



CHORUS

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

General-relativistic neutron star evolutions with the discontinuous Galerkin method

François Hébert, Lawrence E. Kidder, and Saul A. Teukolsky

Phys. Rev. D **98**, 044041 — Published 27 August 2018

DOI: [10.1103/PhysRevD.98.044041](https://doi.org/10.1103/PhysRevD.98.044041)

General-relativistic neutron star evolutions with the discontinuous Galerkin method

François Hébert,^{1,2,*} Lawrence E. Kidder,¹ and Saul A. Teukolsky^{1,2}

¹*Center for Astrophysics and Planetary Science, Cornell University, Ithaca, New York 14853, USA*

²*Theoretical Astrophysics 350-17, California Institute of Technology, Pasadena, CA 91125, USA*

Simulations of relativistic hydrodynamics often need both high accuracy and robust shock-handling properties. The discontinuous Galerkin (DG) method combines these features — a high order of convergence in regions where the solution is smooth, and shock-capturing properties for regions where it is not — with geometric flexibility, and is therefore well-suited to solve the PDEs describing astrophysical scenarios. We present here evolutions of a general-relativistic neutron star with the DG method. In these simulations, we simultaneously evolve the spacetime geometry and the matter on the same computational grid, which we conform to the spherical geometry of the problem. To verify the correctness of our implementation, we perform standard convergence and shock tests. We then show results for evolving, in 3D, a Kerr black hole, a neutron star in the Cowling approximation (holding the spacetime metric fixed), and, finally, a neutron star where the spacetime and matter are both dynamical. The evolutions show long-term stability, good accuracy, and improved rate of convergence versus a comparable-resolution finite volume method.

I. INTRODUCTION

Numerical simulations are a crucial tool in the study of core-collapse supernovae, compact binary mergers, accretion disks with relativistic jets, and other energetic astrophysical sources. In these events, the dynamics are governed by the high-density matter and its coupling to the strong gravitational field. Nuclear reactions, neutrino physics, and magnetic fields can also play significant roles. Because of the highly nonlinear nature of the underlying general-relativistic hydrodynamics (GR-hydro), simulations are necessary to obtain observable predictions from physics models. Achieving sufficient accuracy in the simulation outputs (e.g. gravitational waveforms, ejected masses, nucleosynthesis products) remains a challenge, however. High resolution is needed to resolve multi-scale fluid flows, and the presence of shocks in the matter reduces the accuracy of the numerical schemes.

The standard approach taken in present-day GR-hydro codes is to cast the partial differential equations into conservative form, and discretize them using a finite-volume (FV) method (see reviews [1–3] for an overview and history). FV methods are favored for their robustness and the various “shock-capturing” schemes that enable them to handle fluid shocks and stellar surfaces. The Einstein equations for the spacetime geometry are typically solved with a finite-difference method on the same grid, but can instead be solved with a spectral method on a different computational grid [4]. Over the past decade, the application of improved high-resolution shock-capturing schemes (e.g. PPM [5, 6], WENO [7]) and higher-order difference schemes has led to significant advances in the accuracy and stability of the numerical results (e.g. for core-collapse supernovae [8, 9], binary mergers [10–12], and accretion flows [13–15]). In spite of these successes, FV methods have inherent limitations when used as high-

order methods: the large stencils required for the differencing and shock-capturing schemes make it difficult to adapt the grid to the problem geometry, and can also lead to challenges in efficiently parallelizing the algorithm.

In the pursuit of improved accuracy and efficiency, discontinuous Galerkin (DG) methods have recently emerged as a promising contender for astrophysical problems. DG methods share properties with both spectral methods and FV methods — they inherit the high-order accuracy of the former for smooth solutions while maintaining the robust shock-handling properties of the latter. They are geometrically flexible, enabling the use of grids adapted to the problem geometry. They are well suited to *hp*-adaptivity, where the grid resolution can be set either by adjusting the order of the polynomial approximation within an element (*p*-refinement), or by adjusting the size of the element (*h*-refinement). Finally, DG methods are locally formulated, enabling efficient parallelization and good scaling.

The application of DG methods to problems in relativistic astrophysics is recent and remains exploratory in nature. With several of these explorations focussing on the evolution of the spacetime geometry, different formulations of Einstein’s equations have been investigated. In an early study, Zumbush [16] obtained a space-time DG scheme for the linearized vacuum Einstein equations in harmonic gauge. For the commonly used Baumgarte-Shapiro-Shibata-Nakamura (BSSN) formulation of the Einstein equations, Field *et al.* [17] and Brown *et al.* [18] developed DG schemes in spherical symmetry. More recently, Miller and Schnetter [19] developed an operator-based (vs. the typical differential equation-based) DG discretization of the BSSN equations, and showed success in evolving 3D test problems. Using a new first-order form (FO-CCZ4) of the constraint-damping Z4 formulation, Dumbser *et al.* [20] evolved a single BH spacetime using a puncture and showed a short-timescale “proof of concept” evolution of a binary BH system.

Efforts on the hydrodynamics side began with Radice and Rezzolla [21], who presented a formulation of DG for

* fhebert@caltech.edu

the evolution of fluids in curved spacetimes and evolved a neutron star (NS) in spherical symmetry. In their work, the spacetime is treated self-consistently by satisfying a radial constraint equation. In [22], Zhao and Tang implemented DG with a WENO shock-capturing scheme for special-relativistic hydrodynamics in 1D and 2D. Bugner *et al.* [23] were the first to apply DG to a 3D astrophysical fluid problem, evolving a NS in the Cowling approximation (i.e. fixed background metric). In a DG code using a task-based parallelism paradigm (SpECTRE), Kidder *et al.* [24] showed special-relativistic magneto-hydrodynamic tests in 2D and 3D. Anninos *et al.* [25] and Fambri *et al.* [26] (see also [27]) implemented DG schemes with adaptive mesh refinement (AMR) for applications to special- and (fixed-background) general-relativistic magneto-hydrodynamics, and showed results in 2D and 3D.

In this paper we use a DG method to evolve a NS in coupled GR-hydro in 3D (prior efforts in this direction are the subject of theses by F.H. [28] and Bugner [29]). As tests of our implementation, we also evolve a NS in the Cowling approximation and a Kerr black hole (BH). In these simulations, we investigate the use of cubed-sphere grids conforming to the spherical geometry of the BH and NS problems. We adopt the DG formulation described by Teukolsky [30], using the generalized harmonic formulation of Einstein’s equations [31–33] and the València formulation [1] of the general-relativistic hydrodynamics.

We implement our DG code in the framework of the Spectral Einstein Code [34] (SpEC). SpEC combines a multi-domain penalty spectral method to evolve binary BH spacetimes [35–37] with a FV method to evolve the matter in BH-NS [4] and NS-NS [12, 38] systems. Our DG GR-hydro code is independent from SpEC’s FV component, and is instead built on the algorithms from SpEC’s vacuum spectral code: spectral bases and differentiations, domain mappings, communication, etc.

There are two main goals of this work:

1. Explore the DG method as a means of solving the GR and hydrodynamics equations simultaneously. As we will see below, the equations of the two theories take fundamentally different forms (conservative vs. non-conservative), so it is not a priori obvious that solving them on the same grid with the same technique will work.
2. Explore the use of conforming grids for BH and NS applications. In these grids, cubical elements are mapped to match the spherical geometry of an excision boundary inside the BH, or the spherical boundary at large distances from the BH or NS.

This paper is organized as follows. We first summarize the formulation of our DG method in Sec. II. We discuss our use of geometrically adapted grids, “manual” mesh refinement, and limiters in Sec. III. We detail the GR-hydro equations and associated algorithms of our numerical implementation in Sec. IV. To validate our code,

we perform standard test cases; we show these in Sec. V. We present our results — NS evolutions using the DG method — in Sec. VI, before concluding in Sec. VII.

II. DISCONTINUOUS GALERKIN FORMULATION

Our code uses a DG method to solve conservation laws in curved spacetimes, and also to evolve the spacetime itself. We express the spacetime metric $g_{\mu\nu}$ using the standard 3+1 form

$$ds^2 = g_{\mu\nu} dx^\mu dx^\nu = -\alpha^2 dt^2 + \gamma_{ab}(dx^a + \beta^a dt)(dx^b + \beta^b dt), \quad (1)$$

where α is the lapse function, β^a is the shift vector, and γ_{ab} is the spatial metric (with determinant γ) on hypersurfaces of constant time t . Our index convention is as follows: Greek indices (μ, ν, \dots) refer to spacetime components and range from 0 to d in d spatial dimensions. Latin indices (a, b, \dots) refer to spatial components and range from 1 to d . Repeated indices are summed over. We denote by \mathbf{x} the spatial point with coordinates x^a . We use units where $G, c = 1$. We additionally set $M_\odot = 1$ for the NS simulations in Sec. VI.

A conservation law in this curved spacetime can be written as a 4-divergence $\nabla_\mu F^\mu = s$, where ∇_μ is the covariant derivative, F^μ encodes the conserved quantity $u = F^0$ and its corresponding spatial flux vector $F^a(u)$, and s is the source term for u . Separating the time and spatial components gives the more common form

$$\frac{1}{\sqrt{\gamma}} \partial_t (\sqrt{\gamma} u) + \frac{1}{\sqrt{\gamma}} \partial_a (\sqrt{\gamma} F^a) = s, \quad (2)$$

which we aim to solve for $\sqrt{\gamma} u(\mathbf{x}, t)$ given initial conditions $\sqrt{\gamma} u(\mathbf{x}, 0)$ and suitable boundary conditions. When solving a system of conservation laws (e.g. for mass, energy and momentum in hydrodynamics), u is a vector of several conserved quantities and F^a is a vector of flux vectors.

We numerically solve the conservation law¹ using a strong-form, nodal DG method on square/cube elements. In this section we summarize the method, and give the specifics of our implementation. We follow the formulation given by Teukolsky [30], in which greater detail may be found.

A. Representing the solution

We divide the spatial domain into K elements. On each element we expand the quantities u, F^a, s , etc. over

¹ The conservation law is discretized (see Sec. II B) and solved for a numerical approximation to the true solution u . We do not make the distinction between the approximate and true solutions.

a set of polynomial basis functions ϕ_i , e.g.

$$u(\mathbf{x}, t) = \sum_i u_i(t) \phi_i(\mathbf{x}). \quad (3)$$

We adopt a nodal representation: we evolve the values $u_i(t) = u(\mathbf{x}_i, t)$ at the nodes \mathbf{x}_i of the computational grid, and the ϕ_i interpolate between these grid nodes. Below we define these quantities; more detailed discussion can be found in textbooks [39, 40].

The partition into elements is chosen so that each element is a mapping of a topologically simple reference element: a cube (in 3D), square (in 2D), or interval (in 1D). The mapping from the reference element coordinates $\bar{\mathbf{x}}$ to the computational coordinates $\mathbf{x} = \mathbf{x}(\bar{\mathbf{x}})$ of each element has a Jacobian matrix

$$\mathbf{J} = \frac{\partial x^a}{\partial x^{\bar{a}}} \quad (4)$$

and Jacobian $J = \det \mathbf{J}$.

In each direction, the $x^{\bar{a}}$ coordinate spans the interval $[-1, 1]$, and on this interval we place the nodes $x_i^{\bar{a}}$ of a Gauss-Legendre-Lobatto quadrature. The 1D Lagrange interpolation polynomials $\ell_j(x^{\bar{a}})$ are defined on these nodes, and satisfy $\ell_j(x_i^{\bar{a}}) = \delta_{ij}$. In the full d dimensions, we construct a tensor-product grid — we obtain the grid nodes $\bar{\mathbf{x}}_i$ from the direct product of the $x_i^{\bar{a}}$ and the basis functions ϕ_i from the product of the $\ell_i(x^{\bar{a}})$, e.g. (with some abuse of indices to indicate the tensor product)

$$\phi_i(\bar{\mathbf{x}}) \rightarrow \phi_{ijk}(\bar{\mathbf{x}}) = \ell_i(x^{\bar{1}}) \ell_j(x^{\bar{2}}) \ell_k(x^{\bar{3}}). \quad (5)$$

With N_p nodes in the $x^{\bar{a}}$ coordinate, $\ell_i(x^{\bar{a}})$ is a polynomial of degree $N = N_p - 1$. When N is the same in all directions, we say we have an N^{th} -order DG element.

We will occasionally use a modal representation in which the solution is expanded over a basis of orthonormal polynomials, e.g.

$$u(\bar{\mathbf{x}}, t) = \sum_i \hat{u}_i(t) \psi_i(\bar{\mathbf{x}}). \quad (6)$$

The \hat{u}_i are the expansion weights and the ψ_i are obtained from the tensor product of orthonormal 1D basis functions, the Legendre polynomials P_l . The Vandermonde matrix $\mathcal{V}_{ij} = P_j(x_i)$ gives the transformation between the nodal and modal representations,

$$u_i = \sum_j \mathcal{V}_{ij} \hat{u}_j. \quad (7)$$

B. DG for conservation laws

We impose the conservation law (2) in a Galerkin sense, by integrating the equation against each basis function ϕ_i on each element. We integrate over proper volume $\sqrt{\gamma} d^3 x$, giving

$$\int [\partial_t(\sqrt{\gamma}u) + \partial_a(\sqrt{\gamma}F^a) - \sqrt{\gamma}s] \phi_i(\mathbf{x}) d^3 x = 0. \quad (8)$$

To establish the flow of information between neighboring elements, we integrate the flux divergence term by parts, and apply Gauss's law to the resulting boundary term (see [30]),

$$\int \partial_a(\sqrt{\gamma}F^a) \phi_i(\mathbf{x}) d^3 x = - \int \sqrt{\gamma} F^a \partial_a \phi_i(\mathbf{x}) d^3 x + \oint F^a n_a \phi_i(\mathbf{x}) d^2 \Sigma. \quad (9)$$

Here $d^2 \Sigma$ is the proper surface element on the element's boundary, and n_a is the outward-directed unit normal.

The flux vector F^a is double-valued on the boundary because of the local (i.e. discontinuous) nature of the solution. However, for the scheme to be conservative, a unique flux must cross the boundary between two adjacent elements — this is the so-called *numerical flux* F^{a*} . The numerical flux is computed from the data on both sides of the boundary, and so requires the communication of boundary data between nearest-neighbor elements. We substitute $F^a \rightarrow F^{a*}$ in the last term of (9).

We now undo the integration by parts, using (9) to eliminate the second (i.e. $\partial_a \phi_i$) term (this time, however, we do not substitute in the numerical flux), and obtain

$$\int \partial_a(\sqrt{\gamma}F^a) \phi_i(\mathbf{x}) d^3 x \rightarrow \int \partial_a(\sqrt{\gamma}F^a) \phi_i(\mathbf{x}) d^3 x + \oint (F^{a*} - F^a) n_a \phi_i(\mathbf{x}) d^2 \Sigma. \quad (10)$$

The surface integral term provides a boundary condition on the element, and serves to connect the solution between neighboring elements of the domain. Defining $F = (F^{a*} - F^a) n_a$ and putting the terms back together, we get the DG equation in integral form,

$$\int [\partial_t(\sqrt{\gamma}u) + \partial_a(\sqrt{\gamma}F^a) - \sqrt{\gamma}s] \phi_i(\mathbf{x}) d^3 x = - \oint F \phi_i(\mathbf{x}) d^2 \Sigma. \quad (11)$$

To obtain a form more suitable for computation, we first expand each term of (11) using the nodal expansion (3). We rewrite the integrals in the reference coordinates $\bar{\mathbf{x}}$, where $d^3 x \rightarrow J d^3 \bar{x}$ and $d^2 \Sigma \rightarrow \sqrt{^{(2)}\gamma} d^2 \bar{x}$, with $^{(2)}\gamma$ the determinant of 2D metric induced by γ_{ab} on the surface. Finally, we evaluate the integrals with a Gauss-Lobatto quadrature rule. By using the grid nodes $\bar{\mathbf{x}}_i$ as the quadrature nodes we can use the identity $\ell_i(x_j^{\bar{1}}) = \delta_{ij}$ to greatly simplify the scheme. The tradeoff is that the quadrature rule will not be exact — especially when a non-trivial Jacobian J multiplies the integrand — and this can lead to aliasing and introduce numerical instabilities that require filtering.

Finally, after simplifying the geometric factors on the boundary terms (see [30], Appendix A) and dividing through by common factors, we arrive at the evolution equation,

$$\begin{aligned}
& \frac{d(\sqrt{\gamma}u)_{ijk}}{dt} + \left[\frac{\partial x^{\bar{1}}}{\partial x^a} \Big|_{ijk} \sum_l D_{il}^{\bar{1}} (\sqrt{\gamma}F^a)_{ljk} + \frac{\partial x^{\bar{2}}}{\partial x^a} \Big|_{ijk} \sum_m D_{jm}^{\bar{2}} (\sqrt{\gamma}F^a)_{imk} + \frac{\partial x^{\bar{3}}}{\partial x^a} \Big|_{ijk} \sum_n D_{kn}^{\bar{3}} (\sqrt{\gamma}F^a)_{ijn} \right] - (\sqrt{\gamma}s)_{ijk} \\
&= -\frac{\sqrt{\gamma_{Njk}^{\bar{1}\bar{1}}}}{w_N} (\sqrt{\gamma}F)_{Njk} \delta_{iN} - \frac{\sqrt{\gamma_{iNk}^{\bar{2}\bar{2}}}}{w_N} (\sqrt{\gamma}F)_{iNk} \delta_{jN} - \frac{\sqrt{\gamma_{ijN}^{\bar{3}\bar{3}}}}{w_N} (\sqrt{\gamma}F)_{ijN} \delta_{kN} \\
&+ \frac{\sqrt{\gamma_{0jk}^{\bar{1}\bar{1}}}}{w_0} (\sqrt{\gamma}F)_{0jk} \delta_{i0} + \frac{\sqrt{\gamma_{i0k}^{\bar{2}\bar{2}}}}{w_0} (\sqrt{\gamma}F)_{i0k} \delta_{j0} + \frac{\sqrt{\gamma_{ij0}^{\bar{3}\bar{3}}}}{w_0} (\sqrt{\gamma}F)_{ij0} \delta_{k0}.
\end{aligned} \tag{12}$$

Here $D_{il}^{\bar{1}}$ is the differentiation matrix along the $x^{\bar{1}}$ direction, given by

$$D_{il}^{\bar{1}} = \partial_{\bar{1}} \ell_l(x^{\bar{1}}) \Big|_i. \tag{13}$$

Although our derivation and resulting evolution equation (12) are given for the 3D case, restricting to a lower-dimensional problem is straightforward. For instance, in a 2D problem, the 3rd tensor-product index on each term is dropped (e.g. $u_{ijk} \rightarrow u_{ij}$), as are the “3” terms of the flux derivative and flux boundary terms.

C. DG for the Einstein equations

We use a formulation of the Einstein equations, detailed in the next section, that cannot be written in conservative form. These equations are instead in hyperbolic form,

$$\partial_t u + A^a \partial_a u = s, \tag{14}$$

where the matrices A^a and the vector s may be functions of u , but not of derivatives of u . To obtain the corresponding DG algorithm, we again multiply by a basis function ϕ_i and integrate over the proper volume element. We integrate by parts twice, substituting the numerical flux after the first integration, to obtain the integral form akin to (11),

$$\begin{aligned}
& \int [\partial_t u + A^a \partial_a u - s] \phi_i(\mathbf{x}) \sqrt{\gamma} d^3x = \\
& - \oint [(A^a u)^* - (A^a u)] n_a \phi_i(\mathbf{x}) d^2\Sigma.
\end{aligned} \tag{15}$$

Evaluating the integrals as before, we find

$$\begin{aligned}
& \frac{du_{ijk}}{dt} + A_{ijk}^a \left[\frac{\partial x^{\bar{1}}}{\partial x^a} \Big|_{ijk} \sum_l D_{il}^{\bar{1}} u_{ljk} + \dots \right] - s_{ijk} \\
&= -\frac{\sqrt{\gamma_{Njk}^{\bar{1}\bar{1}}}}{w_N} ((A^a u)^* - (A^a u))_{Njk} \delta_{iN} + \dots
\end{aligned} \tag{16}$$

This result is analogous to (12), so we have reproduced here only one term of each type.

1. Comparison with SpEC’s penalty spectral algorithm

SpEC solves the Einstein equations using a multi-domain penalty pseudospectral method (see, e.g. [41]). This method is closely related to our nodal DG method: the DG boundary term represents a particular type of penalty term, one chosen to enforce conservation via the numerical flux. Indeed, the spectral method in SpEC takes the form of (16) with an upwind flux, differing only in the numerical prefactor multiplying the boundary flux term. Where our DG method has a prefactor of $1/w_N$ arising from the Legendre Gauss-Lobatto quadrature rule, the SpEC penalty method instead uses the prefactor derived for stability of a Chebyshev penalty method [42]. In numerical experiments (not reported in this paper), we observe a higher order of convergence under h -refinement from the DG method (order $N+1$) than from SpEC’s spectral method (order N).

III. APPROACH TO GRID STRUCTURE, MESH REFINEMENT, AND LIMITING

Early applications of the DG method to problems in astrophysics have used uniform grids. We adopt a different philosophy and take advantage of the DG method’s geometric flexibility to tailor our grid to the problem being solved. This approach was also taken by [25, 27], who use a 2D wedge-shaped domain when evolving gas flows in a BH spacetime. We discuss here our choice of grid structures, mesh refinement, and limiting.

A. Grid structure and mesh refinement

It is well known that constructing the computational grid to mirror the underlying symmetries of the problem can greatly increase the accuracy of a numerical method. In astrophysical problems, the symmetry is often spherical, reflecting the gravitational potential of a star or BH. The use of a conforming spherical grid comes with a loss of generality: the grid must remain centered on the astrophysical body. This is especially important when taking advantage of the spherical grid to excise the singularity inside a BH. With the use of moving grids [35] and control

systems [37], however, conforming grids can be successfully used in simulations of binary mergers.

The evolutions shown in this paper make use of two basic types of grid structures.

1. **Cartesian grids**, obtained by a straightforward affine mapping (a translation and a scaling) of the reference element. These grids are used in several standard test problems.
2. **Cubed sphere grids**, obtained by conforming several cube-like elements to the surface of a sphere, using mappings detailed in Appendix A and illustrated in, e.g., Fig. 3. These grids are used for problems with spherical geometry such as single BH or NS evolutions. The cubed sphere grid may cover a hollow spherical shell, allowing for excision of the spacetime region inside the BH’s event horizon, or a filled ball, for evolution of the full NS. As we consider isolated systems at rest, moving grids are not needed.

To further take advantage of the geometric flexibility of the DG method, we use *hp*-adaptivity to vary the spatial resolution across the simulation domain. The AMR infrastructure of SpEC is designed to operate under a restricted set of conditions, and is not general enough to handle the shocks and surfaces encountered in the hydrodynamic evolutions. We instead manually set up fixed mesh refinement, where we initially assign the size and order of the DG elements based on a priori knowledge of the solution. When constructing the grid for a NS evolution, for instance, we use larger, higher-order elements inside the star, and smaller, lower-order elements at the surface. We use “higher-order” (“lower-order”) as a qualitative description of a DG element, typically referring to elements with $N \gtrsim 3$ ($N \lesssim 2$).

The SpEC framework, designed and optimized for evolutions on $\mathcal{O}(10\text{--}100)$ spectral elements, scales poorly to the large number of elements often used in DG simulations. In spite of several improvements to the data structures, we find that the code’s memory usage and parallelism become inefficient when the domain approaches $\mathcal{O}(10^4)$ elements. We therefore stay below this threshold in most of the tests presented. This restriction on the maximum number of elements would be problematic for a typical DG implementation, in which the domain is split into a regular grid of many small cubical elements. As we instead conform our grids to the problem geometry, we obtain satisfactory accuracy using many fewer elements.

B. Limiting

In DG elements containing a shock or surface in the fluid, the solution is susceptible to spurious oscillations (Gibbs phenomenon) and overshoots. If unaddressed, these overshoots can lead to unphysical fluid states (e.g.

negative densities) in which the fluid equations are no longer solvable. A limiter controls these oscillations and overshoots by modifying the solution in a way that is conservative and — ideally — does not overly degrade the accuracy of the method.

Typical DG implementations apply the limiter agnostically across the uniform grid. A “troubled-cell” detector identifies cells containing spurious oscillations and applies the limiter to those cells. While this is the most general way to set up the problem, finding a general troubled-cell detector that does not misidentify smooth extrema in the solution can be challenging. This can lead to problems, such as a smearing out of the density maximum at the center of a star.

In the context of an *hp*-adaptive DG method, however, the AMR criteria can also be used to inform the troubled-cell detector. When the solution is not smooth (i.e. the modal coefficients do not fall off rapidly enough) the AMR algorithm will reduce the order N of the element and trigger *h*-refinement. High-order elements, then, have smooth solutions and do not require limiting. In our manually refined grid, we apply the limiter only to elements with $N \leq 2$ in any spatial direction.

While our choices of grid setup and limiter application are not fully general, they are representative of the outcome from a more general AMR DG code. Our results are an exploration and will serve to inform the choices made in a future AMR update to SpECTRE (the new DG code mentioned in Sec. I).

IV. EVOLUTION OF GR AND HYDRO

A. Spacetime geometry

1. Generalized harmonic equations

We evolve the spacetime geometry using the generalized harmonic formulation of Einstein’s equations [31–33]. We use a first-order representation of the system [43] in which the evolved variables are the spacetime metric $g_{\mu\nu}$, its spatial first derivatives $\Phi_{i\mu\nu} = \partial_i g_{\mu\nu}$, and its first derivative $\Pi_{\mu\nu} = -t^\sigma \partial_\sigma g_{\mu\nu}$ along the (timelike, future-directed) normal t^σ to the constant- t hypersurface. The complete equations for $\partial_t g_{\mu\nu}$, $\partial_t \Phi_{i\mu\nu}$, and $\partial_t \Pi_{\mu\nu}$ ² in a vacuum spacetime can be found in [43]; when coupling the spacetime to matter, we add the source term

$$\partial_t \Pi_{\mu\nu} = \left(\begin{array}{c} \text{vacuum} \\ \text{terms} \end{array} \right) - 2\alpha \left(T_{\mu\nu} - \frac{1}{2} g_{\mu\nu} T^{\rho\sigma} g_{\rho\sigma} \right). \quad (17)$$

The DG method for this system of equations takes the form (16). The characteristic variables and speeds for the

² Where we use $g_{\mu\nu}$, the cited papers use $\psi_{\mu\nu}$ to denote the spacetime metric.

system, used in the upwind numerical flux shown below, are also given in [43].

For the cases we present in this paper, the natural coordinates of the initial data are well-suited to prolonged time evolution. The generalized harmonic gauge function H_σ , which specifies the coordinates, is therefore independent of time. Its precise form will depend on the data being evolved. The constraint-damping parameters γ_0 and γ_2 , which constrain the evolution of the coordinates and the growth of short-wavelength perturbations respectively, are also problem-dependent. Following [43], we fix the parameter γ_1 to -1 because this makes the generalized harmonic system linearly degenerate.

2. Upwind flux

As the solutions to the Einstein equations are smooth, we use an upwind numerical flux, which sets the flux through the boundary according to the propagation direction of each characteristic variable. The characteristic decomposition of the system is given by

$$A^a n_a u = S \Lambda S^{-1} u, \quad (18)$$

where S diagonalizes the product $A^a n_a$ ³, i.e. the i^{th} column of S is the right eigenvector of $A^a n_a$, with eigenvalue λ_i . Physically, the $S^{-1}u$ are the characteristic variables of the system, and λ_i are the associated propagation speeds with respect to the normal n_a . The diagonal matrix $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n)$ holds these eigenvalues, and can be separated by the sign of the eigenvalues, $\Lambda = \Lambda^+ + \Lambda^-$. At a boundary with two edge states u^L , u^R and a normal n_a directed towards the R state, the upwind numerical flux takes the form

$$(A^a n_a u)^{\text{upwind}} = S (\Lambda^+ S^{-1} u^L + \Lambda^- S^{-1} u^R), \quad (19)$$

so that characteristic variables propagating left-to-right (in the direction of n_a , with $\lambda_i > 0$) are set from the u^L state, whereas variables propagating right-to-left (with $\lambda_i < 0$) are set from u^R .

B. Hydrodynamics

1. Relativistic fluid equations

We treat the matter as a perfect fluid. Its stress-energy tensor takes the form

$$T_{\mu\nu} = \rho h u_\mu u_\nu + p g_{\mu\nu}, \quad (20)$$

where ρ is the fluid's rest-frame mass density, p the pressure, and $h = 1 + \epsilon + p/\rho$ the relativistic specific enthalpy, with ϵ the specific internal energy density. From the fluid's 4-velocity $u^\mu = W(1, v^i)$, we define the lower 3-velocity components $v_i = \gamma_{ij} v^j$ and the Lorentz factor $W = \alpha u^0 = 1/\sqrt{1 - v_i v^i}$. An equation of state (EOS) relates p , ρ , and ϵ ; we use an ideal-gas EOS $p = (\Gamma - 1)\rho\epsilon$, with Γ the adiabatic index. In the absence of shocks, this is equivalent to a polytropic EOS where $p = \kappa\rho^\Gamma$, with κ some constant.

The dynamics of the fluid are governed by the relativistic Euler equations. We use the València form of these equations [1], with conserved quantities $\{D, S_i, \tau\}$: the mass-energy density, momentum density, and internal energy, as measured by a generalized Eulerian observer. These are given by

$$\sqrt{\gamma} u = \begin{pmatrix} \tilde{D} \\ \tilde{S}_i \\ \tilde{\tau} \end{pmatrix} = \begin{pmatrix} \sqrt{\gamma} W \rho \\ \sqrt{\gamma} W^2 \rho h v_i \\ \sqrt{\gamma} (W^2 \rho h - p - W \rho) \end{pmatrix}. \quad (21)$$

We follow the convention of using tildes to indicate “densitized” variables, $\tilde{X} \equiv \sqrt{\gamma} X$. The corresponding flux vector and source term are

$$\sqrt{\gamma} F^a = \begin{pmatrix} \tilde{D} v_{\text{tr}}^a \\ \tilde{S}_i v_{\text{tr}}^a + \sqrt{\gamma} \alpha p \delta_i^a \\ \tilde{\tau} v_{\text{tr}}^a + \sqrt{\gamma} \alpha p v^a \end{pmatrix} \quad (22)$$

$$\sqrt{\gamma} s = \begin{pmatrix} 0 \\ (\alpha/2) \tilde{S}^{lm} \partial_i \gamma_{lm} + \tilde{S}_k \partial_i \beta^k - \tilde{E} \partial_i \alpha \\ \alpha \tilde{S}^{lm} K_{lm} - \tilde{S}^l \partial_l \alpha \end{pmatrix}. \quad (23)$$

Here $v_{\text{tr}}^a = \alpha v^a - \beta^a = u^a/u^0$ is the transport velocity relative to the coordinates, S^{lm} and E are components of the stress-energy,

$$\tilde{S}^{lm} = \sqrt{\gamma} T^{lm} = \sqrt{\gamma} \rho h W^2 v^l v^m + \sqrt{\gamma} p \gamma^{lm} \quad (24)$$

$$\tilde{E} = \sqrt{\gamma} n^\mu n^\nu T_{\mu\nu} = \sqrt{\gamma} \rho h W^2 - \sqrt{\gamma} p, \quad (25)$$

and K_{lm} is the usual extrinsic curvature of the constant- t hypersurface. The system of equations is evolved according to the discretized form (12), with the densitized conserved variables $\{\tilde{D}, \tilde{S}_i, \tilde{\tau}\}$ serving as the primary variables in the code. The characteristic speeds, used in the numerical fluxes shown below, are given in [44].

Solving for the primitive variables $\{\rho, v_i, \epsilon\}$ from $\{D, S_i, \tau\}$ requires root finding, and may additionally require “atmosphere fixing” in regions of low density where the solve may be numerically poorly behaved. We follow the inversion and fixing procedure of [45], Appendix C. This fixing procedure takes grid points where the low-density state $\{D, S_i, \tau\}$ does not correspond to a physical state $\{\rho, v_i, \epsilon\}$, and alters the conserved variables to recover a physical state. Additionally, a small (i.e. dynamically negligible) floor ρ_{atmo} is set on the fluid density, ensuring that round-off level errors are controlled. In the test problems of Sec. V, fixing is not needed; we set ρ_{atmo} to 0. For the NS evolutions of Sec. VI, fixing is necessary outside the star; we give the parameters of the fixing within that section.

³ At each point, we treat the background spacetime (i.e. $A^a n_a$) as constant, and compute the wave decomposition of the state vector u by treating it as a perturbation.

2. Numerical fluxes

For the fluid, we use a numerical flux chosen to approximately solve the Riemann problem corresponding to the discontinuity between elements. As before, we label the two states at the boundary as u^L , u^R , and the normal n_a points towards the R state. A popular choice of numerical flux, because of its robustness and simplicity, is the local Lax-Friedrichs (LLF) flux. This flux is computed according to

$$(F^{a*}n_a)_{\text{LLF}} = \frac{F^a(u^L)n_a + F^a(u^R)n_a}{2} - \frac{C}{2}(u^R - u^L), \quad (26)$$

where $C = \max(|\lambda_i(u^L)|, |\lambda_i(u^R)|)$ is the largest speed across the interface. The speeds λ_i are again the eigenvalues of the flux Jacobian (see the upwind flux discussion above, with $A^a \rightarrow \partial F^a / \partial u$). We maximize over the λ_i on both sides of the interface, but independently at each interface grid point.

A more sophisticated numerical flux, which includes an approximate treatment of the system's underlying wave structure, is given by Harten, Lax, and van Leer (HLL) [46, 47],

$$(F^{a*}n_a)_{\text{HLL}} = \frac{c_{\max}F^a(u^L)n_a + c_{\min}F^a(u^R)n_a}{c_{\max} - c_{\min}} - \frac{c_{\max}c_{\min}}{c_{\max} - c_{\min}}(u^R - u^L). \quad (27)$$

Here c_{\min} , c_{\max} are estimates for the fastest left- and right-moving signal speeds, respectively. We use the simple estimates [48], computed point-wise,

$$\begin{aligned} c_{\min} &= \min(\lambda_i(u^L), \lambda_i(u^R), 0) \\ c_{\max} &= \max(\lambda_i(u^L), \lambda_i(u^R), 0). \end{aligned} \quad (28)$$

Note that the HLL flux reduces to upwinding when all λ_i share the same sign, i.e. all characteristic variables are propagating in the same direction.

We find that the LLF and HLL fluxes give very similar results in most of the cases we tested (for an exception, see the supersonic accretion flow test in Sec. VB), and conclude that the use of an approximate solution to the Riemann problem does not introduce significant error in these problems. The results presented in this paper are computed using the HLL flux.

3. Limiters

In this work we use and compare two limiters. The first is the simple, but also low-order, minmod-based slope limiter [39, 49], which we will refer to simply as minmod. This limiter computes several estimates for the slope of the solution on each element, and then, in elements where these estimates indicate the presence of oscillations in the solution, it acts to reduce the slope. Taking the 1D case

as example, we write the solution u^k on the k^{th} element as a series expansion

$$u^k = \bar{u}^k + u_1(x - x_0) + \mathcal{O}(x - x_0)^2, \quad (29)$$

where \bar{u}^k is the element-averaged mean of u^k , u_1 is the mean slope, and x_0 is the center of the element. The minmod limiter's slope estimates are

$$a_1 = u_1, \quad a_2 = \frac{\bar{u}^{k+1} - \bar{u}^k}{h/2}, \quad a_3 = \frac{\bar{u}^k - \bar{u}^{k-1}}{h/2}, \quad (30)$$

where h is the width of the element. The limiter selects the estimate with the smallest absolute value (or 0, if the three estimates differ in sign). If the selected estimate is not the original slope u_1 , the limiter activates by reducing the slope u_1 to the selected estimate (or 0), and discarding any higher-order terms in the approximation. On elements with order $N > 1$, we use the generalization of the limiter described in [39]. We do not use the ‘‘total variation bound’’ (TVB) generalization, which sets a scale below which oscillations are tolerated, since we find that it is not robust at star surfaces.

In higher dimensions, the 1D limiter is applied to each direction in turn. After this process, the limited solution may occasionally correspond to a non-physical state. When this occurs, we further reduce the slope until the following are satisfied throughout the element: $\min(D) > \rho_{\text{atmo}}$, $\min(\tau) > 0$, and $S^2 < \tau(\tau + 2D)$.

For evolutions on deformed grids, we apply the 1D limiter along each direction of the reference $\bar{\mathbf{x}}$ coordinates. This choice of coordinates leads to a straightforward computation of the minmod slope estimates, because the series representation of u^k takes a simple form, and the element's ‘‘upper’’ and ‘‘lower’’ neighbors in each direction are well-defined. However, the choice introduces a (relatively small) violation of conservation: the limiter will conserve the means \bar{u}^k w.r.t. the reference $\bar{\mathbf{x}}$ coordinates, but the means w.r.t. the ‘‘global’’ \mathbf{x} coordinates will in general be modified after the limiter activates. This can be understood by noting that the means in the two coordinate systems are differently sensitive to the shape of the function u^k :

$$\bar{u}^k|_{\bar{\mathbf{x}}} = \frac{\int u^k d^3\bar{\mathbf{x}}}{\int d^3\bar{\mathbf{x}}} \quad \text{vs.} \quad \bar{u}^k|_{\mathbf{x}} = \frac{\int u^k d^3\mathbf{x}}{\int d^3\mathbf{x}} = \frac{\int u^k J d^3\bar{\mathbf{x}}}{\int J d^3\bar{\mathbf{x}}}. \quad (31)$$

We explored two simple corrections to the minmod limiter that restore conservation in the \mathbf{x} coordinates — limiting the Jacobian-weighted solution Ju^k instead of u^k , or shifting the post-limiting solution $u^k \rightarrow u^k + \delta u^k$ with δu^k a constant computed to restore the pre-limiting mean \bar{u}^k . Both of these corrections successfully restore the limiter's conservative properties, but we found that they also introduced long-timescale instabilities at the surface of the star — we note that Radice and Rezzolla [21] also found poor behavior when using similar corrections with the simple minmod limiter. Consequently, we do not use these corrections in our simulations. Instead

we will quantify the error in maintaining conservation when presenting our results.

The second limiter we consider is that of Moe *et al.* [50], henceforth MRS. This limiter acts by scaling the conserved variables u about their means \bar{u} ,

$$u \rightarrow \bar{u} + \theta(u - \bar{u}), \quad (32)$$

with $\theta \in [0, 1]$ determined from analysis of the minima and maxima of the solution in the immediate neighborhood of the element. A tolerance function $\alpha(h)$ sets the scale below which oscillations are tolerated; we use the function $\alpha(h) = 100h^{3/2}$ for the cases presented in this paper, as it performs well on many different test problems.

We obtain best results when computing θ from the primitive variables, as MRS recommend. However, care must be taken when computing the primitive variables, as the fluid state may be unphysical until limited. We “pre-limit” by applying an additional scaling of the form (32) to the conserved variables. The steps below restore a physical state and ensure the inversion procedure is well-posed:

1. if $\min(D) < \rho_{\text{atmo}}$ or $\min(\tau) < 0$, scale to fix these violations.
2. if $S_i S^i > \tau(\tau + 2D)$ at any grid point, scale to fix this violation. This requires solving a quadratic equation for θ .
3. if the inversion to primitive variables encounters any of the errors outlined in [45], Appendix C (this is rare), scale again with $\theta = 1/2$.

This procedure is conservative by construction, and we find it to be robust. After this “pre-limiting” step, we compute the primitive variables and limit according to the MRS prescription. We handle deformed grids as for the minmod limiter, by computing the means in the reference $\bar{\mathbf{x}}$ coordinates and incurring some error due to loss of conservation. As with minmod, attempts to reformulate the limiter to restore conservation (we tried the simple approach of computing the MRS means directly w.r.t. the \mathbf{x} coordinates, as well as the same two reformulations described for minmod) were not stable at the star surface.

We apply the limiter to the hydro variables at the end of each time-stepper substep. As described in Sec. III B, we may not apply the limiter to every element, choosing instead to mimic an AMR scheme in which high-order elements are known to be smooth. The use of more complex, higher-order, limiters, e.g. sub-cell methods [21, 23] or the compact-stencil WENO [51] and HWENO [52] limiters, will be the subject of future investigation.

C. Combined GR-hydro system

For self-consistent NS evolutions, the equations of the GR and hydro systems are each treated as described

above, and are evolved in parallel. The two systems couple via their respective source terms and the geometry terms in the hydro flux $F^a(u)$. We compute the characteristic speeds independently for each system, leaving out the cross-coupling arising from the off-diagonal $\partial F_{\text{hydro}}^a / \partial u_{\text{GR}}$ flux Jacobian terms. When the hydro variables require limiting, the limiter is applied to the hydro variables only, and the spacetime variables are left unmodified.

D. Filtering

The use of inexact quadratures to obtain an efficient DG scheme may result in numerical instabilities caused by aliasing. Where these numerical instabilities exist, we address them by filtering the higher modes in the solution’s modal representation. We use an exponential filter, e.g. in 1D

$$\hat{u}_i \rightarrow F(i)\hat{u}_i = \exp\{-\alpha(i/N)^s\}\hat{u}_i, \quad (33)$$

where α controls the strength of the filter’s effect, and s is an even integer controlling how many modes are affected. In $d > 1$ dimensions, we take advantage of the tensor-product basis to apply the filter dimension-by-dimension; this gives d exponentials. On deformed grids, we filter the Jacobian-weighted solution Ju and then divide by J , so that the operation remains conservative. We apply the filter at the end of each complete time step to the components of u and on the elements that show numerical instability.

E. Time stepping

We use the 3rd-order strong stability-preserving Runge Kutta scheme of [53] for the time integration. Given the solution u^n at time t^n , the solution u^{n+1} at time $t^{n+1} = t^n + \Delta t$ is computed as

$$\begin{aligned} u^{(1)} &= u^n + \Delta t L(u^n) \\ u^{(2)} &= \frac{3}{4}u^n + \frac{1}{4}\left[u^{(1)} + \Delta t L(u^{(1)})\right] \\ u^{n+1} &= \frac{1}{3}u^n + \frac{2}{3}\left[u^{(2)} + \Delta t L(u^{(2)})\right] \end{aligned} \quad (34)$$

Here $L(u) = du/dt$ is computed from expressions (12) for the hydro variables or (16) for the GR variables.

In all cases presented, the initial $t = 0$ data is computed by pointwise evaluation of a known state. The limiter is applied to the initial data and at the end of every subsequent substep. Filtering is done at the end of full time steps.

V. CODE TESTS

In this section we present a selection of benchmark tests that we use to validate our implementation of the

DG method within SpEC.

We first show tests of vacuum spacetime evolution. From a family of gauge wave evolutions at varying resolutions, we verify that the method converges to the exact solution at the expected rate. Next, by evolving a Kerr (i.e. isolated and spinning) BH over long timescales, we show the stability of the algorithm.

We then show our tests of the hydrodynamics implementation. We again verify the convergence rate of the errors, now with a generalized Bondi problem in which the fluid undergoes spherically symmetric accretion onto a Schwarzschild BH. This test verifies the fluid equations as well as the sourcing of the fluid by the spacetime curvature. We then show standard shock tests in 1D and 2D, comparing the effectiveness of the implemented limiters.

In these tests, whenever possible, we compare the numerical solution to an exact solution, and we use their difference as an error measure. We report a normalized error $\text{err}[X]$ in a quantity X , defined as

$$\text{err}[X] = \|X - X_{\text{exact}}\| / \|X_{\text{exact}}\|. \quad (35)$$

Here $\|X\|$ is the L^2 -norm, evaluated pointwise by direct summation over every node of the computational grid,

$$\|X\|^2 = \frac{1}{N_{\text{nodes}}} \sum_{i=0}^{N_{\text{nodes}}} X_i^2. \quad (36)$$

When X is a vector or tensor quantity, we compute a **component-wise norm** $\|X\|^2 = \|X_0\|^2 + \|X_1\|^2 + \dots$, rather than the physical norm $X_a X^a$. When $X_{\text{exact}} = 0$ so that we cannot define the normalized error, we instead use $\|X\|$ as our error measure.

A. Spacetime tests

1. Gauge wave test

The spacetime of the ‘‘apples to apples’’ gauge wave test [54], obtained via a nonlinear, plane-wave transformation of Minkowski space, takes the form

$$ds^2 = -(1+a)dt^2 + (1+a)dx^2 + dy^2 + dz^2, \quad (37)$$

with

$$a = A \sin[2\pi(x-t)]. \quad (38)$$

We show results for a wave of amplitude $A = 0.1$ on a unit-cube domain with extents $[0, 1]^3$. As the gauge wave is harmonic, the generalized harmonic gauge function H_σ is zero. We set the constraint-damping parameters $(\gamma_0, \gamma_1, \gamma_2)$ to $(1, -1, 1)$, values that give stable evolutions over long timescales (up to at least $t_{\text{fin}} = 1000$, or 1000 crossing times). For the convergence study, however, we measure the error in the spacetime metric $g_{\mu\nu}$ at a final time $t_{\text{fin}} = 10$, after evolution with time steps of size $\Delta t = 10^{-4}$. This time step corresponds




FIG. 1. The error in $g_{\mu\nu}$ as a function of the number of elements (h -refinement) for the gauge wave test of Einstein’s equations. The symbols indicate the measured error norms for methods of order $N = \{2, 3, 4\}$. The dashed lines, normalized to the $K = 16$ data, indicate the expected error scaling for third, fourth, and fifth-order convergence.

to $\Delta t / \Delta x_{\text{min}} \simeq 0.074$ for the highest-resolution case in the convergence study ($K = 128$, $N = 4$).

We show in Fig. 1 the convergence under h -refinement, measured for elements of order $N = 2, 3, 4$. As a base resolution we partition the unit-cube domain into 16 elements along the x direction; we h -refine by further splitting each element along x , reducing the element’s width h in half each time. We do not split in y or z — the anisotropic refinement is chosen to match the x -only dependence in the problem. For each order N of the DG method, we compare our measurements to the theoretical scaling of the error (see, for instance, [39]),

$$\text{err}[g_{\mu\nu}] \leq Ch^{N+1} \propto 1/K^{N+1}, \quad (39)$$

for some constant C . We find excellent agreement between the measured and expected convergence rates. The highest-resolution case ($K = 128$ and $N = 4$) has slightly larger error, having reached the round-off level error in the derivatives of the spacetime.

In Fig. 2 we show the convergence under p -refinement, obtained by increasing the order N of the DG method while maintaining the base resolution of 16 elements. We expect the errors to decrease exponentially with the order N , and recover this behavior in our measurements. This result demonstrates the spectral convergence of the DG method for smooth solutions.

2. Kerr black hole

We next evolve the spacetime of a Kerr BH, described by the Kerr metric in Kerr-Schild coordinates [55]. The




FIG. 2. The error in $g_{\mu\nu}$ as a function of the order of approximation (p -refinement) for the gauge wave test of Einstein’s equations. The number of elements is fixed at $K = 16$. The dots indicate the measured errors; the dashed line is a fit demonstrating the exponential decrease in error with N .

BH has spin $\vec{a} = (0.1, 0.2, 0.3)M_{\text{BH}}$ with magnitude $a \approx 0.374M_{\text{BH}}$, not aligned with any grid symmetries. We use units where $M_{\text{BH}} = 1$.

The domain is a hollow spherical shell that excises the singularity within the BH. In terms of the coordinate radius r , the domain extends from $r_{\text{in}} = 1.8$ (just inside the event horizon) to $r_{\text{out}} = 32$. At the inner boundary, all the characteristics of the system are outgoing (i.e., leaving the domain, towards the singularity) so no boundary condition needs to be imposed. Physically, no information enters the simulation from the interior of the BH. At the outer boundary, we impose the analytic solution as a Dirichlet boundary condition⁴. We choose constraint damping parameters

$$\gamma_0 = 3 \exp[-(r/8)^2/2] + 0.1 \quad (40)$$

$$\gamma_1 = -1 \quad (41)$$

$$\gamma_2 = \exp[-(r/8)^2/2] + 0.1. \quad (42)$$

The generalized harmonic gauge function $H_\sigma = \Gamma_\sigma \equiv g^{\mu\nu}\Gamma_{\sigma\mu\nu}$ is the trace of the Christoffel symbols of the Kerr-Schild metric; it is constant in time.

We set up a cubed-sphere grid on this domain, using the mappings from Appendix A. The wedges of the cubed sphere are split radially into 5 concentric shells located between the surfaces $r = \{1.8, 3.2, 5.7, 10, 18, 32\}$, and then tangentially into 2×2 angular portions, for a total of 120 elements. The tangential coordinates of each wedge are mapped to obtain an equiangular grid, as this

⁴ We do not use the constraint-preserving boundary conditions typically used in SpEC simulations, because these are boundary conditions on $\partial_t u$, rather than u , and so would require a modification of the DG formulation.

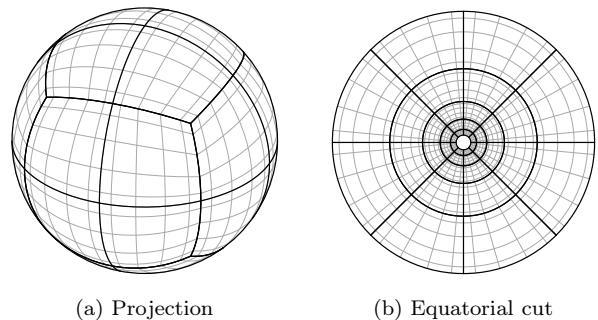


FIG. 3. The grid structure for the Kerr BH evolution test. Shown are (a) a projected view, and (b) an equatorial cut. The black lines show the element boundaries, the light grey lines show the Gauss Legendre Lobatto grid within each element for order $N = 5$.

is a more optimal distribution for the grid points on the spherical surface. We show in Fig. 3 two views of this grid: on the left a projected view showing the angular structure on a constant-radius surface, on the right an equatorial cut showing the radial structure. The clustering of the grid points towards the center helps to resolve the spacetime curvature near the BH.

In Fig. 4 we show the stability of the Kerr BH evolution by monitoring the simulation errors over a duration of $10^4 M_{\text{BH}}$. We carry out the simulation using elements of order $N = 5, 6$, and 7 ; the time-step size is $\Delta t = 10^{-2}$, giving $\Delta t/\Delta x_{\text{min}} \simeq 0.15$ for the $N = 7$ case. The figure’s top panel shows the error $\text{err}[g_{\mu\nu}]$ in the spacetime metric, a measure of the solution’s drift from the exact value. The bottom panel shows the dimensionless norm $\|\mathcal{C}\|$ of the generalized harmonic energy constraint [43], a measure of how well the numerical solution at each constant- t slice satisfies Einstein’s equations. After a rapid settling of the solution to its numerical equilibrium, we see clear convergence in the error quantities. We conclude that the method is convergent and stable up to at least $t = 10^4 M_{\text{BH}}$, and, we presume, forever.

B. Relativistic hydrodynamic tests

1. Spherical accretion onto black hole

In the relativistic Bondi problem, an ideal gas accretes radially onto a non-rotating BH. The feedback from the fluid onto the spacetime is ignored: the BH mass is constant and the spacetime is Schwarzschild. We use Kerr-Schild coordinates, and again we set $M_{\text{BH}} = 1$. The analytic profile for the fluid flow is presented by Michel [56]; following [44], we pick a solution for a $\Gamma = 5/3$ ideal gas with the sonic point and mass accretion rate given by $r_{\text{crit}} = 200$ and $\dot{M} = 10^{-3}$. We measure the error in the conserved relativistic density \tilde{D} at a final time $t_{\text{fin}} = 100$, after evolution with time steps




FIG. 4. The errors during the Kerr BH evolution test. The top panel shows the error in the spacetime metric $g_{\mu\nu}$ for three different orders of the DG method. The lower panel shows the dimensionless norm of the generalized harmonic energy constraint at the same three orders.

of size $\Delta t = 5 \times 10^{-3}$. This time step corresponds to $\Delta t/\Delta x_{\min} \simeq 0.15$ for the highest-resolution case in the convergence study ($K = 120 \times 4^3$, $N = 4$).

We evolve the fluid in a hollow spherical shell extending from $r_{\text{in}} = 1.8$ (just inside the event horizon), to $r_{\text{out}} = 12$. The sonic point in the accretion flow is located outside this region, so the flow is smooth and supersonic throughout the simulation domain. In this test problem we obtain significantly more accurate results when using the HLL numerical flux (vs. LLF), as the supersonic flow is best represented by the HLL upwinding limit. At the inner boundary, the characteristics of the fluid system are outgoing (i.e., leaving the domain into the BH), so no boundary condition needs to be applied. At the outer boundary, we impose the analytic solution as boundary condition.

We use a cubed-sphere grid similar to that of the Kerr BH test above. At the base resolution, we divide the domain into 5 spherical shells between the surfaces located at radii $r = \{1.8, 2.7, 4, 6, 9, 12\}$, and we split each wedge into 2×2 angular portions. The tangential coordinates are again mapped to obtain an equiangular grid.

We show in Fig. 5 the convergence under h -refinement of this grid, for elements of order $N = 2, 3, 4$. We h -refine by splitting each element into 2^3 smaller ele-




FIG. 5. The error in the conserved density \tilde{D} as a function of the number of elements (h -refinement) for the spherical accretion test. The symbols indicate the measured error norms for methods of order $N = \{2, 3, 4\}$. The dashed lines, normalized to the $K = 120$ data, indicate the expected error scaling for third, fourth, and fifth-order convergence.

ments: we split geometrically in radius according to $r_{\text{split}} = \sqrt{r_{\text{lower}} r_{\text{upper}}}$, and linearly in the tangential directions. As the elements are not uniform, this choice of radial split is not unique, but we find it gives reduced error compared to a linear split $r_{\text{split}} = (r_{\text{lower}} + r_{\text{upper}})/2$. We again see the errors converging at the expected rate.

In Fig. 6 we show the convergence under p -refinement. Again we use the base configuration of elements, and increase the order N of the method from 2 to 7. We confirm that for this smooth fluid evolution problem, the errors decrease exponentially with the order of the method.

2. 1D shock tube test

We perform a standard 1D relativistic shock test problem, in which a high density and pressure fluid expands into a low density and pressure fluid. Following [4], we take a $\Gamma = 5/3$ ideal gas initially split at $x = 0.5$ into left and right states characterized by

$$(\rho, v_x, p) = \begin{cases} (10, 0, 40/3), & x < 0.5 \\ (1, 0, 0), & x > 0.5 \end{cases} \quad (43)$$

The simulation domain is an interval $x \in [0, 1]$, which we divide into $K = 160$ elements of order $N = 2$. We evolve the shock until a final time $t_{\text{fin}} = 0.4$, with time steps $\Delta t = 4 \times 10^{-3}$ ($\Delta t/\Delta x_{\min} = 0.128$).

In Fig. 7 we show the profiles of ρ , v_x , and p at the final state, comparing the minmod and MRS limiters. Both limiters capture the features of the shock profile. The minmod limiter produces a larger overshoot at the main




FIG. 6. The error in the conserved density \tilde{D} as a function of the order of approximation (p -refinement) for the spherical accretion test. The number of elements is fixed at $K = 120$. The dots indicate the measured errors; the dashed line is a fit demonstrating the exponential decrease in error with N .




FIG. 7. Snapshot of the fluid variables in the shock tube test. The fluid pressure p , the rest-mass density ρ , and the velocity v_x (scaled up $10\times$) are plotted after evolutions using the minmod and MRS limiters. The mean value on each element is shown. The exact solution to the problem is given by Centrella & Wilson [58], and is plotted here in the solid line.

shock front and increased oscillation at the front end of the rarefaction fan, a known behavior when applying this limiter to the conserved variables (rather than characteristic variables [57]).

3. 2D Riemann shock interaction test

We next study a standard 2D Riemann problem in which two shocks and two contact discontinuities inter-

act. As in the 1D shock test, the fluid is a $\Gamma = 5/3$ ideal gas. The initial conditions for the problem were first generalized from Newtonian to relativistic hydrodynamics by Del Zanna & Bucciantini [59] and later modified by Mignone & Bodo [6] to give a cleaner wave structure. The initial condition divides the computational domain $[-1, 1]^2$ into four quadrants, each of which holds a constant fluid state,

$$(\rho, v_x, v_y, p) = \begin{cases} (0.5, 0, 0, 1), & x < 0, y < 0 \\ (0.1, 0, 0.99, 1), & x > 0, y < 0 \\ (0.1, 0.99, 0, 1), & x < 0, y > 0 \\ (\rho_1, 0, 0, p_1), & x > 0, y > 0 \end{cases}, \quad (44)$$

where the low-density state in the upper-right quadrant is defined by $\rho_1 = 5.477875 \times 10^{-3}$ and $p_1 = 2.762987 \times 10^{-3}$. We partition the domain into 200×200 elements of order $N = 2$, and we evolve until a final time $t_{\text{fin}} = 0.8$ with time steps $\Delta t = 10^{-3}$ ($\Delta t / \Delta x_{\text{min}} = 0.2$).

In Fig. 8 we show contour plots of the density ρ at the final state. We interpolate the evolved ρ onto a high-resolution uniform grid on which the contours are computed. The results in the top panel are computed with a minmod limiter, those in the bottom panel with MRS. We find, qualitatively, an excellent agreement between the results from the two limiters; only the jet feature (in the lower-left quadrant) shows a clear difference in resolution, with the MRS limiter producing a cleaner structure.

VI. NEUTRON STAR EVOLUTIONS

Having verified the convergence and shock-capturing properties of our code, we now present our main results: evolutions of an isolated, spherical NS using the DG method. We first evolve the NS under the Cowling approximation, i.e. keeping the background spacetime fixed to the TOV solution. This remains a challenging test of the hydro code's ability to handle the discontinuity at the stellar surface. We then evolve the NS self-consistently using the coupled GR-hydro system.

The initial data for the NS fluid and spacetime are found by integrating the Tolman-Oppenheimer-Volkoff (TOV) equations [55, 60, 61] for the mass-energy density $\rho_E(R) \equiv \rho(R)(1 + \epsilon(R))$, enclosed ADM mass $m(R)$, and metric potential $\phi(R)$ in terms of the areal radius R . The spacetime metric is given by

$$ds^2 = -e^{2\phi} dt^2 + \left(1 - \frac{2m}{R}\right)^{-1} dR^2 + R^2 d\Omega^2. \quad (45)$$

In computing the TOV solution, we describe the NS matter by a polytropic EOS. When time-evolving the solution, we return to the corresponding ideal-gas EOS.

The results presented throughout this section are for a star with $\kappa = 100$ and $\Gamma = 2$. The star has central mass density $\rho_c = 1.28 \times 10^{-3}$, giving a stable, non-rotating TOV solution with ADM mass $M_{\text{NS}} \simeq 1.4M_{\odot}$ and areal

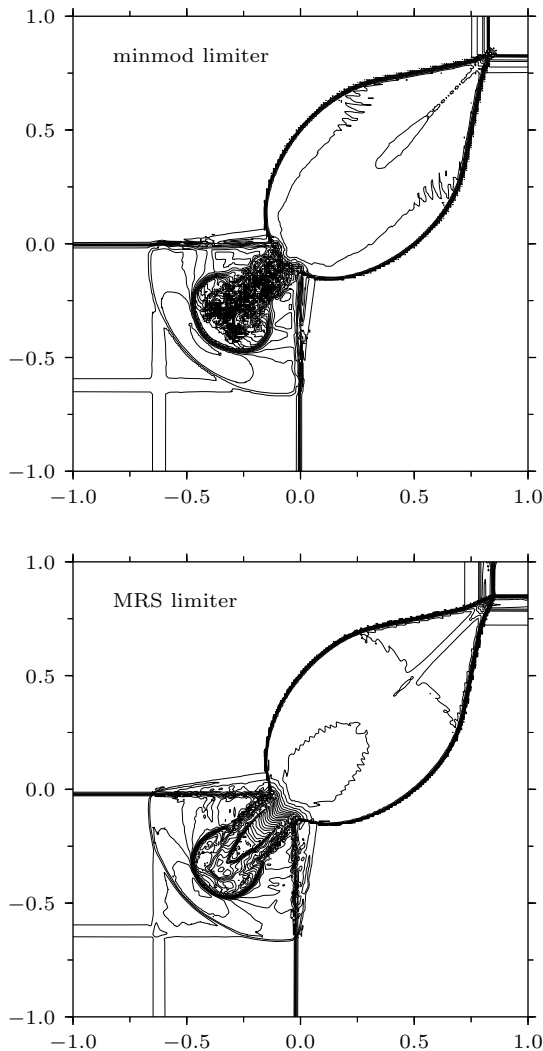


FIG. 8. The density ρ in the 2D Riemann problem. The top panel is computed with the minmod limiter, the bottom panel with the MRS limiter. The plots each show 30 contour lines, equally spaced in $\log \rho$.

radius $R_{\text{NS}} \simeq 9.6M_{\odot} \simeq 14$ km. Its radius in the isotropic coordinates used during evolution is $r_{\text{NS}} \simeq 8.125M_{\odot}$. In this section, we use units where $M_{\odot} = 1$.

For the NS evolutions, we use the atmosphere fixing from [45], Appendix C. We set the density cutoff $\rho_{\text{atmo}}^{\text{cut}} = 10^{-15}$ so as to resolve 12 orders of magnitude in density. Where the density falls below this cutoff, we set the fluid to the “atmosphere” state where $\rho = \rho_{\text{atmo}} = 10^{-16}$, $v_i = 0$, and $\epsilon = 0$. Elsewhere, we constrain the specific internal energy to the range $\kappa\rho \leq \epsilon \leq 100\kappa\rho$, with κ from the polytrope describing the initial conditions. These bounds serve to control the fluid entropy in the region around the star surface, by preventing numerical errors from causing an entropy decrease, and allowing heating only within a reasonable range. To check that our results are not influenced by the choice of these atmosphere

TABLE I. The structure of the spherically symmetric NS grids I1, I2, and I1R. For each grid, the parameters defining the elements in the interior, right-side surface, and right-side exterior regions are given; the elements in the left-side surface and left-side exterior regions are obtained by symmetry. The interior and exterior regions of I2 are identical to those of I1.

	extents	K_{region}	N_{region}	
I1	interior	$[-7.5, 7.5]$	25	3
	surface (right side)	$[7.5, 10]$	10	1
	exterior (right side)	$[10, 24]$	7	3
I2	surface (right side)	$[7.5, 10]$	5	2
I1R	interior	$[-8, 8]$	101	3
	surface (right side)	$[8, 9]$	20	1
	exterior (right side)	$[9, 24]$	30	3

fixing thresholds, we evolved a few comparison cases in which we increased the densities $\rho_{\text{atmo}}^{\text{cut}}$ and ρ_{atmo} by a factor of 10. These comparison evolutions deviated only slightly from the primary evolutions, confirming that our atmosphere treatment does not strongly impact the neutron star simulations.

A. Cowling neutron star in spherical symmetry

We begin with 1D evolutions in spherical symmetry. For these simulations, we rewrite the conservation law (2) and the relativistic Euler equations in terms of spherical coordinates $\{r, \theta, \phi\}$. The DG formulation takes a form similar to (12) in 1D, but with a spherical divergence $\partial_r(r^2 u^r)/r^2$ replacing the Cartesian divergence $\partial_x u^x$. The fluid equations pick up an additional momentum source term: $s(S_r) = s(S_x) + \alpha p(g^{rr} \partial_r g_{rr} + 2/r)$. To avoid the coordinate singularity at $r = 0$, we set up a symmetric domain on the interval $[-24, 24]$ and use a staggered grid so that no nodes are located at the origin.

On this domain we consider three grids with different resolutions. The first two, which we name I1 and I2, have comparable resolutions to the grids of our 3D simulations. These two grids differ in the order of the DG elements near the surface of the star: linear elements are used in I1 vs. quadratic elements in I2. The third grid, I1R⁵, has a higher resolution and is more aggressively refined around the surface of the star. In all three grids, we divide the domain into five regions: the interior of the star, the surface on the left/right, and the exterior on the left/right. We use larger, higher-order elements in the interior and exterior regions, and smaller,

⁵ The grid names are structured as follows: the first letter encodes the domain’s topology (“I”: interval; “B”: ball), the integer gives the (radial) order of approximation of the elements near the surface of the star, and a final “R” indicates a refined, higher-resolution grid.




FIG. 9. The density errors in the spherically symmetric Cowling NS evolution. The four cases correspond to different choices of grid (I1 or I2) and limiter (minmod or MRS) at the star surface. As some of the curves are highly oscillatory, we plot for each case the mean error in a solid line, and the envelope as a light-colored shaded region. The mean is computed by applying a gaussian smoothing to the data, with half-width $\sigma \simeq 5$; the envelope min/max are computed in bins of width $\Delta \simeq 15$.

lower-order elements in the neighborhood of the star’s surface. The number and order of the elements within each region are listed in Table I. We evolve the system until $t = 10^4 \simeq 50$ ms. On the lower-resolution grids I1 and I2, we use time steps $\Delta t = 0.04$ corresponding to $\Delta t/\Delta x_{\min} = 0.29$. On the higher-resolution grid I1R, we use time steps $\Delta t = 0.025$ with $\Delta t/\Delta x_{\min} = 0.57$.

We now compare evolutions of the spherically symmetric NS for different choices of grid and limiter — specifically the I1 or I2 grids, the minmod or MRS limiters. We plot in Fig. 9 the normalized density error $\text{err}[\tilde{D}]$ for each case over the duration of the simulation. We first examine the two minmod cases. Here the data reveal two components in the dynamics: a short-period oscillatory behavior, and a gradual drift as the star settles to its numerical equilibrium configuration on much longer timescales. The I2 case has a much higher initial error and increased dissipation, as indicated by the more rapid decay of the oscillatory component. The increased error and dissipation occur because the minmod limiter linearizes the solution on the quadratic elements at the surface, resulting in the loss of information. In the two MRS cases, we find very different behavior: the density error is roughly an order of magnitude larger (vs. minmod) and grows over time, and the high-frequency oscillations are not damped on the timescale of the simulation. Although the MRS evolutions are stable (on this timescale), the star does not settle to an equilibrium.

To better understand this difference in behavior be-




FIG. 10. The rest-frame density error vs. radius, at time $t = 2000$, in the spherically symmetric Cowling NS evolution. The shaded region shows the min/max errors, at each radius, from the left and right sides of the symmetric domain. We plot the mean error in the solid line. The vertical dotted line indicates the location of the TOV star surface at $r_{\text{NS}} \simeq 8.125$.

tween the minmod and MRS limiters, we now look at the distribution of the errors across the star. In Fig. 10 we show the error in the fluid density vs. the stellar radius, at time $t = 2000$. The solid lines in this plot show the angle-averaged errors (i.e. the average of the left and right-side data), and the lighter filled region shows the spread in error values at fixed radius. From this plot we make two observations. First, while the minmod limiter maintains excellent symmetry across the star, the MRS case shows a large spread in the error values, indicating a loss of spherical (i.e. reflection) symmetry. Second, while the density and velocity errors in the minmod case are largest at the surface of the star, denoted by a vertical dotted line in the figure, the fluid remains confined within the true surface of the star. When using MRS, the star instead extends significantly beyond the true surface: matter with non-negligible densities and large ($v > 0.01$) velocities is present out to $r \simeq 15$. Our interpretation is that the MRS limiter provides insufficient damping of small-scale fluctuations in the atmosphere near the star⁶. These slowly grow, leading to the expansion of the star beyond its true surface and the contamination of the solution inside the star.

We conclude from these comparisons that the MRS limiter — although effective at handling shocks — is poorly suited to the task of controlling a stellar surface on the

⁶ Moe *et al.* [50] discuss interpolating the solution to a finer grid (to better sample its shape) before computing the maxima and minima used by the limiter. We found no significant improvement in behavior when using this interpolation.

conforming grids that we are using. In the remainder of this paper we therefore only show results obtained with the minmod limiter. We will also restrict to results from grids with linear-order elements at the star surface, as these are much less dissipative given our use of this low-order limiter.

In Sec. V, we showed that our DG implementation has the expected convergence properties for test problems with smooth solutions. For the NS problem, the expected convergence behavior is less clear: the convergence rate will be degraded by the discontinuity at the stellar surface, and, furthermore, our use of a geometrically adapted grid with elements of different sizes and different orders makes it more difficult to quantify the resolution. In spite of this, we examine the convergence behavior with a series of short evolutions in which we successively h -refine the I1 grid, and the timestep, by factors of 2 and 4 (in the interior region of the grid we refine from 25 to 49, then 97, elements; this maintains a grid that straddles the origin). From these evolutions, we measure the error in \tilde{D} at time $t = 1000$; we expect the error decrease to lie between a fourth-order convergence (corresponding to the case where the error is dominated by the $N = 3$ interior) and first-order convergence (corresponding to the case where the error is dominated by the discontinuity at the surface). We find (but do not show) that as the grid is refined, the error decrease is consistent with a third-order convergence.

To investigate the degradation in accuracy caused by the stellar surface, we plot in Fig. 11 the spatial and temporal variation of a locally defined convergence order. At each point on the plot, the convergence order is computed from Eq. (39) by measuring the decrease in the element-averaged error in \tilde{D} between simulations on two computational grids: the I1 grid and the denser grid obtained by a $4\times$ refinement of I1. The interior of the star initially shows the expected fourth-order convergence, but during the first $50 M_\odot$ the order quickly decreases to roughly second-order as lower-accuracy data from the stellar surface propagates into the interior. We find, in agreement with [23, 62], that as the star settles to its numerical equilibrium, the order of convergence increases again to roughly third-order convergence. This shows that the degradation in convergence caused by the stellar surface is limited, and so the high-order qualities of the DG method, to a large degree, continue to apply.

We now take a second, closer look at the spherically symmetric NS evolution. In Fig. 12 we compare evolutions of the spherically symmetric NS on the I1 and I1R grids. We show, in the top two panels, the errors in the conserved quantities \tilde{D} and \tilde{S}_r over the first $4000 M_\odot \simeq 20$ ms of evolution time. The errors are lower by one or two orders of magnitude in the I1R case. In the bottom panel, we plot the time-dependence of the central density ρ_c as a fractional error with respect to the initial central density $\rho_{c,0}$. We see in ρ_c a qualitative difference between the two evolutions, with the I1R case showing a clearly periodic structure corresponding to the cross-




FIG. 11. The local convergence order in the spherically symmetric Cowling NS evolution. The convergence order is measured from the decrease in the element-average of $\text{err}[\tilde{D}]$ between simulations on two computational grids: the I1 grid and a denser grid obtained by refining each I1 element into 4. Lighter colors correspond to a higher order of convergence, i.e. a more rapid decrease in error.

ing time for perturbations seeded at the surface of the star. In the full evolution to $t = 10^4$, not shown in the figure, the high-resolution evolution maintains its equilibrium, with the remaining oscillations in \tilde{S}_r and ρ_c slowly decaying. In the low-resolution evolution, the density error $\text{err}[\tilde{D}]$ asymptotes to roughly 10^{-3} , the oscillations in $\|\tilde{S}\|$ slowly decay, and the central density continues to slowly drop, reaching a 0.05% deficit at $t = 10^4$.

While the DG method is fundamentally conservative, we have discussed (in Sec. IV B 3) how our use of limiters on deformed elements can violate this property. The spherically symmetric simulations are also affected, even though the 1D elements are themselves undeformed, because the limiter does not account for the spherical Jacobian $4\pi r^2$ that takes the 1D volume element to the spherical volume element. The corrections to restore conservation explored in Sec. IV B 3 are no more effective in the 1D case: conservation is restored at the expense of stability near the star surface. We quantify the conservation error by tracking the NS's baryon mass $M_b \equiv \int \tilde{D} 4\pi r^2 dr$ during the evolution, as this should be a conserved quantity. We find (but do not show) that M_b slowly grows. On the I1 grid, the relative error in M_b (w.r.t. its initial value) reaches roughly 10^{-4} at $t = 10^4$, with over 90% of this growth occurring over the initial $4000 M_\odot$ as the star settles towards equilibrium. On the I1R grid, the error grows to about 4×10^{-7} , again mostly over the initial portion of the evolution.

We now reconsider Fig. 12, and focus on the oscillatory behavior seen in the different quantities. These oscillations are triggered by errors from two sources: truncation errors from the evaluation of the exact TOV solution on the finite-resolution numerical grid, and the action of the limiter which modifies the initial solution near the star's surface. These errors seed perturbations of the various radial eigenmodes of the star, each of which subsequently resonates with its corresponding eigenfrequency. A common test of NS evolution codes is to compare the fre-




FIG. 12. The errors in the spherically symmetric Cowling NS evolution. The top (middle) panel shows the error in the conserved density \tilde{D} (conserved momentum \tilde{S}_i) for evolutions using the minmod limiter on the grids I1 and I1R. The bottom panel shows the evolution of the central density ρ_c as a fractional error with respect to its initial value $\rho_{c,0}$. The inset in the bottom panel zooms in to better show the initial ρ_c evolution in the I1R case; the I1 curve is omitted from the inset for visual clarity.

quency spectrum of the simulated star against the eigenfrequencies obtained from linear theory.

To make this comparison, we compute the frequency spectrum from the central rest-mass density during the first $4000M_\odot$ of evolution time. After subtracting the initial density offset $\rho_{c,0}$, we apply a Hanning window to the time interval and compute the discrete Fourier transform. We plot in Fig. 13 the absolute value of the Fourier coefficients against frequency. The dotted vertical lines indicate the (Cowling) NS’s radial eigenmode frequencies, as listed in Table I of [63]. The evolution on the I1 grid resolves few of the star’s eigenmodes: the spectrum has sharp peaks corresponding to the fundamental mode and the first harmonic only. Modes with higher frequencies (i.e. shorter wavelengths) are not spatially resolved by




FIG. 13. The Fourier transform of the central rest-mass density ρ_c from the spherically symmetric Cowling NS evolutions. The data from evolutions on the I1 and I1R grids are shown — the I1R curve is shifted downwards on the plot, by a factor of 10, for visual clarity. The vertical dotted lines indicate the frequencies of the fundamental normal mode and the first six harmonics. The units of the vertical axis are arbitrary.

this computational grid, and so the power they contain aliases into the lower-frequency modes. The evolution on I1R, on the other hand, reproduces very clearly the fundamental mode frequency and the first four harmonic frequencies. At higher frequencies, the peaks are still identifiable, though they become broader and less precisely centered. We note the presence of intermediate peaks in the spectrum, shaped in a manner suggestive of sidebands; however, these features are too noisy for unambiguous identification.

The evolution on I1R, with roughly 210 points across the NS’s radius, has a similar resolution to the 75-element case presented by Radice and Rezzolla [21]. While these two simulations are not directly comparable (the 1D star in the cited work self-consistently treats the gravity, and uses a uniform grid), we see a qualitative agreement in the number of resolved modes and the precision with which they are resolved.

We conclude the discussion of 1D evolutions by noting that while the I1R grid has significantly reduced error, the lower-resolution I1 grid, representative of the 3D resolution, is sufficient to resolve the important features in the evolution. The lower-resolution case remains stable on long timescales, and the oscillations as the star settles to its numerical equilibrium correctly reflect the low-frequency eigenmodes from linearized theory.




FIG. 14. The B1 grid used in the 3D NS evolutions. The thick dotted circle indicates the location of the star’s surface. The black lines show the element boundaries, the light grey lines show the Gauss Legendre Lobatto grid within each element. The details of the grid mappings and structure are given in Appendices A and B respectively.

B. Cowling neutron star in 3D

The simulation domain for the 3D star is a filled ball extending to $r_{\max} = 24$. We consider two different cubed-sphere grids on this domain, constructed using the mappings detailed in Appendix A. As in the spherically symmetric case, we adapt these grids to the geometry by using larger, higher-order elements in the star’s interior as well as outside the star. In the region near the surface, the grids are composed of thin cubed-sphere shells with a linear basis in the radial direction. The first grid, B1, has a similar resolution across the radius of the star to the I1 grid used in the 1D evolutions; an equatorial cut through this grid is shown in Fig. 14. The second grid, B1R, is obtained by adaptively refining B1: the large elements in the interior and exterior regions are p -refined; the thin elements near the surface of the star are h -refined, i.e. are radially split into thinner shells. The complete description of these two grid structures is given in Appendix B. In these 3D evolutions, we apply as before the minmod limiter to the surface region of the grid only. We evolve the hydro system until $t = 10^4 \simeq 50$ ms, with time steps $\Delta t = 0.04$ on the B1 grid ($\Delta t/\Delta x_{\min} \simeq 0.61$), and $\Delta t = 0.025$ on the B1R grid ($\Delta t/\Delta x_{\min} \simeq 0.59$).

When evolving the star on these cubed-sphere grids, we find a numerical instability in the conserved momentum \tilde{S}_i that leads to an exponential growth of this quantity on $\mathcal{O}(100M_{\odot})$ timescales. This numerical instability is caused by aliasing of the spectral modes as a result of an insufficiently resolved quadrature rule in the DG method (see paragraph below (11)); we therefore filter the \tilde{S}_i variable. In the central (as described in Appendix B) portion

of the grid, the filter takes the form (33) with $\alpha = 36$ and $s = 6^7$. In the cubed-sphere shells that make up the bulk of the interior, the numerical instability is weaker, and is controlled by a milder filter with $\alpha = 36$ and $s = 12$. With these filters, the NS evolutions remain stable until at least $t = 10^4$. As the stars presented in this paper have a rest state with no velocity, i.e. $\tilde{S}_i = 0$, with dynamics that consist primarily of short-timescale oscillations while the system settles to the rest state, the filters have only a minor effect on the long-term evolution. For stars undergoing rotation or pronounced dynamics, the filtering should not qualitatively affect the results, but would reduce the method’s order of convergence.

We plot, as before, the errors in \tilde{D} , \tilde{S}_i , and ρ_c during the first $4000M_{\odot}$ of evolution time in Fig. 15. The errors in the B1 simulation closely match those seen in the spherically symmetric case (cf. the I1 curves in Fig. 12) — this is expected, given the comparable resolution and the spherically symmetric nature of the problem. We note that the gradual decrease in central density seen in the 1D simulation is not observed in 3D, and the central density instead approaches a constant as the star settles. The errors in the B1R case are reduced by about an order of magnitude. Note that this case cannot be directly compared to the high-resolution 1D case I1R, which has a substantially higher resolution across the NS radius. In the full evolution to $t = 10^4$, for both grids, the oscillations damp away and the errors tend towards a constant equilibrium value. We again check the conservation error by tracking the baryon mass, which we find to slowly grow during the evolution. The relative error in M_b in the B1 evolution is comparable to the I1 case, reaching about 10^{-4} at $t = 10^4$. For the B1R case the relative error is about 6×10^{-6} . Again, most of these drifts are accumulated during the initial settling of the star.

We compute the frequency spectrum of the stellar oscillations from ρ_c , using the procedure described for the spherically symmetric case. The results are shown in Fig. 16. Comparing the B1 spectrum to the I1 spectrum from Fig. 13, we see good agreement: the first two resonant frequencies are clearly resolved, and additional peaks at higher frequencies are suggestive but not conclusive. Going from B1 to B1R we see improvement in the mode resolution, with a third and fourth peak appearing in the frequency spectrum. These new peaks are increasingly shifted towards higher frequencies, which indicates that the corresponding modes are not yet fully resolved. The intermediate peaks seen in 1D are still visible in 3D, but remain close to the noise level.

We also performed (but do not show) simulations of the 3D Cowling NS using the same grid and limiter configurations that, in the 1D study, were found to be problematic. These configurations are: using a grid where the

⁷ With these values, the filter reduces power in approximately the upper half of the modes. The power in the highest mode is reduced to round-off.

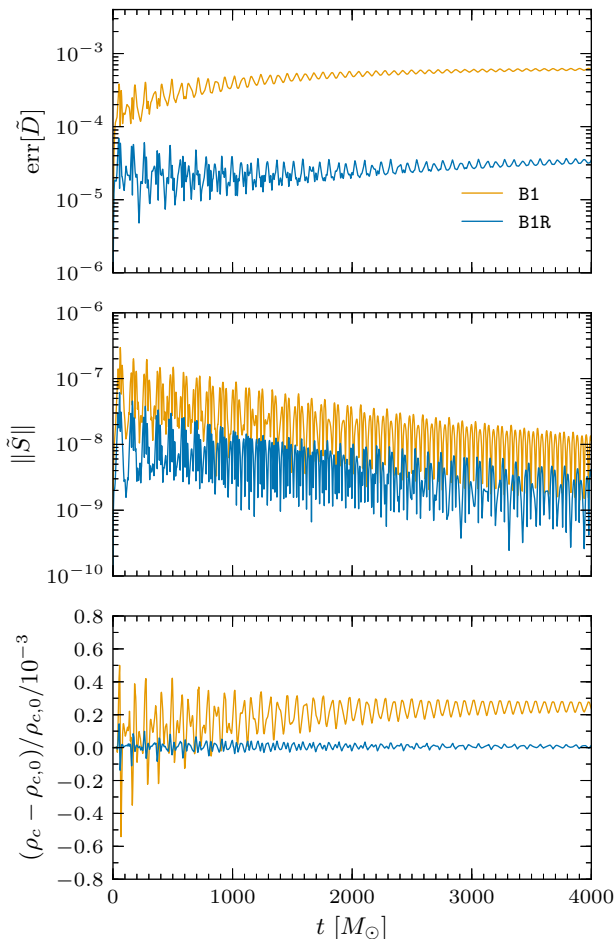


FIG. 15. The errors in the 3D Cowling NS evolution. The top (middle) panel shows the error in the conserved density \tilde{D} (conserved momentum \tilde{S}_i) for evolutions using the minmod limiter on the grids **B1** and **B1R**. The bottom panel shows the evolution of the central density ρ_c as a fractional error with respect to its initial value $\rho_{c,0}$.

surface elements have a quadratic radial basis, and/or using the MRS limiter near the NS surface. For the grid check, we employed a third cubed-sphere grid on the domain, **B2**, that is similar to **B1** but uses thicker shells with a quadratic radial basis (comparable to the **I2** grid in 1D; see Appendix B for details). We found, as in the 1D case, that evolutions on this grid using the low-order minmod limiter are stable over long timescales, but with high dissipation and increased error. For the MRS limiter check, we found, in contrast to the 1D case, that the 3D simulations are unstable on $\mathcal{O}(1000M_\odot)$ timescales: the high-frequency oscillations seen in the 1D case grow in 3D, presumably due to the limiter’s inability to control the additional tangential basis modes, and the star rapidly becomes unstable.




FIG. 16. The Fourier transform of the central rest-mass density ρ_c from the 3D Cowling NS evolution. The data from evolutions on the **B1** and **B1R** grids are shown — the **B1R** curve is shifted downwards on the plot, by a factor of 100, for visual clarity. The vertical dotted lines indicate the frequencies of the fundamental normal mode and the first six harmonics. The units of the vertical axis are arbitrary.

C. GR-hydro neutron star

For the coupled GR-hydro evolutions, we again use the two grids **B1** and **B1R** from above. The hydrodynamics are treated as for the Cowling star, with a minmod limiter at the star surface. We additionally evolve the spacetime geometry, with the constraint damping parameters set to

$$\gamma_0 = 0.1 \exp[-(r/12)^2/2] + 0.01 \quad (46)$$

$$\gamma_1 = -1 \quad (47)$$

$$\gamma_2 = 3 \exp[-(r/12)^2/2] + 0.01. \quad (48)$$

The gauge function H_σ is computed, as for the Kerr BH evolution, from the contraction of the Christoffel symbols of the exact metric; it is constant in time. We evolve the combined system until $t = 10^4 \simeq 50$ ms, with time steps $\Delta t = 0.04$ on the **B1** grid ($\Delta t/\Delta x_{\min} \simeq 0.61$), and $\Delta t = 0.025$ on the **B1R** grid ($\Delta t/\Delta x_{\min} \simeq 0.59$).

We show in Fig. 17 the errors in \tilde{D} , \tilde{S}_i , and ρ_c for the self-consistent NS evolution. Comparing the results from the **B1** grid to the Cowling results of Fig. 15, we see that the self-consistent NS is more dissipative than the Cowling one — the oscillations decay quickly and become negligible by $t \sim 3000$. Additionally, we see that the star settles to a different equilibrium, because the gravity responds to the fluid rather than providing a fixed potential well. The equilibrium central density is higher than the TOV value, indicating that in its numerical equilibrium, the star has compressed slightly. The errors using the higher-resolution grid **B1R** are, as in the Cowling case, significantly reduced as compared to the **B1** grid. In the




FIG. 17. The errors in the coupled GR-hydro NS evolution. The top (middle) panel shows the error in the conserved density \tilde{D} (conserved momentum \tilde{S}_i) for evolutions using the minmod limiter on the grids **B1** and **B1R**. The bottom panel shows the evolution of the central density ρ_c as a fractional error with respect to its initial value $\rho_{c,0}$.

full evolution to $t = 10^4$, the **B1R** case exhibits a slowly-growing error component at late times: from $t \simeq 7000$ onwards the errors increase by order 10%. This growing error is consistent with a weak numerical instability over $\mathcal{O}(10^4 M_\odot)$ timescales, and could presumably be addressed by improved filtering. We again check the conservation error by tracking the baryon mass during the evolution. The relative error in M_b for the **B1** case is about 10^{-4} at $t = 10^4$, as in the Cowling case, and as before is mostly accumulated during the initial settling of the star. In the **B1R** case, however, the relative error reaches about 2×10^{-5} , or three times the value from the Cowling case. Here the error is accumulated in two phases: first during the initial settling of the star, and again at the end of the evolution when the slowly-growing errors begin to affect the computation of M_b .

We compute once more the frequency spectrum of the stellar oscillations from ρ_c , and we show in Fig. 18 the re-




FIG. 18. The Fourier transform of the central rest-mass density ρ_c from the coupled GR-hydro NS evolutions. Results are shown for evolutions on the **B1** and **B1R** grids — the **B1R** curve is shifted downwards on the plot, by a factor of 100, for visual clarity. The vertical dotted lines indicate the frequencies of the fundamental normal mode and the first six harmonics. The units of the vertical axis are arbitrary.

sults from the evolutions on both grids. We also indicate the first 7 radial eigenmode frequencies from linear theory⁸ by the vertical dotted lines. In the lower-resolution **B1** case, we see clear peaks in the spectrum from the fundamental mode up through the 6th harmonic. The first three of these peaks are sharpest, indicating well-resolved modes; the subsequent peaks become gradually less prominent and increasingly shifted towards higher frequencies. The **B1R** case is qualitatively similar — we see the same 7 peaks in the spectrum, and although they are more prominent than in the **B1** case because the noise floor is lower, the shift towards high frequencies persists. Compared to the Cowling case, more modes are resolved. We also note that the intermediate peaks seen in the Cowling case are no longer prominent in the full GR-hydro results.

We conclude our analysis by comparing the accuracy of the DG and FV methods for the NS problem. We use the SpEC hydro code — a FV code that takes a dual-grid approach for coupled GR-hydro problems — to perform additional evolutions of the NS. The spacetime is evolved on a high-resolution grid of nested spherical shells using a pseudo-spectral penalty method, closely related to the DG method presented in this paper. The matter is evolved on a Cartesian grid covering the interval $[0, 12]$ in each direction (octant symmetry is imposed), using a 4th-order finite difference scheme with a WENO

⁸ These eigenfrequencies were kindly provided to us by David Radice; see the discussion of Fig. 11 of [21] for details.




FIG. 19. The central-density error in the coupled GR-hydro NS, for evolutions using the DG and FV methods. For each method, two resolutions are shown: the DG method uses the grids B1 (base) and B1R (refined, with $2\times$ as many grid points); the FV method uses the grids FV (base) and FVR (refined, with $8\times$ as many grid points).

reconstructor. We consider two resolutions for the FV grid. For the base resolution, we require that the FV grid have the same number of grid points within the volume of the star as the B1 grid of the DG evolution. This corresponds to a grid of 51^3 points on $[0, 12]^3$. The high-resolution grid uses 101^3 points — far more than B1R. These two cases are labelled FV and FVR respectively.

In Fig. 19 we compare the central density errors in evolutions with the DG and FV methods. The DG results make use of the grids B1 and B1R (a $2\times$ increase in the number of grid points), and the FV results make use of FV and FVR (an $8\times$ increase) described above. Comparing the results, we find a few differences between the DG and FV evolutions. First, the DG method is more dissipative than the FV method used, with the star’s oscillations damping away by $t \sim 3000$. A contributing factor to the increased dissipation is the use of a low-order shock-capturing scheme in the surface regions for the DG evolutions vs. the high-order reconstruction scheme of the FV method. Second, the error in the central density is greatly reduced in the DG evolution, primarily because of the negligible drift rate after the star has settled to its numerical equilibrium. Finally, in going to the refined B1R and FVR grids, we find that the error decreases more rapidly in the DG case even though the resolution change is smaller. This is because the DG method has higher order in the bulk of the star’s interior, so that p -refinement leads to rapid convergence. Precise statements about the order of convergence for the DG results are difficult to make, however, because we use geometrically adapted grids with elements of different order.

VII. CONCLUSIONS

In this paper we have presented 3D evolutions using a DG method of (a) a Kerr BH, and (b) a general-relativistic NS treated self-consistently. We adopted the DG formulation of Teukolsky [30] to solve the generalized harmonic formulation of Einstein’s equations and the València formulation of general-relativistic hydrodynamics. We used conforming grids to take advantage of the problem geometries, and we evolved the spacetime and matter simultaneously on these grids. We implemented the DG method in the SpEC framework and showed convergence and shock-capturing tests for our code. We also evolved a NS under the Cowling approximation (fixed spacetime metric) in spherical symmetry and in 3D.

With the 3D Kerr BH evolution, we showed that the DG method is accurate and stable for long-timescale spacetime evolutions. By adapting the grid to the (nearly) spherical geometry of the BH spacetime, we were able to excise the singularity from the domain — a promising result for the future use of the DG method in compact-object binary simulations. The success of the DG method here draws on previous successes of the (closely-related) spectral penalty method for the BH problem.

For the NS, we again showed long and stable evolutions and we additionally recovered the eigenfrequencies from linearized theory. By using domains conforming to the star’s spherical geometry and adapted to resolve the surface, we were able to obtain good accuracy with comparatively few elements and a low-order shock-capturing scheme. We compared the DG evolution to a FV evolution and found significantly lower errors and improved rate of convergence from the DG case.

One of the advantages of the DG method over the FV method is that it is easier to scale the algorithm on large machines. However, we were not able to show scaling results from our implementation in SpEC. As discussed, the SpEC framework scales poorly to large numbers of elements. For the NS results shown, the domains are composed of over 5000 elements, enough for SpEC’s scaling to break down and for timing measurements to lose their significance. We do note that our DG method, which uses the same grid for the spacetime geometry and the matter, solves the Einstein equations on a denser grid of points than the dual-grid SpEC hydro code. This adds a significant computational cost for the runs presented in this paper. The added cost would be reduced in the context of a science-producing simulation with a spacetime grid extending to large radii, as the addition of some extra grid points in the central portion of the domain would be less significant.

Improvements to our work will include the adoption of higher-order shock-capturing schemes (e.g. WENO) to lower the errors in the treatment of the star surface. The development of an adaptive mesh-refinement scheme will allow geometrically adapted grids to be used in systems with reduced symmetry and/or dynamics. These im-

provements are planned for implementation in the SpECTRE code, where they will enable evolutions with the DG method of dynamical systems such as rotating or dynamically unstable stars.

ACKNOWLEDGMENTS

We thank Andy Bohn, Mike Boyle, Nils Deppe, Matt Duez, Francois Foucart, Jan Hesthaven, Curran Muhlberger, and Will Thrope for many helpful conversations through the course of this work. We gratefully acknowledge support for this research from the Sherman Fairchild Foundation, from NSF grants PHY-1606654 and AST-1333129 at Cornell, and from NSF grants PHY-1404569, PHY-1708212, and PHY-1708213 at Caltech. F.H. acknowledges support by the NSF Graduate Research Fellowship under Grant No. DGE-1144153. Computations were performed at Caltech on the Zwicky cluster, which is supported by the Sherman Fairchild Foundation and by NSF award PHY-0960291, and on the Wheeler cluster, which is supported by the Sherman Fairchild Foundation.


Appendix A: Cubed-sphere mappings

In simulations of systems with spherical geometries, we use grids based on the cubed sphere [64]. The cubed sphere is obtained by projecting the faces of a cube onto its circumscribed sphere, thereby defining a grid on the sphere composed of six deformed Cartesian grid patches. The radial direction is introduced by tensor product, giving a grid on a hollow spherical shell composed of six mapped cubes; we call each of these mapped cubes a wedge of the spherical shell. For our NS simulations, however, we require a filled ball topology, rather than a hollow spherical shell.

To obtain a grid on the filled ball, one possibility is to insert a cube-shaped element at the center of the grid and deform the inner surface of the spherical shell so that it conforms to this cube. This is shown in panel (a) of Fig. 20. In numerical experiments, we find that this grid configuration often suffers from large errors along the diagonal edges where three of the wedges meet, (e.g. the line $x = y = z$) because of the large grid distortions at these locations. This source of error can be reduced by inserting a “rounded” cube, which reduces the grid distortion in the wedges, as shown in panel (b) of Fig. 20. As we are not aware of previous uses of such a grid configuration, we show here the mappings used.

1. Wedges

The geometry of a cubed-sphere wedge is specified by its inner and outer surfaces. Each of these surfaces is described by two parameters — its curvature c and its



(a) Undeformed cube

(b) Rounded cube

FIG. 20. Two grids on the filled ball, constructed from a cubed sphere with (a) an undeformed cube as central element, and (b) a rounded cube as central element. Both panels show an equatorial cut through the grid. The grids are obtained from the mappings described in Appendix A, with parameters $c_{\min} = 0$ in panel (a), and $c_{\min} = 0.66$ in panel (b); in both panels $c_{\max} = 1$, $x_{\min} = 0.75$, and $x_{\max} = 2$. The black lines show the element boundaries, the light grey lines show the Gauss Legendre Lobatto grid within each element for order $N = 5$.

position x . The surface’s curvature $c \in [0, 1]$ controls the shape: when $c = 0$ the surface is flat (i.e. the 6 wedges together form a cube); when $c = 1$ the surface is spherical. The surface’s position is the “radius” from the origin to the center of the surface, i.e. the point where the surface intersects the $x/y/z$ -axis. The positions x_{\min} and x_{\max} of the inner and outer surface satisfy $0 < x_{\min} < x_{\max}$.

The mapping from the reference element to each cubed-sphere wedge is a radial interpolation between the wedge’s inner and outer surfaces, and is computed by composing four transformations,

$$\mathbf{x}(\bar{\mathbf{x}}) = (\mathbf{x}_{\text{rot}} \circ \mathbf{x}_{\text{cs}} \circ \mathbf{x}_{\text{tan}} \circ \mathbf{x}_{\text{affine}})(\bar{\mathbf{x}}). \quad (\text{A1})$$

The actions of these transformations are:

1. $\mathbf{x}_{\text{affine}}$ shifts and scales the reference cube along the $+x$ axis to obtain a parallelepiped spanning $0 < x_{\min} \leq x \leq x_{\max}$. The y and z coordinates are unaffected.
2. \mathbf{x}_{tan} maps the tangential coordinates y and z according to $y \rightarrow \tan(\pi y/4)$, and likewise for z . This shifts the grid point distribution tangentially inwards to produce a more uniform, equiangular grid when the destination surface is spherical. This transformation is optional; we use it for the spherical-shell grids of the Kerr BH and spherical accretion tests, but elsewhere, we omit it.
3. \mathbf{x}_{cs} deforms the parallelepiped into one wedge of the cubed sphere, intersecting the $+x$ axis at x_{\min} and x_{\max} . It is computed with the intermediate

steps

$$a = 1/\sqrt{1 + \bar{y}^2 + \bar{z}^2} \quad (\text{A2})$$

$$b_{\min} = x_{\min} [1 + c_{\min}(a - 1)] \quad (\text{A3})$$

$$b_{\max} = x_{\max} [1 + c_{\max}(a - 1)] \quad (\text{A4})$$

$$\xi = b_{\min} + (b_{\max} - b_{\min}) \frac{\bar{x} - x_{\min}}{x_{\max} - x_{\min}} \quad (\text{A5})$$

$$\mathbf{x}_{\text{cs}}(\bar{\mathbf{x}}) = (\xi, \xi\bar{y}, \xi\bar{z}). \quad (\text{A6})$$

4. \mathbf{x}_{rot} rotates the wedge to its location on the sphere, corresponding to one of the axes $+x$, $-x$, $+y$, $-y$, $+z$, or $-z$.

Fig. 20 shows two (filled) cubed-sphere grids where the outer surfaces are spherical and the inner surfaces have $c = 0$ and 0.66. Fig. 3 shows a cubed-sphere grid where both surfaces are spherical, and each wedge is divided radially and tangentially into several elements. This is achieved by dividing the unit cube into the corresponding elements before applying the chain of maps in (A1).

2. Rounded central cube

The mapping from the reference element to the rounded central cube is chosen to conform to the inner boundary of the cubed-sphere wedges. The cube is therefore parametrized by the x_{\min} and c_{\min} that give the inner boundary of the wedges, and by whether the equiangular transformation is applied. The mapping is again obtained by composition,

$$\mathbf{x}(\bar{\mathbf{x}}) = (\mathbf{x}_{\text{rc}} \circ \mathbf{x}_{\text{tan}})(\bar{\mathbf{x}}), \quad (\text{A7})$$

with \mathbf{x}_{rc} , the transformation that deforms the cube, given by,

$$a = 1/\sqrt{1 + \bar{x}^2\bar{y}^2 + \bar{x}^2\bar{z}^2 + \bar{y}^2\bar{z}^2 - \bar{x}^2\bar{y}^2\bar{z}^2} \quad (\text{A8})$$

$$b_{\min} = x_{\min} [1 + c_{\min}(a - 1)] \quad (\text{A9})$$

$$\mathbf{x}_{\text{rc}}(\bar{\mathbf{x}}) = (b_{\min}\bar{x}, b_{\min}\bar{y}, b_{\min}\bar{z}). \quad (\text{A10})$$

Inverting this mapping for $\bar{\mathbf{x}} = \mathbf{x}_{\text{rc}}^{-1}(\mathbf{x})$ requires root finding and is done numerically.

The right panel of Fig. 20 shows a cubed-sphere grid with a rounded central cube. Fig. 14 shows a rounded cube as used in the NS simulation grids; just as for the wedges, the division of the central cube into several elements is achieved by dividing the unit cube prior to applying the chain of maps in (A7).

Appendix B: Neutron star simulation grids

The simulation domain for the 3D NS evolutions is a filled ball extending to $r_{\max} = 24$. We use three different cubed-sphere grids on this domain: B1, B2, and B1R. Here we define each of these grids.

TABLE II. The radial structure of the 3D NS grids, B1, B2, and B1R. For each region of each grid, the location and curvature of the surfaces that bound the cubed-sphere elements are given. Duplicated information is omitted: the unspecified regions of B2 are identical to those of B1.

	x_i	c_i	N_r	
B1	center	1.3, 1.9, 2.5	0.55, 0.85, 1	4
	int.	2.5, 3.0, 3.6, 4.33, 5.2, 6.24, 7.5	1	4
	surf.	7.5, 7.75, 8, ..., 9.5, 9.75, 10	1	1
	ext.	10, 12, 15, 18, 21, 24	1	3
B2	surf.	7.5, 8, 8.5, 9, 9.5, 10	1	2
B1R	center	(see B1)	(see B1)	5
	int.	(see B1)	(see B1)	5
	surf.	7.5, 7.625, 7.75, ..., 9.875, 10	1	1
	ext.	(see B1)	(see B1)	4

The B1 grid is shown in Fig. 14. For the bulk of the stellar interior, the grid is composed of nested, spherical cubed-sphere shells containing higher-order elements. In the center of the domain, the grid transitions to a rounded cube using the mappings described in Appendix A. In the region near the surface, we use thinner elements with fewer points; for B1 there are 10 shells of thickness $\Delta r = 0.25$, each of which contains elements whose basis in the radial direction is linear (we denote this as $N_r = 1$). Outside the star, the grid is again made up of larger, higher-order elements. The details of this radial structure are given in Table II, which lists the parameters of the cubed-sphere shells that make up the grid. The angular structure of B1 is obtained by tangentially splitting each wedge into 6×6 elements, each of which has a basis of order $N_{\text{tan}} = 3$ in the two tangential directions. The resolution of the central rounded cube is set by conforming to the angular grid of the shell. The equiangular tangent mapping is *not* applied — omitting this mapping gives a more optimal resolution of the cube in the center of the star. The B1 grid has a total of 5184 elements, with $\Delta x_{\min} \simeq 0.0657$.

The B2 grid differs from B1 in the radial resolution of the surface region. Where B1 uses 10 shells of linear order, B2 instead uses 5 spherical shells, of thickness $\Delta r = 0.5$, each of which contains elements whose basis in the radial direction is quadratic (i.e., $N_r = 2$). The B2 grid has a total of 4104 elements, and the same Δx_{\min} as B1.

The B1R grid is obtained from B1 by selectively refining to further take advantage of the hp -adaptivity of the DG method: h -refinement is used in the neighborhood of the surface where the solution is not smooth, and p -refinement is used in the smooth interior and exterior regions. The radial parameters are again given in Table II; the angular parameters are as for B1 but with $N_{\text{tan}} = 4$. This grid has 7344 elements, with $\Delta x_{\min} \simeq 0.0447$, and has roughly twice as many grid points inside the NS ($r \lesssim 8.125$) as the B1 grid.

-
- [1] J. A. Font, *Living Rev. Rel.* **11**, 7 (2008).
- [2] J. M. Martí and E. Müller, *Living Rev. Comput. Astrophys.* **1**, 3 (2015).
- [3] D. S. Balsara, *Living Rev. Comput. Astrophys.* **3**, 2 (2017), arXiv:1703.01241 [astro-ph.IM].
- [4] M. D. Duez, F. Foucart, L. E. Kidder, H. P. Pfeiffer, M. A. Scheel, and S. A. Teukolsky, *Phys. Rev. D* **78**, 104015 (2008), arXiv:0809.0002 [gr-qc].
- [5] P. Colella and P. R. Woodward, *J. Comput. Phys.* **54**, 174 (1984).
- [6] A. Mignone, T. Plewa, and G. Bodo, *Astrophys. J. Suppl. Ser.* **160**, 199 (2005), astro-ph/0505200.
- [7] G.-S. Jiang and C.-W. Shu, *J. Comput. Phys.* **126**, 202 (1996).
- [8] S. W. Bruenn, A. Mezzacappa, W. R. Hix, E. J. Lentz, O. E. B. Messer, E. J. Lingerfelt, J. M. Blondin, E. Endeve, P. Marronetti, and K. N. Yakunin, *Astrophys. J. Lett.* **767**, L6 (2013), arXiv:1212.1747 [astro-ph.SR].
- [9] P. Mösta, C. D. Ott, D. Radice, L. F. Roberts, R. Haas, and E. Schnetter, *Nature* **528**, 376 (2015), arXiv:1512.00838 [astro-ph.HE].
- [10] K. Kiuchi, K. Kyutoku, Y. Sekiguchi, M. Shibata, and T. Wada, *Phys. Rev. D* **90**, 041502 (2014), arXiv:1407.2660 [astro-ph.HE].
- [11] S. Bernuzzi and T. Dietrich, *Phys. Rev. D* **94**, 064062 (2016), arXiv:1604.07999 [gr-qc].
- [12] F. Foucart, R. Haas, M. D. Duez, E. O'Connor, C. D. Ott, L. Roberts, L. E. Kidder, J. Lippuner, H. P. Pfeiffer, and M. A. Scheel, *Phys. Rev. D* **93**, 044019 (2016), arXiv:1510.06398 [astro-ph.HE].
- [13] J. C. McKinney, A. Tchekhovskoy, A. Sadowski, and R. Narayan, *Mon. Not. Roy. Astr. Soc.* **441**, 3177 (2014), arXiv:1312.6127 [astro-ph.CO].
- [14] O. Porth, H. Olivares, Y. Mizuno, Z. Younsi, L. Rezzolla, M. Moscibrodzka, H. Falcke, and M. Kramer, *Computational Astrophysics and Cosmology* **4**, 1 (2017), arXiv:1611.09720 [gr-qc].
- [15] B. J. Kelly, J. G. Baker, Z. B. Etienne, B. Giacomazzo, and J. Schnittman, *Phys. Rev. D* **96**, 123003 (2017), arXiv:1710.02132 [astro-ph.HE].
- [16] G. Zumbusch, *Class. Quantum Grav.* **26**, 175011 (2009), arXiv:0901.0851 [gr-qc].
- [17] S. E. Field, J. S. Hesthaven, S. R. Lau, and A. H. Mroue, *Phys. Rev. D* **82**, 104051 (2010), arXiv:1008.1820 [gr-qc].
- [18] J. D. Brown, P. Diener, S. E. Field, J. S. Hesthaven, F. Herrmann, A. H. Mroué, O. Sarbach, E. Schnetter, M. Tiglio, and M. Wagman, *Phys. Rev. D* **85**, 084004 (2012), arXiv:1202.1038 [gr-qc].
- [19] J. M. Miller and E. Schnetter, *Class. Quantum Grav.* **34**, 015003 (2017), arXiv:1604.00075 [gr-qc].
- [20] M. Dumbser, F. Guercilena, S. Köppel, L. Rezzolla, and O. Zanotti, *Phys. Rev.* **D97**, 084053 (2018), arXiv:1707.09910 [gr-qc].
- [21] D. Radice and L. Rezzolla, *Phys. Rev. D* **84**, 024010 (2011), arXiv:1103.2426 [gr-qc].
- [22] J. Zhao and H. Tang, *J. Comput. Phys.* **242**, 138 (2013).
- [23] M. Bugner, T. Dietrich, S. Bernuzzi, A. Weyhausen, and B. Brügmann, *Phys. Rev. D* **94**, 084004 (2016), arXiv:1508.07147 [gr-qc].
- [24] L. E. Kidder, S. E. Field, F. Foucart, E. Schnetter, S. A. Teukolsky, A. Bohn, N. Deppe, P. Diener, F. Hébert, J. Lippuner, J. Miller, C. D. Ott, M. A. Scheel, and T. Vincent, *J. Comput. Phys.* **335**, 84 (2017), arXiv:1609.00098 [astro-ph.HE].
- [25] P. Anninos, C. Bryant, P. C. Fragile, A. M. Holgado, C. Lau, and D. Nemergut, *Astrophys. J. Suppl. Ser.* **231**, 17 (2017), arXiv:1706.09939 [astro-ph.IM].
- [26] F. Fambri, M. Dumbser, S. Köppel, L. Rezzolla, and O. Zanotti, *Mon. Not. Roy. Astr. Soc.* (2018), 10.1093/mnras/sty734, arXiv:1801.02839 [physics.comp-ph].
- [27] S. Köppel, ArXiv e-prints (2017), arXiv:1711.08221 [gr-qc].
- [28] F. Hébert, *Exploring Relativistic Gravity with Numerical Simulations*, Ph.D. thesis, Cornell University (2017).
- [29] M. Bugner, *Discontinuous Galerkin Methods for General Relativistic Hydrodynamics*, Ph.D. thesis, Friedrich-Schiller-Universität Jena (2018).
- [30] S. A. Teukolsky, *J. Comput. Phys.* **312**, 333 (2016), arXiv:1510.01190 [gr-qc].
- [31] H. Friedrich, *Commun. Math. Phys.* **100**, 525 (1985).
- [32] F. Pretorius, *Class. Quantum Grav.* **22**, 425 (2005), gr-qc/0407110.
- [33] C. Gundlach, G. Calabrese, I. Hinder, and J. M. Martin-Garcia, *Class. Quantum Grav.* **22**, 3767 (2005), gr-qc/0504114.
- [34] <http://www.black-holes.org/SpEC.html>.
- [35] M. A. Scheel, H. P. Pfeiffer, L. Lindblom, L. E. Kidder, O. Rinne, and S. A. Teukolsky, *Phys. Rev. D* **74**, 104006 (2006), gr-qc/0607056.
- [36] B. Szilágyi, L. Lindblom, and M. A. Scheel, *Phys. Rev. D* **80**, 124010 (2009), arXiv:0909.3557 [gr-qc].
- [37] D. A. Hemberger, M. A. Scheel, L. E. Kidder, B. Szilágyi, G. Lovelace, N. W. Taylor, and S. A. Teukolsky, *Class. Quantum Grav.* **30**, 115001 (2013), arXiv:1211.6079 [gr-qc].
- [38] R. Haas, C. D. Ott, B. Szilágyi, J. D. Kaplan, J. Lippuner, M. A. Scheel, K. Barkett, C. D. Muhlberger, T. Dietrich, M. D. Duez, F. Foucart, H. P. Pfeiffer, L. E. Kidder, and S. A. Teukolsky, *Phys. Rev. D* **D93**, 124062 (2016), arXiv:1604.00782 [gr-qc].
- [39] J. Hesthaven and T. Warburton, *Nodal Discontinuous Galerkin Methods: Algorithms, Analysis, and Applications* (Springer, Berlin, New York, 2008).
- [40] D. A. Kopriva, *Implementing spectral methods for partial differential equations: Algorithms for scientists and engineers* (Springer, Berlin, New York, 2009).

- [41] J. Hesthaven, S. Gottlieb, and D. Gottlieb, *Spectral Methods for Time-Dependent Problems* (Cambridge University Press, Cambridge, UK, 2007).
- [42] D. Gottlieb and J. S. Hesthaven, *J. Comput. Appl. Math.* **128**, 83 (2001).
- [43] L. Lindblom, M. A. Scheel, L. E. Kidder, R. Owen, and O. Rinne, *Class. Quantum Grav.* **23**, S447 (2006), gr-qc/0512093.
- [44] F. Banyuls, J. A. Font, J. M. Ibáñez, J. M. Martí, and J. A. Miralles, *Astrophys. J.* **476**, 221 (1997).
- [45] F. Galeazzi, W. Kastaun, L. Rezzolla, and J. A. Font, *Phys. Rev. D* **88**, 064009 (2013), arXiv:1306.4953 [gr-qc].
- [46] A. Harten, P. D. Lax, and B. van Leer, *SIAM Review* **25**, 35 (1983).
- [47] E. F. Toro, *Riemann solvers and Numerical Methods for Fluid Dynamics: A Practical Introduction* (Springer, Berlin, New York, 2013).
- [48] S. F. Davis, *SIAM Journal on Scientific and Statistical Computing* **9**, 445 (1988).
- [49] B. Cockburn, *J. Comput. Appl. Math.* **128**, 187 (2001).
- [50] S. A. Moe, J. A. Rossmann, and D. C. Seal, ArXiv e-prints (2015), arXiv:1507.03024 [math.NA].
- [51] X. Zhong and C.-W. Shu, *J. Comput. Phys.* **232**, 397 (2013).
- [52] J. Zhu, X. Zhong, C.-W. Shu, and J. Qiu, *Communications in Computational Physics* **19**, 944 (2016).
- [53] S. Gottlieb, C.-W. Shu, and E. Tadmor, *SIAM Review* **43**, 89 (2001).
- [54] M. Alcubierre *et al.*, *Class. Quant. Grav.* **21**, 589 (2004), arXiv:gr-qc/0305023 [gr-qc].
- [55] C. W. Misner, K. S. Thorne, and J. A. Wheeler, *Gravitation* (Freeman, New York, New York, 1973).
- [56] F. C. Michel, *Astrophysics and Space Science* **15**, 153 (1972).
- [57] B. Cockburn, S.-Y. Lin, and C.-W. Shu, *J. Comput. Phys.* **84**, 90 (1989).
- [58] J. Centrella and J. R. Wilson, *Astrophys. J. Suppl. Ser.* **54**, 229 (1984).
- [59] L. Del Zanna and N. Bucciantini, *Astron. Astrophys.* **390**, 1177 (2002), astro-ph/0205290.
- [60] R. C. Tolman, *Phys. Rev.* **55**, 364 (1939).
- [61] J. R. Oppenheimer and G. M. Volkoff, *Phys. Rev.* **55**, 374 (1939).
- [62] A. Schoepe, D. Hilditch, and M. Bugner, ArXiv e-prints (2017), arXiv:1712.09837 [gr-qc].
- [63] J. A. Font, T. Goodale, S. Iyer, M. Miller, L. Rezzolla, E. Seidel, N. Stergioulas, W.-M. Suen, and M. Tobias, *Phys. Rev. D* **65**, 084024 (2002), gr-qc/0110047.
- [64] C. Ronchi, R. Iacono, and P. S. Paolucci, *J. Comput. Phys.* **124**, 93 (1996).