Simulating magnetized neutron stars with discontinuous Galerkin methods

Nils Deppe,[1,*] François Hébert,[1] Lawrence E. Kidder,[2] William Throwe,[2] Isha Anantpurkar,[2] Cristóbal Armaza,[2] Gabriel S. Bonilla,[2,3] Michael Boyle,[2] Himanshu Chaudhary,[1] Matthew D. Duez,[4] Nils L. Vu,[5] Francois Foucart,[6] Matthew Giesler,[2] Jason S. Guo,[2] Yoonsoo Kim,[1] Prayush Kumar,[2,7] Isaac Legred,[1] Dongjun Li,[1] Geoffrey Lovelace,[3] Sizheng Ma,[1] Alexandra Macedo,[3] Denyz Melchor,[3] Marlo Morales,[3] Jordan Moxon,[1] Kyle C. Nelli,[1] Eamonn O'Shea,[2] Harald P. Pfeiffer,[5] Teresita Ramirez,[3] Hannes R. Rüter,[8] Jennifer Sanchez,[3] Mark A. Scheel,[1] Sierra Thomas,[3] Daniel Vieira,[2,9] Nikolas A. Wittek,[5] Tom Wlodarczyk,[5] and Saul A. Teukolsky[2,1]

[1]*Theoretical Astrophysics 350-17, California Institute of Technology, Pasadena, California 91125, USA*
[2]*Cornell Center for Astrophysics and Planetary Science, Cornell University, Ithaca, New York 14853, USA*
[3]*Nicholas and Lee Begovich Center for Gravitational-Wave Physics and Astronomy, California State University Fullerton, Fullerton, California 92834, USA*
[4]*Department of Physics & Astronomy, Washington State University, Pullman, Washington, D.C. 99164, USA*
[5]*Max Planck Institute for Gravitational Physics (Albert Einstein Institute), Am Mühlenberg 1, D-14476 Potsdam, Germany*
[6]*Department of Physics & Astronomy, University of New Hampshire, Durham, New Hampshire 03824, USA*
[7]*International Centre for Theoretical Sciences, Tata Institute of Fundamental Research, Bangalore 560089, India*
[8]*CFisUC, Department of Physics, University of Coimbra, 3004-516 Coimbra, Portugal*
[9]*I. Physikalisches Institut, Universität zu Köln, Zülpicher Straße 77, 50937 Köln, Germany*

(Received 29 September 2021; revised 6 May 2022; accepted 1 June 2022; published 27 June 2022)

Discontinuous Galerkin methods are popular because they can achieve high order where the solution is smooth, because they can capture shocks while needing only nearest-neighbor communication, and because they are relatively easy to formulate on complex meshes. We perform a detailed comparison of various limiting strategies presented in the literature applied to the equations of general relativistic magnetohydrodynamics. We compare the standard minmod/$\Lambda\Pi^N$ limiter, the hierarchical limiter of Krivodonova, the simple WENO limiter, the HWENO limiter, and a discontinuous Galerkin-finite-difference hybrid method. The ultimate goal is to understand what limiting strategies are able to robustly simulate magnetized Tolman-Oppenheimer-Volkoff stars without any fine-tuning of parameters. Among the limiters explored here, the only limiting strategy we can endorse is a discontinuous Galerkin-finite-difference hybrid method.

DOI: 10.1103/PhysRevD.105.123031

I. INTRODUCTION

Many of the most energetic phenomena in the universe involve matter under extreme gravitational conditions. These phenomena include neutron-star binary mergers, accretion onto black holes, and supernova explosions. For many of these systems, the motion of this matter is expected to generate extremely strong magnetic fields. The matter and magnetic fields in these systems are governed by the equations of general relativistic magnetohydrodynamics (GRMHD). These equations admit a rich variety of solutions, which often include large-scale relativistic flows and small-scale phenomena such as shocks and turbulence.

High-resolution shock capturing (HRSC) finite-difference (FD) methods are the current standard methods of choice for numerically evolving these solutions since they are able to robustly handle shocks. Unfortunately, HRSC FD methods have significant computational overhead and are less efficient than spectral-type methods like discontinuous Galerkin (DG) where the solution is smooth. Additionally, achieving better than second-order convergence is generally difficult, with recent results presented in [1,2]. The common use of second-order methods means that the current generation of GRMHD codes is not accurate enough to provide useful predictions for many extreme systems. Increasing simulation resolutions can improve this, but is very computationally expensive. More appealing is using numerical methods with higher convergence orders, which can increase accuracy with significantly less cost than a similar improvement from

*ndeppe@caltech.edu

resolution. Unfortunately, while higher-order methods handle smooth solutions very well, they are generally poor at handling discontinuities, such as fluid shocks, losing accuracy and sometimes failing completely.

Additionally, the time required to run simulations is already too long for most interesting astrophysical cases. The performance of individual computational processors has stagnated over the past decade, so to improve computational speed, codes must be parallelized over more processors. New supercomputer clusters will soon routinely have millions of cores. Codes designed for running on thousands of processors generally scale poorly to massively parallel setups, however. As problems are divided up into an increasing numbers of parts, the amount of communication required during the simulation can become prohibitive, particularly for high-order methods.

Discontinuous Galerkin methods [3–8], together with a task-based parallelization strategy, have the potential to deal with these problems. DG methods offer high-order accuracy in smooth regions, with the potential for robust shock capturing by some nonlinear stabilization technique. The methods are also well suited for parallelization: Their formulation in terms of local, nonoverlapping elements requires only nearest-neighbor communication regardless of the scheme's order of convergence. Additionally, these features allow for comparatively straightforward hp-adaptivity/adaptive mesh refinement and local time stepping, enabling better load distribution across a large number of cores.

Despite extensive success in engineering and applied mathematics communities over the past two decades, applications of DG in relativity [9–13] and astrophysics [14–17] have typically been exploratory or confined to simple problems. However, recently there have been significant advances toward production codes for nonrelativistic [18] and relativistic [19–21] hydrodynamics, special relativistic magnetohydrodynamics [16,22–26], the Einstein equations [13,27], and relativistic hydrodynamics coupled to the Einstein equations [28]. Most of these codes use the Message Passing Interface to implement a data parallelism strategy, though [20,23] use task-based parallelism.

In this paper we present a detailed comparison of various different limiting and shock capturing strategies for DG methods in the context of demanding GRMHD test problems. Specifically, we compare the classical limiters $\Lambda\Pi^N$ limiter [29], Krivodonova limiter [30], and WENO-based limiters [31,32], as well as a DG-FD hybrid method similar to that of [33,34]. The ultimate goal is to simulate a magnetized and nonmagnetized Tolman-Oppenheimer-Volkoff (TOV) star in the Cowling approximation. To the best of our knowledge this is the first time a magnetized TOV star has been simulated using DG methods. This paper presents a crucial first step to being able to apply DG methods to simulations of binary neutron star mergers, differentially rotating magnetized single neutron stars, and magnetized accretion disks.

While generally the classical limiters produce the best results when applied to the characteristic variables, these are not known analytically for GRMHD. Even though most test cases in this paper are in special relativity, we intentionally apply the limiters to the conserved variables to evaluate their performance in the form they need to be used for GRMHD. Since FD methods are also known to be less dissipative when applied to the characteristic variables, this choice does not put any of the limiters at a disadvantage.

The paper is organized as follows. Section II describes the formulation of GRMHD used in the problems presented here. Section III describes the algorithms used by our open-source code SpECTRE [35] to solve these equations. Results of the evolutions of a variety of GRMHD problems are presented in Sec. IV comparing the different shock capturing strategies.

II. EQUATIONS OF GRMHD

We adopt the standard $3+1$ form of the spacetime metric (see, e.g., [36,37]),

$$\begin{aligned} ds^2 &= g_{ab}dx^a dx^b \\ &= -\alpha^2 dt^2 + \gamma_{ij}(dx^i + \beta^i dt)(dx^j + \beta^j dt), \end{aligned} \tag{1}$$

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

SpECTRE currently solves equations in flux-balanced and first-order hyperbolic form. The general form of a flux-balanced conservation law in a curved spacetime is

$$\partial_t U + \partial_i F^i = S, \tag{2}$$

where U is the state vector, F^i are the components of the flux vector, and S is the source vector.

We refer the reader to the literature [36,38,39] for a detailed description of the equations of general relativistic magnetohydrodynamics (GRMHD). If we ignore self-gravity, the GRMHD equations constitute a closed system that may be solved on a given background metric. We denote the rest-mass density of the fluid by ρ and its 4-velocity by u^a, where $u^a u_a = -1$. The dual of the Faraday tensor F^{ab} is

$${}^*F^{ab} = \frac{1}{2}\epsilon^{abcd}F_{cd}, \tag{3}$$

where ϵ^{abcd} is the Levi-Civita tensor. Note that the Levi-Civita tensor is defined here with the convention [40] that in flat spacetime $\epsilon_{0123} = +1$. The equations governing the evolution of the GRMHD system are

$$\nabla_a(\rho u^a) = 0 \quad \text{(rest-mass conservation)}, \tag{4}$$

$$\nabla_a T^{ab} = 0 \quad \text{(energy-momentum conservation)}, \tag{5}$$

$$\nabla_a {}^*F^{ab} = 0 \quad \text{(homogeneous Maxwell equation)}. \tag{6}$$

In the ideal MHD limit the stress tensor takes the form

$$T^{ab} = (\rho h)^* u^a u^b + p^* g^{ab} - b^a b^b, \tag{7}$$

where

$$b^a = -{}^*F^{ab} u_b \tag{8}$$

is the magnetic field measured in the comoving frame of the fluid, and $(\rho h)^* = \rho h + b^2$ and $p^* = p + b^2/2$ are the enthalpy density and fluid pressure augmented by contributions of magnetic pressure $p_{\rm mag} = b^2/2$, respectively.

We denote the unit normal vector to the spatial hypersurfaces as n^a, which is given by

$$n^a = (1/\alpha, -\beta^i/\alpha)^T, \tag{9}$$

$$n_a = (-\alpha, 0, 0, 0). \tag{10}$$

The spatial velocity of the fluid as measured by an observer at rest in the spatial hypersurfaces ("Eulerian observer") is

$$v^i = \frac{1}{\alpha}\left(\frac{u^i}{u^0} + \beta^i\right), \tag{11}$$

with a corresponding Lorentz factor W given by

$$W = -u^a n_a = \alpha u^0 = \frac{1}{\sqrt{1 - \gamma_{ij} v^i v^j}} \tag{12}$$

$$= \sqrt{1 + \gamma^{ij} u_i u_j} = \sqrt{1 + \gamma^{ij} W^2 v_i v_j}. \tag{13}$$

The electric and magnetic fields as measured by a Eulerian observer are given by

$$E^i = F^{ia} n_a = \alpha F^{0i}, \tag{14}$$

$$B^i = -{}^*F^{ia} n_a = -\alpha {}^*F^{0i}. \tag{15}$$

Finally, the comoving magnetic field b^a in terms of B^i is

$$b^0 = \frac{W}{\alpha} B^i v_i, \tag{16}$$

$$b^i = \frac{B^i + \alpha b^0 u^i}{W}, \tag{17}$$

while $b^2 = b^a b_a$ is given by

$$b^2 = \frac{B^2}{W^2} + (B^i v_i)^2. \tag{18}$$

We now recast the GRMHD equations in a $3+1$ split by projecting them along and perpendicular to n^a [38]. One of the main complications when solving the GRMHD equations numerically is preserving the constraint

$$\partial_i(\sqrt{\gamma} B^i) = 0, \tag{19}$$

where $\gamma = \det(\gamma_{ij})$ is the determinant of the spatial metric. Analytically, initial data evolved using the dynamical Maxwell equations are guaranteed to preserve the constraint. However, numerical errors generate constraint violations that need to be controlled. We opt to use the generalized Lagrange multiplier or divergence cleaning method [41] where an additional field Φ is evolved in order to propagate constraint violations out of the domain. Our version is very close to the one in Ref. [42]. The augmented system can still be written in flux-balanced form, where the conserved variables are

$$U = \sqrt{\gamma}\begin{pmatrix} D \\ S_j \\ \tau \\ B^j \\ \Phi \end{pmatrix} = \begin{pmatrix} \tilde{D} \\ \tilde{S}_j \\ \tilde{\tau} \\ \tilde{B}^j \\ \tilde{\Phi} \end{pmatrix} = \sqrt{\gamma}\begin{pmatrix} \rho W \\ (\rho h)^* W^2 v_j - \alpha b^0 b_j \\ (\rho h)^* W^2 - p^* - (\alpha b^0)^2 - \rho W \\ B^j \\ \Phi \end{pmatrix}, \tag{20}$$

with corresponding fluxes

$$F^i = \begin{pmatrix} \tilde{D} v^i_{\rm tr} \\ \tilde{S}_j v^i_{\rm tr} + \alpha\sqrt{\gamma} p^* \delta^i_j - \alpha b_j \tilde{B}^i/W \\ \tilde{\tau} v^i_{\rm tr} + \alpha\sqrt{\gamma} p^* v^i - \alpha^2 b^0 \tilde{B}^i/W \\ \tilde{B}^j v^i_{\rm tr} - \alpha v^j \tilde{B}^i + \alpha\gamma^{ij}\tilde{\Phi} \\ \alpha\tilde{B}^i - \tilde{\Phi}\beta^i \end{pmatrix}, \tag{21}$$

and corresponding sources

$$S = \begin{pmatrix} 0 \\ (\alpha/2)\tilde{S}^{kl}\partial_j\gamma_{kl} + \tilde{S}_k\partial_j\beta^k - \tilde{E}\partial_j\alpha \\ \alpha\tilde{S}^{kl}K_{kl} - \tilde{S}^k\partial_k\alpha \\ -\tilde{B}^k\partial_k\beta^j + \Phi\partial_k(\alpha\sqrt{\gamma}\gamma^{jk}) \\ \alpha\tilde{B}^k\partial_k \ln\alpha - \alpha K\tilde{\Phi} - \alpha\kappa\tilde{\Phi} \end{pmatrix}. \tag{22}$$

The transport velocity is defined as $v^i_{\rm tr} = \alpha v^i - \beta^i$ and the generalized energy $\tilde{E}$ and source $\tilde{S}^{ij}$ are given by

$$\tilde{E} = \tilde{\tau} + \tilde{D}, \tag{23}$$

$$\tilde{S}^{ij} = \sqrt{\gamma}[(\rho h)^* W^2 v^i v^j + p^* \gamma^{ij} - \gamma^{ik}\gamma^{jl} b_k b_l]. \tag{24}$$

The $3+1$ GRMHD divergence cleaning evolution equations analytically preserve the constraint (19), while numerically constraint-violating modes will be damped at a rate κ. We typically choose $\kappa \in [0, 10]$, but will specify the exact value used for each test problem. We note that the divergence cleaning method was shown to be strongly hyperbolic in Ref. [43], a necessary condition for a well-posed evolution problem. The primitive variables of the GRMHD system are ρ, v_i, B^i, Φ, and the specific internal energy ϵ.

Approximate Riemann solvers use the characteristic speeds, which in the GRMHD case require solving a nontrivial quartic equation for the fast and slow modes. Instead, we use the approximation [44]

$$\lambda_1 = -\alpha - \beta_n, \tag{25}$$

$$\lambda_2 = \alpha\Lambda^- - \beta_n, \tag{26}$$

$$\lambda_{3,4,5,6,7} = \alpha v_n - \beta_n, \tag{27}$$

$$\lambda_8 = \alpha\Lambda^+ - \beta_n, \tag{28}$$

$$\lambda_9 = \alpha - \beta_n, \tag{29}$$

where β_n and v_n are the shift and spatial velocity projected along the normal vector in the direction that we want to compute the characteristic speeds along, and

$$\Lambda^\pm = \frac{1}{1 - v^2 c_s^2}\left[v_n(1 - c_s^2) \pm c_s\sqrt{(1 - v^2)(1 - v^2 c_s^2 - v_n^2(1 - c_s^2))}\right], \tag{30}$$

where c_s is the sound speed given by

$$c_s^2 = \frac{1}{h}\left[\left(\frac{\partial p}{\partial \rho}\right)_\epsilon + \frac{p}{\rho^2}\left(\frac{\partial p}{\partial \epsilon}\right)_\rho\right]. \tag{31}$$

III. METHODS

A. The discontinuous Galerkin method

We briefly summarize the nodal discontinuous Galerkin (DG) method for curved spacetimes [19] in d spatial dimensions. We decompose the computational domain into k elements, each with a reference coordinate system $\{\xi, \eta, \zeta\} \in [-1, 1]$. We denote the ith element by Ω_i, so our computational domain $\Omega = \cup_{i=1\ldots k}\,\Omega_i$. In this work we consider only dimension-by-dimension affine maps. We expand the solution in each element over a tensor product basis $\phi_{\check{s}}$ of 1d Lagrange polynomials $\ell_{\check{i}}$,

$$U(\xi) = \sum_{\check{s}} U_{\check{s}}(t)\phi_{\check{s}}(\xi) = \sum_{\check{i}}\sum_{\check{j}}\sum_{\check{k}} U_{\check{i}\check{j}\check{k}}(t)\ell_{\check{i}}(\xi)\ell_{\check{j}}(\eta)\ell_{\check{k}}(\zeta), \tag{32}$$

where ξ, η, and ζ are the logical (or reference) coordinates. We use Legendre-Gauss-Lobatto collocation points, though SpECTRE also supports Legendre-Gauss points. We denote a DG scheme with 1d basis functions of degree N by P_N. A P_N scheme is expected to converge at order $\mathcal{O}(\Delta x^{N+1})$ for smooth solutions [4], where Δx is the 1d size of an element.

A spatial discretization is obtained by integrating the evolution equations (2) against the basis functions $\phi_{\check{s}}$,

$$0 = \int_{\Omega_i}[\partial_t U + \partial_i F^i - S]\phi_{\check{s}}(\mathbf{x})d^3x = \int_{\Omega_i}[\partial_t U + \partial_i F^i - S]\phi_{\check{s}}(\xi)Jd^3\xi, \tag{33}$$

where J is the Jacobian determinant of the map from the reference coordinates ξ to the coordinates $\mathbf{x}$. Denoting the normal covector to the spatial boundary of the element as n_i, integrating the flux divergence term by-parts, replacing $F^i n_i$ with a boundary correction/numerical flux G, and undoing the integration by-parts, we obtain

$$\int_{\Omega_i}\partial_t(U)\,Jd^3\xi = \oint_{\partial\Omega_i}(G - F^i n_i)\phi_{\check{s}}(\xi)\,d^2\Sigma + \int_{\Omega_i}[\partial_i F^i + S]\phi_{\check{s}}(\xi)J\,d^3\xi, \tag{34}$$

where $d^2\Sigma$ is the area element on the surface of the element. The area element in the $+\zeta$ direction is given by [19]

$$d^2\Sigma = \frac{\sqrt{^{(2)}\gamma}}{\sqrt{\gamma}}d\xi^1 d\xi^2 = J\sqrt{\delta_{\hat{3}\hat{i}}\frac{\partial\xi^{\hat{i}}}{\partial x^i}\gamma^{ij}\frac{\partial\xi^{\hat{j}}}{\partial x^j}\delta_{\hat{3}\hat{j}}}d\xi^1 d\xi^2. \tag{35}$$

Note that the normalization of the normal vectors in the $G - F^i n_i$ term does not cancel out with the term in (35), as stated in [19]. This is because both the inverse spatial metric and the Jacobian may be different on each side of the boundary. Specifically, when the spacetime is evolved, each element normalizes the normal vector using its local inverse spatial metric.

Finally, the semidiscrete evolution equations are obtained by expanding U, F^i, and S in terms of the basis functions and evaluating the integrals by Gaussian quadrature. Our nodal

DG code uses the mass lumping approximation[1] when Gauss-Lobatto points are employed.

B. Numerical fluxes

One of the key ingredients in conservative numerical schemes is the approximate solution to the Riemann problem on the interface. We use the Rusanov solver [46] (also known as the local Lax-Friedrichs flux), and the solver of Harten, Lax, and van Leer (HLL) [47,48]. While both the Rusanov and the HLL solver are quite simple, their use is standard in numerical relativity. The Rusanov solver is given by

$$G^{\text{Rusanov}} = \frac{1}{2}(F^{k,+}n_k^+ + F^{k,-}n_k^-) - \frac{C}{2}(U^+ - U^-), \tag{36}$$

where $C = \max(|\lambda_i(U^+)|, |\lambda_i(U^-)|)$, and $\lambda_i(U)$ is the set of characteristic speeds. Quantities superscripted with a plus sign are on the exterior side of the boundary between an element and its neighbor, while quantities superscripted with a minus sign are on the interior side. In this section n_k is the outward pointing unit normal to the element.

The HLL solver is given by

$$G^{\text{HLL}} = \frac{\lambda_{\min}F^{k,+}n_k^+ + \lambda_{\max}F^{k,-}n_k^-}{\lambda_{\max} - \lambda_{\min}} - \frac{\lambda_{\max}\lambda_{\min}}{\lambda_{\max} - \lambda_{\min}}(U^+ - U^-), \tag{37}$$

where $\lambda_{\min}$ and $\lambda_{\max}$ are estimates for the fastest left- and right-moving signal speeds, respectively. We compute the approximate signal speeds pointwise using the scheme presented in Ref. [49]. Specifically,

$$\lambda_{\min} = \min(\lambda_i(U^+), \lambda_i(U^-), 0),$$
$$\lambda_{\max} = \max(\lambda_i(U^+), \lambda_i(U^-), 0). \tag{38}$$

C. Time stepping

SpECTRE supports time integration using explicit multistep and substep integrators. The results presented here were obtained using either a strong stability-preserving third-order Runge-Kutta method [4] or a self-starting Adams-Bashforth method. SpECTRE additionally supports local time stepping when using Adams-Bashforth schemes [50], but that feature was not used for any of these problems. The maximum admissible time step size for a P_N scheme is [51]

$$\Delta t \leq \frac{c}{d(2N+1)}\frac{\Delta x}{\lambda_{\max}}, \tag{39}$$

where c is a time-stepper-dependent constant, d is the number of spatial dimensions, Δx is the minimum 1d size (along each Cartesian axis) of the element, and $\lambda_{\max}$ is the maximum characteristic speed in the element.

D. Limiting

Near shocks, discontinuities, and stellar surfaces, the DG solution may exhibit spurious oscillations (i.e., Gibbs phenomenon) and overshoots. These oscillations can lead to a nonphysical fluid state (e.g., negative densities) at individual grid points and prevent stable evolution of the system. To maintain a stable scheme, some nonlinear limiting procedure is necessary. In general, we identify elements where the solution contains spurious oscillations (we label these elements as "troubled cells") and we modify the solution on these elements to reduce the amount of oscillation.

In this work we consider limiters that preserve the order of the DG solution while maintaining a compact (nearest-neighbor) stencil. The compact stencil greatly simplifies communication patterns, but, in order to provide the limiter with sufficient information to preserve the order of the scheme, it becomes necessary to send larger amounts of data from each element for each limiting step. We specifically consider

(i) the $\Lambda\Pi^N$ limiter of [29];
(ii) the hierarchical limiter of Krivodonova [30];
(iii) the simple WENO limiter of [31] (based on weighted essentially nonoscillatory, often abbreviated as WENO, finite volume methods);
(iv) the Hermite WENO (HWENO) limiter of [32];
(v) a DG-finite-difference hybrid scheme similar to that of [33,34].

Note that we do not use the limiter of Moe, Rossmanith, and Seal [52] because our experiments show that it is not very robust for the kinds of problems we study here.

Below we summarize the action of these limiters. Note that because computing the characteristic variables of the GRMHD system is complicated, we apply the limiters to the evolved (i.e., conserved) variables. However, we do not limit the divergence-cleaning field Φ, as it is not expected to form any shocks. The limiters are applied at the end of each time step when using an Adams-Bashforth method, and at the end of each substep when using a Runge-Kutta method.

1. $\Lambda\Pi^N$

The $\Lambda\Pi^N$ limiter [6,29,51,53] works by reducing the spatial slope of each variable U if the data look like they may contain oscillations. Specifically, if the slope exceeds a simple estimate based on differencing the cell average of U vs the neighbor elements' cell averages of U, then the limiter will linearize the solution and reduce its slope in a conservative manner. We use the total variation bounding (TVB) version of this limiter, which only activates if the

[1]"Mass lumping" is the term that describes using the diagonal approximation for the mass matrix. See [45] for more details.

slope is above mh^2, where m is the so-called TVB constant and h is the size of the DG element. This procedure is repeated independently for each variable component U being limited. While quite simple and robust, this limiter is very aggressive and can cause significant smearing of shocks and flattening of smooth extrema.

2. Krivodonova limiter

The Krivodonova limiter [30] works by limiting the coefficients of the solution's modal representation, starting with the highest coefficient then decreasing in order until no more limiting is necessary. This procedure is repeated independently for each variable component U being limited. Although the algorithm is only described in one or two dimensions, the limiting algorithm is straightforwardly generalized to our 3d application. We expand U over a basis of Legendre polynomials P_i,

$$U^{l,m,n} = \sum_{i,j,k=0,0,0}^{N,N,N} c_{i,j,k}^{l,m,n} P_i(\xi) P_j(\eta) P_k(\zeta), \tag{40}$$

where the $c_{i,j,k}^{l,m,n}$ are the modal coefficients, with the superscript $\{l,m,n\}$ representing the element indexed by l, m, n, and the upper bound N is the number of collocation points minus one in each of the ξ, η, ζ directions.

Each coefficient is limited by comparison with the coefficients of U in neighboring elements. The new value $\tilde{c}_{i,j,k}^{l,m,n}$ of $c_{i,j,k}^{l,m,n}$ is computed according to

$$\begin{aligned}\tilde{c}_{i,j,k}^{l,m,n} = \text{minmod}(&c_{i,j,k}^{l,m,n}, \alpha_i(c_{i-1,j,k}^{l+1,m,n} - c_{i-1,j,k}^{l,m,n}),\\ &\alpha_i(c_{i-1,j,k}^{l,m,n} - c_{i-1,j,k}^{l-1,m,n}), \alpha_j(c_{i,j-1,k}^{l,m+1,n} - c_{i,j-1,k}^{l,m,n}),\\ &\alpha_j(c_{i,j-1,k}^{l,m,n} - c_{i,j-1,k}^{l,m-1,n}), \alpha_k(c_{i,j,k-1}^{l,m,n+1} - c_{i,j,k-1}^{l,m,n}),\\ &\alpha_k(c_{i,j,k-1}^{l,m,n} - c_{i,j,k-1}^{l,m,n-1})),\end{aligned} \tag{41}$$

where minmod is the minmod function defined as

$$\begin{aligned}&\text{minmod}(a,b,\ldots)\\ &= \begin{cases} \text{sgn}(a)\,\min(|a|,|b|,\ldots), & \text{if } \text{sgn}(a) = \text{sgn}(b) = \text{sgn}(\ldots)\\ 0, & \text{otherwise,}\end{cases}\end{aligned} \tag{42}$$

and the α_i, α_j, α_k set the strength of the limiter. In all cases shown in this paper, we set $\alpha_i = 1$, at the least dissipative end of the range for these parameters.[2]

[2]Whereas Krivodonova [30] changes normalization convention for the Legendre polynomials in going from one to two dimensions, our convention matches their 1d convention in all cases, so that the range of the α_i parameters is given by Eq. (14) in the reference.

The algorithm for limiting from highest to lowest modal coefficient is as follows. We first compute $\tilde{c}_{N,N,N}$ (we drop the element superscripts here). If this is equal to $c_{N,N,N}$, no limiting is done. Otherwise, we update $c_{N,N,N} = \tilde{c}_{N,N,N}$, and compute the trio of coefficients $\tilde{c}_{N,N,N-1}$, $\tilde{c}_{N,N-1,N}$, $\tilde{c}_{N-1,N,N}$. If *all* of these are unchanged, the limiting stops. Otherwise, we update each coefficient and proceed to limiting all coefficients given by index permutations such that $i+j+k = 3N-2$, then $i+j+k = 3N-3$, etc. up to the three index permutations of $c_{1,0,0}$. Finally, the limited modal coefficients are used to recover the limited nodal values of the function U. Note that by not modifying $c_{0,0,0}$ the cell average is maintained.

3. Simple WENO

For the two WENO limiters, we use a troubled-cell indicator based on the TVB minmod limiter [6,51,53] to determine whether limiting is needed. When needed, each limiter uses a standard WENO procedure to reconstruct the local solution from several different estimated solutions.

In the simple WENO limiter [31], each variable component U being limited is checked independently: if it is flagged for slope reduction by the minmod limiter, then this component is reconstructed. This limiter uses several different estimated solutions for U on the troubled element labeled by k. The first estimate is the unlimited local data U^k. Each neighbor n of k also provides a "modified" solution estimate U^{k_n}; in the case of the simple WENO limiter, this estimate is simply obtained by evaluating the neighbor's solution U^n on the grid points of the element k. We follow the standard WENO algorithm of reconstructing the solution from a weighted sum of these estimates,

$$U_{\text{new}}^k = \omega_k U^k + \sum_n \omega_n U^{k_n}, \tag{43}$$

where the ω_i are the weights associated with each solution estimate, and satisfy the normalization $\sum_i \omega_i = 1$.

The weights are obtained by first computing an oscillation indicator (also called a smoothness indicator) σ_i for each $U^i = \{U^k, U^{k_n}\}$, which measures the amount of oscillation in the data. We use an indicator based on Eq. (23) of [54], but adapted for use on square or cubical grids,

$$\begin{aligned}\sigma_i = \sum_{\alpha=0}^{N}\sum_{\beta=0}^{N}\sum_{\substack{\gamma=0\\ \alpha+\beta+\gamma>0}}^{N} \int 2^{2(\alpha+\beta+\gamma)-1}\\ \times \left(\frac{\partial^{\alpha+\beta+\gamma}}{\partial\xi^\alpha \partial\eta^\beta \partial\zeta^\gamma} U^i\right)^2 d\xi\, d\eta\, d\zeta.\end{aligned} \tag{44}$$

Here the restriction on the sum avoids the term that has no derivatives of U^i, and the powers of two come from the interval width in the reference coordinates. From

the oscillation indicators, we compute the nonlinear weights:

$$\bar{\omega}_i = \frac{\gamma_i}{(\epsilon + \sigma_i)^2}. \tag{45}$$

Here the γ_i are the linear weights that give the relative weight of the local and neighbor contributions before accounting for oscillation in the data, and ϵ is a small number to avoid the denominator vanishing. We use standard values from the literature for both—we take $\gamma_{k_n} = 0.001$ for the neighbor contributions (then $\gamma_k = 0.994$ for an element with six neighbors; in general γ_k is set by the requirement that all the γ_i sum to unity), and $\epsilon = 10^{-6}$. Finally, the normalized nonlinear weights that go into the WENO reconstruction are given by

$$\omega_i = \frac{\bar{\omega}_i}{\sum_i \bar{\omega}_i}. \tag{46}$$

Note that the simple WENO limiter is not conservative since the neighboring elements' polynomials do not have the same element average as the element being limited.

4. HWENO

Our implementation of the HWENO limiter [32] follows similar steps. Note that we again use the TVB minmod limiter as troubled-cell indicator, whereas the reference uses the troubled-cell indicator of [55]. But, in keeping with the HWENO algorithm, we check the minmod indicator on all components of U being limited, and if any component is flagged for slope reduction, then the element is labeled as troubled and every variable being limited is reconstructed using the WENO procedure.

The HWENO modified solution estimates from the neighboring elements are computed as a least-squared fit to U across several elements. This broader fitting reduces oscillations as compared to the polynomial extrapolation used in the simple WENO estimates, and this improves robustness near shocks. The HWENO reconstruction uses a differently weighted oscillation indicator, computed similarly to Eq. (44) but with the prefactor in the integral being instead $(2^{2(\alpha+\beta+\gamma)-1})/((\alpha+\beta+\gamma)!)^2$. The HWENO algorithm explicitly guarantees conservation by constraining the reconstructed polynomials to have the same element-average value.

5. DG-finite-difference hybrid method

To the best of our knowledge the idea of hybridizing efficient spectral-type methods with robust high-resolution shock-capturing finite difference (FD) or finite volume (FV) schemes was first presented in [33]. However, our implementation is more similar to that of [34]. The basic idea is that after a time step or substep we check that the unlimited DG solution is satisfactory. If it is not, we mark the cell as troubled and retake the time step using standard FD methods. In this paper we use monotized-central reconstruction and the same numerical flux/boundary correction as the DG scheme uses. Our DG-FD hybrid method is also similar to that used in [21]. However, [21] did not attempt to run the method in 3d because of memory overhead. We have not done a detailed comparison of memory overhead between different limiting strategies, but have not noticed any significant barriers with the DG-FD hybrid scheme. We present a detailed description of our DG-FD hybrid method in a companion paper [56]. Our DG-FD hybrid method is not strictly conservative at boundaries where one element uses DG and another uses FD. This is because on the DG element we use the boundary correction of the reconstructed FD data, rather than the reconstructed boundary correction computed on the FD grid. In practice we have not found any negative impact from this choice.

E. Primitive recovery

One of the most difficult and expensive aspects of evolving the GRMHD equations is recovering the primitive variables from the conserved variables. Several different primitive recovery schemes are compared in [57]. We use the recently proposed scheme of Kastaun *et al.* [58]. If this scheme fails to recover the primitives, we try the Newman-Hamlin scheme [59]. If the Newman-Hamlin scheme fails, we use the scheme of Palenzuela *et al.* [60], and if that fails we terminate the simulation. Note that we have not yet incorporated all the fixing procedures to avoid recovery failure that are presented in [58].

F. Variable fixing

During the evolution the conserved and primitive variables can become nonphysical or enter regimes where the evolution is no longer stable (e.g., zero density). When limiting the solution does not remove these unphysical or bad values, a pointwise fixing procedure is used—at any grid points where the chosen conditions are not satisfied, the conserved variables are adjusted. The fixing procedures are generally not conservative and are used only as a fallback to ensure a stable evolution. In SpECTRE we currently use two fixing algorithms: The first applies an "atmosphere" in low-density regions, while the second adjusts the conserved variables in an attempt to guarantee primitive recovery.

Our atmosphere treatment is similar to that of [61–63]. We define values $\rho_{\rm atm}$ and $\rho_{\rm cutoff}$, where $\rho_{\rm atm} \leq \rho_{\rm cutoff}$. For any point where $\rho < \rho_{\rm cutoff}$ we set

$$\rho = \rho_{\rm atm}, \quad v^i = 0, \quad W = 1. \tag{47}$$

When $\rho_{\rm cutoff} < \rho \leq 10\rho_{\rm atm}$ we require that $v^i v_i < 10^{-4}$. After the primitive variables are set to the atmosphere we recompute the conserved variables from the primitive ones.

Our fixing of the conserved variables is based on that of Refs. [63,64]. We define $D_{\min}$ and D_{cutoff} and adjust $\tilde{D}$ if $D < D_{\text{cutoff}}$. Specifically, we set $\tilde{D} = \sqrt{\gamma} D_{\min}$. We adjust $\tilde{\tau}$ such that $\tilde{B}^2 \leq 2\sqrt{\gamma}(1-\epsilon_B)\tilde{\tau}$, where ϵ_B is a small number typically set to 10^{-12}.

Finally, we adjust $\tilde{S}_i$ such that $\tilde{S}^2 \leq \tilde{S}^2_{\max}$, where $\tilde{S}^2_{\max}$ is defined below. We define variables

$$\hat{\tau} = \frac{\tilde{\tau}}{\tilde{D}}, \tag{48}$$

$$\hat{B}^2 = \frac{\tilde{B}^2}{\sqrt{\gamma}\tilde{D}}, \tag{49}$$

$$\hat{\mu} = \begin{cases} \frac{\tilde{S}_i \tilde{B}^i}{\sqrt{\tilde{B}^2 \tilde{S}^2}}, & \tilde{B}^2 > \tilde{D} \times 10^{-16} \quad \text{and} \quad \tilde{S}^2 > \tilde{D}^2 \times 10^{-16} \\ 0, & \text{otherwise.} \end{cases} \tag{50}$$

The Lorentz factor is bounded by

$$\max(1, 1 + \hat{\tau} - \hat{B}^2) \leq W \leq 1 + \hat{\tau}, \tag{51}$$

and is determined by finding the root of

$$\begin{aligned} g(W) = {} & (W + \hat{B}^2 - \hat{\tau} - 1)[W^2 + \hat{B}^2\hat{\mu}^2(\hat{B}^2 + 2W)] \\ & - \frac{\hat{B}^2}{2}[1 + \hat{\mu}^2(W^2 + 2W\hat{B}^2 + \hat{B}^4 - 1)]. \end{aligned} \tag{52}$$

Using the Lorentz factor W obtained by solving (52) we define $\tilde{S}_{\max}$ as

$$\begin{aligned} \tilde{S}_{\max} = \tilde{S} \min\Bigg(1, \\ \sqrt{\frac{(1-\epsilon_S)(W+\hat{B}^2)^2(W^2-1)\tilde{D}^2}{(\tilde{S}^2 + \tilde{D}^2 \times 10^{-16})[W^2 + \hat{\mu}^2\hat{B}^2(\hat{B}^2+2W)]}}\Bigg), \end{aligned} \tag{53}$$

where ϵ_S is a small number typically set to 10^{-12}. We apply the check on the conserved variables after each time or substep before a primitive recovery is done.

Implementing root finding for Eq. (52) in a manner that is well behaved for floating point arithmetic is important. Specifically, we solve

$$\begin{aligned} g(W) = {} & \left(\frac{1}{2}\hat{B}^2 - \hat{\tau}\right)(1 + 2\hat{B}^2\hat{\mu}^2 + \hat{B}^4\hat{\mu}^2) \\ & + (W-1)[2(\hat{B}^2 - \hat{\tau})(1 + \hat{B}^2\hat{\mu}^2) + \hat{B}^2\hat{\mu}^2 + 1] \\ & + (W-1)^2\left(\hat{B}^2 - \hat{\tau} + \frac{3}{2}\hat{B}^2\hat{\mu}^2 + 2\right) \\ & + (W-1)^3 \end{aligned} \tag{54}$$

for $W - 1$ when the lower bound for W is 1 and

$$\begin{aligned} g(W) = {} & -\frac{1}{2}\hat{B}^2[1 + \hat{\mu}^2\hat{\tau}(\hat{\tau}+2)] \\ & + [W - (1 + \hat{\tau} - \hat{B}^2)][1 + \hat{B}^2\hat{\mu}^2 \\ & + (\hat{\tau} - \hat{B}^2)(\hat{B}^2\hat{\mu}^2 + \hat{\tau} - \hat{B}^2 + 2)] \\ & + [W - (1 + \hat{\tau} - \hat{B}^2)]^2\left[2(\hat{\tau} - \hat{B}^2) + \frac{3}{3}\hat{B}^2\hat{\mu}^2 + 2\right] \\ & + [W - (1 + \hat{\tau} - \hat{B}^2)]^3 \end{aligned} \tag{55}$$

for $W - (1 + \hat{\tau} - \hat{B}^2)$ when the lower bound for W is $1 + \hat{\tau} - \hat{B}^2$.

We also have a flattening algorithm inspired by [65] that reduces oscillations of the conserved variables if the solution is unphysical. Unlike the pointwise fixing, the flattening algorithm is conservative. In particular, we reduce the oscillations in $\tilde{D}$ if it is negative at any point in the cell, and we rescale $\tilde{\tau}$ to satisfy $\tilde{B}^2 \leq 2\sqrt{\gamma}\tilde{\tau}$. Finally, if the primitive variables cannot be recovered we reset the conserved variables to their mean values.

IV. NUMERICAL RESULTS

For all test problems we use the less dissipative HLL boundary correction. In many cases one of the limiting strategies fails. This failure usually occurs during the primitive recovery. However, this is a symptom of the DG and limiting procedure producing a bad state rather than a poor primitive recovery algorithm. All simulations are performed using SpECTRE v2022.04.04 [35] and the input files used are provided alongside the arXiv version.

A. 1d smooth flow

We consider a simple 1d smooth flow problem to test which of the limiters and troubled-cell indicators are able to solve a smooth problem without degrading the order of accuracy. A smooth density perturbation is advected across the domain with a velocity v^i. The analytic solution is given by

$$\rho = 1 + 0.7\,\sin[k^i(x^i - v^i t)], \tag{56}$$

$$v^i = (0.8, 0, 0), \tag{57}$$

$$k^i = (1, 0, 0), \tag{58}$$

$$p = 1, \tag{59}$$

$$B^i = (0, 0, 0), \tag{60}$$

and we close the system with an adiabatic equation of state,

$$p = \rho\epsilon(\Gamma - 1), \tag{61}$$

where Γ is the adiabatic index, which we set to 1.4. We use a domain given by $[0, 2\pi]^3$ and apply periodic boundary conditions in all directions. The time step size is $\Delta t = 2\pi/5120$ so that the spatial discretization error is larger than the time stepping error for all resolutions we use.

We perform a convergence test using the different limiting strategies and present the results in Table I. We show both the L_2 norm of the error and the convergence order. The L_2 norm is defined as

$$L_2(u) = \sqrt{\frac{1}{M}\sum_{i=0}^{M-1} u_i^2}, \tag{62}$$

where M is the total number of grid points and u_i is the value of u at grid point i and the convergence order is given by

$$L_2 \text{ order} = \log_2\left[\frac{L_2(\mathcal{E}_{N_x/2})}{L_2(\mathcal{E}_{N_x})}\right]. \tag{63}$$

We see that the troubled-cell indicator for the $\Lambda\Pi^N$, HWENO, and simple WENO limiters does not flag any cells as troubled and the full order of accuracy of the DG scheme is preserved. For these simulations we used a TVB constant of 1. The Krivodonova limiter completely flattens the solution and shows no convergence. The reason is that the Krivodonova limiter is unable to preserve a smooth solution if the flow is constant in an orthogonal direction. This can be understood from the minmod algorithm being applied to the neighboring coefficients. The smooth flow solution is constant in the y and z directions, and so the Krivodonova limiter effectively zeros all higher moments. The DG-FD P_5 scheme switches to FD when we use only two elements, but from four to 16 elements it uses DG. The order of convergence is so large for the $N_x = 4$ case because in addition to doubling the resolution, the code also switches from using second-order FD to sixth-order DG, causing a very large decrease in the errors. Using a higher-order or adaptive-order FD scheme is expected to preserve the accuracy much better when the hybrid scheme is using FD, while still being able to capture shocks robustly and accurately.

TABLE I. The errors and local convergence order for the smooth flow problem using different limiting strategies. Note that the limiter is not applied if the troubled-cell indicator determines the DG solution to be valid. Except for the Krivodonova limiter, which is nonconvergent, we observe the expected convergence order except when the solution is under-resolved because too few elements are used.

Limiter	N_x	$L_2(\mathcal{E}(\rho))$	L_2 order
$\Lambda\Pi^N$	2	2.22282×10^{-3}	
	4	2.23822×10^{-5}	6.63
	8	3.18504×10^{-7}	6.13
	16	5.08821×10^{-9}	5.97
HWENO	2	2.22282×10^{-3}	
	4	2.23822×10^{-5}	6.63
	8	3.18504×10^{-7}	6.13
	16	5.08821×10^{-9}	5.97
Simple WENO	2	2.22282×10^{-3}	
	4	2.23822×10^{-5}	6.63
	8	3.18504×10^{-7}	6.13
	16	5.08821×10^{-9}	5.97
Krivodonova	2	3.92346×10^{-1}	
	4	4.94975×10^{-1}	−0.34
	8	4.94975×10^{-1}	0.00
	16	4.73294×10^{-1}	0.06
DG-FD P_5	2	3.45679×10^{-1}	
	4	$2.23822e \times 10^{-5}$	13.91
	8	3.18504×10^{-7}	6.13
	16	5.08821×10^{-9}	5.97

B. 1d Riemann problems

One-dimensional Riemann problems are a standard test for any scheme that must be able to handle shocks. We will focus on the first Riemann problem (RP1) of [66]. The setup is given in Table II. While not the most challenging Riemann problem, it gives a good baseline for different limiting strategies. We perform simulations using an SSP-RK3 method with $\Delta t = 5 \times 10^{-4}$. In Fig. 1 we show the rest mass density ρ at $t_f = 0.4$ for simulations using the simple WENO, HWENO, $\Lambda\Pi^N$, and Krivodonova limiters, as well as a run using the DG-FD hybrid scheme. The thin black curve is the analytic solution obtained using the Riemann solver of [67]. All simulations use 128 elements in the x direction with a P_2 (third-order) DG scheme, and an ideal fluid equation of state, Eq. (61).

While all five limiting strategies evolve to the final time, the DG-FD scheme is the least oscillatory and is also able to resolve the discontinuities much more accurately. The HWENO scheme is slightly less oscillatory if linear neighbor weights of $\gamma_k = 0.01$ are used instead of $\gamma_k = 0.001$. However, the simple WENO limiter fails to evolve the solution with $\gamma_k = 0.01$ and such sensitivity to parameters in the algorithm is not desirable when solving realistic problems. Going to higher order has proven to be especially challenging. While both the $\Lambda\Pi^N$ and the Krivodonova complete the evolution when using a P_5

TABLE II. The initial conditions for Riemann problem 1 of [66]. The domain is $x \in [-0.5, 0.5]$, the final time is $t_f = 0.4$, and an ideal fluid equation of state is used with an adiabatic index of 2.

	ρ	ρ	v^i	B^i
$x < 0$	1.0	1.0	(0,0,0)	(0.5, 1,0)
$x \geq 0$	0.125	0.1	(0,0,0)	(0.5, −1, 0)

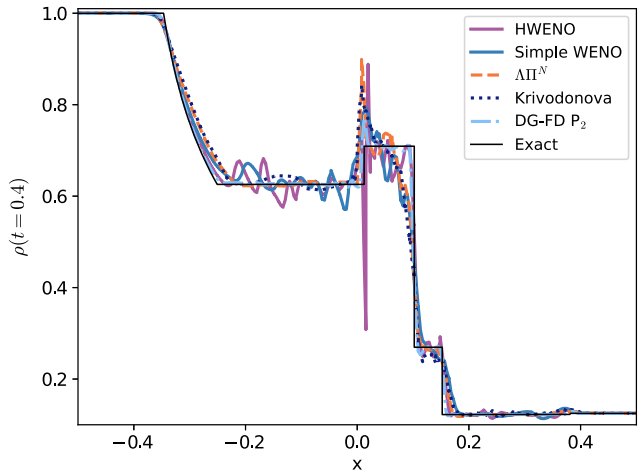

(a) Riemann Problem 1 comparison

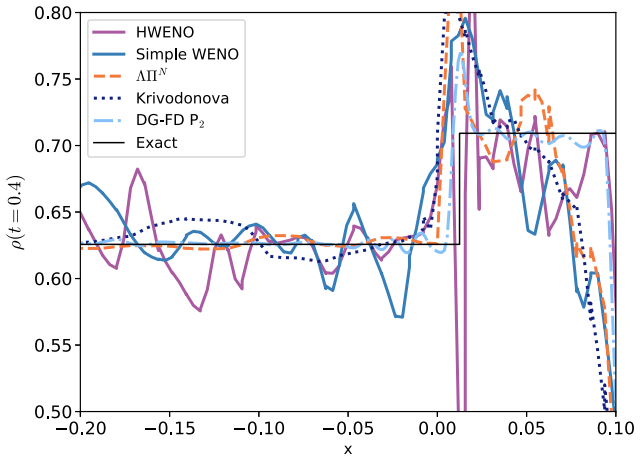

(b) Zoom in of Riemann Problem 1

FIG. 1. A comparison of different limiters used to stabilize the evolution of the Riemann problem 1 of [66]. The problem is solved using 128 third-order (P_2) elements. The DG-FD hybrid scheme significantly outperforms the other limiters both in robustness and accuracy.

scheme with 64 elements (simple WENO and HWENO fail), additional spurious oscillations are present. In comparison, the DG-FD hybrid scheme actually has fewer oscillations when going to higher order. In Fig. 2 we plot the error of the numerical solution using a P_2 DG-FD scheme with 128 elements and a P_5 DG-FD scheme with 64 elements. We see that the P_5 hybrid scheme actually has fewer oscillations than the P_2 scheme, while resolving the discontinuities equally well. We attribute this to the troubled-cell indicators actually triggering earlier when a higher polynomial degree is used since discontinuities entering an element rapidly dump energy into the high modes. While we will compare the different limiting strategies for 2d and 3d problems below, it is already quite apparent that the DG-FD hybrid scheme is by far the most robust and accurate method.

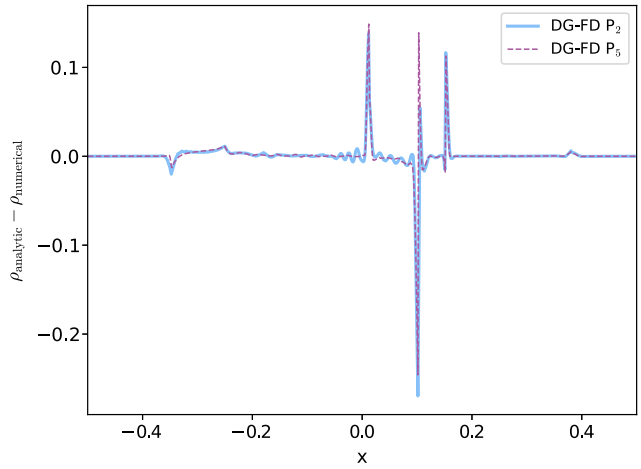

FIG. 2. The difference between the analytic and numerical solution of the Riemann problem 1 of [66] at $t = 0.4$ for the DG-FD P_2 scheme (solid light blue curve) and the DG-FD P_5 scheme (dashed purple curve). The P_5 scheme is able to resolve the discontinuities just as well as the P_2 scheme, while also admitting fewer unphysical oscillations away from the discontinuities.

C. 2d cylindrical blast wave

A standard test problem for GRMHD codes is the cylindrical blast wave [68,69], where a magnetized fluid initially at rest in a constant magnetic field along the x axis is evolved. The fluid obeys the ideal fluid equation of state (61) with $\gamma = 4/3$. The fluid begins in a cylindrically symmetric configuration, with hot, dense fluid in the region with cylindrical radius $r < 0.8$ surrounded by a cooler, less dense fluid in the region $r > 1$. The initial density ρ and pressure p of the fluid are

$$\begin{aligned} \rho(r < 0.8) &= 10^{-2}, \\ \rho(r > 1.0) &= 10^{-4}, \\ p(r < 0.8) &= 1, \\ p(r > 1.0) &= 5 \times 10^{-4}. \end{aligned} \tag{64}$$

In the region $0.8 \leq r \leq 1$, the solution transitions continuously and exponentially (i.e., transitions such that the logarithms of the pressure and density are linear functions of r). The fluid begins threaded with a uniform magnetic field with Cartesian components

$$(B^x, B^y, B^z) = (0.1, 0, 0). \tag{65}$$

The magnetic field causes the blast wave to expand nonaxisymmetrically. For all simulations we use a time step size $\Delta t = 10^{-2}$ and an SSP RK3 time integrator.

We evolve the blast wave to time $t = 4.0$ on a grid of $64 \times 64 \times 1$ elements covering a cube of extent $[-6, 6]^3$ using a DG P_2 scheme, a comparable resolution to what FD code tests use. We apply periodic boundary conditions in all directions, since the explosion does not reach the outer boundary by $t = 4.0$. Figure 3 shows the logarithm of the rest-mass density at time $t = 4.0$, at the end of

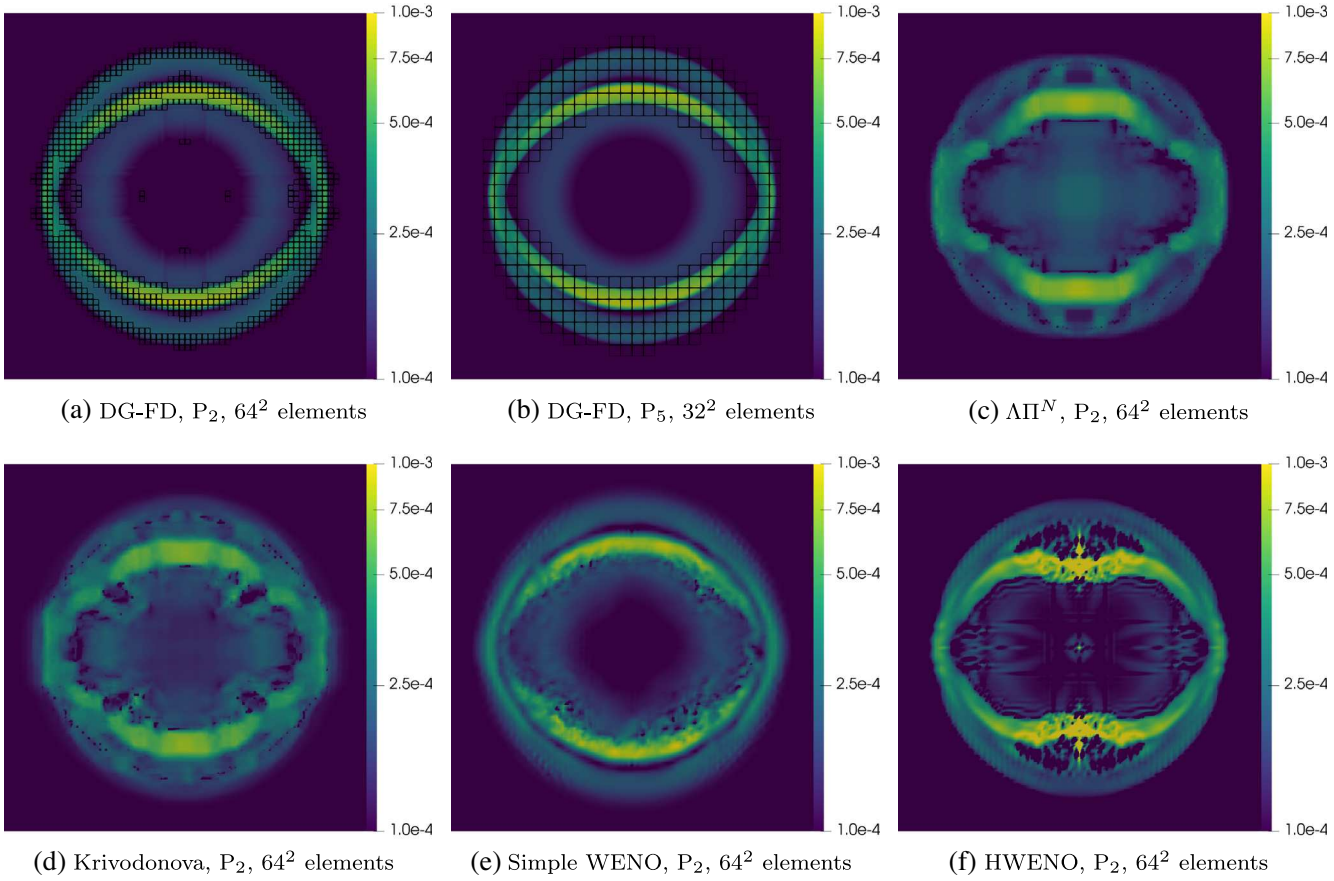

(a) DG-FD, P_2, 64^2 elements (b) DG-FD, P_5, 32^2 elements (c) $\Lambda\Pi^N$, P_2, 64^2 elements

(d) Krivodonova, P_2, 64^2 elements (e) Simple WENO, P_2, 64^2 elements (f) HWENO, P_2, 64^2 elements

FIG. 3. Cylindrical blast wave ρ at $t = 4$ comparing the DG-FD hybrid scheme, the $\Lambda\Pi^N$, Krivodonova, simple WENO, and HWENO limiters using P_2 DG, as well as the DG-FD scheme using P_5 DG. There are 192 degrees of freedom per dimension, comparable to what is used when testing FD schemes. We see that only the DG-FD hybrid scheme really resolves the features to an acceptable level, and the $\Lambda\Pi^N$ and Krivodonova smear out the solution almost completely. In the plots of the DG-FD hybrid scheme the regions surrounded by black squares have switched from DG to FD at the final time.

evolutions using the different limiting strategies. We see from Figs. 3(c) and 3(d) that the $\Lambda\Pi^N$ and Krivodonova limiters result in a very poorly resolved solution. The simple WENO evolution, Fig. 3(e) is much better but still not nearly as good as a FD method with the same number of degrees of freedom. The HWENO limiter, Fig. 3(f), suffers from various spurious artifacts. The DG-FD hybrid scheme, however, again demonstrates its ability to robustly handle discontinuities, while also resolving smooth features with very high order. Figure 3(a) shows the result of a simulation using a P_2 DG-FD scheme and Fig. 3(b) using a P_5 DG-FD scheme with half the number of elements. The increased resolution of a high-order scheme is clear when comparing the P_2 and P_5 solutions in the interior region of the blast wave. We conclude that the DG-FD hybrid scheme is the most robust and accurate method/limiting strategy for solving the cylindrical blast wave problem.

D. 2d magnetic rotor

The second two-dimensional test problem we study is the magnetic rotor problem originally proposed for nonrelativistic MHD [70,71] and later generalized to the relativistic case [72,73]. A rapidly rotating dense fluid cylinder is inside a lower density fluid, with a uniform pressure and magnetic field everywhere. The magnetic braking will slow down the rotor over time, with an approximately 90 degree rotation by the final time $t = 0.4$. We use a domain of $[-0.5, 0.5]^3$ and a time step size $\Delta t = 10^{-3}$ and an SSP RK3 time integrator. An ideal fluid equation of state with $\Gamma = 5/3$ is used, and the following initial conditions are imposed:

$$\begin{aligned} p &= 1 \\ B^i &= (1, 0, 0) \\ v^i &= \begin{cases} (-y\Omega, x\Omega, 0), & \text{if } r \le R_{\text{rotor}} = 0.1 \\ (0, 0, 0), & \text{otherwise,} \end{cases} \\ \rho &= \begin{cases} 10, & \text{if } r \le R_{\text{rotor}} = 0.1 \\ 1, & \text{otherwise,} \end{cases} \end{aligned} \tag{66}$$

with angular velocity $\Omega = 9.95$. The choice of Ω and $R_{\text{rotor}} = 0.1$ guarantees that the maximum velocity of the fluid (0.995) is less than the speed of light.

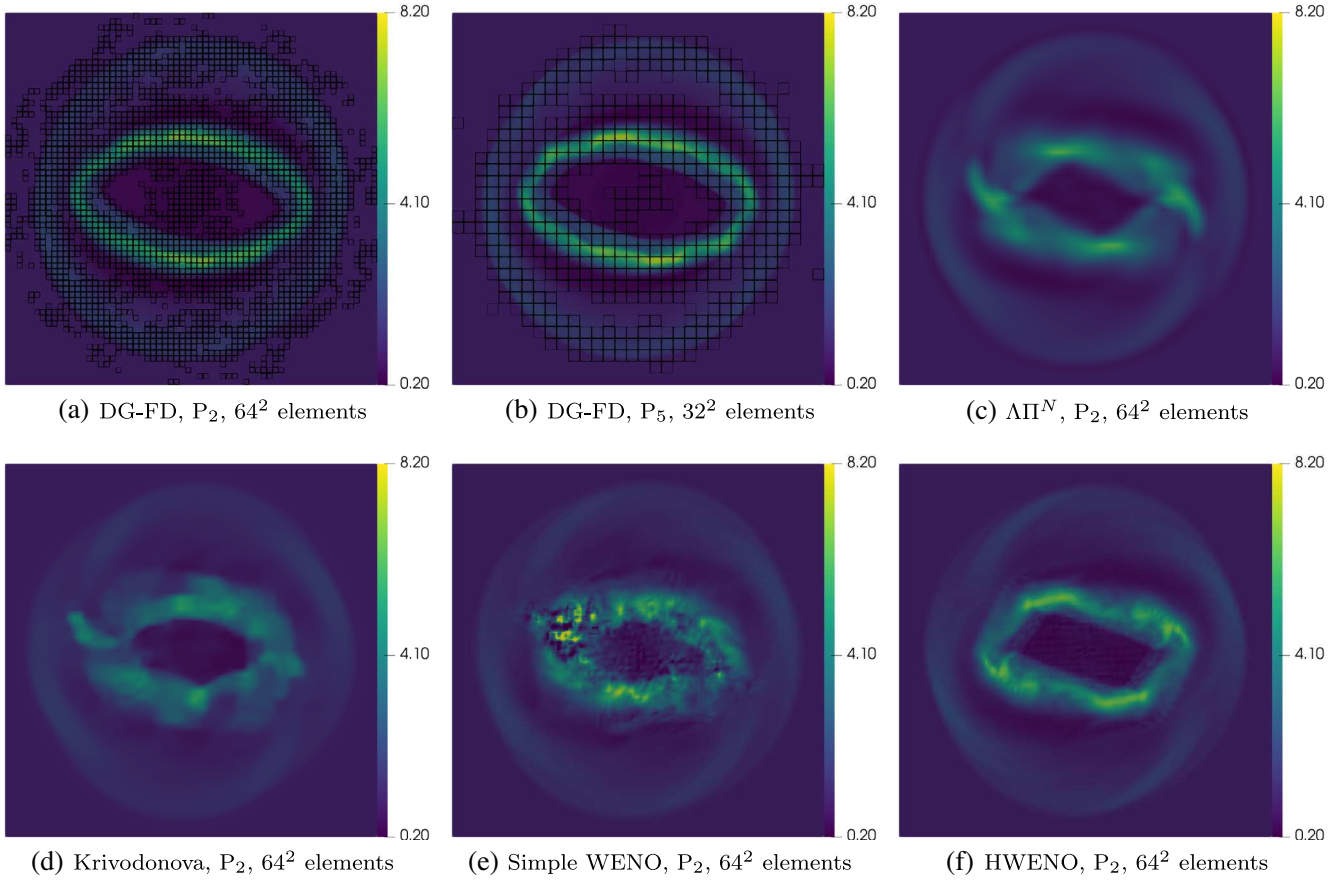

(a) DG-FD, P_2, 64^2 elements (b) DG-FD, P_5, 32^2 elements (c) $\Lambda\Pi^N$, P_2, 64^2 elements

(d) Krivodonova, P_2, 64^2 elements (e) Simple WENO, P_2, 64^2 elements (f) HWENO, P_2, 64^2 elements

FIG. 4. Magnetic rotor ρ at $t = 0.4$ comparing the DG-FD hybrid scheme, the $\Lambda\Pi^N$, Krivodonova, simple WENO, and HWENO limiters using P_2 DG, as well as the DG-FD scheme using P_5 DG. There are 192 degrees of freedom per dimension, comparable to what is used when testing FD schemes. We see that only the DG-FD hybrid scheme really resolves the features to an acceptable level, and the $\Lambda\Pi^N$ and Krivodonova smear out the solution almost completely. The simple WENO limiter fails to solve the problem. In the plots of the DG-FD hybrid scheme the regions surrounded by black squares have switched from DG to FD at the final time.

We show the results of our evolutions in Fig. 4, which are all done with 192 grid points and periodic boundary conditions. Figures 4(c) and 4(d) show results using the $\Lambda\Pi^N$ and Krivodonova limiter which both severely smear out the solution. The simple WENO limiter suffers from spurious artifacts [Fig. 4(d)], while the HWENO limiter does a reasonable job [Fig. 4(e)]. The DG-FD hybrid scheme is most robust and accurate, but a fairly large number of cells end up being marked as troubled in this problem and switched to FD. While ideally fewer cells would be switched to FD, it is better to have a scheme that is capable of solving a large array of problems without fine-tuning than to have a slightly different fine-tuned scheme for each test problem.

E. 2d magnetic loop advection

The third two-dimensional test problem we study is the magnetic loop advection problem [74]. A magnetic loop is advected through the domain until it returns to its starting position. We use an initial configuration very similar to [42,75–77], where

$$\begin{aligned}
\rho &= 1 \\
p &= 3 \\
v^i &= (1/1.2, 1/2.4, 0) \\
B^x &= \begin{cases} -A_{\rm loop} y/R_{\rm in}, & \text{if } r \le R_{\rm in} \\ -A_{\rm loop} y/r, & \text{if } R_{\rm in} < r < R_{\rm loop} \\ 0, & \text{otherwise,} \end{cases} \\
B^y &= \begin{cases} A_{\rm loop} x/R_{\rm in}, & \text{if } r \le R_{\rm in} \\ A_{\rm loop} x/r, & \text{if } R_{\rm in} < r < R_{\rm loop} \\ 0, & \text{otherwise,} \end{cases}
\end{aligned} \tag{67}$$

with $R_{\rm loop} = 0.3$, $R_{\rm in} = 0.001$, and an ideal gas equation of state with $\Gamma = 5/3$. The computational domain is $[-0.5, 0.5]^3$ with $64 \times 64 \times 1$ elements and periodic boundary conditions being applied everywhere. The final time for one period is $t = 2.4$. For all simulations we use a time step size $\Delta t = 10^{-3}$ and an SSP RK3 time integrator.

In Fig. 5 we plot the magnetic field component B^x at $t = 0$ on the left half of each plot and after one period

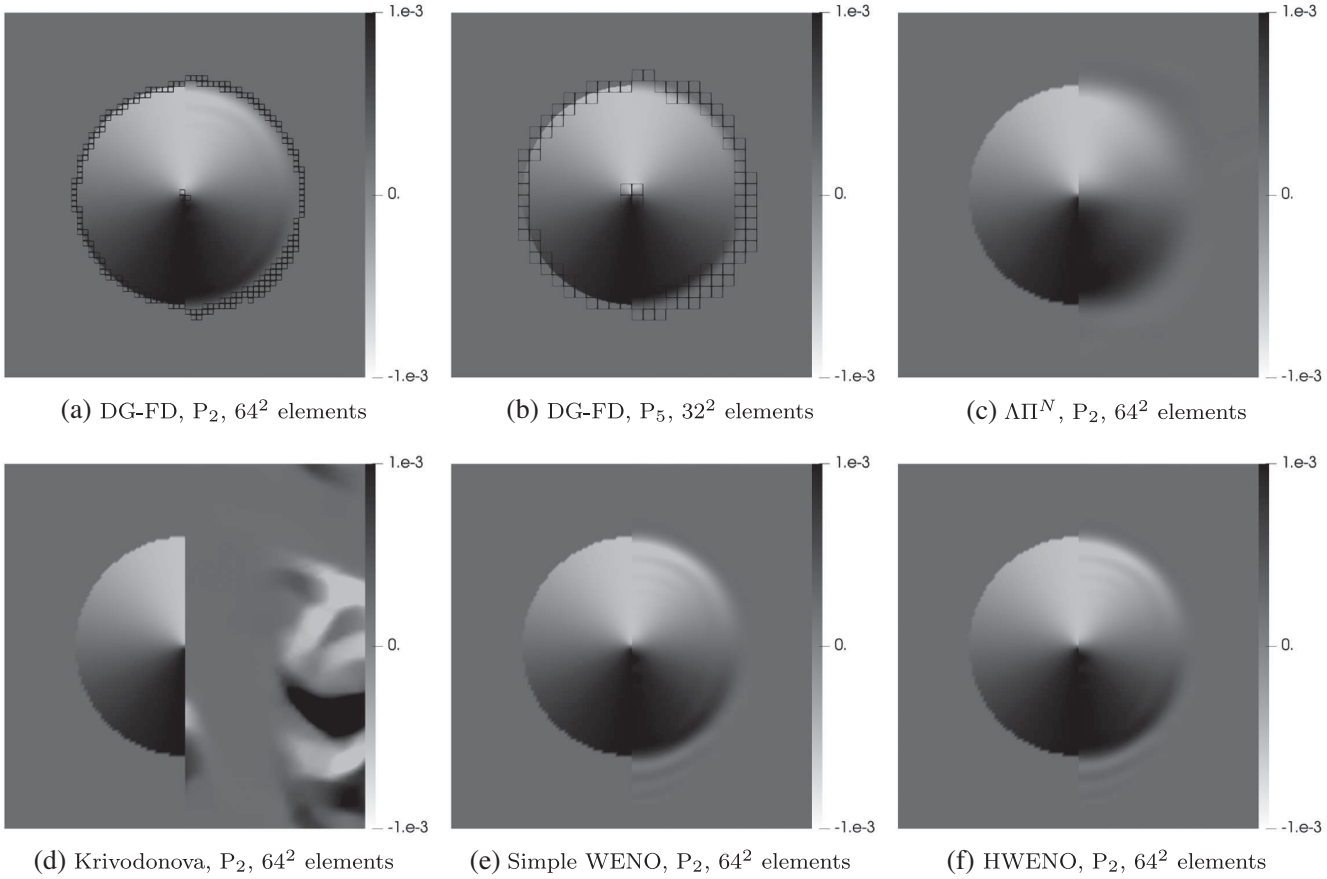

(a) DG-FD, P_2, 64^2 elements (b) DG-FD, P_5, 32^2 elements (c) $\Lambda\Pi^N$, P_2, 64^2 elements

(d) Krivodonova, P_2, 64^2 elements (e) Simple WENO, P_2, 64^2 elements (f) HWENO, P_2, 64^2 elements

FIG. 5. B^x for the magnetic loop advection problem. The left half of each plot is at the initial time, while the right half is after one period ($t = 2.4$). We compare the DG-FD hybrid scheme, the $\Lambda\Pi^N$, Krivodonova, simple WENO, and HWENO limiters using P_2 DG, as well as the DG-FD scheme using P_5 DG. There are 192 degrees of freedom per dimension, comparable to what is used when testing FD schemes. In the plots of the DG-FD hybrid scheme the regions surrounded by black squares have switched from DG to FD at the final time.

$t = 2.4$ on the right half of each plot for results using various limiting strategies. We use a TVB constant of 5 for the $\Lambda\Pi^N$, simple WENO, and HWENO limiters, and use neighbor weights $\gamma_k = 0.001$ for the simple WENO and HWENO limiters. The Krivodonova limiter completely destroys the solution and only remains stable because of our conservative variable fixing scheme. Both WENO limiters work quite well, maintaining the shape of the loop with only some oscillations being generated. The DG-FD hybrid scheme again performs the best. In Fig. 5(a) we show the result using a P_2 DG-FD scheme and in Fig. 5(b) using a P_5 DG-FD scheme. The P_5 scheme resolves the smooth parts of the solution more accurately than the P_2 scheme, as is to be expected. The DG-FD hybrid scheme also does not generate the spurious oscillations that are present when using the WENO limiters. While the spurious oscillations may be reduced by fine-tuning the TVB constant and the neighbor weights, this type of fine-tuning is not possible for complex physics simulations and so we do not spend time searching for the "optimal" parameters.

Since we are using hyperbolic divergence cleaning, violations of the $\partial_i B^i = 0$ constraint occur. In Fig. 6 we plot the divergence cleaning field Φ at the final time $t = 2.4$. The simple WENO, HWENO, and DG-FD hybrid schemes all have $|\Phi| \sim 5 \times 10^{-6}$, while the $\Lambda\Pi^N$ limiter has Φ approximately 1 order of magnitude larger. For the magnetic loop advection problem we find that all classical limiters perform comparably, except the Krivodonova limiter completely destroys the solution and remains stable only because of our conservative variable fixing scheme. Nevertheless, the DG-FD hybrid scheme is better than the classical limiters, and we conclude that the DG-FD hybrid scheme is both the most robust and accurate method/limiting strategy for solving the magnetic loop advection problem.

F. 2d magnetized Kelvin-Helmholtz instability

The last two-dimensional test problem we study is the magnetized Kelvin-Helmholtz (KH) instability, similar to [78]. The domain is $[0, 1]^3$ and we use the following initial conditions [18]:

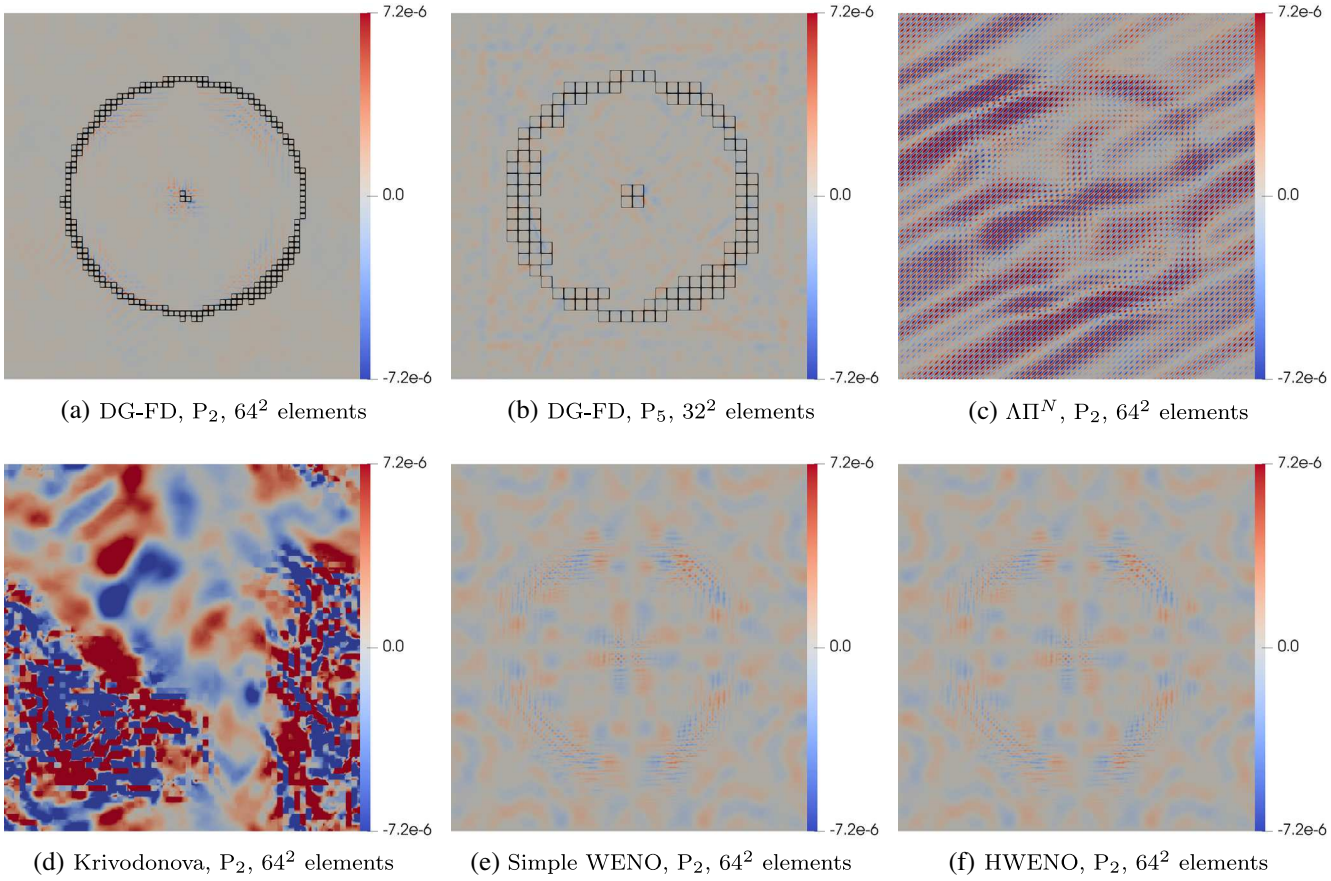

(a) DG-FD, P_2, 64^2 elements (b) DG-FD, P_5, 32^2 elements (c) $\Lambda\Pi^N$, P_2, 64^2 elements

(d) Krivodonova, P_2, 64^2 elements (e) Simple WENO, P_2, 64^2 elements (f) HWENO, P_2, 64^2 elements

FIG. 6. The divergence cleaning field Φ for the magnetic loop advection problem after one period ($t = 2.4$) comparing the DG-FD hybrid scheme, the $\Lambda\Pi^N$, Krivodonova, simple WENO, and HWENO limiters using P_2 DG, as well as the DG-FD scheme using P_5 DG. There are 192 degrees of freedom per dimension, comparable to what is used when testing FD schemes. In the plots from the DG-FD hybrid scheme the regions surrounded by black squares have switched from DG to FD at the final time.

$$\rho = \begin{cases} 1, & |y - 0.5| < 0.25 \\ 10^{-2}, & \text{otherwise,} \end{cases} \tag{68}$$

$$p = 1.0, \tag{69}$$

$$v^x = \begin{cases} 0.5, & |y - 0.5| < 0.25 \\ -0.5, & \text{otherwise,} \end{cases} \tag{70}$$

$$v^y = 0.1 \sin(4\pi x)\left[\exp\left(-\frac{(y-0.75)^2}{0.0707^2}\right) + \exp\left(-\frac{(y-0.25)^2}{0.0707^2}\right)\right], \tag{71}$$

$$v^z = 0.0, \tag{72}$$

$$B^x = 10^{-3}, \tag{73}$$

$$B^y = B^z = 0.0. \tag{74}$$

We use an ideal gas equation of state with $\Gamma = 4/3$, a final time $t_f = 1.6$, a time step size of $\Delta t = 10^{-3}$, an SSP RK3 time integrator, and $[64 \times 1 \times 64]$ P_2 elements for the classical limiters. For the DG-FD hybrid method we use both $[64 \times 1 \times 64]$ P_2 elements and $[32 \times 1 \times 32]$ P_5 elements. We use a TVB constant of 1 for all the limiters. Using the flattening algorithm is crucial for the results obtained here, while for other test problems it is significantly less important.

In Fig. 7 we plot the density at the final time comparing the different limiting strategies. From Figs. 7(e) and 7(d) we see that the simple WENO and Krivodonova limiters destroy the solution almost completely. The $\Lambda\Pi^N$ limiter [Fig. 7(c)] retains some hints of the expected flow pattern, but also nearly completely destroys the solution. The HWENO limiter is plotted in Fig. 7(e) and does by far the best of the classical limiters. Ultimately, only the DG-FD hybrid method [Fig. 7(a) for P_2 and Fig. 7(b) for P_5] is able to produce the expected vortices and flow patterns.

G. TOV star

A rigorous 3d test case in general relativity is the evolution of a static, spherically symmetric star. The TOV solution [79,80] describes such a setup. In this section

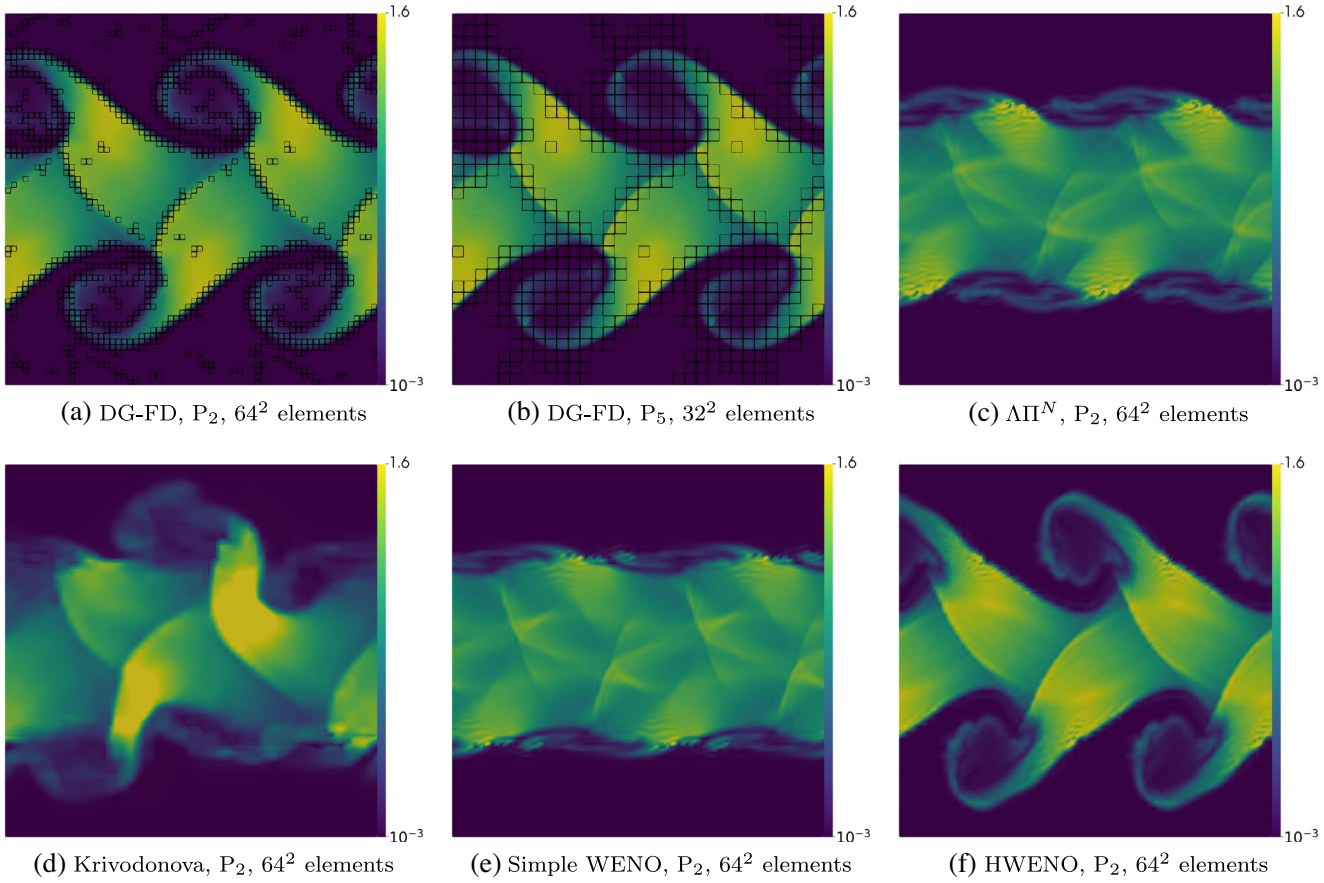

(a) DG-FD, P_2, 64^2 elements (b) DG-FD, P_5, 32^2 elements (c) $\Lambda\Pi^N$, P_2, 64^2 elements

(d) Krivodonova, P_2, 64^2 elements (e) Simple WENO, P_2, 64^2 elements (f) HWENO, P_2, 64^2 elements

FIG. 7. Magnetized Kelvin-Helmholtz instability ρ at $t = 1.6$ comparing the DG-FD hybrid scheme, the $\Lambda\Pi^N$, Krivodonova, simple WENO, and HWENO limiters using P_2 DG, as well as the DG-FD scheme using P_5 DG. There are 192 degrees of freedom per dimension, comparable to what is used when testing FD schemes. Only the DG-FD hybrid scheme and the HWENO limiter produce reasonable results, while the $\Lambda\Pi^N$ limiter has very low effective resolution, and the Krivodonova and simple WENO limiters smear out the solution almost completely. In the plots of the DG-FD hybrid scheme the regions surrounded by black squares have switched from DG to FD at the final time.

we study evolutions of both nonmagnetized and magnetized TOV stars. We adopt the same configuration as in [81]. Specifically, we use a polytropic equation of state,

$$p(\rho) = K\rho^{\Gamma}, \tag{75}$$

with the polytropic exponent $\Gamma = 2$, polytropic constant $K = 100$, and a central density $\rho_c = 1.28 \times 10^{-3}$. When considering a magnetized star we choose a magnetic field given by the vector potential,

$$A_\phi = A_b(x^2 + y^2)\max(p - p_{\rm cut}, 0)^{n_s}, \tag{76}$$

with $A_b = 2500$, $p_{\rm cut} = 0.04 p_{\rm max}$, and $n_s = 2$. This configuration yields a magnetic field strength in CGS units,

$$|B_{\rm CGS}| = \sqrt{b^2} \times 8.352 \times 10^{19}\ {\rm G}, \tag{77}$$

of $|B_{\rm CGS}| = 1.03 \times 10^{16}$ G. The magnetic field is only a perturbation to the dynamics of the star, since $(p_{\rm mag}/p)(r = 0) \sim 5 \times 10^{-5}$. However, evolving the field stably and accurately can be challenging. The magnetic field corresponding to the vector potential in Eq. (76) in the magnetized region is given by

$$\begin{aligned} B^x &= \frac{1}{\sqrt{\gamma}}\frac{xz}{r}A_b n_s (p - p_{\rm cut})^{n_s - 1}\partial_r p, \\ B^y &= \frac{1}{\sqrt{\gamma}}\frac{yz}{r}A_b n_s (p - p_{\rm cut})^{n_s - 1}\partial_r p, \\ B^z &= -\frac{A_b}{\sqrt{\gamma}}\left[2(p - p_{\rm cut})^{n_s} \right. \\ &\quad \left. + \frac{x^2 + y^2}{r} n_s (p - p_{\rm cut})^{n_s - 1}\partial_r p\right], \end{aligned} \tag{78}$$

and at $r = 0$ is

$$B^x = 0, \quad B^y = 0, \quad B^z = -\frac{A_b}{\sqrt{\gamma}} 2(p - p_{\rm cut})^{n_s}. \tag{79}$$

We perform all evolutions in full 3d with no symmetry assumptions and in the Cowling approximation, i.e., we do not evolve the spacetime. To match the resolution usually

used in FD/FV numerical relativity codes we use a domain $[-20, 20]^3$ with a base resolution of 6 P_5 DG elements and 12 P_2 DG elements. This choice means we have approximately 32 FD grid points covering the star's diameter at the lowest resolution, 64 when using 12 P_5 elements, and 128 grid points when using 24 P_5 elements. In all cases we set $\rho_{\text{atm}} = 10^{-15}$ and $\rho_{\text{cutoff}} = 1.01 \times 10^{-15}$.

In Fig. 8 we show the normalized maximum rest mass density over the grid for the nonmagnetized TOV star. The six-element simulation uses FD throughout the interior of the star and so there is no grid point at $r = 0$. This is the reason the data is shifted compared to 12- and 24-element simulations, where the unlimited P_5 DG solver is used throughout the star interior and so there is a grid point at the center of the star. The increased "noise" in the 12- and 24-element data actually stems from the higher oscillation modes in the star that are induced by numerical error. In Fig. 9 we plot the power spectrum using data at the three different resolutions. The six-element simulation only has one mode resolved, while 12 elements resolve two modes well, and the 24-element simulation resolves three modes well. In Fig. 10 we show the normalized maximum rest mass density over the grid for the best two cases using the classical limiters. The simple WENO and HWENO limiters performed similarly and were only stable for P_2. The Krivodonova limiter only succeeded at three of the 16 resolutions we attempted, and its best result is noticeably noisier than the other limiters. Note that our experience is consistent with that of Ref. [21], which was unable to achieve stable evolutions of a 3d TOV star using the simple WENO limiter.

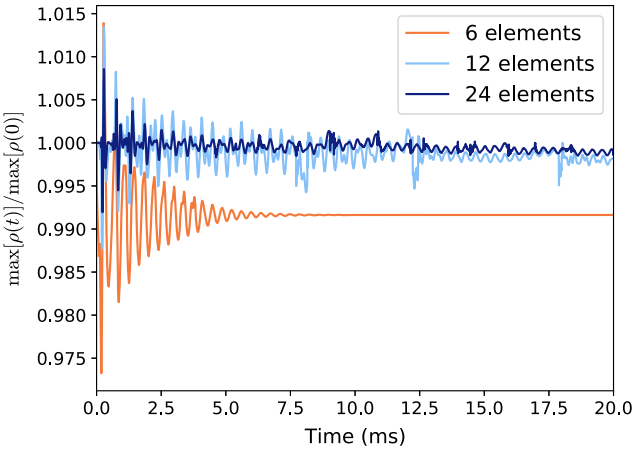

FIG. 8. The maximum density over the grid $\max(\rho)$ divided by the maximum density over the grid at $t = 0$ for three different resolutions for the nonmagnetized TOV star simulations. The six-element simulation uses FD throughout the interior of the star, while 12- and 24-element simulations use DG. The increased high-frequency content in 12- and 24-element simulations occurs because the high-order DG scheme is able to resolve higher oscillation modes in the star. The maximum density in the six-element case drifts down at early times because of the low resolution and the relatively low accuracy of using FD at the center.

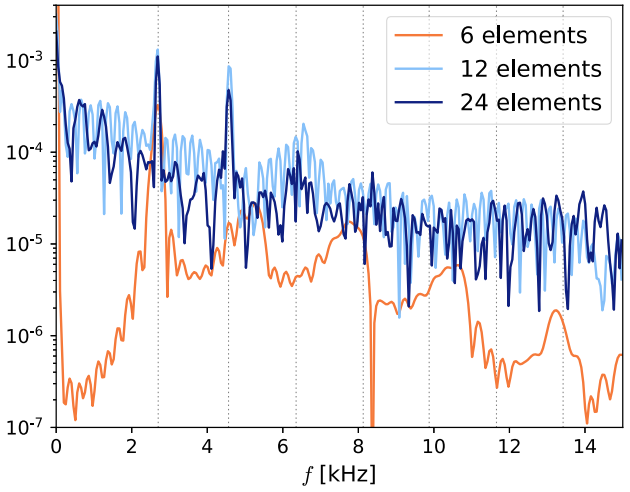

FIG. 9. The power spectrum of the maximum density for three different resolutions for the nonmagnetized TOV star simulations. The six-element simulation uses FD throughout the interior of the star, while the 12- and 24-element simulations use DG. When the high-order DG scheme is used, more oscillation frequencies are resolved. The vertical dashed lines correspond to the known frequencies in the Cowling approximation [82].

We show the normalized maximum rest mass density over the grid for the magnetized TOV star in Fig. 11. Overall the results are nearly identical to the nonmagnetized case. One notable difference is the decrease in the 12-element simulation between 7.5 and 11 ms, which

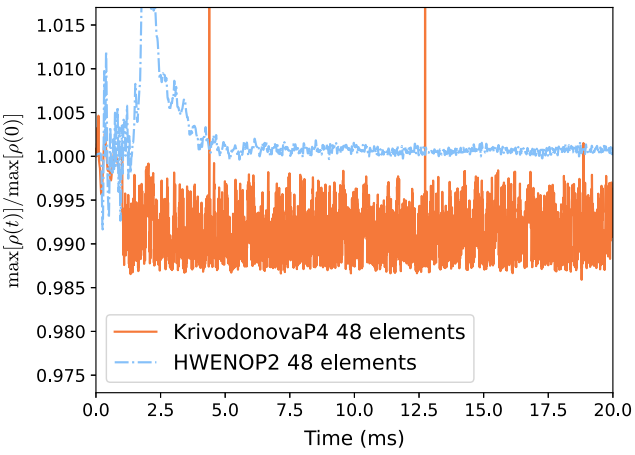

FIG. 10. The maximum density over the grid $\max(\rho)$ divided by the maximum density over the grid at $t = 0$ for the best two cases using classical limiters for the nonmagnetized TOV star simulations. The HWENO limiter is only stable for a P_2 DG solver. Simple WENO (not plotted) gives similar results. The Krivodonova limiter only succeeded at some resolutions (three of the 16 attempted runs) and the shown best result is noticeably noisier than the subcell limiter.

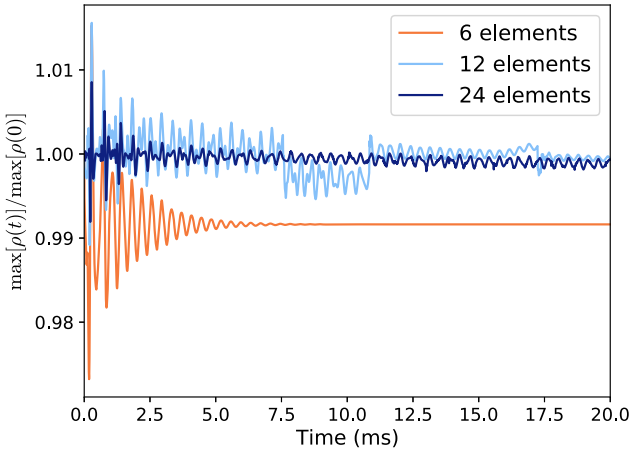

FIG. 11. The maximum density over the grid $\max(\rho)$ divided by the maximum density over the grid at $t = 0$ for three different resolutions for the magnetized TOV star simulation. The six-element simulation uses FD throughout the interior of the star, while 12- and 24-element simulations use DG. The increased high-frequency content in 12- and 24-element simulations occurs because the high-order DG scheme is able to resolve higher oscillation modes in the star. The maximum density in the six-element case drifts down at early times because of the low resolution and the relatively low accuracy of using FD at the center.

occurs because the code switches from DG to FD at the center of the star at 7.5 ms and back to DG at 11 ms. Nevertheless, the frequencies are resolved just as well for the magnetized star as for the nonmagnetized case, as can be seen in Fig. 12 where we plot the power spectrum. Specifically, we are able to resolve the three largest modes

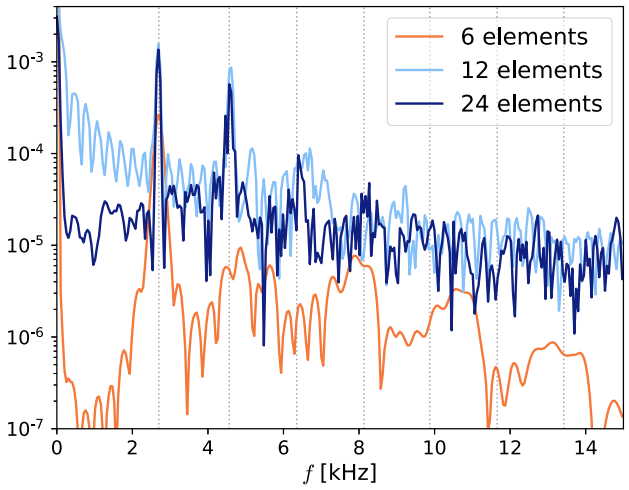

FIG. 12. The power spectrum of the maximum density for three different resolutions of the magnetized TOV star simulations. The six-element simulation uses FD throughout the interior of the star, while the 12- and 24-element simulations use DG. When the high-order DG scheme is used, more oscillation frequencies are resolved. The vertical dashed lines correspond to the known frequencies in the Cowling approximation.

with our P_5 DG-FD hybrid scheme. To the best of our knowledge, these are the first simulations of a magnetized neutron star using high-order DG methods.

V. CONCLUSIONS

We compare various shock capturing strategies to stabilize the DG method applied to the equations of general relativistic magnetohydrodynamics in the presence of discontinuities and shocks. We use the open source numerical relativity code SpECTRE [35] to perform the simulations. We compare the classic $\Lambda\Pi^N$ method [29], the hierarchical limiter of Krivodonova [30], the simple WENO limiter [31], the HWENO limiter [32], and a DG-FD hybrid approach that uses DG where the solution is smooth and HRSC FD methods where the solution contains discontinuities [56]. While many of the limiting strategies appear promising in the Newtonion hydrodynamics case, we have found stable and accurate simulations of GRMHD to be a much more challenging problem. This is in part because limiting the characteristic variables is difficult since the characteristic variables are not known analytically for the GRMHD system.

In the Newtonian hydrodynamics case, the literature advocates for using the classical limiters ($\Lambda\Pi^N$, Krivodonova, simple WENO, HWENO) on the characteristic variables of the evolution system to reduce oscillations, for using more detailed troubled-cell indicators like that of [55], and for supplementing the limiting with flattening schemes to further correct any unphysical values remaining after limiting. We have found these techniques do somewhat improve the accuracy and robustness of the limiters in the Newtonian case, but not enough to avoid the need for problem-dependent tuning of parameters, or to obtain truly robust behavior. Since these techniques do not all easily generalize to the relativistic magnetohydrodynamics case we consider here, we use the classical limiters in their simplest configuration. Our experience with limiters in Newtonian hydrodynamics suggests that limiting characteristic variables with specialized troubled-cell indicators and flatteners will likely still be problematic in the more complicated GRMHD case.

A further challenge with the classical limiters lies in extending the DG method to higher orders. With all these limiters, we consistently find large oscillations and a corresponding loss of accuracy with P_4 or higher-order DG schemes, both in Newtonian and relativistic hydrodynamics evolutions. The difficulty in robustly applying these limiters to higher-order DG schemes gives further motivation to favor the DG-FD hydrid method for scientific applications.

We find that only the DG-FD hybrid method is able to maintain stability when using a sixth-order DG scheme. The other methods are unstable or in the case of the $\Lambda\Pi^N$ limiter fall back to a linear approximation everywhere. The classical limiters all work on only some subset of the

test problems, and even there some tuning of parameters is required. While it is certainly conceivable that with enough fine-tuning each limiter could simulate most or all of the test problems, this does not make the limiter useful in scientific applications where a wide variety of different types of shock interactions and wave patterns appear. A realistic limiting strategy cannot require any fine-tuning for different problems. The only method that presents such a level of robustness is the DG-FD hybrid scheme. As a result, the DG-FD hybrid method is the only method with which we are able to simulate both magnetized and nonmagnetized TOV stars. To the best of our knowledge this paper presents the first simulations of a magnetized TOV star where DG is used.

While the DG-FD hybrid scheme is certainly the most complicated approach for shock capturing in a DG code, our results demonstrate that such complexity unfortunately seems to be necessary. We are not optimistic that any classical limiting strategy can be competitive with the DG-FD hybrid scheme since none of the methods presented in the literature are able to resolve discontinuities within a DG element. This means that discontinuities are at best only able to be resolved at the level of an entire DG element. Thus, at discontinuities the classical limiting strategies effectively turn DG into a finite volume scheme with an extremely stringent time step restriction. Switching the DG scheme to a classical WENO finite-volume-type scheme was actually the only way Ref. [21] was able to evolve a nonmagnetized TOV star.

It is unclear to us how discontinuities could be resolved inside a DG element since the basis functions are polynomials. By switching to FD, the hybrid scheme increases resolution and is able to resolve discontinuities inside an element. This can also be thought of as instead of solving the partial differential equations governing the fluid dynamics, we want to solve as many Rankine-Hugoniot conditions as possible to resolve the discontinuities as cleanly as possible.

Alternatively, we can view the hybrid scheme as a FD method where in smooth regions the solution is compressed to a high-order spectral representation to increase efficiency. The DG-FD hybrid scheme reduces the number of grid points per dimension roughly in half, and so in theory a speedup of approximately eight is expected in 3d. With the current code, we see more moderate speedups of approximately two, so there is certainly room for optimizations in SpECTRE.

In the future we plan to evolve the coupled generalized harmonic and GRMHD system together as one monolithic coupled system, generalize the DG-FD hybrid scheme to curved meshes, and use more robust positivity-preserving adaptive-order FD schemes to achieve high-order accuracy even in regions where the FD scheme is being used.

ACKNOWLEDGMENTS

Charm++/Converse [83] was developed by the Parallel Programming Laboratory in the Department of Computer Science at the University of Illinois at Urbana-Champaign. The figures in this article were produced with MATPLOTLIB [84,85], TIKZ [86] and ParaView [87,88]. Computations were performed with the Wheeler cluster at Caltech. This work was supported in part by the Sherman Fairchild Foundation and by NSF Grants No. PHY-2011961, No. PHY-2011968, and No. OAC-1931266 at Caltech, and NSF Grants No. PHY- 1912081 and No. OAC-1931280 at Cornell. P. K. acknowledges support of the Department of Atomic Energy, Government of India, under Project No. RTI4001, and by the Ashok and Gita Vaish Early Career Faculty Fellowship at the International Centre for Theoretical Sciences. M. D. acknowledges support from the NSF through Grant No. PHY-2110287. F. F. acknowledges support from the DOE through Grant No. DE-SC0020435, from NASA through Grant No. 80NSSC18K0565 and from the NSF through Grant No. PHY-1806278. G. L. is pleased to acknowledge support from the NSF through Grants No. PHY-1654359 and No. AST-1559694 and from Nicholas and Lee Begovich and the Dan Black Family Trust. H. R. R. acknowledges support from the Fundação para a Ciência e Tecnologia (FCT) within the Projects No. UID/04564/2021, No. UIDB/04564/2020, No. UIDP/04564/2020 and No. EXPL/FIS-AST/0735/2021.

[1] E. R. Most, L. J. Papenfort, and L. Rezzolla, Beyond second-order convergence in simulations of magnetized binary neutron stars with realistic microphysics, Mon. Not. R. Astron. Soc. **490**, 3588 (2019).

[2] F. Cipolletta, J. V. Kalinani, E. Giangrandi, B. Giacomazzo, R. Ciolfi, L. Sala, and B. Giudici, Spritz: General relativistic magnetohydrodynamics with neutrinos, Classical Quantum Gravity **38**, 085021 (2021).

[3] W. Reed and T. Hill, Triangular mesh methods for the neutron transport equation, *National topical meeting on mathematical models and computational techniques for analysis of nuclear systems* (1973), https://www.osti.gov/biblio/4491151.

[4] J. Hesthaven and T. Warburton, *Nodal Discontinuous Galerkin Methods: Algorithms, Analysis, and Applications* (Springer-Verlag, New York, 2008).

[5] B. Cockburn, Devising discontinuous Galerkin methods for non-linear hyperbolic conservation laws, J. Comput. Appl. Math. **128**, 187 (2001).
[6] B. Cockburn and C.-W. Shu, The Runge-Kutta discontinuous Galerkin method for conservation laws V. Multidimensional systems, J. Comput. Phys. **141**, 199 (1998).
[7] B. Cockburn, An introduction to the discontinuous Galerkin method for convection-dominated problems, in *Advanced Numerical Approximation of Nonlinear Hyperbolic Equations: Lectures given at the 2nd Session of the Centro Internazionale Matematico Estivo (C.I.M.E.) held in Cetraro, Italy, 1997*, edited by A. Quarteroni (Springer Berlin Heidelberg, Berlin, Heidelberg, 1998), pp. 150–268.
[8] B. Cockburn, G. E. Karniadakis, and C.-W. Shu, The development of discontinuous Galerkin methods, in *Discontinuous Galerkin Methods* (Springer Berlin Heidelberg, Berlin, Heidelberg, 2000), pp. 3–50.
[9] S. E. Field, J. S. Hesthaven, S. R. Lau, and A. H. Mroue, Discontinuous Galerkin method for the spherically reduced BSSN system with second-order operators, Phys. Rev. D **82**, 104051 (2010).
[10] J. D. Brown, P. Diener, S. E. Field, J. S. Hesthaven, F. Herrmann, A. H. Mroue, O. Sarbach, E. Schnetter, M. Tiglio, and M. Wagman, Numerical simulations with a first order BSSN formulation of Einstein's field equations, Phys. Rev. D **85**, 084004 (2012).
[11] S. E. Field, J. S. Hesthaven, and S. R. Lau, Discontinuous Galerkin method for computing gravitational waveforms from extreme mass ratio binaries, Classical Quantum Gravity **26**, 165010 (2009).
[12] G. Zumbusch, Finite element, discontinuous Galerkin, and finite difference evolution schemes in spacetime, Classical Quantum Gravity **26**, 175011 (2009).
[13] M. Dumbser, F. Guercilena, S. Köppel, L. Rezzolla, and O. Zanotti, Conformal and covariant Z4 formulation of the Einstein equations: Strongly hyperbolic first-order reduction and solution with discontinuous Galerkin schemes, Phys. Rev. D **97**, 084053 (2018).
[14] D. Radice and L. Rezzolla, Discontinuous Galerkin methods for general-relativistic hydrodynamics: Formulation and application to spherically symmetric spacetimes, Phys. Rev. D **84**, 024010 (2011).
[15] P. Mocz, M. Vogelsberger, D. Sijacki, R. Pakmor, and L. Hernquist, A discontinuous Galerkin method for solving the fluid and magnetohydrodynamic equations in astrophysical simulations, Mon. Not. R. Astron. Soc. **437**, 397 (2014).
[16] O. Zanotti, M. Dumbser, A. Hidalgo, and D. Balsara, An ADER-WENO finite volume AMR code for astrophysics, in Proceedings, 8th International Conference on Numerical Modeling of Space Plasma Flows (ASTRONUM 2013), Biarritz, France, 2013 [*ASP Conf. Ser.* 488, 285 (2014)].
[17] E. Endeve, C. D. Hauck, Y. Xing, and A. Mezzacappa, Bound-preserving discontinuous Galerkin methods for conservative phase space advection in curvilinear coordinates, J. Comput. Phys. **287**, 151 (2015).
[18] K. Schaal, A. Bauer, P. Chandrashekar, R. Pakmor, C. Klingenberg, and V. Springel, Astrophysical hydrodynamics with a high-order discontinuous Galerkin scheme and adaptive mesh refinement, Mon. Not. R. Astron. Soc. **453**, 4278 (2015).
[19] S. A. Teukolsky, Formulation of discontinuous Galerkin methods for relativistic astrophysics, J. Comput. Phys. **312**, 333 (2016).
[20] L. E. Kidder *et al.*, SpECTRE: A task-based discontinuous Galerkin code for relativistic astrophysics, J. Comput. Phys. **335**, 84 (2017).
[21] M. Bugner, T. Dietrich, S. Bernuzzi, A. Weyhausen, and B. Brügmann, Solving 3D relativistic hydrodynamical problems with weighted essentially nonoscillatory discontinuous Galerkin methods, Phys. Rev. D **94**, 084004 (2016).
[22] O. Zanotti, F. Fambri, and M. Dumbser, Solving the relativistic magnetohydrodynamics equations with ADER discontinuous Galerkin methods, *a posteriori* subcell limiting and adaptive mesh refinement, Mon. Not. R. Astron. Soc. **452**, 3010 (2015).
[23] M. Dumbser, F. Fambri, M. Tavelli, M. Bader, and T. Weinzierl, Efficient implementation of ADER discontinuous Galerkin schemes for a scalable hyperbolic PDE engine, Axioms **7**, 63 (2018).
[24] F. Fambri, M. Dumbser, S. Köppel, L. Rezzolla, and O. Zanotti, ADER discontinuous Galerkin schemes for general-relativistic ideal magnetohydrodynamics, Mon. Not. R. Astron. Soc. **477**, 4543 (2018).
[25] O. Zanotti, F. Fambri, M. Dumbser, and A. Hidalgo, Space-time adaptive ader discontinuous Galerkin finite element schemes with a posteriori sub-cell finite volume limiting, Comput. Fluids **118**, 204 (2015).
[26] O. Zanotti and M. Dumbser, Efficient conservative ADER schemes based on WENO reconstruction and space-time predictor in primitive variables, Comput. Astrophys. Cosmol. **3**, 1 (2016).
[27] J. M. Miller and E. Schnetter, An operator-based local discontinuous Galerkin method compatible with the BSSN formulation of the Einstein equations, Classical Quantum Gravity **34**, 015003 (2017).
[28] F. Hébert, L. E. Kidder, and S. A. Teukolsky, General-relativistic neutron star evolutions with the discontinuous Galerkin method, Phys. Rev. D **98**, 044041 (2018).
[29] B. Cockburn, Discontinuous Galerkin methods for convection-dominated problems, in *High-Order Methods for Computational Physics*, edited by T. J. Barth and H. Deconinck (Springer Berlin Heidelberg, Berlin, Heidelberg, 1999), pp. 69–224.
[30] L. Krivodonova, Limiters for high-order discontinuous Galerkin methods, J. Comput. Phys. **226**, 879 (2007).
[31] X. Zhong and C.-W. Shu, A simple weighted essentially nonoscillatory limiter for Runge-Kutta discontinuous Galerkin methods, J. Comput. Phys. **232**, 397 (2013).
[32] J. Zhu, X. Zhong, C.-W. Shu, and J. Qiu, Runge-Kutta discontinuous Galerkin method with a simple and compact Hermite WENO limiter, Commun. Comput. Phys. **19**, 944 (2016).
[33] B. Costa and W. S. Don, Multi-domain hybrid spectral-WENO methods for hyperbolic conservation laws, J. Comput. Phys. **224**, 970 (2007).
[34] M. Dumbser, O. Zanotti, R. Loubère, and S. Diot, *A posteriori* subcell limiting of the discontinuous Galerkin finite element method for hyperbolic conservation laws, J. Comput. Phys. **278**, 47 (2014).

[35] N. Deppe, W. Throwe, L. E. Kidder, N. L. Vu, F. Hébert, J. Moxon, C. Armaza, G. S. Bonilla, Y. Kim, P. Kumar, G. Lovelace, A. Macedo, K. C. Nelli, E. O'Shea, H. P. Pfeiffer, M. A. Scheel, S. A. Teukolsky, N. A. Wittek *et al.*, SpECTRE v2022.04.04, 10.5281/zenodo.6412468.

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

[37] L. Rezzolla and O. Zanotti, *Relativistic Hydrodynamics* (Oxford University Press, New York, 2013).

[38] L. Antón, O. Zanotti, J. A. Miralles, J. M. Martí, J. M. Ibáñez, J. A. Font, and J. A. Pons, Numerical 3 + 1 general relativistic magnetohydrodynamics: A local characteristic approach, Astrophys. J. **637**, 296 (2006).

[39] J. A. Font, Numerical hydrodynamics and magnetohydrodynamics in general relativity, Living Rev. Relativity **11**, 7 (2008).

[40] C. W. Misner, K. S. Thorne, and J. A. Wheeler, *Gravitation* (Freeman, New York, 1973).

[41] A. Dedner, F. Kemm, D. Kröner, C.-D. Munz, T. Schnitzer, and M. Wesenberg, Hyperbolic divergence cleaning for the MHD equations, J. Comput. Phys. **175**, 645 (2002).

[42] P. Mösta, B. C. Mundim, J. A. Faber, R. Haas, S. C. Noble, T. Bode, F. Löffler, C. D. Ott, C. Reisswig, and E. Schnetter, GRHydro: A new open source general-relativistic magnetohydrodynamics code for the EINSTEIN Toolkit, Classical Quantum Gravity **31**, 015005 (2014).

[43] D. Hilditch and A. Schoepe, Hyperbolicity of divergence cleaning and vector potential formulations of general relativistic magnetohydrodynamics, Phys. Rev. D **99**, 104034 (2019).

[44] C. F. Gammie, J. C. McKinney, and G. Tóth, HARM: A numerical scheme for general relativistic magnetohydrodynamics, Astrophys. J. **589**, 444 (2003).

[45] S. A. Teukolsky, Short note on the mass matrix for Gauss-Lobatto grid points, J. Comput. Phys. **283**, 408 (2015).

[46] V. Rusanov, The calculation of the interaction of non-stationary shock waves and obstacles, USSR Computational Mathematics and Mathematical Physics **1**, 304 (1962).

[47] A. Harten, P. Lax, and B. Leer, On upstream differencing and Godunov-type schemes for hyperbolic conservation laws, SIAM Rev. **25**, 35 (1983).

[48] E. Toro, *Riemann Solvers and Numerical Methods for Fluid Dynamics: A Practical Introduction* (Springer, Berlin Heidelberg, 2009).

[49] S. Davis, Simplified second-order Godunov-type methods, SIAM J. Sci. Stat. Comput. **9**, 445 (1988).

[50] W. Throwe and S. Teukolsky, A high-order, conservative integrator with local time-stepping, SIAM J. Sci. Comput. **42**, A3730 (2020).

[51] B. Cockburn, S. Hou, and C.-W. Shu, The Runge-Kutta local projection discontinuous Galerkin finite element method for conservation laws. IV. The multidimensional case, Math. Comput. **54**, 545 (1990).

[52] S. A. Moe, J. A. Rossmanith, and D. C. Seal, A simple and effective high-order shock-capturing limiter for discontinuous Galerkin methods, arXiv:1507.03024.

[53] B. Cockburn, S.-Y. Lin, and C.-W. Shu, TVB Runge Kutta local projection discontinuous Galerkin finite element method for conservation laws III: One-dimensional systems, J. Comput. Phys. **84**, 90 (1989).

[54] M. Dumbser and M. Käser, Arbitrary high order non-oscillatory finite volume schemes on unstructured meshes for linear hyperbolic systems, J. Comput. Phys. **221**, 693 (2007).

[55] L. Krivodonova, J. Xin, J.-F. Remacle, N. Chevaugeon, and J. Flaherty, Shock detection and limiting with discontinuous Galerkin methods for hyperbolic conservation laws, Applied Numerical Mathematics **48**, 323 (2004).

[56] N. Deppe, F. Hébert, L. E. Kidder, and S. A. Teukolsky, A high-order shock capturing discontinuous Galerkin-finite-difference hybrid method for GRMHD, arXiv:2109.11645.

[57] D. M. Siegel, P. Mösta, D. Desai, and S. Wu, Recovery schemes for primitive variables in general-relativistic magnetohydrodynamics, Astrophys. J. **859**, 71 (2018).

[58] W. Kastaun, J. V. Kalinani, and R. Ciolfi, Robust recovery of primitive variables in relativistic ideal magnetohydrodynamics, Phys. Rev. D **103**, 023018 (2021).

[59] W. Newman and N. Hamlin, Primitive variable determination in conservative relativistic magnetohydrodynamic simulations, SIAM J. Sci. Comput. **36**, B661 (2014).

[60] C. Palenzuela, S. L. Liebling, D. Neilsen, L. Lehner, O. L. Caballero, E. O'Connor, and M. Anderson, Effects of the microphysical equation of state in the mergers of magnetized neutron stars with neutrino cooling, Phys. Rev. D **92**, 044045 (2015).

[61] F. Foucart, M. B. Deaton, M. D. Duez, L. E. Kidder, I. MacDonald, C. D. Ott, H. P. Pfeiffer, M. A. Scheel, B. Szilagyi, and S. A. Teukolsky, Black-hole-neutron-star mergers at realistic mass ratios: Equation of state and spin orientation effects, Phys. Rev. D **87**, 084006 (2013).

[62] F. Galeazzi, W. Kastaun, L. Rezzolla, and J. A. Font, Implementation of a simplified approach to radiative transfer in general relativity, Phys. Rev. D **88**, 064009 (2013).

[63] C. D. Muhlberger, F. H. Nouri, M. D. Duez, F. Foucart, L. E. Kidder, C. D. Ott, M. A. Scheel, B. Szilágyi, and S. A. Teukolsky, Magnetic effects on the low-T/|W| instability in differentially rotating neutron stars, Phys. Rev. D **90**, 104014 (2014).

[64] F. Foucart, Numerical studies of black hole-neutron star binaries, Ph.D. thesis, Cornell University, 2011.

[65] D. S. Balsara, Self-adjusting, positivity preserving high order schemes for hydrodynamics and magnetohydrodynamics, J. Comput. Phys. **231**, 7504 (2012).

[66] D. Balsara, Total variation diminishing scheme for relativistic magnetohydrodynamics, Astrophys. J. Suppl. Ser. **132**, 83 (2001).

[67] B. Giacomazzo and L. Rezzolla, The exact solution of the Riemann problem in relativistic magnetohydrodynamics, J. Fluid Mech. **562**, 223 (2006).

[68] T. Leismann, L. Antón, M. A. Aloy, E. Müller, J. M. Martí, J. A. Miralles, and J. M. Ibáñez, Relativistic MHD simulations of extragalactic jets, Astron. Astrophys. **436**, 503 (2005).

[69] L. Del Zanna, O. Zanotti, N. Bucciantini, and P. Londrillo, ECHO: A Eulerian conservative high order scheme for general relativistic magnetohydrodynamics and magnetodynamics, Astron. Astrophys. **473**, 11 (2007).

[70] D. S. Balsara and D. S. Spicer, A staggered mesh algorithm using high order Godunov fluxes to ensure solenoidal magnetic fields in magnetohydrodynamic simulations, J. Comput. Phys. **149**, 270 (1999).
[71] G. Tóth, The $\nabla \cdot b = 0$ constraint in shock-capturing magnetohydrodynamics codes, J. Comput. Phys. **161**, 605 (2000).
[72] Z. B. Etienne, Y. T. Liu, and S. L. Shapiro, Relativistic magnetohydrodynamics in dynamical spacetimes: A new adaptive mesh refinement implementation, Phys. Rev. D **82**, 084031 (2010).
[73] L. Del Zanna, N. Bucciantini, and P. Londrillo, An efficient shock-capturing central-type scheme for multidimensional relativistic flows. II. Magnetohydrodynamics, Astron. Astrophys. **400**, 397 (2003).
[74] C. R. DeVore, Flux-corrected transport techniques for multidimensional compressible magnetohydrodynamics, J. Comput. Phys. **92**, 142 (1991).
[75] K. Beckwith and J. M. Stone, A second-order Godunov method for multi-dimensional relativistic magnetohydrodynamics, Astrophys. J. Suppl. Ser. **193**, 6 (2011).
[76] T. A. Gardiner and J. M. Stone, An unsplit Godunov method for ideal MHD via constrained transport, J. Comput. Phys. **205**, 509 (2005).
[77] J. M. Stone, T. A. Gardiner, P. Teuben, J. F. Hawley, and J. B. Simon, ATHENA: A new code for astrophysical MHD, Astrophys. J. Suppl. Ser. **178**, 137 (2008).
[78] K. Beckwith and J. M. Stone, A second order Godunov method for multidimensional relativistic magnetohydrodynamics, Astrophys. J. Suppl. Ser. **193**, 6 (2011).
[79] R. C. Tolman, Static solutions of Einstein's field equations for spheres of fluid, Phys. Rev. **55**, 364 (1939).
[80] J. R. Oppenheimer and G. M. Volkoff, On massive neutron cores, Phys. Rev. **55**, 374 (1939).
[81] F. Cipolletta, J. V. Kalinani, B. Giacomazzo, and R. Ciolfi, SPRITZ: A new fully general-relativistic magnetohydrodynamic code, Classical Quantum Gravity **37**, 135010 (2020).
[82] J. A. Font, T. Goodale, S. Iyer, M. Miller, L. Rezzolla, E. Seidel, N. Stergioulas, W.-M. Suen, and M. Tobias, Three-dimensional numerical general relativistic hydrodynamics. II. Long-term dynamics of single relativistic stars, Phys. Rev. D **65**, 084024 (2002).
[83] L. Kale *et al.*, Uiuc-ppl/charm: CHARM++ version 6.10.2 (2020).
[84] J. D. Hunter, MATPLOTLIB: A 2d graphics environment, Comput. Sci. Eng. **9**, 90 (2007).
[85] T. A. Caswell *et al.*, matplotlib/matplotlib: Rel: v3.3.0, 2020.
[86] T. Tantau, The TIKZ and PGF packages.
[87] U. Ayachit, *The ParaView Guide: A Parallel Visualization Application* (Kitware, Inc., Clifton Park, NY, 2015).
[88] J. Ahrens, B. Geveci, and C. Law, ParaView: An end-user tool for large-data visualization, in *The Visualization Handbook* (Butterworth-Heinemann, Burlington, 2005).