Classical and Quantum
Gravity

Simulating binary black hole mergers using discontinuous Galerkin methods

To cite this article: Geoffrey Lovelace *et al* 2025 *Class. Quantum Grav.* 42 035001

OPEN ACCESS
IOP Publishing
Classical and Quantum Gravity
Class. Quantum Grav. 42 (2025) 035001 (31pp)
https://doi.org/10.1088/1361-6382/ad9f19

Simulating binary black hole mergers using discontinuous Galerkin methods

Geoffrey Lovelace[1,12,*], **Kyle C Nelli**[2,12,*], **Nils Deppe**[3,4,5], **Nils L Vu**[2], **William Throwe**[5], **Marceline S Bonilla**[1], **Alexander Carpenter**[1], **Lawrence E Kidder**[5], **Alexandra Macedo**[1], **Mark A Scheel**[2], **Azer Afram**[1], **Michael Boyle**[5], **Andrea Ceja**[1,6], **Matthew Giesler**[5], **Sarah Habib**[2], **Ken Z Jones**[1], **Prayush Kumar**[7], **Guillermo Lara**[8], **Denyz Melchor**[1,9], **Iago B Mendes**[2,10], **Keefe Mitman**[5], **Marlo Morales**[1,11], **Jordan Moxon**[2], **Eamonn O'Shea**[5], **Kyle Pannone**[1], **Harald P Pfeiffer**[8], **Teresita Ramirez-Aguilar**[1,6], **Jennifer Sanchez**[1,6], **Daniel Tellez**[1], **Saul A Teukolsky**[2,5] and **Nikolas A Wittek**[8]

[1] Nicholas and Lee Begovich Center for Gravitational-Wave Physics and Astronomy, California State University Fullerton, Fullerton, CA 92834, United States of America

[2] Theoretical Astrophysics 350-17, California Institute of Technology, Pasadena, CA 91125, United States of America

[3] Department of Physics, Cornell University, Ithaca, NY 14853, United States of America

[4] Laboratory for Elementary Particle Physics, Cornell University, Ithaca, NY 14853, United States of America

[5] Cornell Center for Astrophysics and Planetary Science, Cornell University, Ithaca, NY 14853, United States of America

[6] Center for Interdisciplinary Exploration and Research in Astrophysics, Northwestern University, Evanston, IL 60201, United States of America

[7] International Centre for Theoretical Sciences, Tata Institute of Fundamental Research, Bangalore 560089, India

[8] Max Planck Institute for Gravitational Physics (Albert Einstein Institute), D-14467 Potsdam, Germany

---

[12] The first two authors contributed equally.

[*] Authors to whom any correspondence should be addressed.

[9] Department of Physics and Astronomy, UCLA, Mani L. Bhaumik Institute for Theoretical Physics, Los Angeles, CA 90095, United States of America
[10] Department of Physics and Astronomy, Oberlin College, Oberlin, OH 44074, United States of America
[11] Department of Physics & Astronomy, Washington State University, Pullman, WA 99164, United States of America

E-mail: glovelace@fullerton.edu and knelli@caltech.edu

Received 2 October 2024; revised 9 December 2024
Accepted for publication 13 December 2024
Published 7 January 2025

Abstract

Binary black holes are the most abundant source of gravitational-wave observations. Gravitational-wave observatories in the next decade will require tremendous increases in the accuracy of numerical waveforms modeling binary black holes, compared to today's state of the art. One approach to achieving the required accuracy is using spectral-type methods that scale to many processors. Using the `SpECTRE` numerical-relativity (NR) code, we present the first simulations of a binary black hole inspiral, merger, and ringdown using discontinuous Galerkin (DG) methods. The efficiency of DG methods allows us to evolve the binary through ~ 18 orbits at reasonable computational cost. We then use `SpECTRE`'s Cauchy Characteristic Evolution (CCE) code to extract the gravitational waves at future null infinity. The open-source nature of `SpECTRE` means this is the first time a spectral-type method for simulating binary black hole evolutions is available to the entire NR community.

Keywords: discontinuous Galerkin, binary black holes, numerical relativity

1. Introduction

Binary black holes are the most abundant source of gravitational-wave observations to date [1]. Realizing the scientific potential of these observations requires accurate models of the emitted gravitational waves as the black holes inspiral, merge, and ring down to a final, stationary state. Building these models requires numerical-relativity (NR) simulations of binary black holes, because analytic approximations (e.g. the post-Newtonian [2] approximation) alone break down near the time of merger.

Since the first breakthrough simulations [3–5], the NR community has developed codes capable of evolving two black holes through inspiral, merger, and ringdown (see [6, 7] for a review). Several groups have used NR codes to build catalogs of gravitational waveforms for applications to gravitational-wave astronomy [8–15]. Today's NR codes are sufficiently accurate for the observations that LIGO and Virgo are making. However, observatories planned for the next decade, including the Einstein Telescope [16] and Cosmic Explorer [17] on Earth and the Laser Interferometer Space Antenna (LISA) [18] in space, will be so sensitive that they will require NR waveforms with a substantial increase in accuracy [19–21].

Spectral-type methods are extremely efficient; this makes them a promising avenue toward the ultimate goal of achieving the needed accuracy for future gravitational-wave observatories. In comparison, almost all current NR codes for evolving binary black holes use finite-difference methods, with numerical errors decreasing as a power law with increasing resolution. However, recent results from the `AthenaK` code [22] show that finite-difference methods using graphics processing units (GPUs) might be another approach to achieving the required accuracy. The spectral Einstein code (`SpEC`) [23] uses a pseudospectral method (see [24] for a review of these methods) to construct and evolve binary-black-hole initial data. With pseudospectral methods, errors decrease exponentially with increasing number of grids points in the computational domain's elements ('p-refinement'). `SpEC`'s exponential convergence makes it highly efficient, but its performance, and therefore the achievable accuracy, is limited by aspects of its design. For instance, because it uses computational domains divided into a small number of high-resolution elements, `SpEC` simulations of binary black holes cannot scale beyond $\mathcal{O}(10^2)$ CPU cores. `SpEC` is also a closed-source code, unavailable to most of the NR community. Other examples of pseudospectral or spectral methods being used for solving the initial value problem are `Elliptica` [25], `FUKA` [26], and `bamps` [27]. In terms of evolving spacetimes, the `Nmesh` [28] code has been used to successfully simulate single black holes using discontinuous Galerkin (DG) methods, and the `bamps` [29, 30] code uses pseudospectral methods to evolve spacetimes with single dynamical black holes with a focus on critical behavior [31–38] but has also simulated boson stars [39]. Recently, [40] performed a 0.5 orbit grazing collision of two black holes (a similar setup to [41]) using a finite volume grid in the strong field region and a DG method in the wave zone.

We present the first simulations of a binary black hole inspiral, merger, and ringdown using a DG method [42] (see [43] for a review of DG). The efficiency of DG methods allows us to evolve the binary through ~ 18 orbits at reasonable computational cost: DG, being a spectral-type method, has exponential convergence with p-refinement. For context, 18 orbits is slightly less than the median (20) for binary-black-hole simulations in the SXS catalog [12] (which also uses a spectral-type method) but larger than almost all of the simulations in the RIT and Maya catalogs [14, 15] (which use finite-difference methods). We chose the length of `SpECTRE`'s first binary-black-hole simulation largely out of convenience, balancing a desire to demonstrate `SpECTRE`'s capability with minimizing turnaround time as we tested and fine-tuned our methods. We expect that using `SpECTRE` to simulate more orbits would be straightforward, without requiring changes to the code, although extending the length beyond 100 orbits would likely require implementing in `SpECTRE` similar techniques as those discussed in [44], which presents a 175-orbit `SpEC` binary-black-hole simulation.

Specifically, in this work, we present `SpECTRE`'s [45] first simulations of ~ 18 orbits of inspiral, merger, and ringdown of an equal-mass, non-spinning binary black hole, using DG methods. We then use `SpECTRE`'s Cauchy Characteristic Evolution (CCE) module [45–47] to evolve the gravitational waves to future null infinity. These results demonstrate that DG methods can provide high-accuracy gravitational waveforms from binary black hole mergers for application to gravitational-wave data analysis. By implementing our approach in `SpECTRE`, an open-source NR code, we are also making a spectral-type binary-black-hole evolution code available to the entire NR community for the first time.

The rest of this paper is organized as follows. In section 2, we discuss the numerical methods used in `SpECTRE`'s binary-black-hole simulations. Then, in section 3, we first test our method's stability with simulations of single black holes before presenting results for simulations of binary black holes with `SpECTRE`. We briefly conclude in section 4.

2. Methods

2.1. Equations of motion

We adopt the standard 3+1 form of the spacetime metric, (see, e.g. [48, 49]),

$$\mathrm{d}s^2 = g_{ab}\mathrm{d}x^a\mathrm{d}x^b = -\alpha^2\mathrm{d}t^2 + \gamma_{ij}\left(\mathrm{d}x^i + \beta^i\mathrm{d}t\right)\left(\mathrm{d}x^j + \beta^j\mathrm{d}t\right), \tag{1}$$

where α is the lapse, β^i the shift vector, and γ_{ij} is the spatial metric.We use the Einstein summation convention, summing over repeated indices. Latin indices from the first part of the alphabet $a, b, c, \ldots$ denote spacetime indices ranging from 0 to 3, while Latin indices $i, j, \ldots$ are purely spatial, ranging from 1 to 3. We work in units where $c = G = 1$.

We evolve the first-order generalized harmonic (FOGH) system, given by [50],

$$\partial_t g_{ab} = (1+\gamma_1)\beta^k\partial_k g_{ab} - \alpha\Pi_{ab} - \gamma_1\beta^i\Phi_{iab} + \gamma_1 v_g^k\left(\partial_k g_{ab} - \Phi_{kab}\right), \tag{2}$$

$$\begin{aligned}\partial_t\Phi_{iab} = {}& \beta^k\partial_k\Phi_{iab} - \alpha\partial_i\Pi_{ab} + \alpha\gamma_2\partial_i g_{ab} + \frac{1}{2}\alpha n^c n^d\Phi_{icd}\Pi_{ab} \\ & + \alpha\gamma^{jk}n^c\Phi_{ijc}\Phi_{kab} - \alpha\gamma_2\Phi_{iab},\end{aligned} \tag{3}$$

$$\begin{aligned}\partial_t\Pi_{ab} = {}& \beta^k\partial_k\Pi_{ab} - \alpha\gamma^{ki}\partial_k\Phi_{iab} + \gamma_1\gamma_2\beta^k\partial_k g_{ab} \\ & + 2\alpha g^{cd}\left(\gamma^{ij}\Phi_{ica}\Phi_{jdb} - \Pi_{ca}\Pi_{db} - g^{ef}\Gamma_{ace}\Gamma_{bdf}\right) \\ & - 2\alpha\nabla_{(a}H_{b)} - \frac{1}{2}\alpha n^c n^d\Pi_{cd}\Pi_{ab} - \alpha n^c\Pi_{ci}\gamma^{ij}\Phi_{jab} \\ & + \alpha\gamma_0\left(2\delta^c{}_{(a}n_{b)} - g_{ab}n^c\right)\mathcal{C}_c - \gamma_1\gamma_2\beta^i\Phi_{iab},\end{aligned} \tag{4}$$

where g_{ab} is the spacetime metric, $\Phi_{iab} = \partial_i g_{ab}$, $\Pi_{ab} = n^c\partial_c g_{ab}$, n^a is the unit normal vector to the spatial slice, γ_0 damps the 1-index or gauge constraint $\mathcal{C}_a = H_a + \Gamma_a$, γ_1 controls the linear degeneracy of the system, γ_2 damps the 3-index constraint $\mathcal{C}_{iab} = \partial_i g_{ab} - \Phi_{iab}$, Γ_{abc} are the spacetime Christoffel symbols of the first kind, $\Gamma_a = g^{bc}\Gamma_{bca}$, and v_g^k is the grid/mesh velocity as discussed in section 2.2.

The gauge source function H_a can be any arbitrary function depending only upon the spacetime coordinates x^a and g_{ab}, but not derivatives of g_{ab}, since that may spoil the strong hyperbolicity of the system [51, 52].

Defining s_i to be the unit normal covector to a 2d surface with $s_a = (0, s_i)$, and $s^a = g^{ab}s_b$, the characteristic fields for the FOGH system are [50]

$$w^g_{ab} = g_{ab}, \tag{5}$$

$$w^0_{iab} = \left(\delta^k_i - s^k s_i\right)\Phi_{kab}, \tag{6}$$

$$w^{\pm}_{ab} = \Pi_{ab} \pm s^i\Phi_{iab} - \gamma_2 g_{ab}, \tag{7}$$

with associated characteristic speeds

$$\lambda_{w^g} = -(1+\gamma_1)\beta^i s_i - (1+\gamma_1)v_g^i s_i, \tag{8}$$

$$\lambda_{w^0} = -\beta^i s_i - v_g^i s_i, \tag{9}$$

$$\lambda_{w^\pm} = \pm\alpha - \beta^i s_i - v_g^i s_i, \tag{10}$$

where we denote the velocity of the grid/mesh as v_g^i (see section 2.2 for details on our moving mesh method). The evolved variables as a function of the characteristic fields are given by

$$g_{ab} = w^g_{ab}, \tag{11}$$

$$\Pi_{ab} = \frac{1}{2}\left(w^{+}_{ab} + w^{-}_{ab}\right) + \gamma_2 w^{g}_{ab}, \tag{12}$$

$$\Phi_{iab} = \frac{1}{2}\left(w^{+}_{ab} - w^{-}_{ab}\right) s_i + w^{0}_{iab}. \tag{13}$$

The constraints for the FOGH system are [50]

$$\mathcal{C}_a = H_a + \Gamma_a, \tag{14}$$

$$\begin{aligned}\mathcal{C}_{ia} = {} & \gamma^{jk}\partial_j\Phi_{ika} - \frac{1}{2}\gamma^{j}_{a}g^{cd}\partial_j\Phi_{icd} + n^b\partial_i\Pi_{ba} - \frac{1}{2}n_a g^{cd}\partial_i\Pi_{cd} \\ & + \partial_i H_a + \frac{1}{2}\gamma^{j}_{a}\Phi_{jcd}\Phi_{ief}g^{ce}g^{df} + \frac{1}{2}\gamma^{jk}\Phi_{jcd}\Phi_{ike}g^{cd}n^e n_a \\ & - \gamma^{jk}\gamma^{mn}\Phi_{jma}\Phi_{ikn} + \frac{1}{2}\Phi_{icd}\Pi_{be}n_a\left(g^{cb}g^{de} + \frac{1}{2}g^{be}n^c n^d\right) \\ & - \Phi_{icd}\Pi_{ba}n^c\left(g^{bd} + \frac{1}{2}n^b n^d\right) + \frac{1}{2}\gamma_2\left(n_a g^{cd} - 2\delta^{c}_{a}n^d\right)\mathcal{C}_{icd},\end{aligned} \tag{15}$$

$$\mathcal{C}_{iab} = \partial_i g_{ab} - \Phi_{iab}, \tag{16}$$

$$\mathcal{C}_{ijab} = \partial_i\Phi_{jab} - \partial_j\Phi_{iab}, \tag{17}$$

and

$$\begin{aligned}\mathcal{F}_a = {} & \frac{1}{2}\gamma^{i}_{a}g^{bc}\partial_i\Pi_{bc} - \gamma^{ij}\partial_i\Pi_{ja} - \gamma^{ij}n^b\partial_i\Phi_{jba} + \frac{1}{2}n_a g^{bc}\gamma^{ij}\partial_i\Phi_{jbc} \\ & + n_a\gamma^{ij}\partial_i H_j + \gamma^{i}_{a}\Phi_{ijb}\gamma^{jk}\Phi_{kcd}g^{bd}n^c - \frac{1}{2}\gamma^{i}_{a}\Phi_{ijb}\gamma^{jk}\Phi_{kcd}g^{cd}n^b \\ & - \gamma^{i}_{a}n^b\partial_i H_b + \gamma^{ij}\Phi_{icd}\Phi_{jba}g^{bc}n^d - \frac{1}{2}n_a\gamma^{ij}\gamma^{mn}\Phi_{imc}\Phi_{njd}g^{cd} \\ & - \frac{1}{4}n_a\gamma^{ij}\Phi_{icd}\Phi_{jbe}g^{cb}g^{de} + \frac{1}{4}n_a\Pi_{cd}\Pi_{be}g^{cb}g^{de} - \gamma^{ij}H_i\Pi_{ja} \\ & - n^b\gamma^{ij}\Pi_{bi}\Pi_{ja} - \frac{1}{4}\gamma^{i}_{a}\Phi_{icd}n^c n^d\Pi_{be}g^{be} + \frac{1}{2}n_a\Pi_{cd}\Pi_{be}g^{ce}n^d n^b \\ & + \gamma^{i}_{a}\Phi_{icd}\Pi_{be}n^c n^b g^{de} - \gamma^{ij}\Phi_{iba}n^b\Pi_{je}n^e - \frac{1}{2}\gamma^{ij}\Phi_{icd}n^c n^d\Pi_{ja} \\ & - \gamma^{ij}H_i\Phi_{jba}n^b + \gamma^{i}_{a}\Phi_{icd}H_b g^{bc}n^d + \gamma_2\left(\gamma^{id}\mathcal{C}_{ida} - \frac{1}{2}\gamma^{i}_{a}g^{cd}\mathcal{C}_{icd}\right) \\ & + \frac{1}{2}n_a\Pi_{cd}g^{cd}H_b n^b - n_a\gamma^{ij}\Phi_{ijc}H_d g^{cd} + \frac{1}{2}n_a\gamma^{ij}H_i\Phi_{jcd}g^{cd}.\end{aligned} \tag{18}$$

While only the gauge constraint (14) and 3-index constraint (16) are damped, all constraints can be monitored to check the accuracy of the numerical simulation. All the constraints can be combined into a scalar, the constraint energy, given by [50]

$$\mathcal{E} = \delta^{ab}\left[\mathcal{C}_a\mathcal{C}_b + \left(\mathcal{F}_a\mathcal{F}_b + \mathcal{C}_{ia}\mathcal{C}_{jb}\gamma^{ij}\right)\right] + \delta^{ab}\delta^{cd}\left(\mathcal{C}_{iac}\mathcal{C}_{jbd}\gamma^{ij} + \mathcal{C}_{ikac}\mathcal{C}_{jlbd}\gamma^{ij}\gamma^{kl}\right). \tag{19}$$

In practice we have also found that it is typically only necessary to monitor violations of the constraints $\mathcal{C}_a$ and $\mathcal{C}_{iab}$, because they typically grow first and the other violations grow as a consequence.

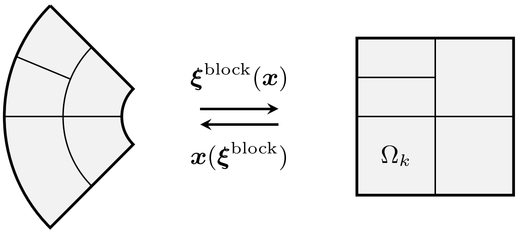

Figure 1. A diagram of the forward and reverse mappings $x^i(\hat{t},\xi^{\hat{\imath}})$ and $\xi^{\hat{\imath}}(t,x^i)$, respectively, from logical (right side of the diagram) to inertial coordinates (left side of the diagram) for an element Ω_k.

2.2. DG Method

SpECTRE uses a DG method for the spatial discretization. We refer readers to [53, 54] and references therein for a detailed discussion of the method and its implementation in SpECTRE; here we summarize the necessary results. The FOGH equations are a first-order strongly hyperbolic system in non-conservative form, which takes the general form

$$\partial_t u_\alpha + B^i_{\alpha\beta}(u_\alpha)\,\partial_i u_\beta = S_\alpha(u_\alpha), \tag{20}$$

where $u_\alpha = \{g_{ab}, \Phi_{iab}, \Pi_{ab}\}$ is the state vector of evolved variables, $B^i_{\alpha\beta}(u)$ is a matrix that depends only on u_α, and $S_\alpha(u)$ are source terms that also only depend on u_α. We denote the logical coordinates of our Legendre–Gauss–Lobatto DG scheme by $\{\hat{t},\xi^{\hat{\imath}}\} = \{\hat{t},\xi,\eta,\zeta\}$ and the inertial coordinates as $\{t=\hat{t}, x^i(\hat{t},\xi^{\hat{\imath}})\}$. We are using a moving mesh as in [53, 55]; therefore, the mapping from logical to inertial coordinates is time dependent. We denote the determinant of the spatial Jacobian of this map as

$$J = \det\left(\frac{\partial x^i}{\partial \xi^{\hat{\imath}}}\right), \tag{21}$$

and the grid or mesh velocity by [53, 55]

$$v^i_g = \partial_{\hat{t}} x^i. \tag{22}$$

SpECTRE decomposes the domain into a set of non-overlapping hexahedra which are deformed using the map $x^i(\hat{t},\xi^{\hat{\imath}})$ illustrated in figure 1. SpECTRE uses a different number of grid points in each logical direction, which we denote by $N_{\breve{\imath}}$ in the ξ direction, $N_{\breve{\jmath}}$ in the η direction, and $N_{\breve{k}}$ in the ζ direction below.

Since we use a moving mesh, we evolve the system

$$\partial_{\hat{t}} u_\alpha + \left[B^i_{\alpha\beta}(u_\alpha) - v^i_g \delta_{\alpha\beta}\right]\partial_i u_\beta = S(u_\alpha). \tag{23}$$

We denote grid point and modal indices with a breve, i.e. $u_{\breve{\imath}\breve{\jmath}\breve{k}}$ is the value of u at the grid point $(\breve{\imath},\breve{\jmath},\breve{k})$. The semi-discrete equations are given by [53, 55]

$$
\begin{aligned}
(\partial_{\hat{t}} u_\alpha)_{\breve{\imath}\breve{\jmath}\breve{k}} &= (S_\alpha)_{\breve{\imath}\breve{\jmath}\breve{k}} - \left(B^i_{\alpha\beta} - v^i_g \delta_{\alpha\beta}\right)_{\breve{\imath}\breve{\jmath}\breve{k}} \left[\left(\frac{\partial \xi^{\hat{1}}}{\partial x^i}\right)_{\breve{\imath}\breve{\jmath}\breve{k}} \sum_{\breve{l}} D^{(\hat{1})}_{\breve{\imath}\breve{l}} (u_\beta)_{\breve{l}\breve{\jmath}\breve{k}} \right. \\
&\quad \left. - \left(\frac{\partial \xi^{\hat{2}}}{\partial x^i}\right)_{\breve{\imath}\breve{\jmath}\breve{k}} \sum_{\breve{l}} D^{(\hat{2})}_{\breve{\jmath}\breve{l}} (u_\beta)_{\breve{\imath}\breve{l}\breve{k}} + \left(\frac{\partial \xi^{\hat{3}}}{\partial x^i}\right)_{\breve{\imath}\breve{\jmath}\breve{k}} \sum_{\breve{l}} D^{(\hat{3})}_{\breve{k}\breve{l}} (u_\beta)_{\breve{\imath}\breve{\jmath}\breve{l}} \right] \\
&\quad - \frac{\delta_{N_{\breve{k}}\breve{k}}}{w_{\breve{k}} J_{\breve{\imath}\breve{\jmath}\breve{k}}} \left\{ \left[J \sqrt{\frac{\partial \xi^{\hat{3}}}{\partial x^i} \gamma^{ij} \frac{\partial \xi^{\hat{3}}}{\partial x^j}} D_\alpha \right]_{\breve{\imath}\breve{\jmath}N_{\breve{k}}} + \left[J \sqrt{\frac{\partial \xi^{\hat{3}}}{\partial x^i} \gamma^{ij} \frac{\partial \xi^{\hat{3}}}{\partial x^j}} D_\alpha \right]_{\breve{\imath}\breve{\jmath}0} \right\} \\
&\quad - \frac{\delta_{N_{\breve{\jmath}}\breve{\jmath}}}{w_{\breve{\jmath}} J_{\breve{\imath}\breve{\jmath}\breve{k}}} \left\{ \left[J \sqrt{\frac{\partial \xi^{\hat{3}}}{\partial x^i} \gamma^{ij} \frac{\partial \xi^{\hat{3}}}{\partial x^j}} D_\alpha \right]_{\breve{\imath}N_{\breve{\jmath}}\breve{k}} + \left[J \sqrt{\frac{\partial \xi^{\hat{3}}}{\partial x^i} \gamma^{ij} \frac{\partial \xi^{\hat{3}}}{\partial x^j}} D_\alpha \right]_{\breve{\imath}0\breve{k}} \right\} \\
&\quad - \frac{\delta_{N_{\breve{\imath}}\breve{\imath}}}{w_{\breve{\imath}} J_{\breve{\imath}\breve{\jmath}\breve{k}}} \left\{ \left[J \sqrt{\frac{\partial \xi^{\hat{3}}}{\partial x^i} \gamma^{ij} \frac{\partial \xi^{\hat{3}}}{\partial x^j}} D_\alpha \right]_{N_{\breve{\imath}}\breve{\jmath}\breve{k}} + \left[J \sqrt{\frac{\partial \xi^{\hat{3}}}{\partial x^i} \gamma^{ij} \frac{\partial \xi^{\hat{3}}}{\partial x^j}} D_\alpha \right]_{0\breve{\jmath}\breve{k}} \right\},
\end{aligned}
\tag{24}
$$

where $w_{\breve{\imath}}$ are the Legendre–Gauss–Lobatto integration weights. We use a method of lines approach to integrate these in time, with the details discussed in section 2.5 below.

For the boundary terms D_α, we use an upwind multi-penalty method [24, 56–58] given by

$$
D_{g_{ab}} = \tilde{\lambda}^{\text{ext}}_{w^g} w^{\text{ext},g}_{ab} - \tilde{\lambda}^{\text{int}}_{w^g} w^{\text{int},g}_{ab}, \tag{25}
$$

$$
\begin{aligned}
D_{\Pi_{ab}} &= \frac{1}{2}\left(\tilde{\lambda}^{\text{ext}}_{w^+} w^{\text{ext},+}_{ab} + \tilde{\lambda}^{\text{ext}}_{w^-} w^{\text{ext},-}_{ab}\right) + \tilde{\lambda}^{\text{ext}}_{w^g} \gamma_2 w^{\text{ext},g}_{ab} \\
&\quad - \frac{1}{2}\left(\tilde{\lambda}^{\text{int}}_{w^+} w^{\text{int},+}_{ab} + \tilde{\lambda}^{\text{int}}_{w^-} w^{\text{int},-}_{ab}\right) - \tilde{\lambda}^{\text{int}}_{w^g} \gamma_2 w^{\text{int},g}_{ab},
\end{aligned}
\tag{26}
$$

$$
\begin{aligned}
D_{\Phi_{iab}} &= \frac{1}{2}\left(\tilde{\lambda}^{\text{ext}}_{w^+} w^{\text{ext},+}_{ab} - \tilde{\lambda}^{\text{ext}}_{w^-} w^{\text{ext},-}_{ab}\right) s^{\text{ext}}_i + \tilde{\lambda}^{\text{ext}}_{w^0} w^{\text{ext},0}_{iab} \\
&\quad - \frac{1}{2}\left(\tilde{\lambda}^{\text{int}}_{w^+} w^{\text{int},+}_{ab} - \tilde{\lambda}^{\text{int}}_{w^-} w^{\text{int},-}_{ab}\right) s^{\text{int}}_i - \tilde{\lambda}^{\text{int}}_{w^0} w^{\text{int},0}_{iab},
\end{aligned}
\tag{27}
$$

where the spatial normal vector to the element interface s^{int}_i is pointing out of the DG element and $\tilde{\lambda} = 0$ if $\lambda > 0$, otherwise $\tilde{\lambda} = \lambda$, i.e. $\tilde{\lambda} = \lambda\Theta(-\lambda)$. Note that we assume s^{ext}_i and s^{int}_i point in the same direction. Also note that these boundary flux terms differ from the multi-penalty approach used in `SpEC` by a factor of 2. That is,

$$
D^{\texttt{SpECTRE}} = 2D^{\texttt{SpEC}}, \tag{28}
$$

ultimately because the lifting terms are different. In `SpEC` and, similarly in `bamps`[29], the penalty term is derived from requiring that the total energy be non-increasing, while in `SpECTRE` the terms come from an integration by parts when deriving the semi-discrete DG equations.

2.3. *Boundary conditions*

At the outer radial boundary, we apply constraint-preserving boundary conditions [50, 59] by adding terms to the time derivative of the characteristic fields and thus also the time derivatives of the evolved variables. We use the characteristic fields and speeds defined in section 2.1.

We define $d_{\hat{t}}w_{\hat{\alpha}}$ as the time derivatives substituted into the transformation equations to the characteristic fields. That is,

$$d_{\hat{t}}w^{g}_{ab} = \partial_{\hat{t}}g_{ab}, \tag{29}$$

$$d_{\hat{t}}w^{0}_{iab} = \left(\delta^{k}_{i} - s^{k}s_{i}\right)\partial_{\hat{t}}\Phi_{kab}, \tag{30}$$

$$d_{\hat{t}}w^{\pm}_{ab} = \partial_{\hat{t}}\Pi_{ab} \pm s^{i}\partial_{\hat{t}}\Phi_{iab} - \gamma_{2}\partial_{\hat{t}}g_{ab}. \tag{31}$$

We also define $D_{\hat{t}}w_{\hat{\alpha}}$ as the characteristic field transformation of the volume right-hand-side, i.e. $\partial_{\hat{t}}u_{\alpha}$ *without* any boundary terms. Finally, for brevity we define the projection tensor $P_{ab} = g_{ab} + n_{a}n_{b} - s_{a}s_{b}$, the inward directed null vector field $k^{a} = (n^{a} - s^{a})/\sqrt{2}$, and the outgoing null vector field $l^{a} = (n^{a} + s^{a})/\sqrt{2}$.

The fields $d_{\hat{t}}w^{g}_{ab}$ and $d_{\hat{t}}w^{0}_{iab}$ are determined solely by the constraint-preserving boundary condition, while the boundary condition for $d_{\hat{t}}w^{-}_{ab}$ is composed of three parts: the constraint preserving part, the physical part, and the gauge part. We denote these as B^{C}_{ab}, B^{P}_{ab} and B^{G}_{ab}. With this, the boundary conditions imposed on the fields are

$$d_{\hat{t}}w^{g}_{ab} = D_{\hat{t}}w^{g}_{ab} + \lambda_{w^{g}}s^{i}\mathcal{C}_{iab}, \tag{32}$$

$$d_{\hat{t}}w^{0}_{kab} = D_{\hat{t}}w^{0}_{ab} + \lambda_{w^{0}}s^{i}P^{j}{}_{k}\mathcal{C}_{ijab}, \tag{33}$$

$$d_{\hat{t}}w^{-}_{ab} = D_{\hat{t}}w^{-}_{ab} + \lambda_{w^{-}}\left[B^{C}_{ab} + B^{P}_{ab} + B^{G}_{ab}\right]. \tag{34}$$

Transforming to the evolved variables we find that the following terms need to be added in order to impose the boundary condition,

$$\partial_{\hat{t}}g_{ab} \to \partial_{\hat{t}}g_{ab} + \lambda_{w^{g}}s^{i}\mathcal{C}_{iab}, \tag{35}$$

$$\partial_{\hat{t}}\Pi_{ab} \to \partial_{\hat{t}}\Pi_{ab} + \frac{1}{2}\lambda_{w^{-}}\left[B^{C}_{ab} + B^{P}_{ab} + B^{G}_{ab}\right] + \gamma_{2}\lambda_{w^{g}}s^{i}\mathcal{C}_{iab}, \tag{36}$$

$$\partial_{\hat{t}}\Phi_{iab} \to \partial_{\hat{t}}\Phi_{iab} - \frac{s_{i}}{2}\lambda_{w^{-}}\left[B^{C}_{ab} + B^{P}_{ab} + B^{G}_{ab}\right] + \lambda_{w^{0}}s^{i}P^{j}{}_{k}\mathcal{C}_{ijab}. \tag{37}$$

We now need to specify the B_{ab} boundary conditions. The constraint-preserving part is

$$B^{C}_{ab} = \sqrt{2}\left(\frac{1}{2}P_{ab}l^{c} + \frac{1}{2}l_{a}l_{b}k^{c} - l_{(a}P_{b)}{}^{c}\right)\left(\mathcal{F}_{c} - s^{k}\mathcal{C}_{kc}\right). \tag{38}$$

The physical boundary conditions are determined by the propagating parts of the Weyl curvature tensor. That is,

$$B^{P}_{ab} = \left(P_{a}{}^{c}P_{b}{}^{d} - \frac{1}{2}P_{ab}P^{cd}\right)\left[C^{-}_{cd} - \gamma_{2}s^{i}\mathcal{C}_{icd}\right], \tag{39}$$

where C^{-}_{ab} is the inward propagating part of the Weyl tensor, given by

$$C^{\pm}_{ab} = \left(P_{a}{}^{c}P_{b}{}^{d} - \frac{1}{2}P_{ab}P^{cd}\right)\left(n^{e} \mp s^{e}\right)\left(n^{f} \mp s^{f}\right)C_{cedf}. \tag{40}$$

For the simulations presented here, we set $C^{-}_{ab} = 0$, though Cauchy-Characteristic Matching (CCM) [60] can be used to prescribe a more physically motivated boundary condition. Recently [61] presented an alternative approach to CCM for providing high-order non-reflecting boundary conditions. Finally, the gauge boundary condition is set using a

Sommerfeld condition on the components not set by the constraint-preserving and physical boundary conditions. The projector for the gauge boundary condition is given by

$$\begin{aligned}\delta_a^c\delta_b^d - P^{C}_{ab}{}^{cd} - P^{P}_{ab}{}^{cd} &= \delta_a^c\delta_b^d - \frac{1}{2}P_{ab}P^{cd} + 2l_{(a}P_{b)}{}^{(c}k^{d)} \\ &\quad - l_a l_b k^c k^d - P_a{}^c P_b{}^d + \frac{1}{2}P_{ab}P^{cd} \\ &= \delta_a^c\delta_b^d + 2l_{(a}P_{b)}{}^{(c}k^{d)} - l_a l_b k^c k^d - P_a{}^c P_b{}^d. \end{aligned} \tag{41}$$

The Sommerfeld condition is

$$B^G_{ab} = \frac{1}{\lambda_{w^-}}\left(2l_{(a}P_{b)}{}^{(c}k^{d)} - 2k_{(a}l_{b)}k^{(c}l^{d)} - k_a k_b l^c l^d\right)\left(\gamma_2 - \frac{1}{r}\right)\partial_t g_{cd}. \tag{42}$$

When evolving spacetimes with black holes, we excise the interior of the black hole as is done in `SpEC` [50]. At excision boundaries, all information is flowing out of the grid and into the black hole, so no boundary condition needs to be applied. However, we monitor the characteristic speeds, (8)–(10), and terminate the code if any of them point into the computational domain. We denote the radius of the excision surfaces by r_{exc}. See section 2.7 for a brief explanation of how we control r_{exc} to avoid any characteristic speed pointing into the computational domain.

2.4. Spectral filter

We use an exponential filter applied to the spectral coefficient c_i in order to eliminate aliasing-driven instabilities. Specifically, for a 1d spectral expansion

$$u(x) = \sum_{\breve{i}=0}^{N} c_{\breve{i}} P_{\breve{i}}(x), \tag{43}$$

where $P_{\breve{i}}(x)$ are the Legendre polynomials, we use the filter

$$c_{\breve{i}} \to c_{\breve{i}} \exp\left[-a\left(\frac{i}{N}\right)^{2b}\right]. \tag{44}$$

We choose the parameters $a = 64$ and $b = 210$ so that only the highest spectral mode is filtered. We apply the filter to all FOGH variables g_{ab}, Φ_{iab} and Π_{ab}. Note that the filter drops the order of convergence for the FOGH variables from $\mathcal{O}(N+1)$ to $\mathcal{O}(N)$ on the DG grid, but is necessary for stability.

2.5. Time integration

We decompose the system using the method of lines and solve the resulting differential equations using a local adaptive time-stepper based on the Adams–Moulton predictor–corrector method [62].The step size in each element is chosen based on an estimate of the truncation error of the time step, using the algorithm described in [63] section 17.2.1. The specific values for the absolute and relative tolerances are given in section 3. As the time-stepping algorithm is more efficient for aligned steps of the same size, the step size in each element is rounded down to a value of the form $0.1M/2^n$ for some non-negative integer n. For the highest-resolution binary-black-hole run in section 3.3, this results in the most-demanding element taking 2^6–2^8 steps for each step on the least demanding element for most of the inspiral. At the time of merger, this can increase to as high as 2^{11} steps for the most-demanding element for each step on the least demanding element.

2.6. Gauge condition

We evolve binary black holes (section 3.3) using the Damped Harmonic gauge condition [52, 64]:

$$H_a = [\mu_{L1} \log(\sqrt{\gamma}/\alpha) + \mu_{L2} \log(1/\alpha)]\, n_a - \mu_S g_{ai} \beta^i/\alpha, \tag{45}$$

using

$$\mu_{L1} = A_{L1} \mathrm{e}^{-(r/\sigma_r)^2} [\log(\sqrt{\gamma}/\alpha)]^{e_{L1}}, \tag{46}$$

$$\mu_{L2} = A_{L2} \mathrm{e}^{-(r/\sigma_r)^2} [\log(\sqrt{\gamma}/\alpha)]^{e_{L2}}, \tag{47}$$

$$\mu_S = A_S \mathrm{e}^{-(r/\sigma_r)^2} [\log(\sqrt{\gamma}/\alpha)]^{e_S}, \tag{48}$$

where r is the coordinate distance from the origin. This condition is designed to drive $\sqrt{\gamma}$ and α to one, while damping out oscillations in the shift. This is because we observe an explosive growth in $\sqrt{\gamma}$ and a rapid collapse in α as the black holes merge. In practice, this ensures coordinates remain sufficiently well behaved throughout inspiral, merger, and ringdown. The amplitudes A_{L1}, A_{L2}, and A_S and exponents e_{L1}, e_{L2} and e_S control the amount of damping, and the spatial decay width σ_r ensures that at large distances, the gauge reduces to harmonic gauge (i.e. to $H_a = 0$). In this paper, we choose $A_{L1} = A_S = 1, A_{L2} = 0, e_{L1} = e_{L2} = e_S = 2$, and $\sigma_r = 100/\sqrt{34.54}$. This choice for σ_r ensures that the spatial decay Gaussian falls to 10^{-15} at a distance $r = 100M$ from the origin.

For some of the single black-hole evolutions (section 3.1), we instead choose H_a to be Γ_a of the analytic initial data. For other single black-hole evolutions, we evolve in harmonic gauge, setting $H_a = 0$ everywhere. For the gauge wave evolution (section 3.2), we set $H_a(t, x^i)$ to the value of $\Gamma_a(t, x^i)$ of the gauge wave analytic solution.

2.7. Control systems

When evolving the FOGH system, if there are black holes, the physical singularities inside of the black holes must be excised from the computational domain. To position the excisions with our moving mesh (described in section 2.2), we use a feedback control system similar to what is presented in [55, 65]. As discussed in section 2.3, the excision surfaces must have all characteristic speeds pointing out of the computational domain, so that no boundary condition must be imposed. In practice this means that the excision surfaces must remain inside the apparent horizons, with the caveat that having them too close to the singularity causes instabilities. In practice the excision surfaces are kept at approximately 95%–99% of the apparent horizons' radii.

Since we *a priori* do not know the motion or shape of the apparent horizons, we use control theory to dynamically update the parameters of the moving mesh periodically during the simulation. The time-dependent coordinate maps of the moving mesh and control signals used to update them are discussed in [65] in sections 4.1–4.3, 4.5, and 5 for the inspiral and section 6 for the ringdown.

The details of how the control systems are implemented within the context of asynchronous task-based parallelism along with the local adaptive time stepping described in section 2.5 are described in [66].

3. Results

In this section, we begin by testing SpECTRE's long-term stability and convergence; first with evolutions of single black holes in different coordinate systems (section 3.1) and then with an evolution of a time-dependent gauge wave on a flat spacetime background (section 3.2). Finally, we present results from a complete simulation of the inspiral, merger, and ringdown of two black holes (section 3.3). The SpECTRE input files used for simulations, including generating the BBH initial data, are provided as ancillary material with the paper.

3.1. Single black hole evolutions

In this section, we use SpECTRE to evolve a single, stationary, black hole that, unless otherwise noted, is non-spinning. We evolve a black hole from analytic initial data corresponding to a black hole at rest centered at the origin with zero spin. We choose the mass of the black hole to be $M = 1$ and work in units of M. In each evolution we use the following values for the FOGH constraint damping parameters,

$$\gamma_0 = \gamma_2 = A_0 \mathrm{e}^{-r^2/w_0^2} + A_1 \mathrm{e}^{-r^2/w_1^2}, \tag{49}$$

$$\gamma_1 = -1, \tag{50}$$

where r is the coordinate distance from the origin, $A_0 = 7.0/M$, $A_1 = 0.1/M$, $w_0 = 2.5M$, and $w_1 = 100.0M$. The computational domain of each evolution covers a spherical shell volume (figure 2) with inner radius $r_{\mathrm{in}} = r_{\mathrm{exc}}$ which differs for our different test cases, and outer boundary coordinate radius $r_{\mathrm{out}} = 1000M$. We apply boundary conditions as described in section 2.3. We use a fourth-order Adams–Moulton predictor–corrector time integrator with absolute and relative time stepper tolerances of 10^{-8} and 10^{-6}, respectively, unless otherwise stated.

3.1.1. Kerr–Schild coordinates. We first evolve a single black hole in Kerr–Schild coordinates from Kerr–Schild initial data. The inner radius of the computational domain is $r_{\mathrm{in}} = r_{\mathrm{exc}} = 1.9M$. In this case there are no coordinate dynamics, so a feedback control system is not necessary, though it is enabled in the simulations presented here. The left panel of figure 3 shows the gauge constraint $\mathcal{C}_a$ and the 3-index constraint $\mathcal{C}_{iab}$ as a function of time for several different resolutions. We evolve the lowest resolution to time $t = 10\,000\,M$ to assess long-term stability, and we evolve the medium and high resolutions to assess convergence. To limit the computational cost of these tests, we choose to evolve the medium and high resolutions only to time $t = 2000\,M$. All simulations are stable to time $t = 2000\,M$, and the lowest resolution remains stable to $t = 10\,000\,M$. The amount of violation of the gauge constraint $\mathcal{C}_a$ and the 3-index constraint $\mathcal{C}_{iab}$ is indicative of the overall constraint violation in the simulation. In this evolution, the constraints remain approximately constant, and they decrease exponentially with increasing p-refinement (that is, increasing points per cell per dimension), as expected. We suspect the transient at $t = 2000\,Mt = 2000\,M$ results from constraint violations reflecting off the outer boundary back to the interior.

The right panel of figure 3 demonstrates long-term stability and convergence for the same setup but with a black hole of dimensionless spin $\chi \equiv S/M^2 = 0.8$ (with $r_{\mathrm{in}} = r_{\mathrm{exc}} = 1.57\,M$).

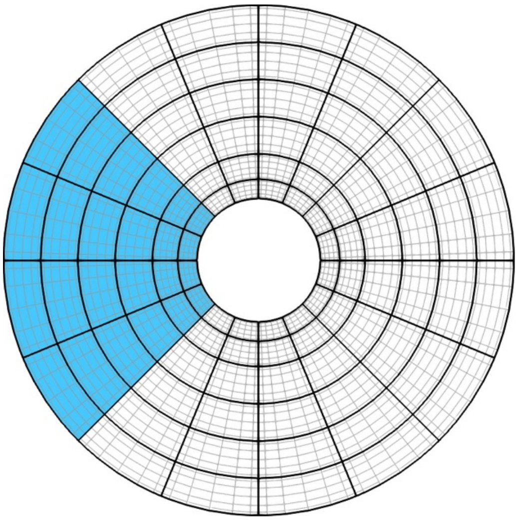

Figure 2. An illustration of a slice through the computational domain used in the evolutions of single black holes described in this paper. Six regions, each in the shape of a deformed cube, combine to cover the volume of a spherical shell, with the inner boundary (the excision surface) inside the black hole's apparent horizon. The six regions are themselves refined radially and angularly into smaller deformed cubes. One of the six regions is shown in blue for clarity.

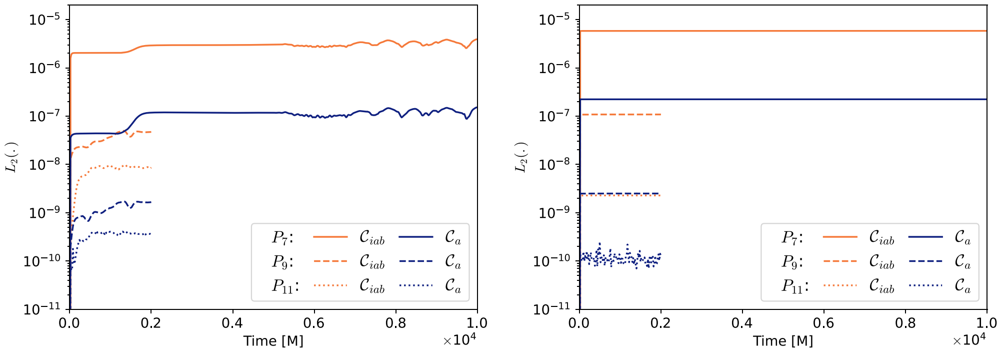

Figure 3. Constraint violations for single black hole evolutions in Kerr–Schild coordinates with Kerr–Schild initial data. We evolve the lowest resolution to time $t = 10\,000\,M$ to demonstrate long term stability and evolve two higher resolutions to time $t = 2000\,M$ to assess convergence with spatial resolution. *Left*: Non-spinning black hole. *Right*: Spinning black-hole with dimensionless spin $\chi = S/M^2 = 0.8$

Again, we see that the constraints remain approximately constant and converge exponentially with increasing p-refinement.

3.1.2. Harmonic coordinates. Next, we evolve a single black hole in harmonic gauge $H_a = 0$ using initial data also in harmonic gauge. Here, $r_{\text{in}} = r_{\text{exc}} = 0.9\,M$. The absolute and relative time stepper tolerances for the highest resolution of this test case are 10^{-10} and 10^{-8}, respectively. Again, since the initial data and evolution use the same gauge there are no gauge

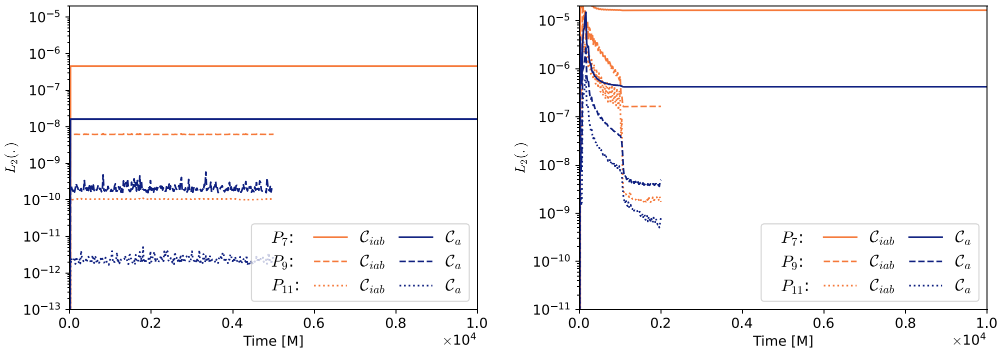

Figure 4. Constraint violations for single black hole evolutions in harmonic coordinates. We evolve the lowest resolution to time $t = 10\,000\,M$ to demonstrate long term stability and evolve two higher resolutions to an earlier time to assess convergence with spatial resolution. *Left*: Analytic initial data is in harmonic coordinates. Higher resolutions evolved to $t = 5000\,M$. The highest resolution has $10\times$ tighter time stepper tolerances. *Right*: Analytic initial data is in Kerr–Schild coordinates. Higher resolutions evolved to $t = 2000\,M$. The difference between the initial data and evolution gauge causes non-trivial dynamics.

dynamics. The left panel of figure 4 shows the gauge constraint $\mathcal{C}_a$ and the 3-index constraint $\mathcal{C}_{iab}$ as a function of time for several different resolutions. We evolve the lowest resolution to time $t = 10\,000\,M$ to assess stability and two higher resolutions to time $t = 5000\,M$ to assess convergence. The constraints again remain constant, and they decrease exponentially with increasing p-refinement.

As a first test of the control system, we evolve a Kerr–Schild black hole in harmonic coordinates. The inner radius of the domain again is $r_{\text{in}} = r_{\text{exc}} = 1.8\,M$. The differing gauge choices in the initial data and evolution create coordinate dynamics that cause the BH horizon to shrink. The control system (section 2.7) must decrease the radius of the excision surface smoothly and precisely to avoid incoming characteristic speeds, so that the problem remains well-posed and the code does not terminate. The right panel of figure 4 shows constraint violations over time for three resolutions. All evolutions are stable, and the constraint violations converge away. However, the constraints remain larger until after time $1000\,M$. We suspect this is caused by initial gauge dynamics, i.e. by time-dependent, outward-moving coordinate effects that travel outward until exiting the domain through the outer boundary at $r_{\text{out}} = 1000\,M$.

3.1.3. Damped harmonic coordinates.

Our final single-black-hole test consists of evolving Kerr–Schild analytic initial data in damped harmonic gauge with $r_{\text{in}} = r_{\text{exc}} = 1.8\,M$. The left panel of figure 5 shows the constraints as a function of time for several different resolutions. Just as in the harmonic gauge case, the non-trivial gauge dynamics cause larger constraint violations until after one light-crossing time to the outer boundary of $r_{\text{out}} = 1000\,M$. The evolutions are stable and converge with increasing resolution. We also repeated this evolution but for a black hole with a dimensionless spin of $\chi = 0.8$ and $r_{\text{in}} = r_{\text{exc}} = 0.57\,M$. We show the constraint violations in the right panel of figure 5. Again we see stable evolutions and exponential convergence with increasing p-refinement.

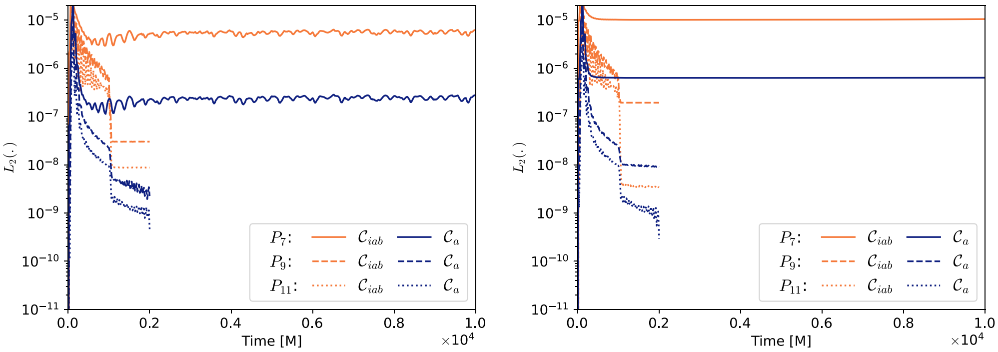

Figure 5. Constraint violations for single black hole evolutions in damped harmonic coordinates with Kerr–Schild initial data. We evolve the lowest resolution to time $t = 10\,000\,M$ to demonstrate long term stability and evolve two higher resolutions to time $t = 2000\,M$ to assess convergence with spatial resolution. *Left*: Non-spinning black hole. *Right*: Spinning black-hole with dimensionless spin $\chi = S/M^2 = 0.8$.

3.2. Gauge wave

As a final test of convergence and stability, we evolve analytic initial data consisting of a gauge wave in flat spacetime, a test conceived in [67] as part of a 'standard testbed' for NR codes. Physically, the solution is equivalent to flat spacetime, but the chosen coordinates include a sinusoidal traveling wave that introduces time-dependence, with a line element given by

$$\mathrm{d}s^2 = -H(t,x)\,\mathrm{d}t^2 + H(t,x)\,\mathrm{d}x^2 + \mathrm{d}y^2 + \mathrm{d}z^2, \tag{51}$$

where

$$H(t,x) = 1 - A\sin\left(\frac{2\pi(x-t)}{d}\right), \tag{52}$$

where A and d are the amplitude and wavelength of the gauge wave, which travels along the x-axis.

We evolve analytic initial data of this solution, using the gauge source function H_a computed directly from the analytic initial data. We set the FOGH constraint damping parameters to $\gamma_0 = \gamma_2 = 1$ and $\gamma_1 = -1$. We evolve on the domain $[0,1]^3$ with two elements in the x-direction, one element in the y- and z-directions. We fix the y and z points per element to 6 (P_5) and perform a convergence test by running three resolutions with 15 (P_{14}), 18 (P_{17}), and 20 (P_{19}) points per element in the x direction. We apply periodic boundary conditions in all directions. We use a sixth-order Adams–Moulton predictor–corrector time integrator. In our simulations we choose $A = 0.1$ and $d = 1$.

Gauge wave simulations are known to be unstable in the BSSN formulation of the Einstein equations [68], but are stable in the Z4 system [69]. Gauge wave simulations are stable in the FOGH system, as we demonstrate with `SpECTRE` in figure 6. The left panel of figure 6 shows the L_2 norm of the 1-index and 3-index constraints as a function of time. While the lowest resolution (P_{14}) simulation has exponentially growing constraints, the higher resolution simulations have constant and convergent constraints. Similarly, the right panel of figure 6 shows the L_2 norm of the error in the evolved variables g_{ab}, Π_{ab}, and Φ_{iab} at the three resolutions. We observe stable and convergent long-term behavior. The highest resolution simulation is close to being limited by the time stepper tolerance.

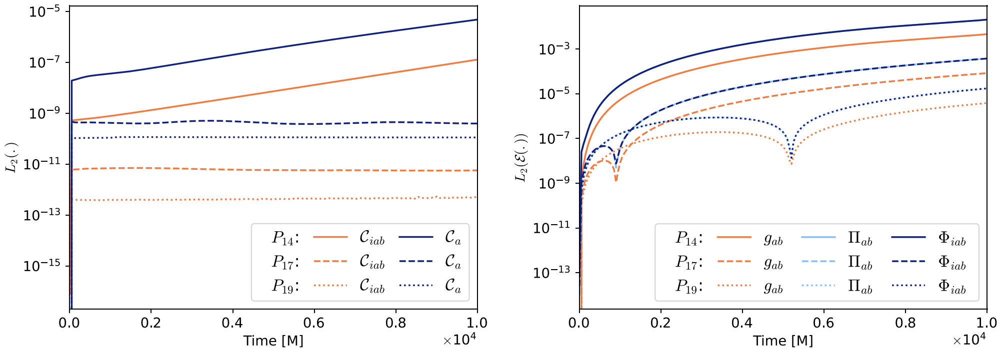

Figure 6. Results from an amplitude 0.1 gauge wave simulation demonstrating long term stable and convergent simulations. We show simulations at three resolutions, 15, 18, and 20 grid points per element in the x-direction. *Left*: The 1-index and 3-index constraints as a function of time. The lowest resolution has exponentially growing constraints in time while the two higher resolutions have constant and convergent constraints in time. *Right*: Difference (error) of the FOGH evolved variables g_{ab}, Π_{ab}, and Φ_{iab} from the analytic solution at the three resolutions. Since these are errors and not constraint violations, we expect the errors will accumulate linearly in time, even for constant constraint violations. We see stable and convergent long-term behavior with the errors decreasing exponentially with increasing resolution. The errors for Π_{ab} and Φ_{iab} lie on top of each other and so only the errors for Φ_{iab} are visible.

3.3. Binary black hole inspiral, merger, and ringdown

In this section, we use SpECTRE to generate binary black hole initial data and evolve the binary through ~ 18 orbits of inspiral, merger, and ringdown to a final, stationary state. We then use SpECTRE's CCE module to evolve the outgoing gravitational waves to future null infinity. We perform the simulations at three different resolutions, which we refer to as 'Lev0', 'Lev1', and 'Lev2' with Lev0 being the lowest resolution and Lev2 being the highest. Each increase in resolution increases the number of points per element per dimension by one. During the inspiral, each simulation uses 4800 elements, with Lev0, Lev1, and Lev2 having ~ 2.6; ~ 3.7; and ~ 5.0 million total grid points. During the ringdown, each simulation uses 7680 elements and ~ 10.2, ~ 13.3, and ~ 16.9 million total grid points, respectively. All simulations use the damped harmonic gauge to prevent the lapse collapsing and $\sqrt{\gamma}$ diverging at merger. All simulations (both inspiral and ringdown) also use a fourth-order Adams–Moulton predictor–corrector time integrator with absolute and relative time stepper tolerances of 10^{-10} and 10^{-8}, respectively.

The evolutions were each performed on 10 compute-nodes in the Resnick High Performance Computing Center at Caltech. Each compute node has two 28-core Intel Cascade Lake CPUs. Our Lev0, Lev1, and Lev2 evolutions cost 58 000; 71 000; and 117 000 core hours, which amounts to 104; 127; and 209 wallclock hours, and an average of 120; 80; and 41 M/hr during the inspiral. In a future paper, we will assess SpECTRE's performance and scaling in more detail; our purpose for this paper is to demonstrate that SpECTRE can evolve binary black holes through inspiral, merger, and ringdown.

3.3.1. Initial data. We begin our evolutions with initial data of two equal-mass and non-spinning black holes in a quasicircular orbit. To generate the initial data we use the SpECTRE initial data module [70–72], which solves the elliptic constraint sector of the Einstein equations

Table 1. Parameters used to generate the binary-black-hole initial data evolved in section 3.3: in terms of the sum of the black holes' Christodoulou masses $M = M_1 + M_2$ of each black hole, the individual masses M_1 and M_2, the dimensionless spin magnitudes $|\chi_1|/M^2$ and $|\chi_2|/M^2$, the orbital eccentricity e, the initial coordinate separation d_0 of the horizon centers, the initial orbital angular velocity Ω_0, and the initial expansion rate $\dot{a}_0$.

M_1/M	M_2/M	$\|\chi_1\|/M^2$	$\|\chi_2\|/M^2$	d_0/M	$M\Omega_0$	e	$\dot{a}_0$
0.5	0.5	3×10^{-8}	3×10^{-8}	15.366	0.0159	7.2×10^{-4}	3.4×10^{-5}

in the extended conformal thin sandwich (XCTS) formalism [73–75]. It uses the superposed Kerr–Schild formalism to construct a conformal background to the XCTS equations based on the weighted superposition of two isolated Kerr–Schild black holes [76, 77]. The black holes are represented as excisions with negative-expansion apparent horizon boundary conditions [78, 79]. The initial data solver uses DG methods similar to those described in this article to achieve scalable and parallelizable solutions to the elliptic equations and is also open source [45, 71].

We choose an initial coordinate separation of the excision centers to be $d_0 = 15.366\,M$ to facilitate future comparison with the family of simulations in [80] and to place the two black holes at ~ 18 orbits before they merge. The masses and spins of the black holes measured on the horizons are driven to the desired values in a control loop that adjusts the initial data parameters, similar to [81]. In a second control loop we performed eccentricity reduction by evolving the initial data for a few orbits and adjusting the initial orbital parameters to iteratively reduce the eccentricity of the orbit [82, 83] to 7.2×10^{-4}. The resulting initial data parameters are summarized in table 1, and a plot of the resulting inertial trajectories for Lev2 is shown in figure 7.

3.3.2. Computational domain decomposition. The numerical evolution of the FOGH system is performed with a DG scheme in which the physical domain of the problem is partitioned into deformed hexahedral elements with conforming boundaries. The boundaries and gridpoint distributions of an element are determined by a continuous and differentiable coordinate map applied to the logical Cartesian coordinates, which we label ξ, η and ζ, of a regular cube spanning $[-1,1]^3$. The maps corresponding to neighboring elements are required to be continuous but are not required to be differentiable at element boundaries. This provides the flexibility necessary to construct the complicated domains needed for binary merger simulations using DG methods. An example of the domain used during the inspiral is shown in figure 8 and an example of the ringdown domain is shown in figure 9. While the coordinate maps provide significant flexibility, we found that instabilities arise if neighboring elements differ by significantly more than a factor of two in size, placing a practical constraint on how quickly the resolution can be reduced as one moves away from the BHs.

Our computational domain is the region of space between the outer spherical boundary and the excision boundaries that remove the black hole singularities from the domain. The excision boundaries are spherical in the comoving coordinates with their sizes and shapes in the inertial coordinates informed by the size and shape of the apparent horizons (section 2.7). At merger the apparent horizons of the inspiraling black holes become enveloped by a single common apparent horizon. We handle the different number of excision boundaries during the inspiral and ringdown by having distinct domains for each. At merger we interpolate data

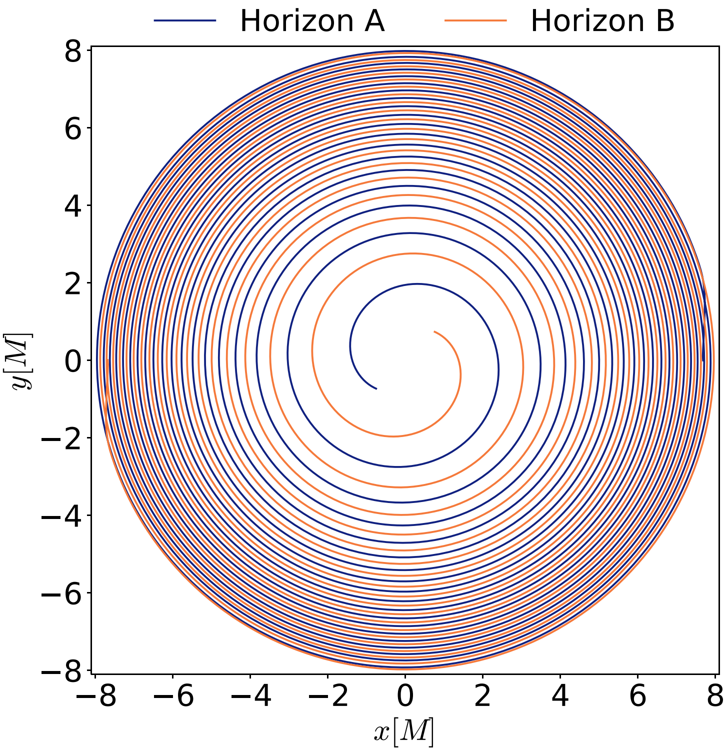

Figure 7. The trajectories of the centers of each apparent horizon during the Lev2 inspiral evolution, in the evolution's asymptotically inertial coordinates. The coordinates x and y indicate the position of the black hole within the orbital plane in units of the total initial mass M of the binary.

from the inspiral domain that has two excision boundaries to the ringdown domain that has one excision boundary. We describe each domain below.

The inspiral domain:

The inspiral domain is more complicated than the ringdown domain. This is because of the complexity of having two excisions. During the inspiral we must tile a two-excision domain with conforming hexahedra. Our solution to this tiling problem makes use of 44 element collections grouped into two radial and two 'biradial' layers. Our description of the domain decomposition starts at the excision surfaces and extends radially outward.

The first layer consists of six wedges composing the spherical shell surrounding each excision. Each wedge is subdivided into multiple elements. The black holes are located on the x-axis at $x \sim \pm 7.683\,M$ with an excision radius of $r_{\mathrm{exc}} \sim 0.792\,M$. The shell around each black hole has an outer radius of 6 M. For a fixed target accuracy, distributing the grid points logarithmically in radius and equiangularly [84, 85] in angles significantly reduces computational cost.

The second layer uses a set of wedges that wrap the shells around each black hole in a cubical shell, as seen in figure 8. A consequence of the decreasing separation between the two black holes during the inspiral is that the size of the excision within each cube grows as the simulation progresses. Since we only deform the region inside the cubes to conform to the

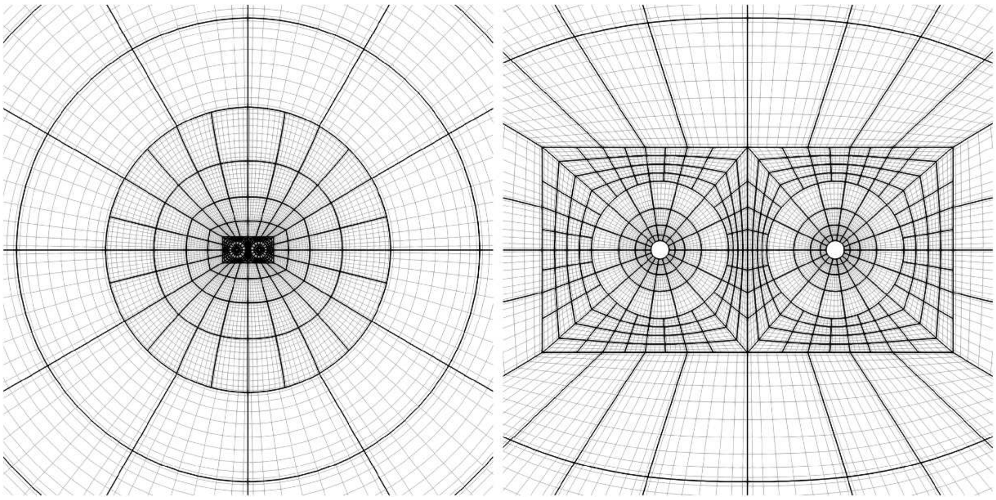

Figure 8. An illustration of the computational grid used during the inspiral. We make use of two excision regions, each region lying inside a black hole's apparent horizon. Each excision is surrounded by a spherical shell partitioned into six deformed cubes as in figure 2. Each spherical shell is then surrounded by another shell of six deformed cubes that transition to a cubical boundary. Then the two cubes themselves are surrounded by a transitionary envelope which becomes spherical. *Left*: The transitionary envelope. *Right*: A close-up of the domain structure around the excisions. The center of each excision is offset from the center of the cube.

shape of the apparent horizons, the simulation will fail if the apparent horizons grow beyond the cube boundaries. We remedy this by decoupling the excision center from the center of the cube (see the right panel of figure 8), effectively increasing the size of the cube relative to the size of the excision by a constant factor that is sufficient to keep the apparent horizons within the cubes throughout the simulation. This generalized map is crucial for robust inspiral and merger simulations.

The third layer consists of the elements surrounding the cubes around each black hole (layers 1 and 2). We refer to this region, which extends from $r \sim 18\,M$ to $r = 100\,M$, as the 'envelope'. This layer serves to transition the grid point distributions from what is used near the black holes to the distribution that is used in the wave zone. We use a logarithmic map in the radial direction and we interpolate between a 'biradial' equiangular map used for two excisions and a 'radial' equiangular map suited for the spherical outer boundary.

The fourth and final layer is a spherical shell extending from the end of the envelope ($r = 100\,M$) to the outer boundary at $r = 600\,M$. This shell uses a linear distribution in the radial direction and an equiangular distribution in the angular directions. Since the GW wavelength is constant in radius, a linear distribution is necessary to avoid under resolving the waves. In production quality simulations we expect to place the outer boundary at $r = 1500\,M$, since in typical `SpEC` simulations we find that this radius is necessary to avoid a center-of-mass drift caused by the gauge boundary condition. Errors in the gauge boundary condition fall off as $\sim 1/r^2$. We find that the drift is larger for longer simulations, but this can be compensated for with a larger outer boundary. Since we use CCE to extract the gravitational waves, a large domain for wave extraction is not required, as it would be if we instead extrapolated the waves to spatial infinity.

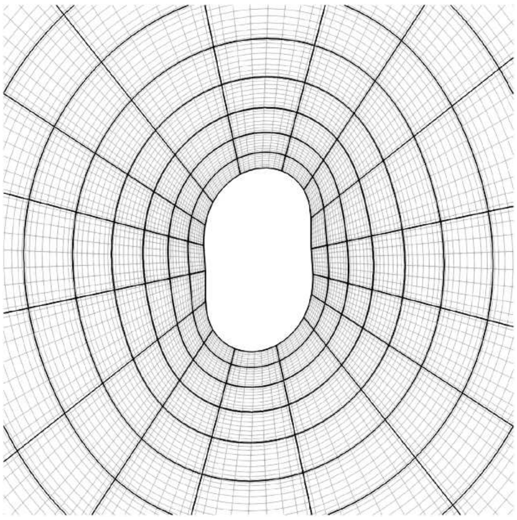

Figure 9. An image of the ringdown domain. After a common horizon is found, the evolution is regridded onto a domain with a single excision as in figure 2. Additional time-dependent maps are used to deform the excision's spherical shape into one that corresponds to the common horizon shape.

In addition to the hexahedral maps that partition the domain, we also globally apply rotation and expansion maps that track the angular, radial, and center of mass motion of the binary system.

To characterize the errors in these coordinate mappings, we introduce the diagnostic quantity (which we refer to as the 'jacobian diagnsotic') as

$$C_{\hat{i}} = 1 - \frac{\sum_i |\partial_{\hat{i}} x^i|}{\sum_i |D_{\hat{i}} x^i|}, \tag{53}$$

following the convention in section 2.2 where hatted quantities are in the logical coordinates and non-hatted quantities are in the inertial coordinates. In the numerator, $\partial_{\hat{i}} x^i$ is the analytic Jacobian provided by the analytic coordinate mappings (both time dependent and time-independent) described in this section. In the denominator, $D_{\hat{i}} x^i$ is the Jacobian computed by taking numerical derivatives of the inertial coordinates in each logical direction. The sums are over all gridpoints. If $C_{\hat{i}} = 0$, the analytic and numerical Jacobians are identical, meaning that our coordinate mappings are perfectly represented by our Legendre–Gauss–Lobatto DG scheme.

The ringdown domain:

The ringdown domain is a single excision domain used after a common apparent horizon has formed. An example of this domain is shown in figure 9. It is similar in structure to the domain used in the single black hole tests in section 3.1. We use a logarithmic radial map from the excision surface to $r = 50\,M$ and use a linear spacing further away to resolve the gravitational waves. We used an equidistant map instead of an equiangular map in the angular directions.

A significant challenge compared to the single black hole evolution is that the time-dependent maps used during ringdown must be initialized from, and matched to, the corresponding time-dependent maps in the inspiral. The rotation and expansion maps are matched and then decay exponentially to being time-independent. Most challenging are the shape and size maps. For the shape map we perform a least squares fit in time to the spherical harmonic coefficients of the common horizon (and their time derivatives) found during the inspiral. We fit to 100 times, and then initialize the shape of the excision by evaluating the fit at the transition time. For the size of the excision, we manually specify an excision radius $r_{\mathrm{exc}} = 1.45$ and gave the excision surface an initial outward radial velocity of 1.0; we adjusted these choices by hand until the ringdown was able to begin successfully, with the excision surface remaining inside the common apparent horizon while having all characteristic characteristic speeds pointing out of the computational domain (cf section 2.3).

For the simulations presented here, we carried out the transition from the inspiral domain to the ringdown domain manually. First, when the coordinate separation of the centers of the black holes' apparent horizons in the asymptotically inertial frame fell to less than $2.38\,M$ (a value chosen by hand to roughly correspond to the time that the common horizon first forms), we chose to have the simulation begin outputting the evolution variables at every gridpoint every $\Delta t = 0.01\,M$. Then, we chose to terminate the inspiral portion of the simulation when the coordinate separation between the two black holes decreased to less than $2.138\,M$. This value was chosen to be late enough that it yielded enough distinct common horizon finds, but also early enough that the constraint energy (cf section 3.3.4) did not grow too large. In the future, we intend to implement the techniques that `SpEC` uses to automate the transition process.

3.3.3. *Constraint damping.*

Based on our experience evolving binary black holes in `SpEC`, we use a superposition of three Gaussians and a constant for γ_0 and γ_2, and a single Gaussian plus a constant for γ_1. See (4) for how the constraint damping terms appear in the evolution equations. The motivation for the different Gaussians is to increase the constraint damping near each black hole, which requires the Gaussians to move with the black holes as they inspiral. In `SpEC` and `SpECTRE` we achieve this by making the Gaussians functions of the comoving 'grid frame' coordinates $x^{\bar{\imath}}$. As the black holes inspiral, their coordinate radius increases in the comoving $x^{\bar{\imath}}$ coordinates, which means the width w of the Gaussian must also increase by the same amount. Increasing the width is achieved by dividing the width by the expansion factor $E(t)$, which starts at 1 and decreases as the black holes inspiral. The specific form of the damping parameters we use is

$$\gamma_0\left(t,x^{\bar{\imath}}\right) = \gamma_2\left(t,x^{\bar{\imath}}\right) = C + \sum_{I=0}^{2} A_I \exp\left[-\left(\frac{\bar{r}_I}{\bar{w}_I/E(t)}\right)^2\right], \tag{54}$$

$$\gamma_1\left(t,x^{\bar{\imath}}\right) = C + A_0 \exp\left[-\left(\frac{\bar{r}_0}{\bar{w}_0}\right)^2\right], \tag{55}$$

where C is a constant, A_I are the amplitudes of the Gaussians, $\bar{r}_I$ are the grid-frame radii from the center of each Gaussian, and $\bar{w}_I$ are the widths of the Gaussians in the grid frame. Table 2 shows the parameters during the inspiral and table 3 shows them during the ringdown. In the grid frame the black holes are always located on the $\bar{x}$-axis so we only specify the grid frame $\bar{x}$-coordinate at which each Gaussian is centered, denoted by $\bar{x}^C$ in tables 2 and 3.

3.3.4. *Constraint violations.*

Figure 10 shows the constraint energy $\mathcal{E}$ (see (19)) as a function of time at each resolution. Experience from `SpEC` suggests that the initial, rapid growth in $\mathcal{E}$ is caused by not resolving the initial data's rapid relaxation and emission of spurious, high-frequency 'junk' gravitational radiation at the start of the simulation. Resolving the junk

Table 2. Parameters for the Gaussians that make up the constraint damping functions during the inspiral. See (54) and (55) for how the coefficients appear in the functional form of the constraint damping parameters.

	C	A_0	$\bar{w}_0$	$\bar{x}_0^C$	A_1	$\bar{w}_1$	$\bar{x}_1^C$	A_2	$\bar{w}_2$	$\bar{x}_2^C$
γ_0	$0.01/M$	$0.75/M$	38.415	0	$8/M$	3.5	7.683	$8/M$	3.5	-7.683
γ_1	$-0.999/M$	$0.999/M$	$10d_0$	—	—	—	—	—	—	—
γ_2	$0.01/M$	$0.75/M$	38.415	0	$8/M$	3.5	7.683	$8/\mathrm{M}$	3.5	-7.683

Table 3. Parameters for the Gaussians that make up the constraint damping functions during the ringdown. See (54) and (55) for how the coefficients appear in the functional form of the constraint damping parameters.

	C	A_0	$\bar{w}_0$	$\bar{x}_0^C$	A_1	$\bar{w}_1$	$\bar{x}_1^C$	A_2	$\bar{w}_2$	$\bar{x}_2^C$
γ_0	$0.01/M$	$1.0/M$	100	0	$7/M$	2.5	0	—	—	—
γ_1	$-0.999\,999/M$	—	—	—	—	—	—	—	—	—
γ_2	$0.001/M$	$0.1/M$	100	0	$7/M$	2.5	0	—	—	—

radiation is computationally expensive and generally not done in NR evolutions of binary black holes. After the initial growth of constraint violation damps away, the constraints converge exponentially with increasing p-refinement. The constraints grow sharply near the time of merger, as the black holes become more distorted by each others' tidal gravity, causing the solution to be less resolved by our fixed computational mesh. We anticipate that future `SpECTRE` simulations using adaptive mesh refinement will improve the behavior of the constraints near the time of merger.

To ensure the errors from our coordinate mappings do not affect our constraint violations, we compare the L2-norms of the constraint energy and the jacobian diagnostic in table 4 at $t = 3000\,M$ during the inspiral. For all our resolutions, the mapping error is below our constraint violations, meaning that our coordinate maps do not affect the results of these simulations.

After the merger, as the remnant black hole rings down, the constraint violations decrease to much smaller values again. When the black hole has settled to its final stationary state, the constraint violations continue to slowly decrease in time. Because the ringdown constraints in the highest resolution are not smaller than those with the medium spatial resolution, we suspect that the numerical error is dominated not by spatial resolution in the ringdown but by some other factor. One possibility is the time-stepping accuracy during the late inspiral and ringdown. We leave a careful study of this, including improvements to the domain decomposition and use of adaptive mesh refinement during the ringdown, to future work.

3.3.5. Apparent horizons. Another way of measuring the accuracy of BBH simulations is to track the masses and spins of the BHs using the apparent horizon surfaces. We find the apparent horizons surfaces using a 'fast flow' approach similar to the one outlined in [86]. There are a number of subtleties that arise when implementing this approach with an asynchronous task-based parallelism setting. These are due to the fact that the 'fast flow' approach is an iterative method and thus needs to store the metric on the full computational domain at a given simulation time until all iterations are complete. A more in-depth explanation of this will be available in [66].

We measure masses and spins on the individual apparent horizons every $0.5\,M$ during the inspiral and merger and on the common apparent horizon every $0.1\,M$ during ringdown. The

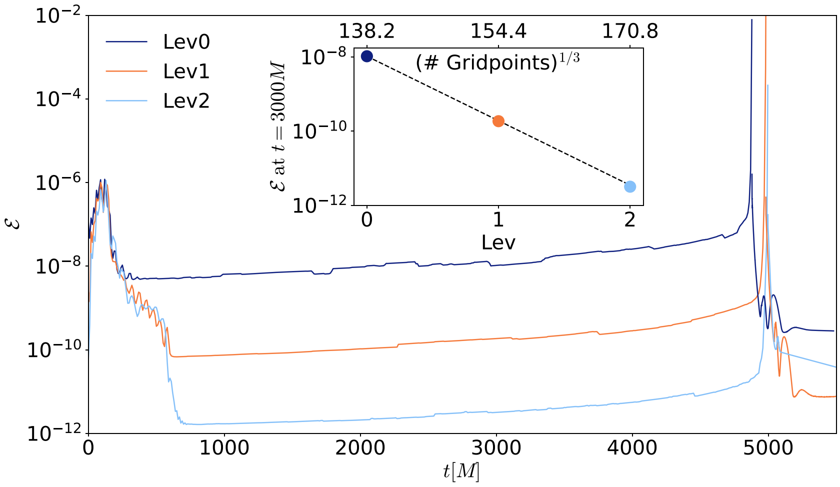

Figure 10. L2-norm of the constraint energy $\mathcal{E}$ (see (19)) over the entire computational domain for evolutions of an equal-mass, non-spinning, binary black hole at low ('Lev0'), medium ('Lev1'), and high ('Lev2') spatial resolution, as a function of time t in units of the initial total mass M. The peak of the constraint energy for each resolution is when we transition to the ringdown grid. We see exponential convergence with resolution except at merger and early ringdown, where we suspect adaptive mesh refinement will be necessary to resolve the additional dynamics. The inset shows this exponential convergence at time $t = 3000\,M$ during the inspiral as a function of the cube root of the total number of grid points (and 'Lev').

Table 4. The L2-norm of the constraint energy $\mathcal{E}$ and jacobian diagnostic $C_{\hat{i}}$ over all gridpoints at $t = 3000\,M$ during the inspiral. For all of our resolutions, the jacobian diagnostic (which is a measure of our mapping error) is below the error of the constraint energy.

L2-norm	Lev0	Lev1	Lev2
$\mathcal{E}$	1.0×10^{-8}	1.9×10^{-10}	3.2×10^{-12}
$C_{\hat{i}}$	8.8×10^{-11}	6.2×10^{-12}	5.7×10^{-13}

irreducible mass, defined as $M_{\mathrm{irr}} \equiv \sqrt{A/16\pi}$, where A is the surface area of the apparent horizon, should be monotonically increasing. Thus, any decreases in M_{irr} can be viewed as a measure of the numerical error in the simulation. Another useful metric is the Christodoulou mass $M_{\mathrm{Ch}} \equiv \sqrt{M_{\mathrm{irr}}^2 + S^2/4M_{\mathrm{irr}}^2}$, which includes both the irreducible mass and rotational kinetic energy. The dimensionless spin $\chi = S/M_{\mathrm{Ch}}^2$ measures the spin in terms of approximate rotational Killing vectors, as discussed in appendix A of [76]. For equal-mass non-spinning simulations M_{irr} and M_{Ch} should remain constant until merger, while χ should remain identically zero. Deviations from this behavior help quantify numerical errors in the simulation.

Figure 11 shows M_{irr}, M_{Ch}, and χ during the inspiral and ringdown. We find that the masses and spins over time are convergent, in the sense that the difference between Lev0 and Lev1 is greater than the difference between Lev1 and Lev2. During the inspiral, after the initial transient, the masses and spins remain more constant in time as resolution increases, until sharp gains near the time of merger as the black holes gain energy and angular momentum. During the ringdown, the masses and spins relax to final, constant values, as expected. The

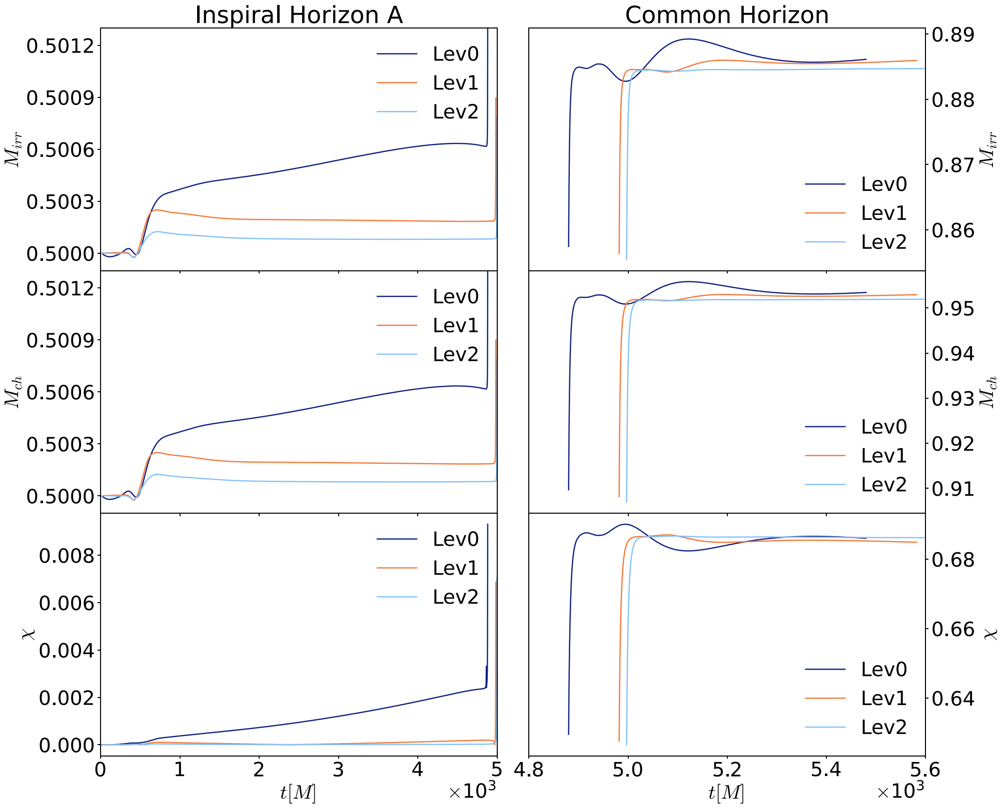

Figure 11. Apparent horizon masses and spins as a function of time. The left column shows the irreducible mass M_{irr}, Christodoulou mass M_{ch}, and dimensionless spin angular momentum $\chi \equiv S/M_{\mathrm{ch}}^2$ for one of the individual apparent horizons during inspiral, as a function of time t in units of the total initial Christodoulou mass M. The right column shows the same quantities during the ringdown for the common apparent horizon. We see convergence with increasing resolution, specifically that the Lev1 and Lev2 evolutions track each other more closely than the Lev0 and Lev1 simulations.

final Christodoulou mass $M_{\mathrm{ch}} = 0.952$ differs from the initial Christodoulou mass ($M_{\mathrm{ch}} = 1$) by 4.8%, and the final spin is 0.686. Both values are consistent with the fitting formulas in [80], tuned using SpEC evolutions of equal-mass, equal-aligned-spin binary black holes.

3.3.6. Gravitational waveforms. We compute gravitational waveforms using Cauchy-Characteristic Evolution (CCE) [46, 87–90], using the SpECTRE implementation of CCE [47]. This method utilizes an additional characteristic evolution code, the one described in [47], that solves the full Einstein equations on a set of outgoing null slices that extend from some inner worldtube all the way to future null infinity. Boundary conditions on the worldtube are provided by the interior Cauchy evolution, in this case also done with SpECTRE. For the characteristic evolution, there is freedom to choose one complex function on the initial null slice, which encodes the initial incoming radiation. We set it according to equation (16) of [47]. From the characteristic evolution, one can compute the gravitational-wave strain and all five

Weyl scalars at future null infinity. Gravitational waveforms computed via CCE are in a well-defined gauge modulo Bondi-van der Burg–Metzner–Sachs (BMS) transformations [91, 92], which are extensions of Poincaré transformations and correspond to symmetries of asymptotically flat spacetimes at future null infinity. The raw output from CCE is in an effectively random BMS frame, so to completely fix the gauge, it is necessary to perform a BMS transformation [93].We choose to transform waveforms into the superrest frame of the inspiral [94, 95], which can be thought of as the BMS extension of a frame in which the binary is at rest during the inspiral. We find the BMS transformation to map to this frame using data from the strain and Weyl scalars over the time window $[1800\,M, 2200\,M]$. See [93] and references therein for an in-depth review of BMS frame fixing.

Figure 12 shows the leading-order $\ell = m = 2$ spin-weighted spherical harmonic mode of the gravitational-wave strain as a function of retarded time t. To estimate the accuracy of the waveforms, we first apply a time shift t_{22}^{peak} so that the peak amplitudes at each spatial resolution occur at the same shifted time. Then, we apply a constant phase offset such that the gravitational-wave phase of the $\ell = m = 2$ mode vanishes at time $t - t_{22}^{\text{peak}} = -4000\,M$. This time was chosen as an early time after most of the initial spurious 'junk' radiation (especially visible in the amplitude at early times) has been emitted. This junk radiation is characteristically different than the junk radiation seen in figure 10. The junk radiation in figure 12 is of a lower frequency and is due to our choice of data on the initial null slice. It is an active area of research to improve data on the initial null slice.

We choose to post-process the data in this way because it enables us to transform each simulation to some reasonable BMS frame without using information from the other simulations. As a result, we can perform more meaningful convergence tests, since the output from each simulation is independent of every other. While one could obtain better agreement between different resolutions by finding the BMS transformation which minimizes the residual between the simulations' waveforms, this would go beyond computing a convergence error, because the frame of one simulation is determined by the other.

Finally, we interpolate the amplitudes and phases at each spatial resolution onto a common set of times and take differences, estimating the numerical error of a spatial resolution in terms of its difference with the next highest spatial resolution. We find that between time $t - t_{22}^{\text{peak}} = -4000\,M$ and merger time $t - t_{22}^{\text{peak}} = 0$, the amplitude and phase differences decrease with increasing resolution, as expected, with the medium and high resolution differing in amplitude by 0.01% at time of merger and about 0.1% throughout the inspiral. During the window between $t - t_{22}^{\text{peak}} = -4000\,M$ and $t = t_{22}^{\text{peak}}$, the medium and high spatial resolutions accumulate 0.03 radians of phase error. After merger time, the fractional amplitude and phase errors grow, which is expected, because the amplitude itself is exponentially falling to zero, making determining the phase accurately increasingly challenging.

4. Conclusion

We present the first inspiral-merger-ringdown simulations of two binary black holes using DG methods. We use the open-source numerical relativity code `SpECTRE` [45] to perform all simulations. These include several long-term stability and accuracy tests, e.g. evolutions of a single black hole in Kerr–Schild, harmonic, and damped harmonic gauge, as well as a long-term stable gauge wave simulation. All simulations demonstrate the expected exponential convergence.

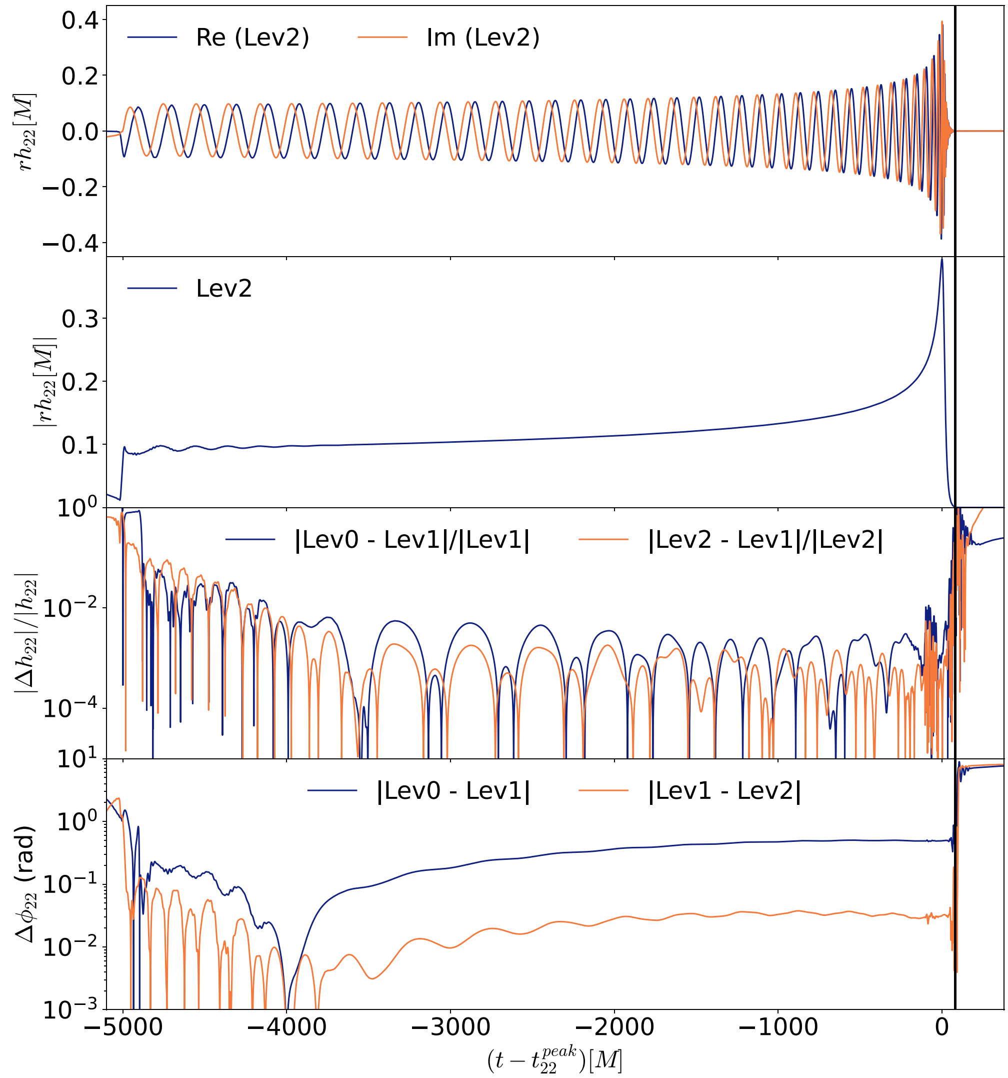

Figure 12. The $\ell = m = 2$ spin-weighted spherical-harmonic mode of the gravitational waves, computed using Cauchy Characteristic Evolution (CCE) and frame-fixed using the Python package `scri`, with a worldtube radius of 200 M, where M is the total initial Christodoulou mass. From top to bottom, as a function of retarded time $t - t_{22}^{\rm peak}$, the panels show (i) the real and imaginary part of the strain rh_{22}/M; ii) the amplitude $|rh_{22}/M|$ at the highest spatial resolution; iii) the fractional amplitude difference $|\Delta h_{22}|/|h_{22}|$, defined as the magnitude of the $\ell = m = 2$ amplitude difference between two spatial resolutions, divided by the magnitude of the higher spatial resolution's amplitude; iv) the difference in phase $\Delta\phi_{22}$, defined as the difference between the $\ell = m = 2$ gravitational-wave phase at two different spatial resolutions, with each phase offset such that the phase vanishes at $t - t_{22}^{\rm peak} = -4000\,M$. The vertical black lines shows where the amplitude goes below $\sim 10^{-3}$.

The binary black hole simulation is of the last 18 orbits before merger of an equal mass non-spinning binary. We extract gravitational waveforms at future null infinity using `SpECTRE`'s CCE module. We observe exponential convergence in the constraint violations with increasing resolution, and demonstrate convergence in amplitude and phase of the $\ell = m = 2$ mode of the gravitational wave strain. The simulations presented here are the first binary merger simulations where the initial data, evolution, and wave extraction are all performed using the open-source code `SpECTRE`.

Our medium and high resolution `SpECTRE` simulations run at 80 and 41 M/hr on ten 56-core Intel Cascade Lake nodes during inspiral. A proper comparison between this performance from `SpECTRE` and an analogous calculation in `SpEC` would be nontrivial, in part because a meaningful comparison should ensure that the `SpEC` and `SpECTRE` simulations have the same accuracy. We plan to perform such a comparison in the future. Here, for an initial ballpark comparison, we simply note that a `SpEC` evolution of analogous initial data, at what we typically would consider high resolution in `SpEC` (approximately 300,000 grid points, with `SpEC`'s adaptive mesh refinement algorithm varying the precise number of points throughout the simulation), ran at approximately 90 M/hr on 32 Intel Sky Lake cores. The `SpECTRE` high-resolution calculation ran about half as fast in wall time, using about 17 times more cores on a grid with about 17 times more gridpoints. `SpECTRE`'s domain has many more elements with fewer points per element, enabling it to scale to more CPU cores than `SpEC`; however, `SpEC`'s choice to use fewer elements with higher number of points (especially spherical shells near the horizons and in the wave zone) mean that `SpEC`'s domain is currently much more efficient.

We expect that `SpECTRE` would outperform `SpEC` in wall time at sufficiently high resolutions, running on enough CPU cores, but that it would be less efficient, requiring more cost for the same accuracy. We also expect that planned future optimizations (discussed in the next paragraph) will greatly reduce the computational cost of a `SpECTRE` binary-black-hole calculation at a given accuracy. Finally, note that in its current state, `SpECTRE` is still efficient enough to perform an ~ 18 orbit inspiral—longer than almost all inspirals published to date using moving-puncture codes—at feasible cost (less than 9 days of wall time on less than 600 compute cores).

While the results here present a milestone for `SpECTRE` simulations, several further advancements are necessary to enable building catalogs for future gravitational wave detectors. These fall in one of four categories: (i) performance improvements like using adaptive mesh refinement, dynamic load balancing, and GPU support; (ii) robustness and parameter space improvements like ensuring that high-spin, high-mass-ratio, and eccentric simulations can be performed robustly without hand tuning; (iii) automation infrastructure that allows a single user to run hundreds of simulations, such as automatically restarting failed simulations, automatically transitioning from inspiral to ringdown, and automatically running CCE after the Cauchy simulation completes; and (iv) improving documentation and tutorials to make the code more accessible to the broader community. Item (v) also includes developing a good understanding of how error diagnostics, like constraint violations, impact gravitational-wave amplitude and phase errors; our experience with `SpEC` suggests that the answer is complicated and will require careful investigation. These will allow `SpECTRE` to outperform our current code `SpEC` and to be more useful to the broader NR community.

Data availability statement

The data cannot be made publicly available upon publication because they are not available in a format that is sufficiently accessible or reusable by other researchers. The data that support the findings of this study are available upon reasonable request from the authors.

Acknowledgment

Charm++/Converse [96] was developed by the Parallel Programming Laboratory in the Department of Computer Science at the University of Illinois at Urbana-Champaign. The figures in this article were produced with `matplotlib` [97, 98], `TikZ` [99], `ParaView` [100, 101], `numpy` [102], `scipy` [103], and `scri` [104–107]. Computations were performed at the Resnick High-Performance Computing Center; a facility supported by the Resnick Sustainability Institute; at Caltech, the mbot cluster at Cornell, the ocean cluster at Cal State Fullerton, and the Urania cluster at the Max Planck Computing and Data Facility. We are pleased to thank Josh Smith for helpful discussions. This work was supported in part by NSF awards PHY-2208014 and AST-2219109, the Dan Black Family Trust, and Nicholas and Lee Begovich at Cal State Fullerton. This material is based upon work supported by the National Science Foundation under Grants Nos. PHY-2407742, PHY-2207342 and OAC-2209655 at Cornell. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. This work was supported by the Sherman Fairchild Foundation at Cornell. Support for this work was provided by NASA through the NASA Hubble Fellowship grant number HST-HF2-51562.001-A awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Incorporated, under NASA contract NAS5-26555.This work was supported in part by the Sherman Fairchild Foundation and by NSF Grants Nos. PHY-2309211, PHY-2309231 and OAC-2209656 at Caltech. P K acknowledges support of the Department of Atomic Energy, Government of India, under Project No. RTI4001, and by the Ashok and Gita Vaish Early Career Faculty Fellowship at the International Centre for Theoretical Sciences.

ORCID iDs

Geoffrey Lovelace https://orcid.org/0000-0002-7084-1070
Kyle C Nelli https://orcid.org/0000-0003-2426-8768
Nils Deppe https://orcid.org/0000-0003-4557-4115
Nils L Vu https://orcid.org/0000-0002-5767-3949
William Throwe https://orcid.org/0000-0001-5059-4378
Marceline S Bonilla https://orcid.org/0000-0003-4502-528X
Alexander Carpenter https://orcid.org/0000-0002-9183-8006
Lawrence E Kidder https://orcid.org/0000-0001-5392-7342
Alexandra Macedo https://orcid.org/0009-0001-7671-6377
Mark A Scheel https://orcid.org/0000-0001-6656-9134
Azer Afram https://orcid.org/0009-0003-2340-4059
Michael Boyle https://orcid.org/0000-0002-5075-5116
Andrea Ceja https://orcid.org/0000-0002-1681-7299
Matthew Giesler https://orcid.org/0000-0003-2300-893X
Sarah Habib https://orcid.org/0000-0002-4725-4978
Ken Z Jones https://orcid.org/0009-0003-1034-0498
Prayush Kumar https://orcid.org/0000-0001-5523-4603
Guillermo Lara https://orcid.org/0000-0001-9461-6292
Denyz Melchor https://orcid.org/0000-0002-7854-1953
Iago B Mendes https://orcid.org/0009-0007-9845-8448
Keefe Mitman https://orcid.org/0000-0003-0276-3856

Marlo Morales https://orcid.org/0000-0002-0593-4318
Jordan Moxon https://orcid.org/0000-0001-9891-8677
Kyle Pannone https://orcid.org/0009-0005-8607-2113
Harald P Pfeiffer https://orcid.org/0000-0001-9288-519X
Teresita Ramirez-Aguilar https://orcid.org/0000-0003-0994-115X
Jennifer Sanchez https://orcid.org/0000-0002-5335-4924
Daniel Tellez https://orcid.org/0009-0008-7784-2528
Saul A Teukolsky https://orcid.org/0000-0001-9765-4526
Nikolas A Wittek https://orcid.org/0000-0001-8575-5450

References

[1] Abbott R *et al* 2023 GWTC-3: compact binary coalescences observed by LIGO and Virgo during the second part of the third observing run *Phys. Rev.* X **13** 041039

[2] Blanchet L 2024 Post-Newtonian theory for gravitational waves *Living Rev. Relativ.* **27** 4

[3] Pretorius F 2005 Evolution of binary black hole spacetimes *Phys. Rev. Lett.* **95** 121101

[4] Campanelli M, Lousto C O, Marronetti P and Zlochower Y 2006 Accurate evolutions of orbiting black-hole binaries without excision *Phys. Rev. Lett.* **96** 111101

[5] Baker J G, Centrella J, Choi D-I, Koppitz M and van Meter J 2006 Gravitational wave extraction from an inspiraling configuration of merging black holes *Phys. Rev. Lett.* **96** 111102

[6] Pfeiffer H P 2012 Numerical simulations of compact object binaries *Class. Quantum Grav.* **29** 124004

[7] Duez M D and Zlochower Y 2019 Numerical relativity of compact binaries in the 21st century *Rep. Prog. Phys.* **82** 016902

[8] Mroue A H *et al* 2013 Catalog of 174 binary black hole simulations for gravitational wave astronomy *Phys. Rev. Lett.* **111** 241104

[9] Jani K, Healy J, Clark J A, London L, Laguna P and Shoemaker D 2016 Georgia tech catalog of gravitational waveforms *Class. Quantum Grav.* **33** 204001

[10] Healy J, Lousto C O, Zlochower Y and Campanelli M 2017 The RIT binary black hole simulations catalog *Class. Quantum Grav.* **34** 224001

[11] Healy J, Lousto C O, Lange J, O'Shaughnessy R, Zlochower Y and Campanelli M 2019 Second RIT binary black hole simulations catalog and its application to gravitational waves parameter estimation *Phys. Rev.* D **100** 024021

[12] Boyle M *et al* 2019 The SXS Collaboration catalog of binary black hole simulations *Class. Quantum Grav.* **36** 195006

[13] Healy J and Lousto C O 2020 Third RIT binary black hole simulations catalog *Phys. Rev.* D **102** 104018

[14] Healy J and Lousto C O 2022 Fourth RIT binary black hole simulations catalog: Extension to eccentric orbits *Phys. Rev.* D **105** 124010

[15] Ferguson D *et al* (Second MAYA) 2023 Catalog of binary black hole numerical relativity waveforms (arXiv:2309.00262)

[16] Punturo M *et al* 2010 The Einstein Telescope: a third-generation gravitational wave observatory *Class. Quantum Grav.* **27** 194002

[17] Evans M *et al* 2021 A horizon study for cosmic explorer: science, observatories, and community (arXiv:2109.09882)

[18] Amaro-Seoane P *et al* 2017 Laser interferometer space antenna (arXiv:1702.00786)

[19] Pürrer M and Haster C-J 2020 Gravitational waveform accuracy requirements for future ground-based detectors *Phys. Rev. Res.* **2** 023151

[20] Ferguson D, Jani K, Laguna P and Shoemaker D 2021 Assessing the readiness of numerical relativity for LISA and 3G detectors *Phys. Rev.* D **104** 044037

[21] Jan A, Ferguson D, Lange J, Shoemaker D and Zimmerman A 2024 Accuracy limitations of existing numerical relativity waveforms on the data analysis of current and future ground-based detectors *Phys. Rev.* D **110** 024023

[22] Zhu H, Fields J, Zappa F, Radice D, Stone J, Rashti A, Cook W, Bernuzzi S, and Daszuta B 2024 Performance-portable numerical relativity with AthenaK (arXiv:2409.10383)

[23] (available at: www.black-holes.org/SpEC.html)

[24] Hesthaven J S, Gottlieb S and Gottlieb D 2007 *Spectral Methods for Time-Dependent Problems* (Cambridge University Press)

[25] Rashti A, Maria Fabbri F, Brügmann B, Vivekanandji Chaurasia S, Dietrich T, Ujevic M and Tichy W 2022 New pseudospectral code for the construction of initial data *Phys. Rev.* D **105** 104027

[26] Jens Papenfort L, Tootle S D, Grandclément P, Most E R and Rezzolla L 2021 New public code for initial data of unequal-mass, spinning compact-object binaries *Phys. Rev.* D **104** 024057

[27] Rüter H R, Hilditch D, Bugner M and Brügmann B 2018 Hyperbolic relaxation method for elliptic equations *Phys. Rev.* D **98** 084044

[28] Tichy W, Ji L, Adhikari A, Rashti A and Pirog M 2023 The new discontinuous Galerkin methods based numerical relativity program Nmesh *Class. Quantum Grav.* **40** 025004

[29] Hilditch D, Weyhausen A and Brügmann B 2016 A pseudospectral method for gravitational wave collapse *Phys. Rev.* D **93** 063006

[30] Bugner M, Dietrich T, Bernuzzi S, Weyhausen A and Brügmann B 2016 Solving 3D relativistic hydrodynamical problems with weighted essentially nonoscillatory discontinuous Galerkin methods *Phys. Rev.* D **94** 084004

[31] Hilditch D, Weyhausen A and Brügmann B 2017 Evolutions of centered Brill waves with a pseudo-spectral method *Phys. Rev.* D **96** 104051

[32] Suárez Fernández I, Vicente R and Hilditch D 2021 Semilinear wave model for critical collapse *Phys. Rev.* D **103** 044016

[33] Bhattacharyya M K, Hilditch D, Rajesh Nayak K, Renkhoff S, Rüter H R and Brügmann B 2021 Implementation of the dual foliation generalized harmonic gauge formulation with application to spherical black hole excision *Phys. Rev.* D **103** 064072

[34] Suárez Fernández I, Renkhoff S, Cors D, Bruegmann B and Hilditch D 2022 Evolution of Brill waves with an adaptive pseudospectral method *Phys. Rev.* D **106** 024036

[35] Renkhoff S, Cors D, Hilditch D and Brügmann B 2023 Adaptive hp refinement for spectral elements in numerical relativity *Phys. Rev.* D **107** 104043

[36] Baumgarte T W *et al* 2023 Critical phenomena in the collapse of gravitational waves *Phys. Rev. Lett.* **131** 181401

[37] Cors D, Renkhoff S, Rüter H R, Hilditch D and Brügmann B 2023 Formulation improvements for critical collapse simulations *Phys. Rev.* D **108** 124021

[38] Marouda K, Cors D, Rüter H R, Atteneder F and Hilditch D 2024 Twist-free axisymmetric critical collapse of a complex scalar field *Phys. Rev.* D **109** 124042

[39] Atteneder F, Rüter H R, Cors D, Rosca-Mead R, Hilditch D and Brügmann B 2024 Boson star head-on collisions with constraint-violating and constraint-satisfying initial data *Phys. Rev.* D **109** 044058

[40] Dumbser M, Zanotti O and Peshkov I 2024 High order discontinuous Galerkin schemes with subcell finite volume limiter and AMR for a monolithic first–order BSSNOK formulation of the Einstein–Euler equations (arXiv:2406.15798)

[41] Alcubierre M, Benger W, Bruegmann B, Lanfermann G, Nerger L, Seidel E and Takahashi R 2001 The 3-D grazing collision of two black holes *Phys. Rev. Lett.* **87** 271103

[42] Reed W H and Hill T R 1973 Triangular mesh methods for the neutron transport equation *Technical Report* LA-UR-73-479 Los Alamos Scientific Lab., N. Mex.(USA)

[43] Hesthaven J S and Warburton T 2008 *Nodal Discontinuous Galerkin Methods: Algorithms, Analysis and Applications* (Springer)

[44] Szilágyi B, Blackman J, Buonanno A, Taracchini A, Pfeiffer H P, Scheel M A, Chu T, Kidder L E and Pan Y 2015 Approaching the post-Newtonian regime with numerical relativity: a compact-object binary simulation spanning 350 gravitational-wave cycles *Phys. Rev. Lett.* **115** 031102

[45] Deppe N *et al* 2024 SpECTRE v2024.09.29 (available at: https://doi.org/10.5281/zenodo.13858965)

[46] Moxon J, Scheel M A and Teukolsky S A 2020 Improved Cauchy-characteristic evolution system for high-precision numerical relativity waveforms *Phys. Rev.* D **102** 044052

[47] Moxon J, Scheel M A, Teukolsky S A, Deppe N, Fischer N, Hébert F, Kidder L E and Throwe W 2023 SpECTRE Cauchy-characteristic evolution system for rapid, precise waveform extraction *Phys. Rev.* D **107** 064013

[48] Baumgarte T W and Shapiro S L 2010 *Numerical Relativity: Solving Einstein's Equations on the Computer* (Cambridge University Press)

[49] Rezzolla L and Zanotti O 2013 *Relativistic Hydrodynamics* (Oxford University Press)

[50] Lindblom L, Scheel M A, Kidder L E, Owen R and Rinne O 2006 A New generalized harmonic evolution system *Class. Quantum Grav.* **23** S447–62

[51] Lindblom L, Matthews K D, Rinne O and Scheel M A 2008 Gauge drivers for the generalized harmonic Einstein equations *Phys. Rev.* D **77** 084001

[52] Szilagyi B, Lindblom L and Scheel M A 2009 Simulations of binary black hole mergers using spectral methods *Phys. Rev.* D **80** 124010

[53] Teukolsky S A 2016 Formulation of discontinuous Galerkin methods for relativistic astrophysics *J. Comput. Phys.* **312** 333–56

[54] Deppe N, Hébert F, Kidder L E and Teukolsky S A 2022 A high-order shock capturing discontinuous Galerkin–finite difference hybrid method for GRMHD *Class. Quantum Grav.* **39** 195001

[55] Scheel M A, Pfeiffer H P, Lindblom L, Kidder L E, Rinne O and Teukolsky S A 2006 Solving Einstein's equations with dual coordinate frames *Phys. Rev.* **D74** 104006

[56] Hesthaven J S 1997 A stable penalty method for the compressible Navier–Stokes equations: II. One-dimensional domain decomposition schemes *SIAM J. Sci. Comput.* **18** 658–85

[57] Hesthaven J S 1999 A stable penalty method for the compressible Navier–Stokes equations: III. Multidimensional domain decomposition schemes *SIAM J. Sci. Comput.* **20** 62–93

[58] Hesthaven J S 2000 Spectral penalty methods *Appl. Num. Math.* **33** 23–41

[59] Rinne O, Lindblom L and Scheel M A 2007 Testing outer boundary treatments for the Einstein equations *Class. Quantum Grav.* **24** 4053–78

[60] Ma S *et al* 2024 Fully relativistic three-dimensional Cauchy-characteristic matching for physical degrees of freedom *Phys. Rev.* D **109** 124027

[61] Buchman L T, Duez M D, Morales M, Scheel M A, Kostersitz T M, and Evans A M 2024 Numerical relativity multimodal waveforms using absorbing boundary conditions (arXiv:2402.12544)

[62] Throwe W and Teukolsky S Local time-stepping with predictor-corrector methods (in preparation)

[63] Press W H, Teukolsky S A, Vetterling W T and Flannery B P 2007 *Numerical Recipes* (*The Art of Scientific Computing*) 3rd edn (Cambridge University Press)

[64] Deppe N, Kidder L E, Scheel M A and Teukolsky S A 2019 Critical behavior in 3D gravitational collapse of massless scalar fields *Phys. Rev.* **D99** 024018

[65] Hemberger D A, Scheel M A, Kidder L E, Szilagyi B, Lovelace G, Taylor N W and Teukolsky S A 2013 Dynamical excision boundaries in spectral evolutions of binary black hole spacetimes *Class. Quantum Grav.* **30** 115001

[66] Nelli K C *et al* 2025 (in preparation)

[67] Alcubierre M *et al* 2004 Toward standard testbeds for numerical relativity *Class. Quantum Grav.* **21** 589–613

[68] Babiuc M C *et al* 2008 Implementation of standard testbeds for numerical relativity *Class. Quantum Grav.* **25** 125012

[69] Alic D, Bona-Casas C, Bona C, Rezzolla L and Palenzuela C 2012 Conformal and covariant formulation of the Z4 system with constraint-violation damping *Phys. Rev.* D **85** 064040

[70] Fischer N L and Pfeiffer H P 2022 Unified discontinuous Galerkin scheme for a large class of elliptic equations *Phys. Rev.* D **105** 024034

[71] Vu N L *et al* 2022 A scalable elliptic solver with task-based parallelism for the SpECTRE numerical relativity code *Phys. Rev.* D **105** 084027

[72] Vu N L 2024 Discontinuous Galerkin scheme for elliptic equations on extremely stretched grids *Phys. Rev.* D **110** 084062

[73] York J W 1999 Conformal 'thin sandwich' data for the initial-value problem of general relativity *Phys. Rev. Lett.* **82** 1350–3

[74] Pfeiffer H P and York J W 2003 Extrinsic curvature and the Einstein constraints *Phys. Rev.* D **67** 044022

[75] Pfeiffer H P 2005 The initial value problem in numerical relativity *J. Hyperbolic Differ. Equ.* **2** 497–520

[76] Lovelace G, Owen R, Pfeiffer H P and Chu T 2008 Binary-black-hole initial data with nearly-extremal spins *Phys. Rev.* D **78** 084017

[77] Lovelace G, Scheel M A and Szilagyi B 2011 Simulating merging binary black holes with nearly extremal spins *Phys. Rev.* D **83** 024010

[78] Cook G B and Pfeiffer H P 2004 Excision boundary conditions for black hole initial data *Phys. Rev.* D **70** 104016

[79] Varma V, Scheel M A and Pfeiffer H P 2018 Comparison of binary black hole initial data sets *Phys. Rev.* D **98** 104011

[80] Hemberger D A *et al* 2013 Final spin and radiated energy in numerical simulations of binary black holes with equal masses and equal, aligned or anti-aligned spins *Phys. Rev.* D **88** 064014

[81] Ossokine S, Foucart F, Pfeiffer H P, Boyle M and Szilágyi B 2015 Improvements to the construction of binary black hole initial data *Class. Quantum Grav.* **32** 245010

[82] Buonanno A, Kidder L E, Mroué A H, Pfeiffer H P and Taracchini A 2011 Reducing orbital eccentricity of precessing black-hole binaries *Phys. Rev.* D **83** 104034

[83] Habib S, Scheel M and Teukolsky S 2024 Eccentricity reduction for quasicircular binary evolutions (arXiv:2410.05531)

[84] Sadourny R 1972 Conservative finite-difference approximations of the primitive equations on quasi-uniform spherical grids *Mon. Wea. Rev.* **100** 136–44

[85] Ronchi C, Iacono R and Paolucci P S 1996 The “cubed sphere”: a new method for the solution of partial differential equations in spherical geometry *J. Comput. Phys.* **124** 93–114

[86] Gundlach C 1998 Pseudospectral apparent horizon finders: an efficient new algorithm *Phys. Rev.* D **57** 863–75

[87] Bishop N T, Gómez R, Lehner L and Winicour J 1996 Cauchy-characteristic extraction in numerical relativity *Phys. Rev.* D **54** 6153–65

[88] Bishop N T and Rezzolla L 2016 Extraction of gravitational waves in numerical relativity *Living Rev. Relativ.* **19** 2

[89] Handmer C J, Szilágyi B and Winicour J 2016 Spectral Cauchy characteristic extraction of strain, news and gravitational radiation flux *Class. Quantum Grav.* **33** 225007

[90] Barkett K, Moxon J, Scheel M A and Szilágyi B 2020 Spectral Cauchy-characteristic extraction of the gravitational wave news function *Phys. Rev.* D **102** 024004

[91] Bondi H, Van der Burg M G J and Metzner A W K 1962 Gravitational waves in general relativity, VII. Waves from axi-symmetric isolated system *Proc. R. Soc.* A **269** 21–52

[92] Sachs R 1962 Gravitational waves in general relativity. VIII. Waves in asymptotically flat spacetime *Proc. R. Soc.* A **270** 103–26

[93] Mitman K *et al* 2024 A review of gravitational memory and BMS frame fixing in numerical relativity *Class. Quantum Grav.* **41** 223001

[94] Mitman K *et al* 2021 Fixing the BMS frame of numerical relativity waveforms *Phys. Rev.* D **104** 024051

[95] Mitman K *et al* 2022 Fixing the BMS frame of numerical relativity waveforms with BMS charges *Phys. Rev.* D **106** 084029

[96] Kale L *et al* 2020 Uiuc-ppl/charm: charm++ version 6.10.2 (available at: https://doi.org/10.5281/zenodo.3972617)

[97] Hunter J D 2007 Matplotlib: A 2d graphics environment *Comput. Sci. Eng.* **9** 90–95

[98] Caswell T A *et al* 2020 matplotlib/matplotlib: Rel: v3.3.0. (available at: https://doi.org/10.5281/zenodo.3948793)

[99] Tantau T The tikz and pgf packages (available at: https://github.com/pgf-tikz/pgf)

[100] Ahrens J, Geveci B and Law C 2005 Paraview: an end-user tool for large-data visualization *Visualization Handbook* (Butterworth-Heinemann)

[101] Ayachit U 2015 *The Paraview Guide: A Parallel Visualization Application* (Kitware, Inc.)

[102] Harris C R *et al* 2020 Array programming with NumPy *Nature* **585** 357–62

[103] Virtanen P *et al* SciPy 10 2020 Fundamental algorithms for scientific computing in Python *Nat. Methods* **17** 261–72

[104] Boyle M 2013 Angular velocity of gravitational radiation from precessing binaries and the corotating frame *Phys. Rev.* D **87** 104006

[105] Boyle M, Kidder L E, Ossokine S, and Pfeiffer H P 2014 Gravitational-wave modes from precessing black-hole binaries (arXiv:1409.4431)

[106] Boyle M 2016 Transformations of asymptotic gravitational-wave data *Phys. Rev.* D **93** 084031

[107] Boyle M, Iozzo D and Stein L C 2020 moble/scri: v1.2. (available at: https://doi.org/10.5281/zenodo.4041972)