remarkable discovery of unexpectedly large recoil velocities from the merger of certain spinning binary systems [7,15–27], investigations into the mapping between the binary black hole initial conditions (individual masses and spins) and the final state of the merged black hole [28–35], and improvements in our understanding of the validity of approximate binary black hole orbital calculations using post-Newtonian methods [36–43].

These first results on binary black hole systems were obtained by several different groups using different codes based on two different formulations of the Einstein equations (Baumgarte-Shapiro-Shibata-Nakamura and generalized harmonic), using two different methods for treating the black hole interiors (moving puncture and excision), and using rather different gauge conditions to fix the space-time coordinates. All of these early results, however, were obtained with codes based on finite difference numerical methods and adaptive mesh refinement for computational efficiency. The Caltech/Cornell Collaboration decided a number of years ago to follow a different path by developing an Einstein evolution code, called SpEC, based on spectral methods. The advantages of spectral methods are their superior accuracy and computational efficiency, and their extremely low numerical dissipation that is ideal for solving wave propagation problems with high precision. The disadvantages are the relative complexity of spectral codes, the lack of any appropriate preexisting spectral code

infrastructure (analogous to CACTUS [44] for example), and the extreme sensitivity of spectral algorithms to developing instabilities when any aspect of the solution method is mathematically ill posed. Consequently it has taken our group somewhat longer to bring SpEC up to the level of the state-of-the-art codes in the field.

We, along with our Caltech/Cornell collaborators, have developed a large number of numerical and analytical tools over the years that make it possible for SpEC to evolve binary black hole systems with greater precision than any other code (at the present time) [45]. These technical developments include the derivation and implementation of constraint preserving and physical (no incoming gravitational wave) boundary conditions [46,47], dual frame evolution methods and feedback and control systems to lock the computational grids onto the location of the black holes [48], and special angular filtering methods needed to cure an instability that occurs when tensor fields are evolved [49]. Using these methods we and our collaborators have performed a number of high precision evolutions of the inspiral portions of binary black hole systems, and have used the gravitational waveforms from these evolutions to calibrate the accuracy of the approximate post-Newtonian waveforms that are widely used in gravitational wave data analysis [50–53].

During the past year our group has begun to have some success in performing the more dynamical and difficult (for spectral codes) merger and ringdown portions of binary black hole evolutions [51,53–57]. The techniques we used to achieve these first merger calculations turned out to be rather nonrobust however, requiring a great deal of hands-on adjustment and fine-tuning for each new case we attempted. In this paper we report on several new numerical and analytical developments that allow us now to perform stable and accurate binary merger and ringdown simulations of a fairly wide range of binary black hole systems. These new methods appear to be quite robust, and no fine-tuning is required: the same basic method works on each of the cases we have attempted.

The three new technical breakthroughs that allow us now to perform successful binary black hole merger and ringdown simulations are described in Sec. II. These major advances include the development of a new gauge condition that works extremely well with the generalized harmonic formulation of the Einstein system used in our code. This new gauge chooses spatial coordinates that are solutions of a damped-wave equation, and chooses the time coordinate in a way that limits the growth of \sqrt{g}/N , here g is the determinant of the spatial metric and N is the lapse function. This new gauge condition is described in some detail in Sec. II A. Another important development, reported in Sec. II B, is the construction of a new grid structure on which our binary black hole evolutions are performed. This new nonoverlapping grid structure, and a certain type of spectral filtering used with the new grid structure, removes a class of numerical instabilities that were a limiting factor in our ability to perform the highly dynamical portions of merger calculations, and allows us to increase resolution near the black holes in a more targeted and computationally efficient way. The third technical development is a new more efficient and robust method to conform the structure of the grid to the shape and size of the black holes in an automatic dynamical way. These developments are described in Sec. II C and the Appendix.

Using these new technical tools we have now performed successful merger and the succeeding ringdown simulations of a fairly wide range of binary black hole systems. We describe in Sec. III merger simulations for six reasonably diverse cases. These include two systems in which the black holes are nonspinning: one has equal-mass black holes, the other has holes with a 2:1 mass ratio. We also describe successful merger and ringdown calculations with two different equal-mass binary systems in which the holes have identical intrinsic spins of magnitude $0.4M^2$: aligned to the orbital angular momentum in one system, and anti-aligned in the other. We have also performed successful merger and ringdown simulations on two fairly generic binary systems. These two systems have black holes with the same mass ratio $M_2/M_1 = 2$, and the same “randomly oriented” initial spins of magnitudes $0.2M_1^2$ and $0.4M_2^2$, respectively. But the initial separations of the holes are different in the two cases: one starts about 8.5 orbits before merger, and another (used to perform careful convergence studies) starts about 1.5 orbits before merger. The same methods were used to perform all of these merger and ringdown simulations, and no particular fine-tuning was required.

II. TECHNICAL DEVELOPMENTS

This section describes the three new technical breakthroughs that allow us now to perform successful binary black hole merger and ringdown simulations. These major advances include the development of a new gauge condi-

tion, described in Sec. II A, in which the spatial coordinates satisfy a damped-wave equation and the time coordinate is chosen in a way that controls the growth of the spatial volume element. Another important development, described in Sec. II B, is a new nonoverlapping grid structure for our binary black hole simulations, and a type of spectral filtering for these new grid structures that improves their accuracy and stability. The third new technical development, described in Sec. II C, is a more efficient and robust method of conforming the structure of the grid to the shape and size of the black holes in a dynamical and automatic way.

A. Damped-wave gauge

Harmonic gauge is defined by the condition that each coordinate x^a satisfies the covariant scalar wave equation:

$$\nabla^c \nabla_c x^a = H^a = 0. \quad (1)$$

Harmonic coordinates have proven to be extremely useful for analytical studies of the Einstein equations, but have found only limited success in numerical problems like simulations of complicated highly dynamical black hole mergers. One reason for some of these difficulties is the wealth of “interesting” dynamical solutions to the harmonic gauge condition itself, Eq. (1). Since all “physical” dynamical fields are expressed in terms of the coordinates, an ideal gauge condition would limit coordinates to those that are simple, straightforward, dependable, and nonsingular; having interesting dynamics of their own is *not* a desirable feature for coordinates. The dynamical range available to harmonic coordinates can be reduced by adding a damping term to the equation: [58,59]

$$\nabla^c \nabla_c x^a = \mu_S t^c \partial_c x^a = \mu_S t^a, \quad (2)$$

where t^a is the future directed unit normal to the constant- t hypersurfaces. Adding such a damping term to the equations for the spatial coordinates x^i tends to remove extraneous gauge dynamics and drives the coordinates toward solutions of the covariant spatial Laplace equation on the time scale $1/\mu$. Choosing $1/\mu$ to be comparable to (or smaller than) the characteristic time scale of a particular problem should remove any extraneous coordinate dynamics on time scales shorter than the physical time scale. The addition of such a damping term in the time-coordinate equation is not appropriate however. Such a damped-wave time coordinate is driven toward a constant value, and therefore toward a state in which it fails to be a useful time coordinate at all. It makes sense then to use the damped-wave gauge condition only for the spatial coordinates:

$$\nabla^c \nabla_c x^i = H^i = \mu_S t^i = -\mu_S N^i / N, \quad (3)$$

where N^i is the shift, and N is the lapse. The appropriate contravariant version of this damped-wave gauge condition is therefore

$$H_a = -\mu_S g_{ai} N^i / N, \quad (4)$$

where g_{ab} is the spatial metric of the constant- t hypersurfaces. (Note that we use Latin letters from the beginning of the alphabet, a, b, c, \dots , to denote four-dimensional spacetime indices, and letters from the middle of the alphabet, i, j, k, \dots , for three-dimensional spatial indices.)

While the damped-wave gauge is a poor choice for the time coordinate, the idea of imposing a gauge that adds dissipation to the gauge dynamics of the time coordinate is attractive. To find the appropriate expression for $t^a H_a$, the component of H_a not fixed by Eq. (4), we note that the gauge constraint $H_a + \Gamma_a = 0$ implies that $t^a H_a$ is given by

$$t^a H_a = t^a \partial_a \log\left(\frac{\sqrt{g}}{N}\right) - N^{-1} \partial_k N^k, \quad (5)$$

where $g = \det g_{ij}$ is the spatial volume element. In our experience, a frequent symptom of the failure of simpler gauge conditions in binary black hole simulations is an explosive growth of g in the spacetime region near the black hole horizons. This suggests that a good use of the remaining gauge freedom would be to attempt to control the growth of the spatial volume element, g . Choosing the gauge condition,

$$t^a H_a = -\mu_L \log\left(\frac{\sqrt{g}}{N}\right), \quad (6)$$

together with Eq. (5) implies the following evolution equation for \sqrt{g}/N :

$$t^a \partial_a \log\left(\frac{\sqrt{g}}{N}\right) + \mu_L \log\left(\frac{\sqrt{g}}{N}\right) = N^{-1} \partial_k N^k, \quad (7)$$

whose solutions tend to suppress growth in \sqrt{g}/N . The discussion of this gauge condition in Ref. [58] shows that it also implies that the lapse N satisfies a damped-wave equation, with the damping factor μ_L . So in this sense, the gauge condition on the time coordinate, Eq. (6), is the natural extension of the spatial-coordinate damped-wave gauge condition, Eq. (4).

Combining this new lapse condition, Eq. (6), with the damped-wave spatial-coordinate condition, Eq. (4), gives the gauge-source function for our full damped-wave gauge condition:

$$H_a = \mu_L \log\left(\frac{\sqrt{g}}{N}\right) t_a - \mu_S N^{-1} g_{ai} N^i. \quad (8)$$

The damping factors $\mu_L \geq 0$ and $\mu_S \geq 0$ can be chosen quite arbitrarily as functions of spacetime coordinates x^a , or even as functions of the spacetime metric ψ_{ab} . The gauge-source function H_a depends only on coordinates and the spacetime metric ψ_{ab} in this case, so these gauge conditions can be implemented directly in the generalized harmonic Einstein system without the need for a gauge driver. Previous studies of this condition (developed ini-

tially as a test of the first-order gauge-driver system [58]) showed it to be quite useful for evolving single black hole spacetimes. In those tests, which included several different evolutions of maximal-slice Schwarzschild initial data with large nonspherical gauge perturbations, the black hole always evolved quickly toward a nonsingular time-independent equilibrium state.

The solutions to the lapse gauge condition, Eq. (7), can be thought of as equating $\log(\sqrt{g}/N)$ to a certain weighted time average of $\partial_k N^k/N$. The time scale associated with this time averaging is set by μ_L , which determines, for example, the rate at which \sqrt{g}/N is driven toward an asymptotic equilibrium state. When μ_L is constant, it is easy to show that \sqrt{g}/N is driven exponentially toward this asymptotic state. In highly dynamical spacetimes, however, we find that \sqrt{g}/N must be driven even faster than exponential in order to prevent the formation of singularities in g . This can be accomplished by making μ_L larger whenever g becomes large. In practice we find that the choice $\mu_L = \mu_0 [\log(\sqrt{g}/N)]^2$, (where μ_0 is a constant or perhaps a function of time) is very effective at suppressing the growth of these singularities. We find that choosing the same damping factor μ_S ,

$$\mu_S = \mu_L = \mu_0 \left[\log\left(\frac{\sqrt{g}}{N}\right) \right]^2, \quad (9)$$

for the spatial part of the gauge condition is also quite effective. The binary black hole merger and ringdown simulations described in Sec. III use these μ_L and μ_S , with μ_0 taken to be an order-unity function of time (to accommodate starting up evolutions from initial data satisfying a different gauge condition).

B. Grid structure and filtering

The pseudospectral numerical methods used by our code, the Spectral Einstein Code (SpEC) [46,48,60], represent dynamical fields on a spatial grid structure that is specially constructed for each problem. For binary black hole simulations, our group has been using grid structures constructed from layers of spherical-shell subdomains centered on each black hole, surrounded and connected by cylindrical, cylindrical-shell, and/or rectangular-block subdomains, all surrounded by spherical-shell subdomains that extend to the outer boundary of the computational domain (located far from the holes). The intermediate-zone cylindrical-shell and/or rectangular-block subdomains must overlap both the inner and outer spherical-shell subdomains in these grid structures to cover the computational domain completely. These grid structures are quite efficient, and have allowed our group to perform long stable inspiral calculations for a variety of simple binary black hole systems [48,50,53,61], and also simulations of the merger and ringdown phases of a few of these simple cases [54,57]. The overlap regions in these grid structures are well behaved in these successful cases. But in more

generic inspiral and merger simulations, these overlap regions become significant sources of numerical error and instability.

The overlap-related instability discussed above is most likely caused by the method we use to exchange information between adjacent subdomains. We set the incoming characteristic fields, whose values are determined by interpolation from the adjacent subdomain, using a penalty method [62–65]. This penalty method for imposing boundary conditions was derived explicitly for the case where adjacent subdomain boundaries touch but do not overlap. One possibility for resolving this instability would be to rederive the appropriate penalty boundary terms for the case of overlapping subdomains. We have chosen the simpler option of constructing new grid structures without overlapping subdomains. We do this by changing the intermediate-zone grid structure from the cylindrical-shells and/or rectangular blocks used previously into a series of deformed rectangular blocks whose boundaries conform to the spherical shells used near the inner and outer boundaries. The new type of grid element needed for this construction is illustrated in Fig. 1, which shows how the 3D volume between a sphere and a concentric cube can be mapped into six deformed cubes. This construction is based on the “cubed-sphere” coordinate representations of the sphere (i.e., the mapping of the six faces of the cube onto the sphere as illustrated in Fig. 1), that is widely used in atmospheric and geophysical modeling [66–69]. Three-dimensional grid structures based on cubed spheres have also been used in other types of simulations, including a few black hole simulations [70,71].




FIG. 1 (color online). Illustrates a basic element of the grid structure used to fill the 3D volume between a sphere and a concentric cube.




FIG. 2 (color online). Grid structure around each black hole. A series of spherical shells (two shells shown) is surrounded by a series of cubed-sphere layers (three shown) to fill up a cube. This figure illustrates one quarter of the cube structure surrounding one of the black holes.

Our new grid structure is illustrated in Figs. 2 and 3. It consists of a series of spherical-shell subdomains and several layers of cubed-sphere subdomains surrounding each black hole, as illustrated in Fig. 2. The two cubic blocks containing the black holes are surrounded by a series of rectangular-block subdomains to fill out the volume of a large cube. The center of this large cube is placed at the center-of-mass of the binary system, as illustrated in Fig. 3 for a system having a 2:1 mass ratio. This large cube is connected to the outer spherical-shell subdomains by several additional layers of cubed-sphere subdomains, also illustrated in Fig. 3.

Previous attempts to use this type of cubed-sphere grid structure in SpEC have always proven unsuccessful, due to




FIG. 3 (color online). Illustrates the intermediate grid structure surrounding the two black holes, consisting of a set of rectangular blocks placed around the cubic block containing each hole. This collection of blocks fills up a cube which is surrounded by a series of cubed-sphere layers (two shown) all surrounded by a series of spherical shells (one shell shown) that are centered on the center of mass of the binary system.

numerical instabilities at the interdomain boundaries. These previous attempts used no spectral filtering, or a two-thirds antialiasing filter, in the cubed-sphere subdomains. It has been shown, however, that an appropriate filter is required for stability (and improved accuracy) of Chebyshev polynomial spectral expansions, such as those used in these cubed-sphere subdomains. Such a filter must satisfy two important conditions. The first criterion is that when represented in physical space, this filter must act like a dissipation term in the evolution equations that vanishes on the subdomain boundaries (to avoid conflicting with the boundary conditions). The second criterion is that it must set the highest spectral coefficient to zero. Adding this type of spectral filtering turns out to be one of the key elements in improving our numerical method enough to make possible the binary black hole simulations described in Sec. III. The needed spectral filter is applied to each field $u(x, t)$ that is expanded as a sum of Chebyshev polynomials:

$$u(x, t) = \sum_k u_k(t) T_k(x). \quad (10)$$

These are used for the radial-coordinate expansions in the spherical-shell subdomains, and for all three spectral-coordinate expansions in the rectangular-block and the cubed-sphere subdomains. Each spectral coefficient u_k in these expansions is filtered according to the expression,

$$\mathcal{F}(u_k) = u_k e^{-\alpha(k/k_{\max})^{2p}}, \quad (11)$$

after each time step. This filter satisfies the first necessary filter criterion by adding a $2p^{\text{th}}$ -order dissipation term to the evolution equations which vanishes on each subdomain boundary [65]. For the binary black hole simulations presented here, we use the filter parameters $\alpha = 36$ and $p = 32$. The choice $\alpha = 36$ guarantees that this filter satisfies the second necessary filter criterion by ensuring that the highest spectral coefficient is set to (double precision) zero, while the choice $p = 32$ insures that this is a very mild filter that only influences the largest few spectral coefficients. This new filter is applied wherever Chebyshev expansions, Eq. (10), are used. We use the same filter we have used in previous binary black hole simulations in the angular directions in the spherical-shell subdomains: we set to zero the top four ℓ -coefficients in the tensor spherical-harmonic representation of each dynamical field [48,49].

C. Adaptive conforming grid

One of the important developments that allowed our group (several years ago) to begin performing successful binary black hole inspiral simulations was the introduction of the dual-coordinate-frame evolution method [48]. This method uses two distinct coordinate frames: One is a non-rotating and asymptotically Cartesian coordinate system used to construct the tensor basis for the components of the

various dynamical fields evolved by our code. The second is a coordinate system chosen to follow (approximately) the motions of the black holes. We fix the computational grid to this second coordinate frame, and solve the evolution equations for the “inertial-frame” tensor field components as functions of these “grid-frame” coordinates. The map \mathcal{M} connecting these grid-frame coordinates \bar{x}^i to the inertial-frame x^i , can be written as the composition of more elementary maps:

$$\mathcal{M} = \mathcal{M}_K \circ \mathcal{M}_S. \quad (12)$$

\mathcal{M}_K represents a “kinematical” map that keeps the centers of the black holes located (approximately) at the centers of the excised holes in our grid structures (cf. Sec. II B), and \mathcal{M}_S is a “shape-control” map that makes the grid conform (approximately) to the shapes of the black holes. The third major technical development (which makes it possible now for us to perform robust merger simulations) consists of improvements in the choice of the shape-control map, \mathcal{M}_S , and improvements in the way this map is adapted to the dynamically changing shapes of the black holes.

The kinematical map \mathcal{M}_K can itself be decomposed into more elementary maps:

$$\mathcal{M}_K = \mathcal{M}_T \circ \mathcal{M}_E \circ \mathcal{M}_R. \quad (13)$$

The maps \mathcal{M}_T , \mathcal{M}_E , and \mathcal{M}_R each move the centers of the black holes, but do not significantly distort their shapes. The map \mathcal{M}_T translates the grid to account for the motion of the center of the system due to linear momentum being exchanged with the near field [55] and being emitted in gravitational radiation. The map \mathcal{M}_E does a conformal rescaling that keeps the coordinate distance between the centers of the two black holes fixed in the grid frame, as they inspiral in the inertial frame. And \mathcal{M}_R rotates the frame so that the centers of the two black holes remain on the \bar{x} axis in the grid frame, as they move along their orbits in the inertial frame. These maps are also used during inspiral simulations in the same way we use them for our merger simulations. So here we focus on the new shape-control map \mathcal{M}_S , which is one of the critical new developments that made the merger simulations reported in Sec. III possible.

The interiors of the black holes are excised in our simulations at the spherical boundaries shown in Figs. 2 and 3. This is possible because, *if this excision boundary is chosen wisely*, the spacetime inside this boundary cannot influence the spacetime region covered by our computational grid. For hyperbolic evolution systems, like the generalized harmonic form of the Einstein equations used in our code [46], boundary conditions must be placed on each incoming characteristic field at each boundary point. Apparent horizons are surfaces that are often used to study black holes, because they can be found numerically in a fairly straightforward way, and because (if they exist at all) they are always located within the true event horizons. If

the excision boundaries for our computational domain were placed exactly on the apparent horizons, then all the characteristic fields of the Einstein system would be outflowing (with respect to the computational domain) since apparent horizons are spacelike (or null) hypersurfaces. Boundary conditions would not be needed on any field at such boundaries. Unfortunately we are not able to place excision boundaries precisely on the apparent horizons. So the best that can be done numerically is to place them slightly inside the apparent horizons (i.e., the apparent horizons must remain within the computational domain), if we are to avoid the need for boundary conditions on these excision surfaces.

However, if an excision boundary is placed *inside* an apparent horizon, the outflow condition is no longer automatic or simple: the condition depends on the shape and the location of the excision boundary, its motion with respect to the horizon, and the gauge. In our simulations, the excision boundaries are kept somewhat inside the apparent horizons, so the outflow condition can—and does—fail if we are not careful. One reason (and probably the principal reason) we need to control the shapes of the computational domains through the map \mathcal{M}_S is to keep pure outflow conditions on all the dynamical fields at the excision boundaries. We do this, as described in more detail below, by requiring that the excision boundaries closely track the shapes and sizes of the horizons.

Another (probably secondary) reason that shape and size control are needed in our numerical simulations is related to finite numerical resolution. If the excision boundaries had different shapes than the horizons, then some points on the boundaries would be located deeper into the black hole interior and hence closer to the spacetime singularity. Higher numerical resolution would be needed at these points, and for fixed resolution, the amount of constraint violation and errors in the solution would be largest there. In many situations we find that numerical instabilities at these points cause our simulations to fail. This mode of failure can be eliminated by keeping the shapes and sizes of the boundaries close to those of the horizons.

We have decomposed the map \mathcal{M} , which connects the grid-frame coordinates \bar{x}^i with inertial-frame coordinates x^i , into kinematical and shape-control parts: $\mathcal{M} = \mathcal{M}_K \circ \mathcal{M}_S$. It will be convenient to have a name for the intermediate coordinate system whose existence is implied by this split. The map \mathcal{M}_K connects the inertial-frame coordinates x^i with coordinates \tilde{x}^i in which the centers of the black holes are at rest (approximately). So it is natural to call these intermediate-frame coordinates, \tilde{x}^i , the “rest-frame” coordinate system. The shape-control map \mathcal{M}_S connects the grid-frame coordinates \bar{x}^i with these rest-frame coordinates \tilde{x}^i .

It is useful to express the shape-control map \mathcal{M}_S as the composition of maps acting on each black hole individually: $\mathcal{M}_S = \mathcal{M}_{S_1} \circ \mathcal{M}_{S_2}$. Such individual black hole

shape-control maps can be written quite generally in the form

$$\tilde{\theta}_A = \bar{\theta}_A, \quad (14)$$

$$\tilde{\phi}_A = \bar{\phi}_A, \quad (15)$$

$$\tilde{r}_A = \bar{r}_A - f_A(\bar{r}_A, \bar{\theta}_A, \bar{\phi}_A) \sum_{\ell=0}^{\ell_{\max}} \sum_{m=-\ell}^{\ell} \lambda_A^{\ell m}(t) Y_{\ell m}(\bar{\theta}_A, \bar{\phi}_A), \quad (16)$$

where $(\bar{r}_A, \bar{\theta}_A, \bar{\phi}_A)$ and $(\tilde{r}_A, \tilde{\theta}_A, \tilde{\phi}_A)$ are, respectively, the grid-frame and rest-frame spherical polar coordinates centered at the (fixed) grid-coordinate location of black hole $A = \{1, 2\}$. The functions $f_A(\bar{r}_A, \bar{\theta}_A, \bar{\phi}_A)$ are fixed functions of space in the grid frame. We note that these maps become the identity whenever $f_A(\bar{r}_A, \bar{\theta}_A, \bar{\phi}_A) = 0$, so making $f_A(\bar{r}_A, \bar{\theta}_A, \bar{\phi}_A)$ vanish as $\bar{r}_A \rightarrow \infty$ ensures that the distortion is limited to the neighborhood of each black hole. The parameters $\lambda_A^{\ell m}(t)$ specify the angular structure of the distortion map. They are determined by a feedback control system (discussed in the Appendix) that dynamically adjusts the shape and overall size of the grid-frame coordinates relative to the rest-frame coordinates.

The Caltech/Cornell collaboration has used shape-control maps of this form, Eqs. (14)–(16), in all of our recent binary black hole simulations [54,55,57]. In those cases the functions $f_A(\bar{r}_A, \bar{\theta}_A, \bar{\phi}_A)$ were taken to be smooth functions of \bar{r}_A alone: roughly constant near each black hole and falling off to zero rapidly enough that \mathcal{M}_{S_1} and \mathcal{M}_{S_2} approximately commute. These functions had large gradients which made it difficult to control the grid distortion near the horizons without introducing additional unwanted distortions elsewhere. As a result, we were never able to achieve very robust shape control using these maps.

One of our major breakthroughs, leading to the successful merger results reported here, is an improvement in our choice of the map functions $f_A(\bar{r}_A, \bar{\theta}_A, \bar{\phi}_A)$. We now use simpler functions $f_A(\bar{r}_A, \bar{\theta}_A, \bar{\phi}_A)$ that have much smaller gradients and that are exactly zero outside (nonintersecting) compact regions surrounding each black hole. This means the new maps \mathcal{M}_{S_1} and \mathcal{M}_{S_2} commute, exactly. These new maps can be defined everywhere between the black holes because of the new nonoverlapping grid structure discussed in Sec. II B, but these improvements are achieved at the expense of smoothness: the new $f_A(\bar{r}_A, \bar{\theta}_A, \bar{\phi}_A)$ are smooth except at subdomain boundaries where they may only be continuous, not differentiable.¹ Fortunately, this lack of smoothness at subdomain boundaries is not a problem for our basic evolution method. In our multidomain code the equations are solved independently on each subdomain, and subdomains communicate only by

¹The idea of introducing nonsmooth time-dependent mappings was also suggested by Larry Kidder.

equating the appropriate characteristic fields at their mutual boundaries. These boundary conditions require no differentiation, so the grid coordinates themselves need not be smooth across these boundaries.

The new $f_A(\bar{r}_A, \bar{\theta}_A, \bar{\phi}_A)$ are chosen to be unity inside a sphere of radius \bar{b}_A centered on black hole $A = \{1, 2\}$, then decrease linearly with the radius along a ray through the center of the black hole, and finally vanish on the surface of a centered cube of size $2\bar{a}_A$. A two-dimensional sketch of this grid-frame domain is shown in Fig. 4, which is merely an abstract version of the type of grid structures we use around each black hole as illustrated in Figs. 1 and 2. This function $f_A(\bar{r}_A, \bar{\theta}_A, \bar{\phi}_A)$ can be expressed analytically as

$$f_A(\bar{r}_A, \bar{\theta}_A, \bar{\phi}_A) = \begin{cases} 1, & \bar{r}_A \leq \bar{b}_A, \\ \frac{\bar{r}_A - \bar{\rho}_A}{\bar{b}_A - \bar{\rho}_A}, & \bar{b}_A \leq \bar{r}_A \leq \bar{\rho}_A, \\ 0, & \bar{\rho}_A \leq \bar{r}_A, \end{cases} \quad (17)$$

where

$$\begin{aligned} \bar{\rho}_A(\bar{\theta}_A, \bar{\phi}_A) \\ = \bar{a}_A [\max(|\sin \bar{\theta}_A \cos \bar{\phi}_A|, |\sin \bar{\theta}_A \sin \bar{\phi}_A|, |\cos \bar{\theta}_A|)]^{-1} \end{aligned} \quad (18)$$

is the value of \bar{r}_A on the surface of the cube of size $2\bar{a}_A$ centered on black hole $A = \{1, 2\}$. The different cases of




FIG. 4. Two-dimensional illustration of the grid-frame coordinate domain on which the function $f_A(\bar{r}_A, \bar{\theta}_A, \bar{\phi}_A)$ is defined in Eq. (17).

the maxima which appear in the denominator of Eq. (18) correspond to the different faces of this cube.

We note that the choice of $f_A(\bar{r}_A, \bar{\theta}_A, \bar{\phi}_A)$ in Eq. (17) implies that the shape-control map \mathcal{M}_S , and hence the map between grid-frame and rest-frame coordinates is not smooth. Since the kinematical map \mathcal{M}_K is smooth, this also implies that the full map \mathcal{M} relating grid-frame to inertial-frame coordinates is not smooth either. This non-smoothness does not cause a problem for our basic evolution method, but it implies that calculations that require smoothness must not be performed in the grid frame. For example, the representation of a smooth apparent horizon in grid coordinates will not be smooth. So apparent horizon finding and other calculations that require smoothness are done in the smooth rest-frame or inertial-frame coordinates.

The angular structures of the shape-control maps defined in Eqs. (14)–(16) are determined by the parameters $\lambda_A^{\ell m}(t)$. These are chosen dynamically to ensure that the shapes of the excision boundaries match the evolving shapes of the apparent horizons. These maps also control the sizes of the excision boundaries by adjusting λ_A^{00} , which are chosen to ensure that the excision boundaries remain close to but safely inside the apparent horizons. Figure 5 illustrates the effect of this distortion map on the structure of the grid surrounding a black hole in one of the generic merger simulations described in Sec. III. The choices of suitable target shapes and sizes for these shape-control maps, and the feedback control system that we use to implement these




FIG. 5 (color online). Illustrates the inertial frame representation of the grid structure around one black hole, just at the time of merger. This grid structure has been distorted in relation to the inertial frame by the shape-control maps \mathcal{M}_S described here.

choices dynamically in our simulations, are also among the critical technical advances that allow us now to perform robust merger and ringdown simulations. The details of exactly how this is done are somewhat complicated, however, so we defer their discussion to the Appendix.

III. MERGER AND RINGDOWN SIMULATIONS

In this section we present simulations of the merger and ringdown of binary black hole systems using our Spectral Einstein Code (SpEC). In Sec. III A we describe the binary black hole initial data sets that we evolve. In Sec. III B we discuss our evolution algorithm, pointing out at which stage we employ the various improvements described in Sec. II and what goes wrong when these improvements are not used. Finally, in Sec. III C we describe six different binary black hole merger and ringdown simulations performed using these new methods. We present snapshots of the shapes and locations of the horizons at various times during the merger and ringdown, and a demonstration of convergence of the constraint violations.

A. Initial data

Our initial data are in quasiequilibrium [72–74] (see also [75,76]), and are built using the conformal thin sandwich formalism [77,78] with the simplifying choices of conformal flatness and maximal slicing. Quasiequilibrium boundary conditions are imposed on spherical excision boundaries for each black hole, with the lapse boundary condition given by Eq. (33a) of Ref. [74]. The spins of the black holes are determined by boundary conditions on the shift vector at each excision surface [73].

This formalism for constructing initial data also requires the initial radial velocity v_r of each black hole toward its partner and an initial orbital angular frequency Ω_0 (chosen to be about the z axis without loss of generality). These parameters determine the orbital eccentricity of the binary, and can be tuned by an iterative method [57,61] to produce data with very low eccentricity. This has been done for most of the initial data sets described here (cases A through D in Table I), but is unnecessary for the purposes of the present paper, since our goal here is simply to document

TABLE I. Black hole binary configurations run through merger and ringdown using the methods presented here. For the initial spin parameters \tilde{S}_1/M_1^2 and \tilde{S}_2/M_2^2 , the \hat{z} direction is parallel to the orbital angular momentum.

Case	M_2/M_1	\tilde{S}_1/M_1^2	\tilde{S}_2/M_2^2	N_{orbits}
A	1	0	0	16
B	2	0	0	15
C	1	$-0.4\hat{z}$	$-0.4\hat{z}$	11
D	1	$0.4\hat{z}$	$0.4\hat{z}$	15
E	2	$0.2(\hat{z} - \hat{x})/\sqrt{2}$	$-0.4(\hat{z} + \hat{y})/\sqrt{2}$	8.5
F	2	$0.2(\hat{z} - \hat{x})/\sqrt{2}$	$-0.4(\hat{z} + \hat{y})/\sqrt{2}$	1.5

our improved methods for evolving binaries through merger. For cases E and F, v_r and Ω_0 are chosen roughly by using post-Newtonian formulas that ignore spins.

Table I shows the mass ratio and initial spins for the black hole binary configurations we evolve here. The initial data and first 15 orbits of inspiral for case A are identical to those presented in [50,54]. The initial data and the first 9 orbits of inspiral for case C are identical to those presented in [57]. Cases E and F are fully generic, with unequal masses and unaligned spins, and exhibit precession and radiation-reaction recoil.

B. Evolution procedure

In this section we summarize our procedure for evolving black hole binary systems through merger and ringdown, concentrating on the improvements discussed in Sec. II. This procedure can be divided into several stages, beginning with the early inspiral and extending through ringdown:

- (1) Evolve early inspiral using “quasiequilibrium” gauge.
- (2) Transition smoothly to damped harmonic gauge.
- (3) Eliminate overlapping subdomains.
- (4) Turn on shape control.
- (5) Replace remaining inner spherical shells with cubed spheres.
- (6) Turn on size control just before holes merge.
- (7) After merger, interpolate to a single-hole grid to run through ringdown.

The above ordering of these stages is not the only possible choice: we have exchanged the order of some of these stages without trouble. For example, for the runs described here we perform stage 3 at the instant we start stage 2. The transitions between these stages are relatively simple, require little or no fine-tuning, and can be automated. We now discuss each of these stages in turn. For several of these stages, we will illustrate the effects of new improvements (Sec. II) for one particular evolution, case F of Table I.

1. Early inspiral

The gauge is chosen early in the evolution following the procedure of Ref. [50]: the initial data are constructed in quasiequilibrium, so we choose the initial H_a to make the time derivatives of the lapse and shift zero. We attempt to maintain this quasiequilibrium condition by demanding that H_a remains constant in time in the grid frame in the following sense: we require

$$\partial_{\tilde{t}} \tilde{H}_a = 0, \quad (19)$$

where \tilde{H}_a is a tensor (note that H_a is not a tensor) defined so that $\tilde{H}_a = H_a$ in the inertial frame. The bars in Eq. (19) refer to the grid-frame coordinates. We refer to this as “quasiequilibrium” gauge, although strictly speaking this

gauge maintains quasiequilibrium only when the spin directions are constant in the grid frame.

For cases A–D, the simulations follow many orbits using this quasiequilibrium gauge. For cases E and F, in which the spins change direction in the grid frame, this gauge is not appropriate and causes difficulties, so for these cases we transition to harmonic gauge $H_a = 0$ very early in the inspiral.

During inspiral, the map \mathcal{M}_E uniformly contracts the grid so that the grid-coordinate distance between the centers of the two black holes remains constant. Because this contraction is uniform, the spherical-shell subdomains inside each apparent horizon shrink relative to the horizon. This means that the excision boundary (the inner boundary of the innermost spherical shell) inside each black hole moves further into the strong-field region in the interior, towards the singularity. If this motion continues unchecked, gradients on the innermost spherical shell increase until the solution is no longer resolved. However, the inner boundary of the *next-to-innermost* spherical shell is also shrinking, and it eventually shrinks enough that the characteristic fields on this boundary all become outflowing (into the black hole); when this happens, the innermost spherical shell no longer influences the exterior solution, so we simply drop that shell from the domain. This shell-dropping process occurs automatically as successive spherical shells shrink relative to the horizons. This process is illustrated in Fig. 6, which shows the minimum of the

(into the hole) characteristic speeds on spherical grid boundaries around one of the black holes. At any given time, the innermost of these boundaries is an excision boundary. All characteristic speeds at all points on the excision boundary (and hence the minimum) must be positive for the excision algorithm to be well posed. The innermost shell is dropped whenever the characteristic speeds are positive on the inner boundary of the next shell. For example, surface A_3 in Fig. 6 is the excision boundary at $t = 24M$, but after the characteristic speeds on A_4 become positive (around $t = 26M$), the innermost shell is dropped and A_4 becomes the excision boundary. Similarly, A_5 becomes the excision boundary at $t = 38$, and so on. It is important that the characteristic speed on surface A_n in Fig. 6 does not become negative before the speed on surface A_{n+1} becomes positive; otherwise this shell-dropping procedure will fail. Shell dropping works well during most of the evolution but it fails in the last stage before merger, which is discussed in Sec. III B 6.

When describing mergers, it is useful to introduce a measure of distance between the black holes. One such measure of distance is indicated on the horizontal axis at the top of Fig. 6. This is the spatial (in each time slice) proper separation between the apparent horizon surfaces, obtained by integrating along the rest-frame \tilde{x} axis (recall that in the rest frame, the centers of the horizons always remain along this axis). This is not the true proper separation between the horizons because we do not minimize over all possible paths between all possible pairs of points on the two surfaces; however, we use this only as an approximate measure of how far the black holes are from merger. Note that the horizons touch each other when (or possibly before) this proper separation measure falls to zero.

It turns out that the improvements discussed in Sec. II are unnecessary during early inspiral for cases A–D and are not yet used at this stage. For instance, the domain decomposition uses overlapping cylinders and spherical shells, and not the cubed-sphere subdomains described in Sec. II B. We find no problems while the holes are far enough apart that they remain roughly in equilibrium. However, the binary eventually becomes extremely dynamical, and the black holes become significantly distorted. Unless the algorithm is modified, the simulation fails (typically because of large constraint violations) before the black holes merge.

Typically, it is possible to extend the quasiequilibrium inspiral stage until the proper separation decreases to about $7M$ or $8M$, but the next stages, Secs. III B 2 and III B 3, can be done sooner if desired. (For example, cases E and F use a nonoverlapping cubed-sphere grid from $t = 0$.) For the simulations presented here, the length of this quasiequilibrium inspiral stage is variable. For instance, for case A in Table I this stage lasts for 15 orbits, but for case F we transition to damped harmonic gauge and eliminate over-




FIG. 6 (color online). The minimum of the outgoing (into the hole) characteristic speeds on the nine innermost spherical grid boundaries (labeled $A_0 \dots A_8$) around the larger black hole, for part of run F.

lapping subdomains starting at $t = 0$ because the black holes are initially very close together.

2. Transition to damped harmonic gauge

When the inspiraling black holes reach a proper separation of about $8M$, we smoothly transition from quasiequilibrium gauge to damped harmonic gauge at $t = t_g$ (g stands for “gauge”) by choosing

$$H_a(t) = \tilde{H}_a(t)e^{-(t-t_g)^4/\sigma_g^4} + \mu_L \log\left(\frac{\sqrt{g}}{N}\right)t_a - \mu_S N^{-1} g_{ai} N^i, \quad (20)$$

where $\tilde{H}_a(t)$ is the value of $H_a(t)$ obtained from the quasiequilibrium gauge condition (19), and σ_g is a constant. The last line in Eq. (20) is the damped harmonic gauge condition, Eq. (8). The values of μ_L and μ_S are set according to Eq. (9), using

$$\mu_0 = \begin{cases} 0, & t < t_d \\ 1 - e^{-(t-t_d)^2/\sigma_d^2}, & t > t_d \end{cases}, \quad (21)$$

where t_d and σ_d are constants. It is necessary to choose σ_g and σ_d sufficiently large or else the gauge becomes unnecessarily dynamical, possibly resulting in incoming characteristic fields on the excision boundary.

For all runs shown here, t_d corresponds approximately to a proper separation of $8M$. We use the values $\sigma_g = 20M$, $\sigma_d = 50M$ for the equal-mass binaries (cases A, C, and D), $\sigma_g = 15M$, $\sigma_d = 100M$ for case B, and $\sigma_g = 25M$, $\sigma_d = 50M$ for case F. For these cases, $t_g = t_d$. For case E, the gauge is rolled off to harmonic early in the inspiral, so at $t = t_d$ we need only to roll on the damped harmonic gauge, and we use $\sigma_d = 40M$. The constants in Eqs. (20) and (21) can be chosen quite flexibly: the only constraint is that σ_g and σ_d must be large enough to avoid significant spurious gauge dynamics, but small enough so that the simulation is using damped harmonic gauge before the black holes approach each other too closely. The σ_g and σ_d can always be made longer by choosing earlier transition times t_g and t_d .

3. Eliminate overlapping subdomains

Although the domain decomposition consisting of overlapping spherical shells, cylinders, and cylindrical shells is quite efficient for the early inspiral, it often suffers from numerical instabilities, particularly as the black holes approach merger. Therefore, we regrid onto the cubed-sphere subdomains described in Sec. II B. Use of these nonoverlapping subdomains eliminates these remaining instabilities, and effectively increases our resolution. For most of the runs described here, we happen to regrid at the same time t_d that we begin the transition to damped harmonic gauge. However, for runs E and F, the entire simulation uses the nonoverlapping cube-sphere subdomains. As far

as we know, there is no reason that regridding cannot occur at times other than t_d or t_g .

4. Shape control

The next stage is to turn on shape control by introducing the map \mathcal{M}_S defined in Eqs. (14)–(18). The parameters $\lambda_A^{\ell m}$ that appear in this map are determined (via a feedback control system) by Eqs. (A3) and (A6) for $\ell > 0$, and λ_A^{00} is set to zero. This map deforms the grid so that the boundaries of the spherical shells inside the horizons, including the excision boundaries, are mapped to closely match the horizons’ shapes.

We do not turn on shape control until time $t_g + 2.5\sigma_g$, which is when the coefficient multiplying the gauge term $\tilde{H}_a(t)$ in Eq. (20) becomes smaller than double-precision round off. This ensures that H_a remains a smooth function of space in inertial coordinates. If shape control were turned on too early, H_a would fail to be everywhere smooth because the quasi-equilibrium-gauge quantity \tilde{H}_a is a smooth function of space in the *grid-frame* coordinates (and its numerical value is constant in time for each grid point), and the deformation is only C^0 across subdomain boundaries.

Shape control is necessary for the shell-dropping procedure described in Sec. III B 1 to remain successful as the holes approach each other. For a shell to be dropped, the *entire* inner boundary of that shell must remain an outflow surface until the *entire* inner boundary of the next shell becomes an outflow surface. If these boundaries do not have the same shape as the horizon, then typically some portion of either boundary moves too close or too far from the horizon, and violates this condition. An example is shown in Fig. 7. Like Figs. 6 and 7 show the minimum characteristic speeds on the boundaries of inner spherical shells for a portion of case F, but Fig. 7 shows two separate evolutions: one with shape control turned on at $t = 62M$ (solid curves) and one with no shape control (dashed curves). Without shape control, the minimum characteristic speed on surface A_7 stops growing around $t = 100M$ and becomes negative soon thereafter. Because A_7 is the excision surface at $t \sim 100M$, the excision algorithm becomes ill posed and the simulation stops at $t = 106M$, well before merger (which occurs at $t \sim 128M$). In the simulation with active shape control, the same boundary A_7 shows no such problems—in fact it is dropped at $t = 102M$ when all speeds on the next shell, A_8 become positive.

5. Eliminate inner spherical shells

Spherical-shell subdomains are extremely efficient when the solution is nearly spherically symmetric, because these subdomains use Y_{lm} basis functions, which are well suited for nearly spherical functions. However, once the holes become sufficiently distorted, it is difficult to resolve the region near each horizon using spherical shells. At this point in the evolution, we replace the inner spherical shells




FIG. 7 (color online). The minimum of the outgoing (into the hole) characteristic speeds on the remaining four spherical-shell boundaries (labeled $A_5 \dots A_8$) around the larger black hole, for a portion of run F. The dashed curves correspond to runs without shape control, and the solid curves correspond to runs with shape control turned on at $t = 62M$. The run without shape control terminates at proper separation $5.2M$ (time $t = 106M$) because the excision algorithm fails on surface A_7 .

with cubed-sphere shells. Figure 8 shows the constraints as a function of time for two evolutions of case F: one with inner spherical shells (dashed curves), and another in which the inner spherical shells were replaced with cubed-sphere shells at $t = 80M$ (solid curves). The quantity plotted is $\|C_{\text{GH}}\|$, the L^2 norm over all the constraint fields of our first-order generalized harmonic (GH) system, normalized by the L^2 norm over the spatial gradients of all dynamical fields [see Eq. (71) of Ref. [46]]. The L^2 norms are taken over the portion of the computational volume that lies outside the apparent horizons. For the spherical-shell case in Fig. 8, the constraints in the spherical shells grow large enough that the simulation terminates at $t = 129.5M$, when the proper separation has fallen to $0.5M$. Note that these simulations typically proceed through merger even if we do not replace spherical shells with cubed-sphere shells. So while this replacement step is not strictly necessary, eliminating spherical shells improves the accuracy of the simulation (as measured by the constraint quantity $\|C_{\text{GH}}\|$) by almost an order of magnitude.

6. Size control

Eventually, even using shape control, shell dropping typically fails before the horizons merge: the outflow




FIG. 8 (color online). Constraint norm $\|C_{\text{GH}}\|$ as a function of time for a portion of case F. The dashed curve corresponds to a run with spherical shells around each hole, and the solid curve corresponds to an identical run with these inner spherical shells replaced by cubed-sphere shells at $t = 80M$. The constraint norm grows larger and begins to grow earlier when spherical shells are used.

condition on the inner boundary of the inner shell fails *before* the outflow condition on the inner boundary of the next shell becomes satisfied. This problem occurs because the map \mathcal{M}_{E} shrinks the entire grid—including regions inside each hole—faster and faster as the holes approach each other. As soon as the velocity of excision boundary towards the center of the hole becomes large enough, the boundary becomes timelike, and the excision algorithm fails. To remedy this, we turn off shell dropping and we turn on size control. Size control slows down the infall of the excision surface (by pulling the excision sphere towards the horizon), thus keeping the characteristic speeds from changing sign. Size control is implemented by changing the map parameters $\lambda_A^{00}(t)$, which previously were set to zero, to values given by Eqs. (A3) and (A10). This causes the size of each excision boundary to be driven towards some fraction η of the size of the appropriate horizon. For the equal-mass runs A, C, and D we use $\eta = 0.8$ for both black holes. For the unequal-mass run B and the generic cases E and F we use $\eta = 0.9$ and $\eta = 0.7$ for the smaller and larger black holes, respectively. These values do not require fine-tuning; for instance, case E still runs through merger and ringdown if we change η to 0.9 for both black holes. The main criterion for choosing this parameter is that a grid sphere of radius \bar{b}_A (see Fig. 4) should not be mapped outside the cube of side $2\bar{a}_A$. If this




FIG. 9 (color online). The minimum of the outgoing (into the hole) characteristic speeds on the two innermost cubed-sphere boundaries (labeled B_0 and B_1) around the smaller black hole, for a portion of run F. The dashed curves correspond to runs without size control, and the solid curves correspond to runs with size control turned on at $t = 120.2M$. Without size control, the innermost cubed sphere is dropped at $t = 124.5M$ and B_1 becomes the excision boundary. However, the characteristic speeds become negative on B_1 at $t = 127M$ and excision fails. With size control, the characteristic speeds on the excision boundary B_0 remain positive until proper separation $0.075M$, well after merger.

occurs, the map \mathcal{M}_S becomes singular and the run fails. Figure 9 shows the minimum characteristic speed at particular domain boundaries for two evolutions of case F: one with size control and one without. Without size control, the outflow condition on the excision boundary fails at proper separation $1.67M$ and time $t = 127.24M$. With size control, the evolution proceeds to proper separation $0.015M$, time $t = 130.17$, well past the formation of a common apparent horizon, which occurs at proper separation $1.4M$ (time $t = 128.23M$).

7. Ringdown

At some time $t = t_m$ shortly after a common apparent horizon forms, we define a new grid, composed only of spherical shells. This grid has only a single excision boundary inside the common horizon. We also define a new map $\mathcal{M}_{\text{ringdown}}: \bar{x}^i \rightarrow x^i$ that can be written as a composition of simpler maps:

$$\mathcal{M}_{\text{ringdown}} := \mathcal{M}_{\text{Tf}} \circ \mathcal{M}_{\text{Ef}} \circ \mathcal{M}_{\text{R}} \circ \mathcal{M}_{\text{S}_3}. \quad (22)$$

This is similar to the ringdown maps described in Refs. [54,57] (which had no translation) and Ref. [55] (which had no rotation). The new translation map \mathcal{M}_{Tf} is tied to the center of the new common horizon, and is not continuous with the old translation map \mathcal{M}_{T} except near the outer boundary where both translation maps are the identity. The new expansion map \mathcal{M}_{Ef} is chosen to be continuous with the old expansion map \mathcal{M}_{E} at the outer boundary, but it is the identity near the merged black hole. This new expansion map \mathcal{M}_{Ef} smoothly becomes constant in time shortly after merger. The rotation map \mathcal{M}_{R} is continuous at merger, and after merger it smoothly becomes constant in time. The map \mathcal{M}_{S_3} has the same form as Eqs. (14)–(17) except the function $\bar{\rho}_A(\bar{\theta}_A, \bar{\phi}_A)$ defined in Eq. (18) is replaced by $\bar{\rho}_3 = \text{constant}$, where the constant value corresponds to a subdomain boundary.

As described in Ref. [54], at $t = t_m$ the parameters $\lambda_A^{\ell m}$ are chosen so that the common apparent horizon is stationary and spherical in the grid frame. Similarly, the map \mathcal{M}_{Tf} is chosen at $t = t_m$ so that the center of the horizon is located at the grid-frame origin and is stationary in the grid frame.

Once we have defined the maps at $t = t_m$, we then interpolate all variables from the old grid to the new grid. Note that the grid frame changes discontinuously in time at $t = t_m$. Because of this, we take care to properly use both the new map $\mathcal{M}_{\text{ringdown}}$ and the old map \mathcal{M} in doing this interpolation, so that the inertial frame and quantities defined in that frame remain smooth. In particular, the gauge-source function H_a is still determined by Eq. (20) during ringdown, and remains smooth at $t = t_m$.

For $t > t_m$, the parameters $\lambda_A^{\ell m}(t)$ are determined by a feedback control system that keeps the common apparent horizon stationary in the grid frame. Likewise, the map \mathcal{M}_{Tf} keeps the common apparent horizon centered at the origin in the grid frame.

C. Results

In this section we present some results from the simulations listed in Table I. We first show snapshots of the inertial-coordinate shapes of the apparent horizons during merger and after ringdown. Figures 10–13 show cross sections of the horizons in the orbital plane for cases A–D; in these cases, the orbital plane is well defined and is constant in time. For all cases, we show the apparent horizons of the individual black holes and the common apparent horizon at the time when the common horizon is first detected (solid curves), and at the time t_m when we transition to a new grid that has a single excision boundary (dotted curves). We also show the apparent horizon of the final remnant black hole after it has reached equilibrium (dashed curves). Figures 14 and 15 show cross sections of the horizons for cases E and F, in which the orbital plane precesses. For these cases we show cross sections in the coordinate plane (defined using the flat metric) that are




FIG. 10 (color online). Three snapshots of the apparent horizons of the nonspinning equal-mass binary black hole merger (case A). Solid curves are the orbital plane cross section when a common horizon is first detected; dotted curves represent the time of transition between the binary merger and the single-hole ringdown evolutions; dashed curve shows the final equilibrium horizon.




FIG. 11 (color online). Apparent horizon snapshots as in Fig. 10, except for a 2:1 mass ratio nonspinning binary black hole merger (case B).




FIG. 12 (color online). Apparent horizon snapshots as in Fig. 10, except for an antialigned-spin equal-mass binary black hole merger (case C).

perpendicular to the instantaneous orbital angular velocity at the time the common horizon is first detected, and that passes through the coordinate center of the common horizon at this time. We show cross sections at three different times: the time of horizon formation, the time t_m , and a time $300M$ after merger, all with respect to the same plane. The remnant black hole in cases E and F have nonzero




FIG. 13 (color online). Apparent horizon snapshots as in Fig. 10, except for an aligned-spin equal-mass binary black hole merger (case D).




FIG. 14 (color online). Three snapshots of the apparent horizons of a generic binary black hole merger (case E). The plane of the cross sections is a coordinate plane perpendicular to the instantaneous orbital axis at the time the common horizon is first detected. The additional degree of freedom is fixed by having this plane go through the coordinate center of the shape of the common horizon at its first detection. Solid curves are the cross section when a common horizon is first detected; dotted curves represent the time of transition between the binary merger and the single-hole ringdown evolutions; dashed curve shows the horizon at a time well after merger. Note that the individual apparent horizons intersect; this was also observed in Ref. [79].




FIG. 15 (color online). Apparent horizon snapshots as in Fig. 14, except for the generic binary black hole merger case F in Table I.




FIG. 16 (color online). Constraint norm $\|C_{\text{GH}}\|$ as a function of time for the generic (case F) binary black hole merger and ringdown, computed with four different numerical resolutions. The small inset graph shows that numerical convergence is maintained (at a reduced rate) even during the most dynamical part of the merger.

linear momentum both because of radiation reaction and because of some nonzero linear momentum present in the initial data; this is why the centers of the horizons shown in Figs. 14 and 15 change with time.

For case F, we also show the constraint norm as a function of time in Fig. 16. We plot the same quantity $\|C_{\text{GH}}\|$ as shown in Fig. 8. This quantity is shown for four numerical resolutions, and is convergent at all times (although the convergence rate is smaller near merger when the solution is most dynamical). The constraints are largest at $t = t_m$, when we transition to a grid with a single excision boundary. Just after $t = t_m$ the constraints decrease discontinuously by a small amount because part of the computational domain has been newly excised. The approximate number of grid points for the four resolutions is $\{79^3, 87^3, 95^3, 103^3\}$ just before merger and $\{44^3, 51^3, 57^3, 63^3\}$ during ringdown.

ACKNOWLEDGMENTS

We thank Jan Hesthaven for helpful discussions. We are grateful to Luisa Buchman, Larry Kidder, and Harald Pfeiffer for the improved kinematical coordinate maps and associated control systems that handle inspirals of precessing and unequal-mass binaries, and we thank Harald Pfeiffer for help preparing the initial data. We also thank Luisa Buchman, Tony Chu, Geoffrey Lovelace, and Harald Pfeiffer for producing some of the BBH inspiral simulations that we here extend through

merger and ringdown. We acknowledge use of the Spectral Einstein Code (SpEC). This work was supported in part by grants from the Sherman Fairchild Foundation, by NSF Grants No. PHY-0601459, No. PHY-0652995, and by NASA Grant No. NNX09AF97G.

APPENDIX: CONTROLLING THE MAPS

The purpose of the shape-control maps \mathcal{M}_{S_A} , with $A = \{1, 2\}$, is to distort the grid structures around the black holes so the excision boundaries of the computational domain are mapped to surfaces lying just inside and having the same basic shapes as the apparent horizons. Choosing the right maps is equivalent to choosing the right target values for the parameters $\lambda_A^{\ell m}(t)$ that define these maps, cf. Equation (16). This Appendix describes in some detail how these parameters are chosen, and what target surfaces are used to fix these maps in the merger simulations described in Sec. III.

The grid structures used in our merger simulations, cf. Sec. II B, have excision boundaries that are grid-frame coordinate spheres. The target surfaces S_A^T to which we want to map these grid-frame spheres can be written in the form

$$\tilde{r}_A = \sum_{\ell=0}^{\ell_{\max}} \sum_{m=-\ell}^{\ell} \Lambda_A^{\ell m}(t) Y_{\ell m}(\tilde{\theta}_A, \tilde{\phi}_A), \quad (\text{A1})$$

where the expansion coefficients $\Lambda_A^{\ell m}(t)$ define the target surface and $(\tilde{r}_A, \tilde{\theta}_A, \tilde{\phi}_A)$ are rest-frame coordinates.

Let $\tilde{r}_A^{\text{target}}$ denote the radii of the grid-frame coordinate spheres that are to be mapped onto the target surfaces S_A^T . The rest-frame representations of these coordinate spheres, using Eqs. (14)–(16), are

$$\tilde{r}_A = \tilde{r}_A^{\text{target}} - \sum_{\ell=0}^{\ell_{\max}} \sum_{m=-\ell}^{\ell} \lambda_A^{\ell m}(t) Y_{\ell m}(\tilde{\theta}_A, \tilde{\phi}_A). \quad (\text{A2})$$

Note that the functions f_A that appear in Eq. (17) are set equal to unity here because we always choose the target grid-spheres $\tilde{r}_A^{\text{target}}$ to be smaller than the outer radii of the largest spherical subdomain layers: $\bar{b}_a \geq \tilde{r}_A^{\text{target}}$. It follows that the target spheres will be mapped to the shapes of S_A^T when

$$\lambda_A^{\ell m}(t) = \begin{cases} -\Lambda_A^{00}(t) + \frac{\tilde{r}_A^{\text{target}}}{Y_{00}}, & \ell = 0, \\ -\Lambda_A^{\ell m}(t), & \ell > 0. \end{cases} \quad (\text{A3})$$

These conditions cannot be imposed directly for a number of reasons, but they can be enforced approximately using a feedback control system. The control system used in the merger calculations described in Sec. III is the same one used to perform our earlier binary black hole simulations [48], so we will not describe it in detail here.

To complete the specification of the shape-control maps we must choose the target grid-spheres $\tilde{r}_A^{\text{target}}$, and the target

surface parameters $\Lambda_A^{\ell m}$. From Eq. (A3) it follows that the choice of $\tilde{r}_A^{\text{target}}$ only effects the $\ell = 0$ size control part of the distortion map. So until size control is activated late in our runs, there is no need to specify what $\tilde{r}_A^{\text{target}}$ actually is. We think of it as being the radius of the suitably scaled apparent horizon, but this fact does not influence the shape-control part of the map in any way. Once size control is activated late in the run, then we require $\tilde{r}_A^{\text{target}}$ to be the radius of the excision boundary:

$$\tilde{r}_A^{\text{target}} = \tilde{r}_A^{\text{ex}}. \quad (\text{A4})$$

Next consider the choice of target surfaces S_A^T , starting with the $\ell > 0$ contributions that control the shape but not the overall scaling of the map. The idea is to choose the target surfaces S_A^T to be similar in shape to the apparent horizons \mathcal{H}_A . The \mathcal{H}_A can be represented as smooth surfaces in the rest-frame coordinate system:

$$\tilde{r}_A = \sum_{\ell=0}^{\ell_{\max}} \sum_{m=-\ell}^{\ell} \tilde{S}_A^{\ell m}(t) Y_{\ell m}(\tilde{\theta}, \tilde{\phi}). \quad (\text{A5})$$

To keep the shapes of the target surfaces similar to the shapes of the apparent horizons, the target surface parameters $\Lambda_A^{\ell m}$ should be made proportional to $S_A^{\ell m}$. We find it is appropriate to scale these coefficients by a factor, $G(\tilde{R}_A)$, which depends on the average radius of the apparent horizon, $\tilde{R}_A = S_A^{00} Y_{00}$:

$$\Lambda_A^{\ell m}(t) = G(\tilde{R}_A) S_A^{\ell m}(t). \quad (\text{A6})$$

The larger the apparent horizon radius in relation to the desired excision radius, the smaller this scaling factor must be to maintain an appropriate shape for the excision boundary. In practice, we find the scale factor

$$G(\tilde{R}_A) = \frac{\bar{a}_A}{\tilde{R}_A} \tanh^p \left[\left(\frac{\tilde{R}_A}{\bar{a}_A} \right)^{1/p} \right], \quad (\text{A7})$$

works quite well for $p = 2$. This scaling factor is near unity when $\tilde{R}_A < \bar{a}_A$ and decreases like $1/\tilde{R}_A$ for larger values of \tilde{R}_A . The $\tanh x$ function is introduced here because it is linear for small values of $|x|$, and approaches 1 for $x \gg 1$. The p dependence maintains these asymptotic forms, but allows the transition between them to occur more quickly than the $p = 1$ case.

We have found no need to introduce the $\ell = 0$ size control parts of these maps, until the final plunge phase just before the black holes merge. So during most of our merger simulations we simply set $\lambda_A^{00}(t) = 0$. This allows the apparent horizons to grow relative to the grid coordinates as the map \mathcal{M}_E contracts the rest-frame relative to the inertial coordinates. This growth in the apparent horizons in the grid frame allows us to drop a number of spherical subdomain layers as the evolution progresses, and this helps remove unwanted constraint violations from the computational domain.

At some point however, we may not be able to drop additional spherical subdomain layers: either because we run out of layers, or because the inner boundary experiences incoming characteristic fields when the apparent horizon expands too quickly during the final plunge. When either of these conditions occurs we turn on the $\ell = 0$ part of the shape-control maps, to keep the radius of the excision boundary \bar{r}_A^{ex} close to the size of the apparent horizon. It will be useful to define

$$\Delta r_A = \eta \tilde{R}_A - \bar{r}_A^{\text{ex}}. \quad (\text{A8})$$

We choose η so that when $\Delta r_A < 0$ there is no need to impose size control, and use η in the range $0.7 \leq \eta \leq 0.9$ for the merger simulations reported in Sec. III. When Δr_A becomes positive we want to turn on size control, but we want to do it in a fairly smooth and continuous way. This is done by defining the function

$$P(\Delta r_A) = \begin{cases} 0, & \Delta r_A < 0 \\ \tanh(\frac{40\Delta r_A}{\tilde{R}_A}), & \Delta r_A \geq 0 \end{cases} \quad (\text{A9})$$

which vanishes for $\Delta r_A \leq 0$, and asymptotically approaches $P(\Delta r_A) \rightarrow 1$ for $\Delta r_A \gg \tilde{R}_A/40$. We find that an appropriate value for the overall scale of the target surfaces \mathcal{S}_A^T , i.e., their $\ell = 0$ components, to be

$$\Lambda_A^{00} = \frac{\bar{r}_A^{\text{ex}} + \Delta r_A P(\Delta r_A)}{Y_{00}}. \quad (\text{A10})$$

Once size control is activated, the target radius is given by Eq. (A4): $\bar{r}_A^{\text{target}} = \bar{r}_A^{\text{ex}}$. This choice implies that the average rest-frame radius of the excision boundary, $\tilde{R}_A^{\text{ex}} = \bar{r}_A^{\text{ex}} - \lambda_A^{00} Y_{00}$, approaches $\eta \tilde{R}_A$ when $\Delta r_A \geq R_A/40$.

To summarize, the appropriate shape and size control of the inner grid structures in our binary black hole merger simulations are determined by the target shape-control surfaces \mathcal{S}_A^T , whose definitions are given in Eqs. (A6) and (A10). The shape-control map needed to distort the grid coordinates to the shape of \mathcal{S}_A^T , is given in Eq. (A3). This mapping is enforced with an appropriate feedback control system.

-
- [1] F. Pretorius, Phys. Rev. Lett. **95**, 121101 (2005).
 - [2] M. Campanelli, C. O. Lousto, P. Marronetti, and Y. Zlochower, Phys. Rev. Lett. **96**, 111101 (2006).
 - [3] J. G. Baker, J. Centrella, D.-I. Choi, M. Koppitz, and J. van Meter, Phys. Rev. Lett. **96**, 111102 (2006).
 - [4] M. Campanelli, C. O. Lousto, and Y. Zlochower, Phys. Rev. D **74**, 041501(R) (2006).
 - [5] M. Campanelli, C. O. Lousto, and Y. Zlochower, Phys. Rev. D **74**, 084023 (2006).
 - [6] M. Campanelli, C. O. Lousto, Y. Zlochower, B. Krishnan, and D. Merritt, Phys. Rev. D **75**, 064030 (2007).
 - [7] F. Herrmann, I. Hinder, D. M. Shoemaker, P. Laguna, and R. A. Matzner, Phys. Rev. D **76**, 084032 (2007).
 - [8] P. Marronetti, W. Tichy, B. Brügmann, J. González, and U. Sperhake, Phys. Rev. D **77**, 064010 (2008).
 - [9] E. Berti, V. Cardoso, J. A. Gonzalez, U. Sperhake, and B. Brügmann, Classical Quantum Gravity **25**, 114035 (2008).
 - [10] M. Campanelli, Classical Quantum Gravity **22**, S387 (2005).
 - [11] F. Herrmann, I. Hinder, D. Shoemaker, and P. Laguna, Classical Quantum Gravity **24**, S33 (2007).
 - [12] J. G. Baker *et al.*, Astrophys. J. **653**, L93 (2006).
 - [13] J. A. Gonzalez, U. Sperhake, B. Brügmann, M. Hannam, and S. Husa, Phys. Rev. Lett. **98**, 091101 (2007).
 - [14] D. Pollney *et al.*, Phys. Rev. D **76**, 124002 (2007).
 - [15] M. Campanelli, C. O. Lousto, Y. Zlochower, and D. Merritt, Phys. Rev. Lett. **98**, 231102 (2007).
 - [16] J. A. Gonzalez, M. D. Hannam, U. Sperhake, B. Brügmann, and S. Husa, Phys. Rev. Lett. **98**, 231101 (2007).
 - [17] B. Brügmann, J. A. González, M. Hannam, S. Husa, and U. Sperhake, Phys. Rev. D **77**, 124047 (2008).
 - [18] F. Herrmann, I. Hinder, D. Shoemaker, P. Laguna, and R. A. Matzner, Astrophys. J. **661**, 430 (2007).
 - [19] D.-I. Choi *et al.*, Phys. Rev. D **76**, 104026 (2007).
 - [20] J. G. Baker *et al.*, Astrophys. J. **668**, 1140 (2007).
 - [21] W. Tichy and P. Marronetti, Phys. Rev. D **76**, 061502(R) (2007).
 - [22] J. D. Schnittman *et al.*, Phys. Rev. D **77**, 044031 (2008).
 - [23] M. Campanelli, C. O. Lousto, Y. Zlochower, and D. Merritt, Astrophys. J. Lett. **659**, L5 (2007).
 - [24] M. Koppitz *et al.*, Phys. Rev. Lett. **99**, 041102 (2007).
 - [25] S. H. Miller and R. Matzner, Gen. Relativ. Gravit. **41**, 525 (2009).
 - [26] J. G. Baker *et al.*, Astrophys. J. **682**, L29 (2008).
 - [27] J. Healy *et al.*, Phys. Rev. Lett. **102**, 041101 (2009).
 - [28] L. Boyle, M. Kesden, and S. Nissanke, Phys. Rev. Lett. **100**, 151101 (2008).
 - [29] L. Boyle and M. Kesden, Phys. Rev. D **78**, 024017 (2008).
 - [30] A. Buonanno, L. E. Kidder, and L. Lehner, Phys. Rev. D **77**, 026004 (2008).
 - [31] W. Tichy and P. Marronetti, Phys. Rev. D **78**, 081501 (2008).
 - [32] M. Kesden, Phys. Rev. D **78**, 084030 (2008).
 - [33] L. Rezzolla, Classical Quantum Gravity **26**, 094023 (2009).
 - [34] E. Barausse and L. Rezzolla, Astrophys. J. **704**, L40 (2009).
 - [35] C. O. Lousto, M. Campanelli, and Y. Zlochower, arXiv:0904.3541.
 - [36] A. Buonanno, G. B. Cook, and F. Pretorius, Phys. Rev. D **75**, 124018 (2007).
 - [37] J. G. Baker, J. R. van Meter, S. T. McWilliams, J. Centrella, and B. J. Kelly, Phys. Rev. Lett. **99**, 181101 (2007).

- [38] Y. Pan *et al.*, Phys. Rev. D **77**, 024014 (2008).
- [39] A. Buonanno *et al.*, Phys. Rev. D **76**, 104049 (2007).
- [40] M. Hannam, S. Husa, J. A. González, U. Sperhake, and B. Brügmann, Phys. Rev. D **77**, 044020 (2008).
- [41] A. Gopakumar, M. Hannam, S. Husa, and B. Brügmann, Phys. Rev. D **78**, 064026 (2008).
- [42] M. Hannam, S. Husa, B. Brügmann, and A. Gopakumar, Phys. Rev. D **78**, 104007 (2008).
- [43] I. Hinder, F. Herrmann, P. Laguna, and D. Shoemaker, arXiv:0806.1037.
- [44] The Cactus Computational Toolkit, <http://www.cactuscode.org>.
- [45] M. Hannam *et al.*, Phys. Rev. D **79**, 084025 (2009).
- [46] L. Lindblom, M. A. Scheel, L. E. Kidder, R. Owen, and O. Rinne, Classical Quantum Gravity **23**, S447 (2006).
- [47] O. Rinne, L. Lindblom, and M. A. Scheel, Classical Quantum Gravity **24**, 4053 (2007).
- [48] M. A. Scheel *et al.*, Phys. Rev. D **74**, 104006 (2006).
- [49] L. E. Kidder, L. Lindblom, M. A. Scheel, L. T. Buchman, and H. P. Pfeiffer, Phys. Rev. D **71**, 064020 (2005).
- [50] M. Boyle *et al.*, Phys. Rev. D **76**, 124038 (2007).
- [51] M. Boyle *et al.*, Phys. Rev. D **78**, 104020 (2008).
- [52] M. Boyle, D. A. Brown, and L. Pekowsky, Classical Quantum Gravity **26**, 114006 (2009).
- [53] A. Buonanno *et al.*, Phys. Rev. D **79**, 124028 (2009).
- [54] M. A. Scheel *et al.*, Phys. Rev. D **79**, 024003 (2009).
- [55] G. Lovelace *et al.*, arXiv:0907.0869.
- [56] M. Cohen, H. P. Pfeiffer, and M. A. Scheel, Classical Quantum Gravity **26**, 035005 (2009).
- [57] T. Chu, H. P. Pfeiffer, and M. A. Scheel, arXiv:0909.1313.
- [58] L. Lindblom and B. Szilágyi, Phys. Rev. D **80**, 084019 (2009).
- [59] M. W. Choptuik and F. Pretorius, arXiv:0908.1780.
- [60] M. Boyle, L. Lindblom, H. P. Pfeiffer, M. A. Scheel, and L. E. Kidder, Phys. Rev. D **75**, 024006 (2007).
- [61] H. P. Pfeiffer *et al.*, Classical Quantum Gravity **24**, S59 (2007).
- [62] J. S. Hesthaven, SIAM J. Sci. Comput. **18**, 658 (1997).
- [63] J. S. Hesthaven, SIAM J. Sci. Comput. **20**, 62 (1998).
- [64] J. S. Hesthaven, Applied Numerical Mathematics **33**, 23 (2000).
- [65] D. Gottlieb and J. S. Hesthaven, J. Comput. Appl. Math. **128**, 83 (2001)0377-0427.
- [66] R. Sadourny, Mon. Weather Rev. **100**, 136 (1972).
- [67] C. Ronchi, R. Iacono, and P. S. Paolucci, J. Comput. Phys. **124**, 93 (1996).
- [68] M. Rancic, R. J. Purser, and F. Mesinger, Q. J. R. Meteorol. Soc. **122**, 959 (1996).
- [69] M. Taylor, J. Tribbia, and M. Iskandarani, J. Comput. Phys. **130**, 92 (1997).
- [70] L. Lehner, O. Reula, and M. Tiglio, Classical Quantum Gravity **22**, 5283 (2005).
- [71] E. Pazos, M. Tiglio, M. D. Duez, L. E. Kidder, and S. A. Teukolsky, Phys. Rev. D **80**, 024027 (2009).
- [72] G. B. Cook, Phys. Rev. D **65**, 084003 (2002).
- [73] G. B. Cook and H. P. Pfeiffer, Phys. Rev. D **70**, 104016 (2004).
- [74] M. Caudill, G. B. Cook, J. D. Grigsby, and H. P. Pfeiffer, Phys. Rev. D **74**, 064011 (2006).
- [75] E.ourgoulhon, P. Grandclément, and S. Bonazzola, Phys. Rev. D **65**, 044020 (2002).
- [76] P. Grandclément, E.ourgoulhon, and S. Bonazzola, Phys. Rev. D **65**, 044021 (2002).
- [77] J. W. York, Phys. Rev. Lett. **82**, 1350 (1999).
- [78] H. P. Pfeiffer and J. W. York, Phys. Rev. D **67**, 044022 (2003).
- [79] B. Szilágyi *et al.*, Classical Quantum Gravity **24**, S275 (2007).