

PAPER

AsterX: a new open-source GPU-accelerated GRMHD code for dynamical spacetimes

To cite this article: Jay V Kalinani *et al* 2025 *Class. Quantum Grav.* **42** 025016

View the [article online](#) for updates and enhancements.

You may also like

- [One-Pot Electrochemical Oxidation of Alcohols to Nitriles Mediated by TEMPO](#)
Zhongquan Fan, Xianjing Yang, Chen Chen et al.
- [Dental restorative materials and halitosis: a preliminary *in-vitro* study](#)
Tuçe Odaba and Osman Tolga Haror
- [Theory of Transport in Highly Concentrated Electrolytes](#)
Max Schammer, Birger Horstmann and Arnulf Latz

AsterX: a new open-source GPU-accelerated GRMHD code for dynamical spacetimes

Jay V Kalinani^{1,*} , Liwei Ji¹ , Lorenzo Ennoggi¹ ,
Federico G Lopez Armengol¹ ,
Lucas Timotheo Sanches² , Bing-Jyun Tsao³ ,
Steven R Brandt⁴ , Manuela Campanelli¹ ,
Riccardo Ciolfi^{5,6} , Bruno Giacomazzo^{7,8,9} ,
Roland Haas^{10,11} , Erik Schnetter^{4,12,13} ,
and Yosef Zlochower¹ 

¹ Center for Computational Relativity and Gravitation, & School of Mathematical Sciences, Rochester Institute of Technology, 85 Lomb Memorial Drive, Rochester, NY 14623, United States of America

² Programa de Engenharia de Sistemas e Computação (PESC/COPPE), Universidade Federal do Rio de Janeiro (UFRJ), Avenida Horácio Macedo 2030, Centro de Tecnologia, Bloco H, sala 319, Rio De Janeiro, RJ 21941-914, Brazil

³ Center for Gravitational Physics, Department of Physics, The University of Texas at Austin, Austin, TX 78712, United States of America

⁴ Center for Computation & Technology, Louisiana State University, Baton Rouge, LA, United States of America

⁵ INAF, Osservatorio Astronomico di Padova, Vicolo dell'Osservatorio 5, I-35122 Padova, Italy

⁶ INFN, Sezione di Padova, Via Francesco Marzolo 8, I-35131 Padova, Italy

⁷ Dipartimento di Fisica G Occhialini, Università di Milano-Bicocca, Piazza della Scienza 3, I-20126 Milano, Italy

⁸ INFN, Sezione di Milano-Bicocca, Piazza della Scienza 3, I-20126 Milano, Italy

⁹ INAF, Osservatorio Astronomico di Brera, Via E Bianchi 46, I-23807 Merate, Italy

¹⁰ National Center for Supercomputing Applications, University of Illinois, 1205 W Clark St, Urbana, IL, United States of America

¹¹ Department of Physics, University of Illinois, 1110 West Green St, Urbana, IL, United States of America

¹² Perimeter Institute for Theoretical Physics, Waterloo, Ontario, Canada

¹³ Department of Physics and Astronomy, University of Waterloo, Waterloo, Ontario, Canada

* Author to whom any correspondence should be addressed.

E-mail: jaykalinani@gmail.com

Received 17 June 2024; revised 1 December 2024

Accepted for publication 9 December 2024

Published 27 December 2024



CrossMark

Abstract

We present **AsterX**, a novel open-source, modular, GPU-accelerated, fully general relativistic magnetohydrodynamic (GRMHD) code designed for dynamic spacetimes in 3D Cartesian coordinates, and tailored for exascale computing. We utilize block-structured adaptive mesh refinement (AMR) through **CarpetX**, the new driver for the **Einstein Toolkit**, which is built on **AMReX**, a software framework for massively parallel applications. **AsterX** employs the Valencia formulation for GRMHD, coupled with the ‘Z4c’ formalism for space-time evolution, while incorporating high resolution shock capturing schemes to accurately handle the hydrodynamics. **AsterX** has undergone rigorous testing in both static and dynamic spacetime, demonstrating remarkable accuracy and agreement with other codes in literature. Using subcycling in time, we find an overall performance gain of factor 2.5–4.5. Benchmarking the code through scaling tests on OLCF’s Frontier supercomputer, we demonstrate a weak scaling efficiency of about 67%–77% on 4096 nodes compared to an 8-node performance.

Keywords: general relativity, magnetohydrodynamics, exascale computing, GPUs

1. Introduction

Recent years have witnessed remarkable strides in multimessenger astrophysics, driven by technological advancements and collaborative efforts among global observatories. This progress was kick-started by the breakthrough detection of gravitational waves (GWs) from a binary neutron star (BNS) merger event ‘GW170817’ [1], accompanied by rich variety of electromagnetic (EM) counterparts including a short gamma-ray burst (SGRB) [2] and an ultraviolet/optical/infrared transient consistent with a radioactively powered kilonova [3] (GRB 170817A and AT2017gfo, respectively), flooding the entire EM spectrum from gamma-rays to radio [4]. The comprehensive data collected from this single event provided critical insights into the origin of heavy elements through r-process nucleosynthesis [3, 5, 6], the nature of short gamma-ray bursts [2, 7–9], and the behavior of matter under extreme conditions [10–13]. Undoubtedly, the plethora of science extracted from this event has been unprecedented. Subsequently, a number of compact binary mergers involving at least one neutron star have been detected through GWs [14–16], making such events a cornerstone of the field. In the next decade or two, next-generation ground-based detectors (3G) [17, 18] are expected to expand our ability to observe these systems. On much larger scales instead, accreting supermassive binary black holes (SMBBHs) continue to be as promising multimessenger sources due to their strong emission of both GWs and EM radiation. While low-frequency GWs have already been detected by pulsar timing arrays [19] and are anticipated to be detected by the Laser Interferometer Space Antenna [20, 21], discovering SMBBH mergers in this data would offer profound insights for astrophysics and cosmology (e.g. [22]). Additionally, the combination of EM and GW data

will allow for a more comprehensive understanding of the environments surrounding SMBBH mergers and provide unique cosmological tests, such as measuring the Universe's expansion history.

General relativistic magnetohydrodynamic (GRMHD) simulations remain an indispensable tool for probing the intricate dynamics of relativistic plasmas in the vicinity of such compact astrophysical sources. For instance, in recent years, various GRMHD simulations of BNS and black hole-neutron star (BHNS) mergers have demonstrated that magnetic fields play a critical role in launching relativistic jets, which can act as a pre-cursor or potentially power SGRBs [23–30]. Additionally, they can have a profound impact on matter ejection, contributing to the kilonova production [31–34]. On the other hand, for accurate modelling of the gas dynamics in SMBBH merger simulations, magnetic fields prove to be fundamental in driving turbulence and accretion processes from the circumbinary disks [35–37] to the minidisks surrounding each SMBH [38–42], as well as in jet generation [41, 43–45].

As the computational demands of these multi-physics simulations are substantial, they are typically limited in time and/or length scales. Existing state-of-the-art GRMHD codes such as WhiskyMHD [46, 47], GRHydro [48], HARM3D [49–51], IllinoisGRMHD [52, 53], MHDuet [54–56], Spritz [57–59], GRAthena++ [60], employed to run such simulations, are designed to work on central processing unit (CPU) architectures that are not efficient enough for massive parallel processing. To tackle these challenges, modern GRMHD codes are currently being redesigned directly or being build on infrastructures that support graphical processing unit (GPU) architectures such as Parthenon [61] (based instead on Kokkos library [62]), and AMReX [63, 64] which also provide block structured adaptive mesh refinement (AMR) capabilities. This is also driven by the revolutionary advancement in latest high-performance computing (HPC) supercomputers that employ GPUs to achieve exascale performance. Availing this advantage, the newest GPU-friendly (GR)MHD codes coupled with either static or dynamical spacetimes such as K-Athena [65], H-AMR [66], GRaM-X [67] have reported a performance speed up of a factor 2–10 in challenging simulations.

In this paper, we present *AsterX*, a GPU-accelerated GRMHD code designed to work with dynamical spacetimes in Cartesian coordinates. *AsterX* is built on top of *CarpetX*, the new driver for the *Einstein Toolkit* [68] that provides AMR capability via the AMReX framework. Developed entirely in C++ from scratch as an open-source code [69], *AsterX* incorporates numerous algorithms derived from the *Spritz* code [57–59]. Like *Spritz*, it also adopts the Valencia formulation for GRMHD equations, and evolves the staggered version of the vector potential to preserve the divergenceless character of the magnetic field. At the moment, the code supplies only analytical equations of state (EOS), and support for tabulated microphysical EOSs is currently underway. For dynamical evolution, we utilize the spacetime solver based on the 'Z4c' formulation [70, 71] of the Einstein's field equations. Here, we also present a series of extensive tests in 1D, 2D, and 3D that validate our implementation, and provide performance benchmarks conducted on OLCF's Frontier supercomputer, scaling up to 4096 nodes.

The paper is organized as follows. In section 2, we briefly describe the theoretical formulations on which our code is based. Numerical implementation of the adopted schemes are outlined in section 3. In section 4, we showcase the results of various tests performed in special and general relativistic regimes. Performance benchmarks are summarized in section 5. Finally, we present our conclusions and future outlook in section 6. Unless stated otherwise, we adopt the space-like signature as $(-, +, +, +)$ and use geometric units $G = c = M_{\odot} = 1$, and follow the standard Einstein's convention for the summation over repeated indices.

2. Background equations

In this section, we briefly provide the theoretical background, describing the formulations adopted in `AsterX`. First, we focus on the *3+1 foliation* technique adopted to arrive at numerical solutions of the Einstein field equations, followed by the *3+1 formulation* of GRMHDs equations adopted for matter evolution.

2.1. 3+1 spacetime

The geometry of the 4-dimensional spacetime is governed by the Einstein's equations

$$G_{\mu\nu} = 8\pi T_{\mu\nu}, \quad (1)$$

where $G_{\mu\nu}$ is the Einstein tensor, and $T_{\mu\nu}$ is the total stress-energy tensor. Our implementation employs the standard decomposition of spacetime into 3-dimensional space-like hypersurfaces associated with the line element of the form

$$ds^2 = g_{\mu\nu} dx^\mu dx^\nu = -\alpha^2 dt^2 + \gamma_{ij} (dx^i + \beta^i dt) (dx^j + \beta^j dt), \quad (2)$$

where γ_{ij} are the spatial components of the spacetime 4-metric $g_{\mu\nu}$, α is the *lapse* function and β^i is the *shift* vector. We then define \mathbf{n} as the unit 4-vector orthonormal to the hypersurfaces Σ_t , which has the following components

$$n^\mu = \frac{1}{\alpha} (1, -\beta^i), \quad n_\mu = (-\alpha, 0, 0, 0). \quad (3)$$

The extrinsic curvature is given by the Lie derivative of the three-metric in the \mathbf{n} direction as

$$K_{\mu\nu} = -\frac{1}{2} \mathcal{L}_{\mathbf{n}} \gamma_{\mu\nu} \quad (4)$$

For the evolution of spacetime, we consider the conformal decomposition of the Z4 formulation, commonly denoted as Z4c [70, 71]. In the Z4c system, the Einstein equations are augmented with an additional 4-vector Z^μ such that in the extended evolution system, the Hamiltonian and momentum constraints are damped to zero. The gauge conditions employed with this system include the $1 + \log$ slicing and the Γ -driver shift. For convenience, the so-call ADM variables α , β^i , γ_{ij} , and K_{ij} are made available whenever the Z4c state vector is updated. Thus, `AsterX` is decoupled from details of the Z4c formulation. For further details of this evolution system, we refer to [70, 71].

2.2. Valencia formulation

Analogous to the 3+1 decomposition of 4D spacetime, we require a 3+1 formulation to evolve the matter field equations involving magnetic fields. This section discusses the 3+1 GRMHD formulation implemented in `AsterX`.

First, we define the 4D rest-mass density current and the energy-momentum tensor for a perfect fluid as

$$J^\mu = \rho u^\mu, \quad T^{\mu\nu} = (\rho + \rho\epsilon + p) u^\mu u^\nu + pg^{\mu\nu} = h\rho u^\mu u^\nu + pg^{\mu\nu}, \quad (5)$$

where p is the gas pressure, ρ stands for rest-mass density, u^μ is the four velocity of the fluid, ϵ is the specific internal energy, and $h = 1 + \epsilon + p/\rho$ denotes the specific relativistic enthalpy.

The equation for the conservation of matter current density i.e. the continuity equation, and for the conservation of energy-momentum are written respectively as

$$\nabla_\mu J^\mu = \nabla_\mu (\rho u^\mu) = 0, \quad \nabla_\mu T^{\mu\nu} = \nabla_\mu [h \rho u^\mu u^\nu + p g^{\mu\nu}] = 0. \quad (6)$$

To include the effect of electric and magnetic fields in the GRHD formalism, we start with the Faraday EM tensor $F^{\mu\nu}$ and its dual ${}^*F^{\mu\nu} = (1/2)\eta^{\mu\nu\sigma\delta}F_{\sigma\delta}$, which are given by the expressions

$$F^{\mu\nu} = \tilde{U}^\mu E^\nu - \tilde{U}^\nu E^\mu - \eta^{\mu\nu\lambda\delta} \tilde{U}_\lambda B_\delta, \quad (7)$$

$${}^*F^{\mu\nu} = \frac{1}{2}\eta^{\mu\nu\lambda\delta}F_{\lambda\delta} = \tilde{U}^\mu B^\nu - \tilde{U}^\nu B^\mu - \eta^{\mu\nu\lambda\delta} \tilde{U}_\lambda E_\delta, \quad (8)$$

where E^μ is the electric field, B^μ stands for the magnetic field, \tilde{U}^μ a generic observer's four-velocity, and $\eta^{\mu\nu\lambda\delta} = \frac{1}{\sqrt{-g}} [\mu\nu\lambda\delta]$ is the Levi-Civita pseudo-tensor representing the volume element. For the evolution of EM fields, we have the well-known Maxwell's equations

$$\nabla_\nu {}^*F^{\mu\nu} = 0, \quad \nabla_\nu F^{\mu\nu} = 4\pi \mathcal{J}^\mu, \quad (9)$$

where \mathcal{J}^μ is the charge-current four-vector, expressed using Ohm's law as

$$\mathcal{J}^\mu = qu^\mu + \sigma F^{\mu\nu} u_\nu, \quad (10)$$

where q is the proper charge density and σ represents the electrical conductivity. Under the assumption of ideal MHD, which considers a perfectly conducting fluid such that $\sigma \rightarrow \infty$ and $F^{\mu\nu} u_\nu = 0$ (implying that the comoving observer measures no electric field), the EM tensor can be re-written solely in terms of the magnetic field b measured in the comoving frame as

$$F^{\lambda\delta} = \eta^{\mu\nu\lambda\delta} b_\mu u_\nu, \quad {}^*F^{\mu\nu} = b^\mu u^\nu - b^\nu u^\mu = \frac{u^\mu B^\nu - u^\nu B^\mu}{W}, \quad (11)$$

whereas the Maxwell equations become

$$\nabla_\nu {}^*F^{\mu\nu} = \frac{1}{\sqrt{-g}} \partial_\nu [\sqrt{-g} (b^\mu u^\nu - b^\nu u^\mu)] = 0, \quad (12)$$

where $W = 1/\sqrt{1 - v^i v_i}$ is the Lorentz factor and b is the magnetic field measured by a comoving observer, which can then be written in terms of the magnetic field measured by an Eulerian observer \tilde{B} as

$$b^0 = \frac{WB^i v_i}{\alpha}, \quad b^i = \frac{B^i + \alpha b^0 u^i}{W}, \quad b^2 \equiv b^\mu b_\mu = \frac{B^2 + \alpha^2 (b^0)^2}{W^2}, \quad (13)$$

with $B^2 \equiv B^i B_i$. Considering $\tilde{B}^i \equiv \sqrt{\gamma} B^i$ where γ is the determinant of γ_{ij} , and $\tilde{v}^i \equiv \alpha v^i - \beta^i$, equation (12) can be separated into time and spatial components, with the time component giving the *divergence-free condition* as

$$\partial_i \tilde{B}^i = 0, \quad (14)$$

while the spatial component giving the magnetic-field *induction equations* as

$$\partial_i \tilde{B}^i = \partial_j (\tilde{v}^i \tilde{B}^j - \tilde{v}^j \tilde{B}^i), \quad (15)$$

Including the magnetic field contribution, the stress-energy tensor can be redefined as

$$T^{\mu\nu} = (\rho h + b^2) u^\mu u^\nu + (p + p_{\text{mag}}) g^{\mu\nu} - b^\mu b^\nu, \quad (16)$$

where $p_{\text{mag}} \equiv \frac{b^2}{2}$ is the magnetic pressure. The set of equation (6) can then be written in the following hyperbolic, first-order flux-conservative form

$$\frac{1}{\sqrt{-g}} \left\{ \partial_t [\sqrt{\gamma} \mathbf{F}^0(\mathbf{U})] + \partial_i [\sqrt{-g} \mathbf{F}^{(i)}(\mathbf{U})] \right\} = \mathbf{S}(\mathbf{U}) , \quad (17)$$

where $\mathbf{F}^{(i)}(\mathbf{U})$ stands for the flux-vector in the i -direction, $\mathbf{S}(\mathbf{U})$ are the source terms, and $\sqrt{-g} = \alpha \sqrt{\gamma}$. Except for the dependence on stress-energy tensor, the source terms depend only on the metric and its derivatives. In order to recast equation (6) into a conservative form (17), the *primitive* magnetohydrodynamical variables $\mathbf{U} \equiv (\rho, v^i, \epsilon, B)$ are mapped into a set of conserved variables $\mathbf{F}^0(\mathbf{U}) \equiv (D, S_i, \tau, \tilde{B})$ which are defined in the following way:

$$\mathbf{F}^0 = \begin{pmatrix} D \\ S_j \\ \tau \\ B^k \end{pmatrix} , \quad \mathbf{F}^i = \begin{pmatrix} D \tilde{v}^i / \alpha \\ S_j \tilde{v}^i / \alpha + (p + b^2/2) \delta_j^i - b_j B^i / W \\ \tau \tilde{v}^i / \alpha + (p + b^2/2) v^i - \alpha b^0 B^i / W \\ B^k \tilde{v}^i / \alpha - B^i \tilde{v}^k / \alpha \end{pmatrix} , \quad (18)$$

$$\mathbf{S} = \begin{pmatrix} 0 \\ T^{\mu\nu} (\partial_\mu g_{\nu j} - \Gamma_{\nu\mu}^\delta g_{\delta j}) \\ \alpha (T^{\mu 0} \partial_\mu \ln \alpha - T^{\mu\nu} \Gamma_{\nu\mu}^0) \\ 0^k \end{pmatrix} , \quad (19)$$

where $0^k = (0, 0, 0)^T$ and

$$D \equiv \rho W \quad (20)$$

$$S_j \equiv (\rho h + b^2) W^2 v_j - \alpha b^0 b_j \quad (21)$$

$$\tau \equiv (\rho h + b^2) W^2 - (p + b^2/2) - \alpha^2 (b^0)^2 - D . \quad (22)$$

Along with equations (18) and (19), we require an equation of state (EOS) which relates the pressure to the rest-mass density ρ and to the specific internal energy ϵ , in order to close the system of equations for hydrodynamics.

2.3. EM gauge conditions

To ensure the divergence-free constraint (14), one common approach is to evolve the EM vector potential instead of magnetic field variables. To do so, we first express magnetic field \vec{B} in terms of the vector potential \vec{A} . Considering ∇ as a purely spatial operator, we can write $\vec{B} = \nabla \times \vec{A}$. Taking the divergence of \vec{B} , we get

$$\nabla \cdot \vec{B} = \nabla \cdot (\nabla \times \vec{A}) = 0 . \quad (23)$$

Therefore, by construction, evolving the vector potential \vec{A} automatically satisfies (23).

We now introduce the four-vector potential as $\mathcal{A}_\nu = n_\nu \Phi + A_\nu$, where A_ν is the purely spatial vector potential, whereas Φ stands for the scalar potential. The magnetic field, as measured by an Eulerian observer, can then be written as

$$B^i = \epsilon^{ijk} \partial_j A_k , \quad (24)$$

and the induction equation (15) can be re-written in terms of A_i as

$$\partial_t A_i = -E_i - \partial_i (\alpha \Phi - \beta^j A_j) , \quad (25)$$

where $\epsilon^{ijk} = n_\nu \epsilon^{\nu ijk}$ is the 3-dimensional spatial Levi–Civita tensor.

However, the gauge freedom admitted by Maxwell's equations renders the choice of the four-vector potential \mathcal{A}^ν as not unique, and allows setting a suitable gauge.

2.3.1. Algebraic gauge. The numerical codes that first employed vector potential evolution equations to perform GRMHD simulations considered using the *algebraic* gauge [72, 73], where the scalar potential is set to $\Phi = \frac{1}{\alpha}(\beta^j A_j)$. This choice leads to a much simplified version of (25), reducing it to $\partial_t A_i = -E_i$, and therefore, does not necessitate the evolution of the scalar potential Φ .

2.3.2. Lorenz gauge. The *Lorenz* gauge is another choice that has been recently adopted in GRMHD simulations [72], which is based on the constraint $\nabla_\nu \mathcal{A}^\nu = 0$, which serves as an advection equation for \mathcal{A} . In order to impose this constraint, this gauge condition also requires solving the following evolution equation for the scalar potential

$$\partial_t(\sqrt{\gamma}\Phi) + \partial_i(\alpha\sqrt{\gamma}A^i - \sqrt{\gamma}\beta^i\Phi) = 0. \quad (26)$$

The Lorenz gauge has proven to be advantageous over the algebraic gauge in GRMHD simulations which utilize AMR, for e.g. BNS and NSBH merger simulations [72]. The use of algebraic gauge in such simulations with AMR leads to generation of static gauge modes which then cause interpolation errors at the refinement boundaries, producing spurious magnetic fields near the boundary regions (see [72] for more details), rendering this gauge choice unsuitable for BNS, NSBH or SMBBH merger simulations, and more generally, for all simulations which involve magnetized matter crossing refinement boundaries. A more robust gauge choice has been introduced in [74] with the name of *generalized Lorenz* gauge, which considers

$$\nabla_\nu \mathcal{A}^\nu = \xi n_\nu \mathcal{A}^\nu, \quad (27)$$

where $\xi n_\nu \mathcal{A}^\nu$ acts like a damping term in the advection equation for \mathcal{A} . When employing this gauge choice, the evolution equation for the scalar potential becomes

$$\partial_t(\sqrt{\gamma}\Phi) + \partial_i(\alpha\sqrt{\gamma}A^i - \sqrt{\gamma}\beta^i\Phi) = -\xi\alpha\sqrt{\gamma}\Phi. \quad (28)$$

3. Numerical implementation

3.1. Storage location for grid variables

CarpetX allows for grid variables to be stored at centers, vertices, faces and edges of the grid cell. All the spacetime quantities including α , β^i , γ_{ij} , and K_{ij} , as well as the stress-energy tensor $T_{\mu\nu}$ live at cell vertices. The hydrodynamic variables including primitives (ρ, p, ϵ, v^i) as well as conservatives (D, τ, S_i) are stored at cell centers. Magnetic vector potential components A_i are staggered and are located at respective cell edges, taking the curl of which conveniently results in the magnetic field components B^i at cell faces. Using the adjacent cell-face values, a linear interpolation is then employed to calculate the magnetic field components at the cell-centers, whenever necessary. The scalar potential Φ is instead stored at the cell vertex. Fluxes are also stored at cell faces. This is highlighted in figure 1.

3.2. Spacetime evolution

As mentioned in section 2.1, we use the Z4c module to evolve the spacetime. In our simulations, we employ the standard gauge and constraint damping parameters, specifically $\kappa_1 = 0.02$ and $\kappa_2 = 0.0$ for constraint damping, along with the lapse parameter $\mu_L = 2/\alpha$, and shift parameters $\mu_S = 1$ and $\eta = 2$ [71].

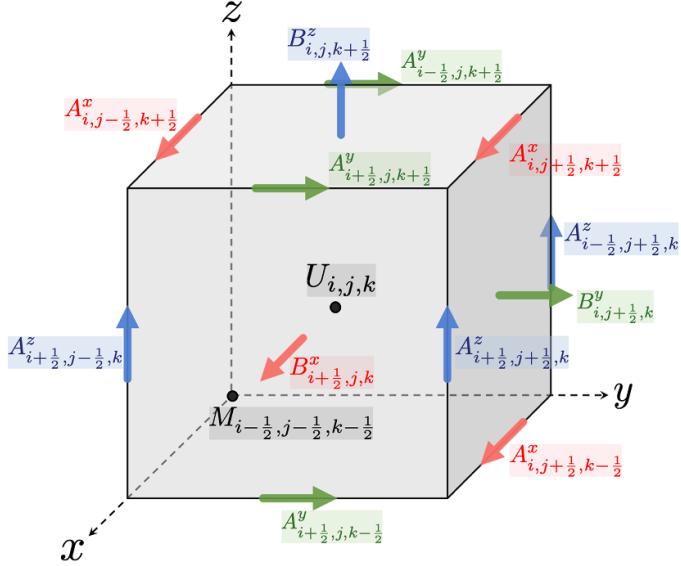


Figure 1. Storage locations of different variables within a grid cell. Both conserved and primitive hydrodynamic quantities live at the cell center, denoted by $U_{i,j,k}$. Magnetic vector potential components A_i are stored at the respective cell edges, whereas the magnetic field components B^i and fluxes are at cell faces. The magnetic scalar potential as well as the spacetime variables are stored at cell vertices, represented by $M_{i-\frac{1}{2},j-\frac{1}{2},k-\frac{1}{2}}$. Figure reproduced with permission from [57].

Since all spacetime quantities are stored at grid vertices, which is different from the hydrodynamical quantities that are located at cell centers and fluxes at cell faces, we use fourth order Lagrange interpolation for interpolating the spacetime variables from vertices to cell centers or faces when evaluating the hydrodynamics sources and fluxes.

3.3. High resolution shock capturing schemes

Accurate and efficient treatment of shocks and discontinuities require state-of-the-art high-resolution shock capturing methods. Two main components of such schemes are (i) the reconstruction algorithms, employed to compute values of primitive variables at the cell interfaces, and (ii) the approximate Riemann solvers, used to compute the fluxes at these cell interfaces.

In AsterX, we have implemented the second order accurate total variation diminishing (TVD) MINMOD scheme [75], as well as the third order accurate piecewise-parabolic method (PPM) [76, 77]. Higher order methods, i.e. WENO-Z [78], and MP5 [79] have also been implemented, and currently are undergoing testing. All the schemes are made available through the ReconX module. Within this module, we also conduct internal unit tests for some of the aforementioned reconstruction algorithms to verify that their specific mathematical properties are satisfied.

Approximate Riemann solvers implemented include Lax–Friedrichs (LxF) [80] and Harten–Lax–van–Leer–Einfeldt (HLL) [81]. For HLL, the numerical fluxes at cell interfaces are computed as

$$\mathbf{F}^i = \frac{c_{\min} \mathbf{F}_r^i + c_{\max} \mathbf{F}_l^i - c_{\max} c_{\min} (\mathbf{F}_r^0 - \mathbf{F}_l^0)}{c_{\max} + c_{\min}}, \quad (29)$$

where the subscript r (l) denotes function computation at the right (left) side of the cell interface, and $c_{\max} \equiv \max(0, c_{+,l}, c_{+,r})$, $c_{\min} \equiv -\min(0, c_{-,l}, c_{-,r})$, where $c_{\pm,r}$ ($c_{\pm,l}$) are the right-going (+) and left-going (−) maximum characteristic wave speeds computed from the primitive variables \mathbf{U} at the right (left) interface. Here, the computed eigenvalues are based on the solutions of a quartic equation, as described in detail for instance, in section 3.1.1 of [46].

Fluxes via LxF can instead be obtained by the expression

$$\mathbf{F}^i = \frac{\mathbf{F}_r^i + \mathbf{F}_l^i - c_{\text{wave}} (\mathbf{F}_r^0 - \mathbf{F}_l^0)}{2}, \quad (30)$$

where $c_{\text{wave}} = \max(c_{\max}, c_{\min})$ [82]. This scheme is more dissipative than HLLE, and can be useful in cases when dealing with strong shocks and interiors of black-hole horizons.

3.4. EM field evolution

In AsterX, we evolve the vector potential A_i instead of the magnetic field variable B^i , and then compute B^i by taking the curl of the vector potential given by the equation (24). For the EM gauge, we have the choice to employ either the algebraic gauge or the generalized Lorenz gauge (27). The latter is adopted by default in our simulations, unless stated otherwise.

When making use of the generalized Lorenz gauge, the damping term ξ in equation (28) is typically set to $1.5/\Delta t_{\max}$, where Δt_{\max} is the time step of the coarsest refinement level, as also adopted in [52, 57].

Recasting the evolution equations of the vector potential and the densitized scalar potential ($\Psi_{\text{mhd}} \equiv \sqrt{\gamma}\Phi$), the update terms for the right-hand side can be written as

$$\begin{aligned} \partial_t A_i &= -E_i - \partial_i (G_A) \\ &= -\epsilon_{ijk} \tilde{v}^j B^k - \partial_i \left(\alpha \frac{\Psi_{\text{mhd}}}{\sqrt{\gamma}} - \beta^j A_j \right), \end{aligned} \quad (31)$$

$$\begin{aligned} \partial_t \Psi_{\text{mhd}} &= -\partial_j (F_{\Psi}^j) - \xi \alpha \Psi_{\text{mhd}} \\ &= -\partial_j (\alpha \sqrt{\gamma} A^j - \beta^j \Psi_{\text{mhd}}) - \xi \alpha \Psi_{\text{mhd}}. \end{aligned} \quad (32)$$

In the above equations, since some of the terms on the right-hand side are stored or staggered in different ways within a cell, they need to be interpolated at a common point. For solving equation (31), we first interpolate A_j to the vertices, where Ψ_{mhd} and the metric components are located. This allows us to compute G_A at the vertices. We apply finite differencing with a stencil that computes the derivative of G_A at cell edges, which gives the second term of the RHS.

To calculate the electric field in equation (31) at cell edges, we have implemented two different schemes. The first approach uses the flux constrained transport method [83], in which the electric field is computed from the magnetic field HLLE fluxes (see, for instance, equation (23) of [84]). However, this scheme might result in generation of numerical instabilities in magnetic field evolution in simulations such as BNS mergers. These instabilities can be curbed by adding an appropriate amount of Kreiss–Oliger dissipation to the magnetic field variables. Our second approach implements the upwind constrained transport (UCT) scheme as described in [84, 85]. To compute, for instance, the E_z component of the electric field on the edge $(i + \frac{1}{2}, j + \frac{1}{2}, k)$, we follow the steps summarized below:

(i) Using equation (57) of [85], the upwind transverse velocities are first computed, for example, on cell-face along x direction, i.e. $(i + \frac{1}{2}, j, k)$ as

$$\bar{v}^j = \frac{a_+^x \tilde{v}^{jL} + a_-^x \tilde{v}^{jR}}{a_+^x + a_-^x}, \quad j = y, z. \quad (33)$$

where L and R stand for left and right upwind states of \tilde{v}^i along an interface, and a_\pm^x is defined as

$$a_\pm^x = \max \{0, \pm \lambda_\pm^x (U^L), \pm \lambda_\pm^x (U^R)\}. \quad (34)$$

where λ_\pm^x are the characteristic wave speeds calculated at both left and right states.

(ii) Then, the quantities $\tilde{v}^x, \tilde{v}^y, B^x, B^y$ are reconstructed from faces $(i + \frac{1}{2}, j, k)$ and $(i, j + \frac{1}{2}, k)$ to the edge $(i + \frac{1}{2}, j + \frac{1}{2}, k)$ using the PPM reconstruction method, following equations (59) and (60) of [85].

(iii) Finally, using the above expressions, the E_z component of the electric field on the edge $(i + \frac{1}{2}, j + \frac{1}{2}, k)$ can be written using equation (61) of [85], as

$$E_z = - \frac{a_+^x \bar{v}^{xL} B^{yL} + a_-^x \bar{v}^{xR} B^{yR} - a_+^x a_-^x (B^{yR} - B^{yL})}{a_+^x + a_-^x} + \frac{a_+^y \bar{v}^{yL} B^{xL} + a_-^y \bar{v}^{yR} B^{xR} - a_+^y a_-^y (B^{xR} - B^{xL})}{a_+^y + a_-^y}. \quad (35)$$

The other components E_x and E_y are also derived in the same way.

For solving equation (32), we only need to interpolate A_j to the vertices, where Ψ_{mhd} and the metric components are located. F_{Ψ^j} is then calculated at vertices. Using finite differencing with a stencil that computes the derivative of F_{Ψ^j} at cell vertices, we obtain all the required RHS terms for equation (32).

3.5. Primitive variables recovery

At each time-step during the evolution, we require values of the primitive variables \mathbf{U} in order to compute the numerical fluxes. However, since we evolve the conservative variables, the system of equation (20) need to be inverted to retrieve the primitives. This must be done numerically, and is one of the most delicate parts of the time evolution step. This process is referred to as *primitive variables recovery* or *conservative-to-primitives* (C2P) inversion.

In **AsterX**, we have implemented two such C2P methods:

- **2D Noble et al scheme** [50, 86]. This method is based on the reduction of (20) to a system of two equations and its subsequent inversion via a Newton–Raphson technique. First, we define

$$z \equiv \rho h W^2, \quad (36)$$

and use expression (13) for b^0 , together with

$$b_i = \frac{B_i}{W} + \alpha b^0 v_i, \quad (37)$$

to write (20) as

$$S_i = -n_\mu T^\mu_i = \alpha T^0_i = (\rho h + b^2) W^2 v_i - \alpha b^0 b_i = (z + B^2) v_i - (B^j v_j) B_i. \quad (38)$$

Next, contracting the above equation with B^i , we get

$$B^i S_j = \rho h W^2 (B^i v_j) = z (B^i v_j). \quad (39)$$

Finally, we use the above relation while contracting (38) with S^i , and recast (20) in terms of z and v^2 to get

$$v^2 (B^2 + z)^2 - \frac{(B^i S_i)^2 (B^2 + 2z)}{z^2} - S^2 = 0 \quad (40a)$$

$$\tau + D - \frac{B^2}{2} (1 + v^2) + \frac{(B^i S_i)^2}{2z^2} - z + p = 0, \quad (40b)$$

where $B^2 \equiv B^i B_i$ and $S^2 \equiv S^i S_i$. The above system is solved for the unknowns z and v^2 via a Newton–Raphson technique, and the primitives are then retrieved in the following order

$$W = \frac{1}{\sqrt{1 - v^2}} \quad (41a)$$

$$\rho = \frac{D}{W} \quad (41b)$$

$$v^i = \frac{\gamma^{ij} S_j}{z + B^2} + \frac{(B^i S_i) B^i}{z(z + B^2)} \quad (41c)$$

$$\epsilon = \frac{1}{\gamma} \left(\frac{z}{\rho W^2} - 1 \right), \quad (41d)$$

where (41c) comes from (38) and (39), while (41d) stems from (36) together with relativistic specific enthalpy h , and pressure p is derived from the evolution EOS.

- **1D Palenzuela et al scheme** [86, 87]. This method is based on the reduction of (20) to a single equation, whose root is found via Brent’s technique [88]. Since the root is bracketed and the algorithm does not attempt to estimate the derivatives of the function whose root is being searched for, this scheme is both more robust and more accurate than Newton–Raphson based techniques (though it is usually slower). The scheme works as follows:

- starting from the conservative variables at the current timestep, define the quantities

$$q \equiv \frac{\tau}{D}, \quad r \equiv \frac{S^2}{D^2}, \quad s \equiv \frac{B^2}{D}, \quad t \equiv \frac{B^i S_i}{D^{\frac{3}{2}}}; \quad (42)$$

- compute hW using the primitives at the previous timestep, and solve

$$f(x) = x - \hat{h} \hat{W} = x - \left(1 + \hat{\epsilon} + \frac{p(\hat{\rho}, \hat{\epsilon}, \hat{Y}_e)}{\hat{\rho}} \right) \hat{W}, \quad (43)$$

via Brent’s method to get the value of the unknown $x \equiv \rho h W^2 / \rho W = hW$ in between the bounds $1 + q - s < x < 2 + 2q - s$, where quantities with a hat in (43) are calculated at every iteration step from x ;

- in terms of (42) and x and using $v^2 = 1 - \frac{1}{W^2}$, equation (40a) can be recast as

$$\frac{1}{\hat{W}^2} = 1 - \frac{rx^2 + t^2(2x + s)}{x^2(x + s)^2}; \quad (44)$$

(iv) compute $\hat{\rho} = \frac{D}{\hat{W}}$;

(v) again, make use of (42), x and $v^2 = 1 - \frac{1}{\hat{W}^2}$ to recast (40b) as

$$\hat{\epsilon} = -1 + \frac{x}{\hat{W}} \left(1 - \hat{W}^2 \right) + \hat{W} \left[1 + q - s + \frac{1}{2} \left(\frac{t^2}{x} + \frac{s}{\hat{W}^2} \right) \right]; \quad (45)$$

(vi) using the evolution EOS, invert $\hat{\epsilon}$ to compute pressure p ;

(vii) noting that (36) holds, and that x is close to hW , the 3-velocity components can be obtained from (41c) setting $z = \rho Wx$.

Our C2Ps are made available through the `Con2PrimFactory` module.

3.6. Atmosphere

One of the limitations of the state-of-the-art GRMHD codes is the difficulty to handle vacuum states. A standard approach imposes an artificial atmosphere, that requires setting a minimum density floor ρ_{atm} . In `AsterX`, during the evolution, if the conserved density ‘ D ’ computed by the C2P falls below the floor density, i.e. $D < \sqrt{\gamma} \rho_{\text{atm}}$, then the primitive ρ is reset to ρ_{atm} , whereas the pressure and specific internal energy are recomputed using a polytropic EOS. Moreover, the fluid’s three-velocity is set to zero. The conserved variables are then reset to atmospheric values using the modified primitives, via equation (20). The magnetic fields are kept unchanged. For the tests presented in this paper, unless stated otherwise, we typically set $\rho_{\text{atm}} = 10^{-7} \rho_{0,\text{max}}$, where $\rho_{0,\text{max}}$ is the initial maximum value of the rest-mass density.

3.7 EOS

In order to close the system of equations (18) and (19), an EOS encoding the thermodynamic properties of the gas under study is needed. To this end, we couple `AsterX` to `EOSX`, a new Cactus module we developed that currently implements the following two analytical EOS

- **ideal-fluid:**

$$P(\rho, \epsilon) = (\Gamma - 1) \rho \epsilon, \quad (46)$$

where Γ is the adiabatic index;

- **polytropic:**

$$P(\rho) = K \rho^\Gamma, \quad (47)$$

$$\epsilon(\rho) = \frac{K \rho^{\Gamma-1}}{\Gamma - 1} = \frac{P(\rho)}{(\Gamma - 1) \rho}, \quad (48)$$

where K is the polytropic constant with $n = \frac{1}{\Gamma-1}$ as the polytropic index.

It is our intention, in future work, to render `EOSX` capable of supporting finite-temperature, tabulated EOS in the form $P(\rho, T, Y_e)$ and/or $P(\rho, \epsilon, Y_e)$.

3.8. Boundary conditions

`CarpetX` provides several types of boundary conditions that can be applied to the numerical domain, namely, ‘Dirichlet’, ‘Neumann’, ‘Robin’, ‘linear extrapolation’ and ‘none’. A

Dirichlet boundary condition sets a fixed value of the function, whereas the Neumann boundary condition imposes a fixed value of the function derivative at the boundary. The Robin boundary condition is a mix of Dirichlet and Neumann conditions, specifying a linear combination of the function value and its derivative at the boundary. There is also support for periodic and reflection (symmetry) boundary conditions. Radiative boundary conditions are also made available via the NewRadX module.

3.9. AMR

CarpetX offers the capability of AMR through the AMReX framework [64]. Support for fixed mesh refinement is provided by the BoxInBox module, which allows for constructing a hierarchical grid structure with nested rectangular boxes. Whereas, the dynamic (non-fixed) block-structured AMR enables higher refinement in localized regions of interest while maintaining a coarser resolution elsewhere, thus optimizing computational resources and enhancing accuracy. In the latter approach, once we have the solution over a rectangular grid, cells that require additional refinement can be identified through user-based criteria, which are then covered with a set of rectangular grids or blocks of higher resolution. AMReX follows the strategy of [89] to determine the most efficient patch layout to cover the cells that have been tagged for refinement (see [64] for more details). In AsterX, we have implemented algorithms that compute either the gradient, first derivative, second derivative, or second derivative norm for one or more MHD variables, the results of which are compared to a user-defined error threshold parameter, to decide cell-tagging for refinement.

3.10. Subcycling

Our latest implementation in CarpetX includes support for subcycling in time, in which grid variables on different levels progress with different time step sizes, thus reducing the computational effort required. This implementation follows a more efficient approach than the one provided in the original Carpet code. Specifically, subcycling in Carpet employs a buffer region with a width four times the number of ghost zones to supply inter-level boundary data at AMR boundaries during an RK4 evolution. Additionally, the buffer zone interpolation is second-order accurate, resulting in loss of accuracy. The new subcycling algorithm in CarpetX, based on [77, 90], does not require any additional buffer zones, and all interpolations are fully fourth-order accurate. While this method will be presented soon in our upcoming paper (Ji *et al* in preparation), we show preliminary results with AsterX employing subcycling in section 5 (see also figure 13).

3.11. Parallelism

As based on the CarpetX driver, AsterX can run on both CPUs and GPUs, and for convenience, code that has been written for GPUs can still execute efficiently on CPUs. We have tested AsterX with both NVIDIA GPUs (CUDA) and AMD GPUs (ROCm). Intel GPUs (oneAPI) are in principle also supported, but this has not been tested yet. This functionality is provided by AMReX.

CarpetX provides several levels of parallelism:

SIMD. SIMD vectorization uses special CPU instructions to execute the same operation on multiple (usually 4 or 8) data elements simultaneously. Under ideal conditions, executing a SIMD instruction takes the same amount of time as executing a scalar instruction,

and a compute kernel might thus run several times faster with this optimization. Support for SIMD instructions is implemented via the *nsimd* library [91]. Future versions of the C++ standard will support SIMD vectorization directly. Unfortunately, SIMD vectorization requires certain smallish changes to the source code, and not all parts of the code have been vectorized. In practice, code that is limited by memory bandwidth instead of compute throughput does not usually benefit from SIMD vectorization.

Multi-threading. Multi-threading is also called ‘shared-memory parallelism.’ Modern CPUs have many cores, and it is beneficial to have multiple cores share data structures (grid functions) because this reduces overhead introduced by ghost zones (see below) that would be necessary by distributed-memory parallelism. AMReX splits grid functions into tiles, and we assign tiles to threads for calculations.

GPUs. GPUs offer multiple levels of parallelism. Using CUDA terminology, a GPU executes *threads* which perform individual operations. Threads are grouped into *warps* which execute identical operations on different data elements. Warps are finally grouped into *blocks* that operate on a common data structure. CarpetX hides these levels from the programmer, who writes compute kernels that are only slightly more complex than a CPU-only compute kernel. As mentioned above, a GPU-enabled compute kernel still runs efficiently on a CPU.

Message passing. This is also called ‘distributed-memory parallelism.’ To split a large calculation onto multiple GPUs or multiple compute nodes we use *MPI* to communicate between independently running processes. CarpetX hides this complexity from the programmer, using *ghost zones* to allow the evaluation of finite differencing stencils and reconstructing fluid states on cell faces. This is also implemented via AMReX.

Overall, the programmer implements compute kernels that look very similar to straightforward C++ code, surrounded by function calls to allow CarpetX to parallelize the code when iterating over tiles.

3.12. AsterX workflow

To facilitate users’ comprehension of the organizational structure of our code, figure 2 illustrates an overview of the workflow within AsterX for computing the right-hand side of the GRMHD equations.

In the figure, elements outlined by green boxes denote GRMHD variables, red boxes represent Einstein’s equations, while grey boxes show functional steps in the GRMHD workflow. Concurrently, blue arrows denote data flow for computing the GRMHD RHS, while those depicted by red boxes and arrows pertain to relevant segments of the spacetime evolution scheme.

The code is initialized with data encompassing both primitive hydrodynamical variables (refer to section 2.2) and spacetime quantities, including lapse, shift vector, spatial metric, and extrinsic curvature. Analytical expressions (20) are then employed to initialize the conserved variables. The code loops over the following steps until the designated final simulation time is reached:

- (i) Utilizing the conserved variables from the previous step and an appropriate EOS, `con2prim` is called to obtain cell averages of the conserved variables.
- (ii) Cell averages of the conserved variables are reconstructed on cell faces while also being used in computing the energy momentum tensor of the evolution.
- (iii) The Riemann problem on each cell face is solved via an approximate Riemann solver in order to obtain the numerical flux between cells. This process also requires the 3-metric

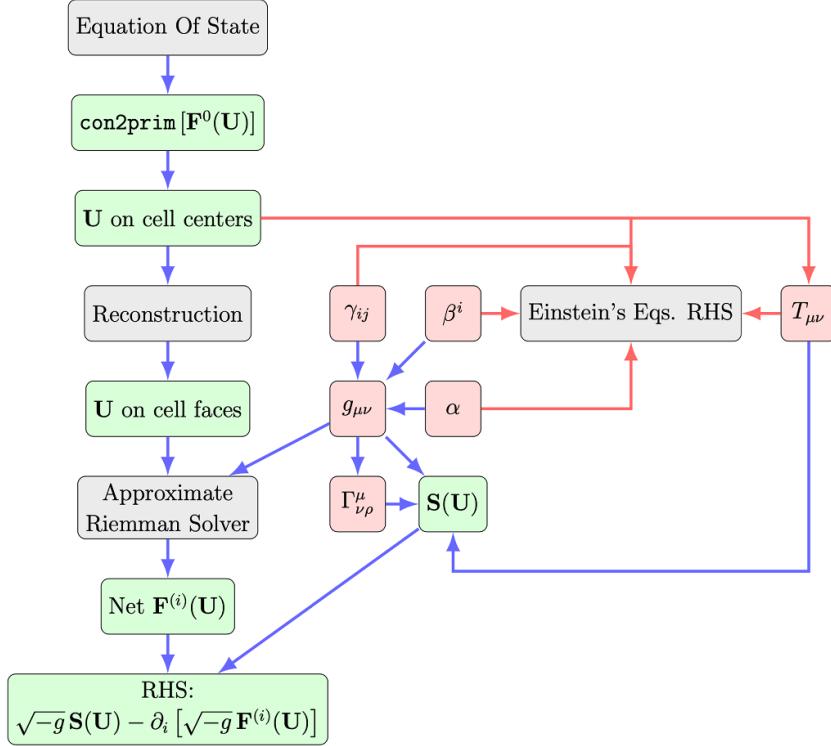


Figure 2. Schematic representation of the AsterX workflow to compute the right-hand side of the conservative equations. See text for more details.

along with lapse and shift (components of the four-metric) from the gravitational evolution part of the code.

- (iv) Once the fluxes have been computed, the 3-metric along with lapse and shift and their derivatives are used to compute source terms, finalizing the computation of the hydrodynamical RHS.

4. Tests

In this section, we report on the results of our extensive testing, including a number of 1-, 2- and 3-dimensional simulations. These simulations include critical tests that have been already considered in the literature in several previous Papers (see, e.g. and references therein), allowing for a direct comparison with other codes.

4.1. Balsara shocktube

In order to check the correctness of the approximate Riemann solvers and other numerical schemes implemented in the code, the first round of tests we performed with AsterX are those involving Riemann problems. In figure 3, we show the results for 1-dimensional (1D) relativistic shock-tube problems based on the testsuite of [92]. Here, our numerical results of

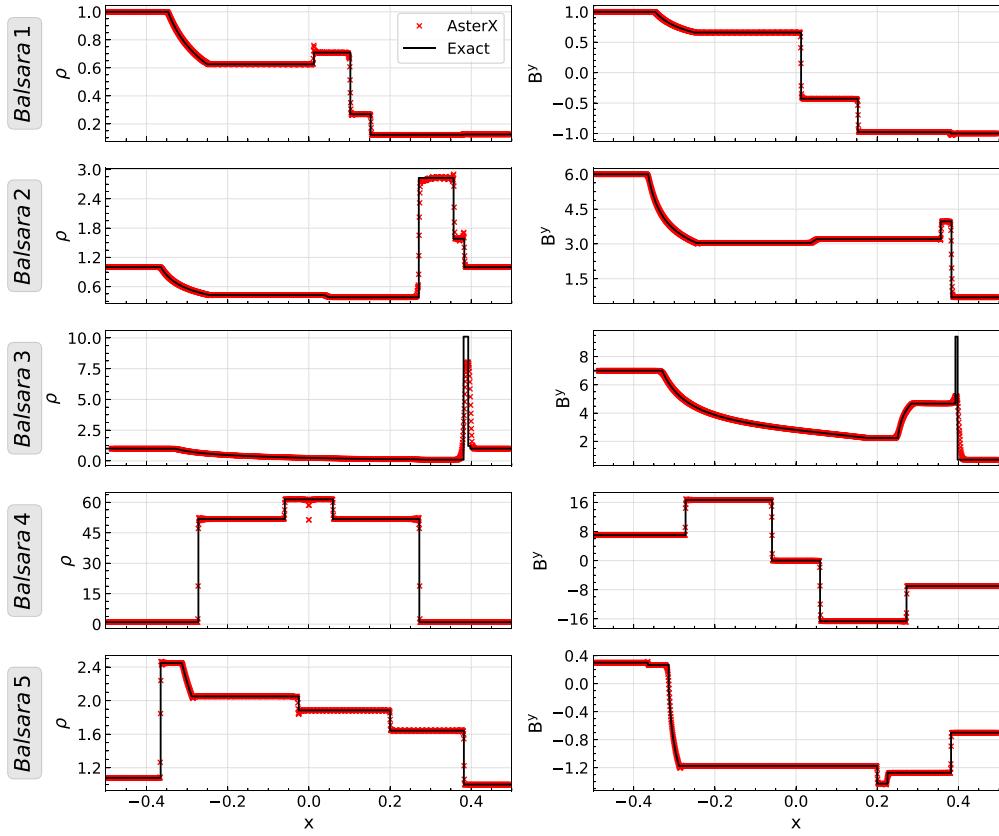


Figure 3. Comparison of numerical results (red dots) and exact solutions (continuous black lines) for the suite of tests of [92]. Left and right columns show respectively the spatial distributions of the rest-mass density and the magnetic field component B_y at the final time of the evolution. Here, the Balsara 1, 2, 4 and 5 tests are performed with the third order PPM method. On the other hand, Balsara 3 is performed with the second order MINMOD method because this test is the most demanding one due to the jump of four orders of magnitude in the initial pressure. This extreme change requires a slightly more dissipative method to succeed.

these tests are directly compared with the exact solutions computed via the code presented in [93]. Initial data for such tests are described in table 1.

For all tests, we employ an ideal fluid EOS, with $\Gamma = 2.0$ for test Balsara 1 and $\Gamma = 5/3$ for the others. The final evolution time is considered as $t = 0.55$ for test Balsara 5 and $t = 0.4$ for the others. We note that all tests show an excellent agreement between the numerical results and the exact solutions.

4.2. Alfvén wave

To determine the convergence order of our code, we perform the conventional Alfvén wave test. This involves advecting a circularly polarized Alfvén wave across the simulation domain. The solutions for such a wave are exact and do not contain shocks, making it an ideal candidate for this test. The initial conditions are adopted from [57, 67], where we set the wave amplitude

Table 1. Initial data for Balsara relativistic shock tube tests.

Test:	1		2		3		4		5	
	L	R	L	R	L	R	L	R	L	R
ρ	1.0	0.125	1.0	1.0	1.0	1.0	1.0	1.0	1.08	1.0
p	1.0	0.1	30.0	1.0	1000.0	0.1	0.1	0.1	0.95	1.0
v_x	0.0	0.0	0.0	0.0	0.0	0.0	0.999	-0.999	0.4	-0.45
v_y	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.3	-0.2
v_z	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.2	0.2
B^x	0.5	0.5	5.0	5.0	10.0	10.0	10.0	10.0	2.0	2.0
B^y	1.0	-1.0	6.0	0.7	7.0	0.7	7.0	-7.0	0.3	-0.7
B^z	0.0	0.0	6.0	0.7	7.0	0.7	7.0	-7.0	0.3	0.5

$A_0 = 1.0$, the fluid rest-mass density $\rho = 1.0$, fluid pressure $p = 0.5$ and the Alfvén wave speed $v_A = 0.5$. The velocity components are initialized as

$$v^x = 0, \quad v^y = -v_A A_0 \cos(kx), \quad v^z = -v_A A_0 \sin(kx), \quad (49)$$

whereas the magnetic field components are initialized as

$$B^x = 1.0, \quad B^y = A_0 B^x \cos(kx), \quad B^z = A_0 B^x \sin(kx). \quad (50)$$

The wave propagates along the x-direction with the wave vector $k = 2\pi/L_x$ and the domain length $L_x = 1$, where the grid along x extends from -0.5 to +0.5. We study convergence based on various different resolutions, varying the number of grid points along x as 8, 16, 32, 64 and 128. For the evolution, we use the ideal gas EOS with $\Gamma = 5/3$, MINMOD reconstruction, HLLC Riemann solver, RK2 for time-stepping with the Courant factor of 0.2, and periodic boundary conditions. The result of the L2 norm of the convergence error in the y-component of the magnetic field B^y as a function of number of grid points is illustrated in figure 4. In particular, we compute the L2 norm based on the initial and final wave profile after 1 full period of evolution. Here, we observe an overall 2nd order convergence rate, an expected outcome that has also been reported, for instance, in [57, 67] for the same test case.

4.3. Cylindrical explosion

The first test considers the evolution of a 2D cross-section of a dense, over-pressured cylinder in a uniformly magnetized environment. Following this blast wave allows a check of whether the code can correctly follow the shock front on the equatorial plane. The initial data is based on the setup described in [57]. For the cylinder, we consider

$$r_{\text{in}} = 0.8, \quad r_{\text{out}} = 1.0, \quad \rho_{\text{in}} = 10^{-2}, \quad p_{\text{in}} = 1.0, \quad (51)$$

while for the surrounding ambient medium, we set

$$\rho_{\text{out}} = 10^{-4}, \quad p_{\text{out}} = 3 \times 10^{-5}. \quad (52)$$

Here, the parameters r_{in} and r_{out} are used for smoothening the density profile (and similarly for the smoothening of the pressure profile) considered in [57], such that

$$\rho(r) = \begin{cases} \rho_{\text{in}} & \text{for } r \leq r_{\text{in}} \\ \exp\left[\frac{(r_{\text{out}}-r) \ln \rho_{\text{in}} + (r-r_{\text{in}}) \ln \rho_{\text{out}}}{r_{\text{out}}-r_{\text{in}}}\right] & \text{for } r_{\text{in}} < r < r_{\text{out}} \\ \rho_{\text{out}} & \text{for } r \geq r_{\text{out}}. \end{cases} \quad (53)$$

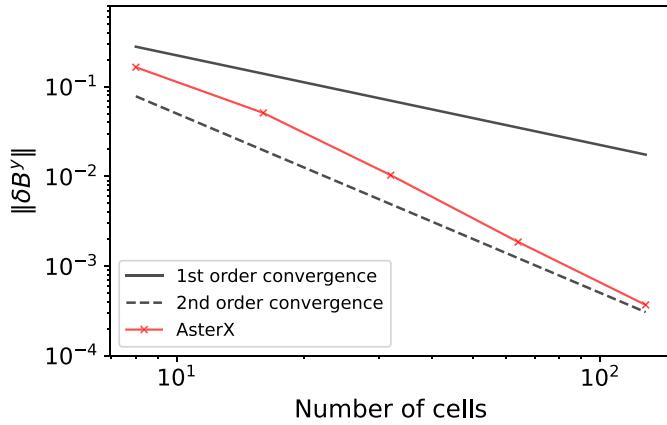


Figure 4. Results of the Alfvén wave test. Here, the L2 norm of error in the y-component of the magnetic field is illustrated (red solid line) as a function of the number of cells used along x-direction. Reference curves for 1st and 2nd order convergence are shown in black solid and dashed lines respectively. With the increase in resolution, our results show good agreement with the expected \sim 2nd order convergence rate as well as with the ones reported in literature [57, 67].

Initially, the fluid 3-velocity is set to zero while the magnetic field is kept uniform with $B^x = 0.1$ and $B^y = B^z = 0$. The grid domain spans the range $[-6, 6]$ along x and y -axes containing 200×200 grid points. For the evolution, we use an ideal fluid EOS with adiabatic index $\Gamma = 4/3$ and set the CFL factor to 0.25. For solving the numerics, the MINMOD reconstruction method is employed along with the HLLE flux solver and the RK4 method for the time-step evolution.

In figure 5, the left panel illustrates the 2D blast wave profiles at the final time $t = 4.0$, representing the distribution of gas pressure p , Lorentz factor W (together with magnetic field lines), and the x - and y -components of the magnetic field, B^x and B^y . Profiles from 1D slices along x - and y -axis are instead shown in the right panel. We find both 1D and 2D profiles to be in good agreement when compared to the results already presented in the literature [57].

4.4. Magnetic rotor

Another 2D test we consider is the magnetic cylindrical rotor, originally introduced for classic MHD in [83, 94] and later employed also for relativistic MHD in [73, 82]. This test consists of a dense, rapidly spinning fluid at the center, surrounded by a static ambient medium starting with a uniform magnetic field and pressure in the entire domain. For initial data, we set the radius of the inner rotating fluid to $r = 0.1$, with inner rest-mass density $\rho_{\text{in}} = 10.0$, uniform angular velocity $\Omega = 9.95$, and therefore the maximum value of the fluid 3-velocity is $v_{\text{max}} = 0.995$. In the outer static ambient medium, we set the rest-mass density to $\rho_{\text{out}} = 1.0$. The initial magnetic field and gas pressure are set to $B^i = (1.0, 0, 0)$ and $p_{\text{in}} = p_{\text{out}} = 1.0$ respectively. The numerical domain is a 400×400 grid with x - and y -coordinates lying in range $[-0.5, 0.5]$. Here, too, we set the CFL factor to 0.25, and consider an ideal fluid EOS with $\Gamma = 5/3$ for evolution. For the numerics, we use the MINMOD reconstruction, the HLLE flux solver, and the RK4 method for time-update.

Figure 6 depicts the 2D profiles of density ρ , gas pressure p , magnetic pressure $p_{\text{mag}} = b^2/2$, and Lorentz factor W along with magnetic field lines at the final time $t = 0.4$. The rotation of the cylinder leads to magnetic winding. This can be seen in the bottom-right panel of left figure 6

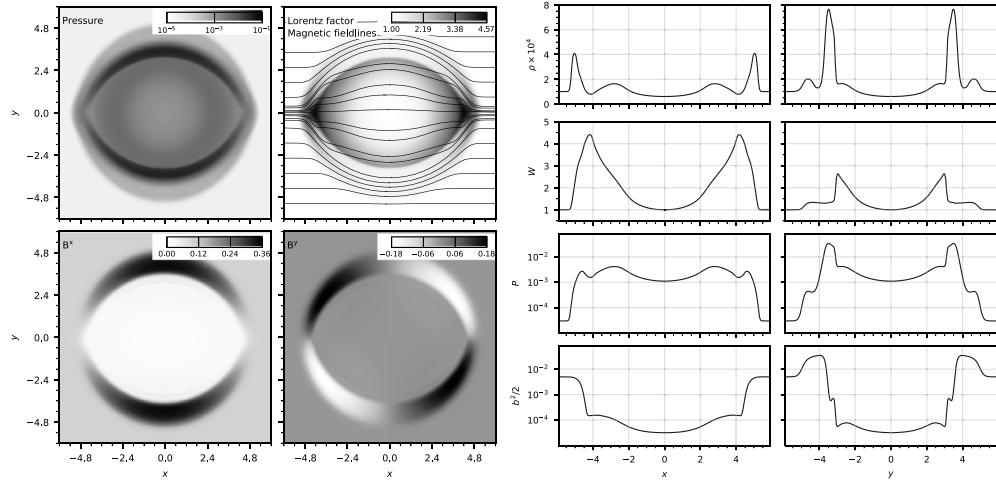


Figure 5. Equatorial snapshots (left panel) and 1D slices along x and y -axis (right-panel) of the cylindrical explosion test at final time $t = 4$. On the left, we have gas pressure P (top-left), Lorentz factor W together with magnetic field lines (top-right), and x - and y -components of the magnetic field, B^x (bottom-left) and B^y (bottom-right). Right panel shows density ρ (top row), lorentz factor W (second row), gas pressure P (third row) and magnetic pressure $b^2/2$ (bottom row).

where, near the central region, the field lines are twisted roughly by $\sim 90^\circ$. This twisting of field lines could eventually slow down the rotation of the cylinder. A decrease in ρ , p , and p_{mag} in the central region is also observed along with the formation of an oblate shell of higher density. Also for this test, the results are in good agreement with the ones in the literature [48, 57, 67, 82].

For a quantitative check, we again take a 1D slice along $x = 0$ and $y = 0$ of the final rotor configuration at $t = 0.4$, as shown in the right panel of 6. Comparing with the results presented in [57, 67], we find the curves including the peak values in agreement with the literature.

To assess the AMR functionality, we conduct an additional simulation by introducing an inner refined grid spanning the x - and y -coordinates from $[-0.32, +0.32]$, with half the grid spacing, utilizing the BoxInBox module. For prolongation at the inner refinement boundary, we interpolate polynomially in vertex centered directions for the spacetime variables, and conserve with third order accuracy along other directions for MHD variables, with a linear fall-back in presence of shocks. Figure 7 presents a comparison of the final magnetic pressure distribution with the results from the uniform grid test. The comparison shows no significant differences nor any noticeable effects at the boundaries between grids, confirming the correct working of our code with the AMR framework.

4.5. Magnetic loop advection

In this test, we simulate an advecting magnetic field loop in 2D, as also considered in [48, 57, 95–97]. This case considers a magnetized circular field loop propagating with a constant

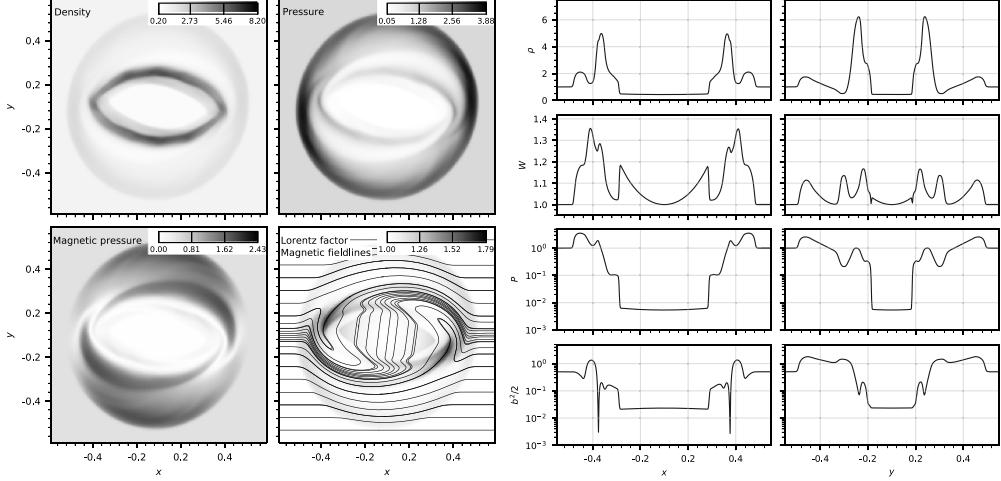


Figure 6. 2D snapshots (left) and 1D slices along x - and y -axis (right) of the Magnetic rotor test for the final evolved time $t = 0.4$. The left panel shows the density ρ (top-left), gas pressure p (top-right), magnetic pressure p_{mag} (bottom-left), and Lorentz factor W together with the magnetic field lines (bottom-right). The right plot illustrates the density ρ , Lorentz factor W , gas pressure p and magnetic pressure $b^2/2$ along the x - and y -axes respectively. The resolution here is $\Delta x = \Delta y = 0.0025$.

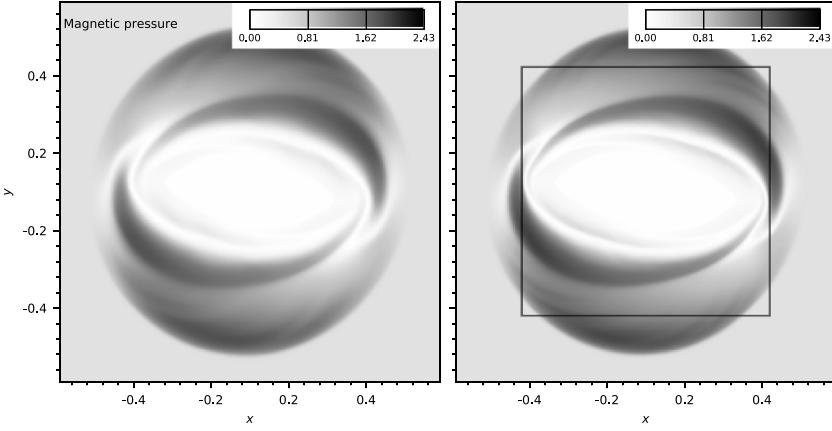


Figure 7. Magnetic rotor test with adaptive mesh refinement (AMR). For comparison, we show the final magnetic pressure configuration at $t = 0.4$ between the uniform grid test (left panel) and the AMR test (right panel) including a refined inner grid (black box) with twice the resolution. We find the two results to be agreement, without any spurious effects at the inner refinement boundary.

velocity in a surrounding non-magnetized ambient medium within a 2D periodic grid. The initial magnetic field is set by the following analytical prescription

$$B^x, B^y = \begin{cases} -A_{\text{loop}}y/r, A_{\text{loop}}x/r & \text{for } r < R_{\text{loop}} \\ 0 & \text{for } r \geq R_{\text{loop}} \end{cases} \quad (54)$$

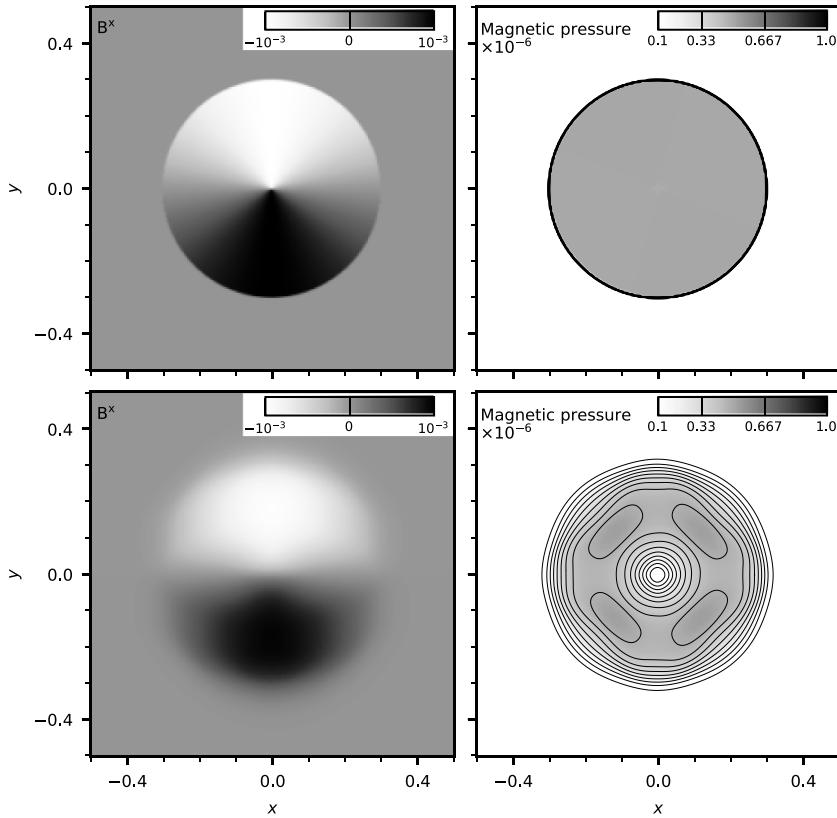


Figure 8. Loop advection test with $v^z = 1/24$. Left and right columns represent the x -component of the magnetic field B^x and the magnetic pressure $p_{\text{mag}} = b^2/2$, respectively. The top row shows the initial data for B^x and its corresponding p_{mag} at $t = 0$ while the bottom row depicts these quantities after one periodic cycle of evolution, i.e. at $t = 24$. Our results exhibit good agreement with those reported in [57].

where $r = \sqrt{x^2 + y^2}$ is the radial coordinate, R_{loop} is the loop radius, A_{loop} sets the amplitude of the magnetic field, and B^z is set to zero. The corresponding vector potential prescription from which equation (54) can be obtained is given by $\vec{A}(r) = (0, 0, \max[0, A_{\text{loop}}(R_{\text{loop}} - r)])$ [96].

As initial data, we set the density and pressure to $\rho = 1.0$ and $p = 3.0$ respectively throughout the computational domain. For the loop, we set $A_{\text{loop}} = 0.001$ and $R_{\text{loop}} = 0.3$. The fluid 3-velocity is initialized to $v^i = (1/12, 1/24, 1/24)$, keeping a non-zero vertical velocity, i.e. $v^z \neq 0$. The test is performed on a 256×256 grid, where the x - and y -components span the range [-0.5, 0.5]. The CFL factor is 0.25 and the adiabatic index for the ideal EOS is set to $\Gamma = 5/3$. Here, too, we use the MINMOD reconstruction method along with the HLLE flux solver and the RK4 method for time-update.

In figure 8, the top row depicts the initial configuration of the magnetic loop for the quantities B^x and $p_{\text{mag}} = b^2/2$ at $t = 0$. After one entire cycle of evolution across the domain, the same quantities are illustrated in the bottom row at $t = 24$. A significant loss of magnetic pressure can be noticed due to numerical dissipation after one evolution cycle, as also reported in [57]. We also perform another simulation using a less dissipative PPM reconstruction, and find similar final outcome profiles as in [57].

4.6. Kelvin–Helmholtz instability

The Kelvin–Helmholtz instability (KHI) is an instability that develops across fluid interfaces in the presence of a tangential shear flow, resulting in the production of fluid turbulences. Here, we performed the KHI test in 2D, to evaluate the block-structured AMR feature provided via `CarpetX`. This test adopts the initial setup similar to the one presented in [98], based on which the simulation domain is $[-0.5, +0.5]^2$ on the xy plane with periodic boundary conditions along x -direction. The central region $|y| < 0.25$ is initialized with density $\rho = 2.0$ and velocity $v^x = 0.5$, whereas the surrounding gas is set up with $\rho = 1.0$ and velocity $v^x = -0.5$. The initial gas pressure is kept constant throughout the domain with value $p = 2.5$. This gives two contact discontinuities or ‘slip’ surfaces at $|y| = 0.25$. To induce the instability, we excite a single mode with wavelength equal to half the domain size by perturbing v^y , given by the expression

$$v^y = \omega_0 \sin(4\pi x) \left(e^{\frac{-(y-0.25)^2}{2\sigma^2}} + e^{\frac{-(y+0.25)^2}{2\sigma^2}} \right), \quad (55)$$

where we set the wave amplitude $\omega_0 = 0.1$, and $\sigma = 0.05/\sqrt{2}$. For simplicity, the magnetic field components are set to zero.

Here, we incorporate 3 levels of AMR with 64^2 cells on the coarsest level. The error threshold for regridding is set to 0.4. At the time of regridding, regions with first derivatives of either of the fluid variables (ρ, p, ϵ, v^i) that surpass this threshold are refined. We evolve the system up to $t = 1.5$ with an ideal gas EOS using $\Gamma = 5/3$, along with the MINMOD reconstruction method, the HLLE flux solver, and RK4 time-stepping with CFL factor 0.5.

In figure 9, the density profile at end is illustrated along with the block-structured AMR mesh. Here, we notice that KHI grows at the slip surfaces, and develops the characteristic wave-like structures. In order to capture the physics at the shear layers, these waves are refined based on our refinement criterion. However, we also note that the refinement criteria are problem-dependent, and can be improved with further experimentation for this case. Our test demonstrates `AsterX`’s ability to utilize the dynamic block-structured AMR, which could be essential for accurately resolving shocks and instabilities in various astrophysical fluid flows. One such application of our code with block-structured AMR could be to better capture KHI in BNS merger simulations, by increasing resolution in regions of interest, i.e. in the shear layer where the KHI develops when the two NS cores touch each other, allowing to effectively model magnetic field amplification.

4.7. TOV star

The Tolman–Oppenheimer–Volkoff (TOV) equations [99, 100] are typically employed to describe static, spherically symmetric stars in general relativity. As the first 3D test to further assess the stability and accuracy of our code, we consider the evolution of a non-rotating stable magnetized TOV configuration in dynamical spacetime. For computing the initial data, we use our module `TOVSolverX`, ported to the new `CarpetX` infrastructure from an earlier version [68]. The initial TOV configuration is generated using a polytropic EOS with adiabatic index $\Gamma = 2.0$, polytropic constant $K = 100$, and initial central rest-mass density $\rho = 1.28 \times 10^{-3}$. This setup is then evolved using an ideal fluid EOS with the same Γ . In addition, we add a magnetic field to the TOV configuration by hand using the analytical prescription of the vector potential A_ϕ given by

$$A_\phi \equiv A_b \varpi^2 \max(p - p_{\text{cut}}, 0)^{n_s}, \quad (56)$$

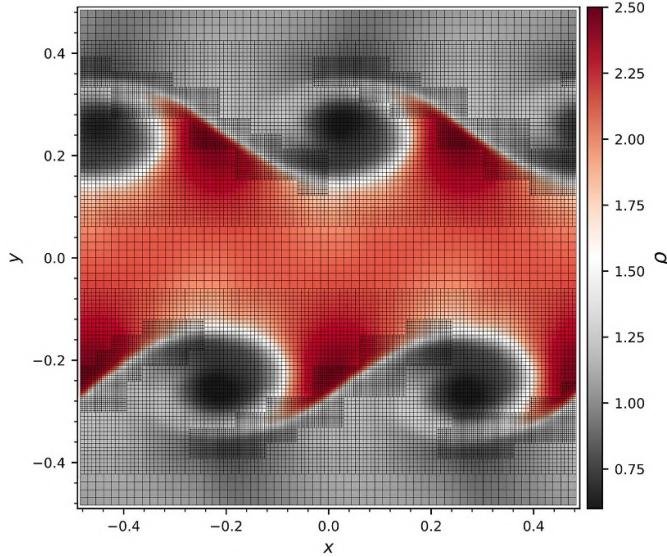


Figure 9. Density profile of the Kelvin–Helmholtz instability test at final time = 1.5, illustrated along with the grid mesh. The wave-like patterns, that develop at the shear layers, are well-captured by the finest mesh refinement blocks.

where ϖ is the cylindrical radius, A_b is a constant, $p_{\text{cut}} = 0.04p_{\text{max}}$ determines the cutoff when the magnetic field goes to zero inside the star, with p_{max} corresponding to the initial maximum gas pressure, and $n_s = 2$ sets the degree of differentiability of the magnetic field strength [47]. The value of A_b is set such that the maximum value of the initial magnetic field strength is set to $\approx 10^{14}$ G. This generates a dipole-like magnetic field confined fully inside the star, which is a standard configuration used for such tests. Exploration of complex magnetic field geometries that include higher-order multipoles and also extend beyond the stellar surface (e.g. [101, 102]), in long-term high-resolution simulations can be non-trivial and is deferred to future work. The spacetime is evolved via Z4c with constraint damping parameters $\kappa 1 = 0.02$ and $\kappa 2 = 0.0$, along with the dissipation coefficient of 0.32. In the atmosphere, we set $\rho_{\text{atm}} = 10^{-11}$.

The tests are run on a box-in-box AMR grid with 6 refinement levels, having the finest level extending from -20 to 20 along x -, y - and z -directions, and the radii along with grid-spacings are increased by a factor of 2 for the outer refinement levels. We perform simulations with three different resolutions, i.e. low, medium and high resolution with finest grid spacing of 0.5 , 0.25 and 0.125 respectively. All test cases are simulated for 8 ms using the PPM reconstruction method, the HLLE flux solver, and the RK4 method for time stepping with a CFL factor of 0.25.

The top-left panel of figure 10 shows the evolution of central rest-mass density ρ_c for the three simulations with different resolutions. In all cases, the star experiences periodic oscillations induced as a result of the truncation errors present in the initial data, whereas the numerical viscosity of the employed finite differencing (FD) scheme are the main cause of dissipation. The low-resolution simulation experiences both large oscillations as well as heavy dissipation. As we move to higher resolution, the oscillations are reduced to nearly 1%, and show convergence of the order ≈ 2.7 , as shown in bottom-left panel of the same figure.

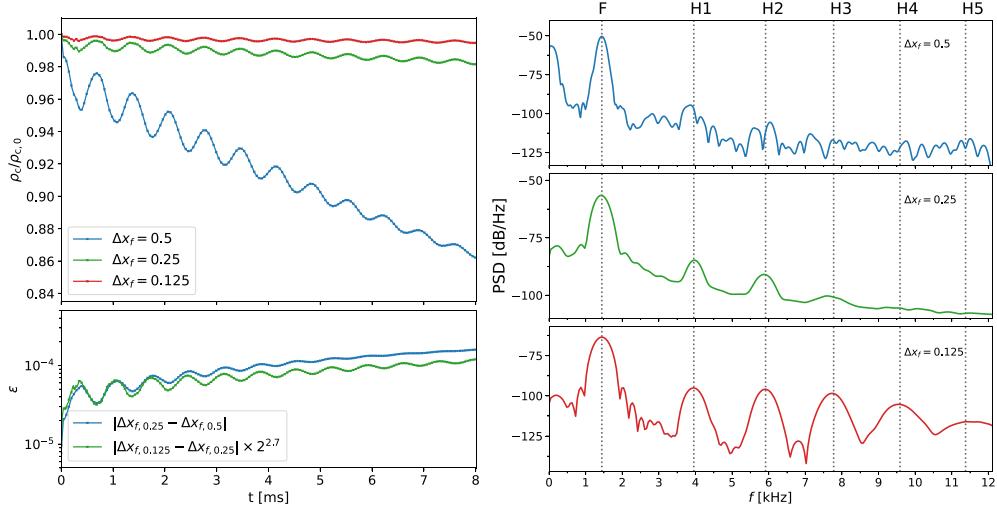


Figure 10. Left: the top panel shows the evolution of the normalised central rest-mass density $\rho_c/\rho_{c,0}$ for the different resolution runs. To quantify convergence with increasing resolution, evolution of the absolute difference ϵ between the different resolutions are shown in the bottom panel. Right: top, middle and bottom plots show the power spectrum of the central rest-mass density evolution for the three different resolutions, i.e. Δx_f equal to 0.5, 0.25 and 0.125 respectively, compared with peak frequencies of the oscillations (vertical dotted lines) computed from perturbation theory [103]. Here, F stands for the fundamental frequency of oscillation whereas (H1, H2, H3, H4, H5) represent the higher harmonics.

As another validation check, we report oscillation frequencies of the three TOV simulations with different resolutions in the right panel of figure 10. For each simulation, the power spectrum density is computed via fast Fourier transform in order to extract the amplitudes and frequencies of the oscillations of the central rest mass density. To compare, in the same figure, we also illustrate the expected peak frequencies of the oscillations taken from [103], which were obtained independently through a 2D perturbative code. Here, we find that the low resolution case reproduces only the fundamental oscillation mode correctly, but as we increase the resolution, the medium resolution run is able to capture at least two higher harmonic peak frequencies (H1, H2) whereas the high resolution case is able to show matching peak frequencies at least up to fourth harmonic (H4). Note that even though we evolve the TOV using ideal gas EOS, which is different from the polytropic EOS used in [103], the peak frequencies still show a good agreement since ideal gas EOS is expected to produce different results from a polytropic one only in presence of shocks, which appear only near the low-density TOV surface, thus having negligible impact on the oscillations at the core. Our results are also in agreement with the ones reported in the literature [57, 67].

5. Performance benchmarks

To evaluate the performance of `AsterX`, we conduct scaling tests on OLCF’s supercomputer Frontier. Each Frontier compute node is equipped with a 64-core AMD CPU having 512 GB of DDR4 memory, alongside 4 AMD MI250X accelerator units, each featuring two Graphic Compute Dies (GCDs). This configuration yields 8 logical GPUs per node, each outfitted with

Table 2. Single-GPU performance tests for our four setups: Setup A (static spacetime + uniform grid), Setup B (dynamical spacetime + uniform grid), Setup C (dynamical spacetime + AMR, without and with subcycling).

	Setup A	Setup B	Setup C	Setup C + subcycling
levels	1	1	8	8
cells/level	240^3	240^3	120^3	120^3
cells GPU^{-1}	240^3	240^3	240^3	240^3
speed (ZCs s^{-1})	0.27×10^8	0.095×10^8	0.072×10^8	0.2×10^8

64 GB of high-bandwidth memory. As noted in [67], the memory available per GPU dictates the problem size, constraining the maximum number of cells per GPU. Furthermore, to maximize GPU utilization, computations should primarily occur within GPU kernels, and CPU-GPU memory transfers should be kept to a minimum.

We perform scaling tests for the following configurations, based on magnetized TOV simulations using an ideal gas EOS for evolution and fourth order Runge-Kutta (RK4) for time-stepping:

- Setup A: static (Cowling) spacetime + uniform grid
- Setup B: dynamic (Z4c) spacetime + uniform grid
- Setup C: dynamic (Z4c) spacetime + 8-level AMR

These tests employ PPM reconstruction, an HLLE Riemann solver, Neumann boundary conditions, and three ghost zones. For benchmarking purposes only, we refrain from writing any output files and exclude the time taken to set up initial conditions. We quantify our performance in terms of ‘zone-cycles s^{-1} ’ and ‘zone-cycles $\text{s}^{-1} \text{GPU}^{-1}$ ’. Each zone-cycle (ZC) corresponds to the evolution of a single grid cell for one full time-step, encompassing all time-integrator (RK4) sub-steps. Thus, ‘ ZCs s^{-1} ’ provides an estimate of the overall simulation efficiency, while ‘ $\text{ZCs s}^{-1} \text{GPU}^{-1}$ ’ measures the code performance per GPU.

By using subcycling, we reduce the number of ZCs required to reach the same time point by allowing larger time step sizes for coarser levels, compared to the non-subcycling case where all levels evolve using the same time step. As a result, less time is needed to reach the same point. For an effective comparison, instead of using the wall time required to reach the same time point, we use the total amount of work divided by the wall time taken. Here, we define the zone cycles required by the non-subcycling simulation as the measure of total work.

5.1. Single GPU benchmarks

To quantify the performance on a single Frontier GPU (i.e. 1 GCD), we consider a compute load of 240^3 cells. The results are summarized in table 2.

For Setup A, the simulation runs with $0.46 \times 10^8 \text{ ZCs s}^{-1}$ when using RK2 integrator, which is comparable to the one reported in [67] for a similar setup, which was run however on a single Summit NVIDIA V100 GPU. Whereas, the performance drops to $0.27 \times 10^8 \text{ ZCs s}^{-1}$ when using RK4 instead. Since RK4 has double the number of sub-steps as compared to RK2, we find the 40% decrease in run speed to be reasonable. Setup B instead yields a speed of $0.095 \times 10^8 \text{ ZCs s}^{-1}$, which indicates that the spacetime solver (Z4c) is roughly twice as

expensive as the GRMHD code¹⁴. For Setup C, we use 120^3 cells per level, which effectively gives a compute load of 240^3 cells in total as well. Without subcycling, we find the performance of 0.072×10^8 ZCs s^{-1} , which is about 3.5 and 1.3 times slower than that of Setup A and B respectively. Inclusion of subcycling in time increases the speed to 0.2×10^8 ZCs s^{-1} , thus proving to be faster by a factor of about 2.8.

We repeated some of the above tests with the NVIDIA Quadro RTX 5000 GPU available on the Frontera cluster, and encounter an overall reduction in performance by a factor of about 8 for uniform grid cases and about 15 times when using AMR.

We also ran Setup B with the aforementioned problem size on a single Frontier node (8 GPUs), and found the parallel efficiency to decrease to 80% compared to a single GPU performance. This drop is likely due to the additional ghost cell communication as well as decrease in compute load per GPU. The same setup was also run using only CPUs on a single Frontier node (64 CPUs), and we find the CPU run to be about an order of magnitude slower than the GPU one.

5.2. Profiling

Profiling is a fundamental tool to pinpoint the bottlenecks that make the code slow and inefficient. To identify the most expensive GPU kernels in `AsterX`, we make use of the HPCToolkit software [104] on Frontier. Specifically, we repeat the simulations for Setup A and B with the compute load of 240^3 cells on a single GPU, but enabling all the necessary options to output the measurement statistics for HPCToolkit. Here, we also ignore the computational costs for initialization routines.

For Setup A, in which we do not evolve spacetime, we find the most expensive kernel to be `AsterX_Fluxes`, which dictates about 70% of the total evolution time. This is expected since this routine involves computations for the reconstruction scheme PPM as well as the approximate Riemann solver. The next two expensive kernels, i.e. `AsterX_SourceTerms` and `AsterX_Con2Prim`, which involve calculations of the source terms and the C2P scheme respectively, consume about 11% and 10% of the evolution time. Instead, for Setup B, the spacetime solver kernels prove to be most expensive, as also reported in [67]. Here, `Z4c_RHS`, which computes the RHSs of the Z4c evolution variables, demands 34%, whereas, `Z4c_ADM2`, which computes the first time derivative of extrinsic curvature and second time derivatives of gauge quantities, consumes 23% of the total evolution. In this case, `AsterX_Fluxes` utilizes 21% whereas computation of $T_{\mu\nu}$ via `AsterX_Tmunu` takes 14% of run time. Since the spacetime and the flux solver kernels deem to be most costly, our future efforts will focus on optimizing them.

5.3. Strong scaling

Strong scaling measures the efficiency of the parallel performance of the code run with an increasing number of processors while keeping a fixed problem size. Ideally, in this case, increasing the number of GPUs should reduce the time to the solution proportionally, however, achieving perfect linear speedup on GPUs can be challenging due to factors such as communication overhead between GPUs, memory bandwidth limitations, and diminishing returns as the number of GPUs increases.

¹⁴ Z4c is hand-written in C++ which makes it difficult to optimize. Other spacetime evolution codes based on automated code generation are in preparation.

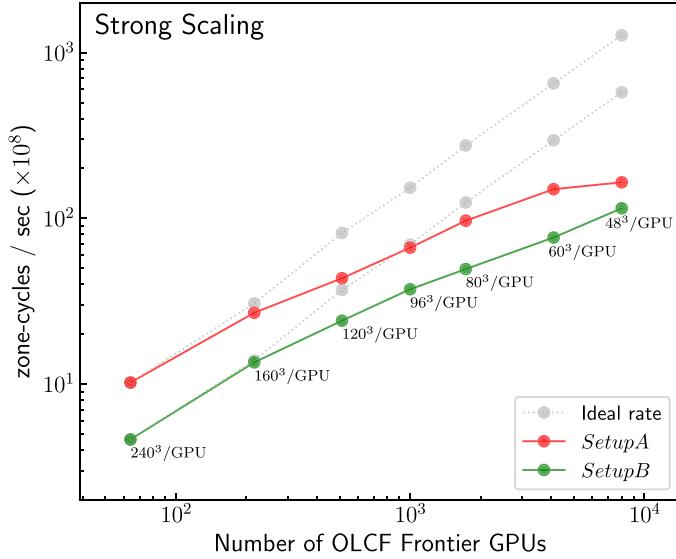


Figure 11. Results of strong scaling test performed for Setup A (static spacetime + uniform grid) in solid red line, and Setup B (dynamical spacetime + uniform grid) in solid green line. Corresponding ideal rates are shown in light grey dotted lines. Here, we vary the number of GPUs (nodes) from 64 (8) to 8000 (1000) as shown along the x -axis, while keeping a fixed problem size of $960 \times 960 \times 960$. Performance is quantified in terms of zone-cycles s^{-1} , shown along the y -axis. As we increase the number of GPUs, the compute load per GPU varies from 240^3 cells GPU^{-1} to 48^3 cells GPU^{-1} . For 512 GPUs, the strong scaling efficiency η_s is maintained at 54% and 65% for Setup A and B, but drops further to 13% and 20% respectively when using 8000 GPUs. One major reason for the decrease in η_s stems from the insufficient compute load per GPU, leading to its ineffective usage.

Table 3. Strong scaling test results for Setup A (static spacetime + uniform grid) and Setup B (dynamical spacetime + uniform grid) with a fixed problem size of $960 \times 960 \times 960$. The strong scaling efficiency η_s is normalized by the 8 node performance.

Nodes	GPUs	cells GPU^{-1}	Setup A		Setup B	
			$ZCs s^{-1}$	η_s	$ZCs s^{-1}$	η_s
8	64	240^3	10.0×10^8	1.0	4.6×10^8	1.0
27	216	160^3	26.9×10^8	0.79	13.5×10^8	0.86
64	512	120^3	43.5×10^8	0.54	24.1×10^8	0.65
125	1000	96^3	66.5×10^8	0.42	37.3×10^8	0.51
216	1728	80^3	96.8×10^8	0.36	49.4×10^8	0.40
512	4096	60^3	150.2×10^8	0.23	76.6×10^8	0.26
1000	8000	48^3	164.6×10^8	0.13	114.8×10^8	0.20

For **AsterX**, the strong scaling test results are illustrated in figure 11, and details presented in table 3. On uniform grid, we consider a problem size of $960 \times 960 \times 960$, and vary the computational load per GPU from 240^3 cells GPU^{-1} to 48^3 cells GPU^{-1} with increasing number of nodes, similar to the study reported in [67]. Due to limited computational resources, we perform each test up to 512 iterations. Our parallel efficiency for strong scaling η_s is normalized

based on the results obtained for 8 nodes (64 GPUs), for which we have 240^3 cells GPU^{-1} . Here, we find that Setup A and B run with about 10×10^8 and 4.6×10^8 ZCs s^{-1} respectively. With an about $\times 16$ increase in node count, η_s drops to 42% and 51% for Setup A and B, and decreases further to 13% and 20% respectively with a $\times 125$ node count increase. This decline in parallel efficiency can be explained by the significantly larger ratio of throughput and memory bandwidth to problem size on GPUs. In other words, we are not efficiently using the GPUs since the problem size per GPU becomes too small, and for the chosen simulation setup, operating below a compute load of about 96^3 cells GPU^{-1} results in significant inefficiency. We also find that for a given number of nodes, the performance in terms of ZCs s^{-1} is roughly twice as faster for Setup A which assumes static spacetime, than that of Setup B, which involves evolving the spacetime equations as well. In contrast, the parallel efficiency η_s is higher by roughly 10%–20% for Setup B when compared to that of Setup A. This is expected since the work load per GPU for Setup B is larger, in terms of solving the number of equations, as compared to Setup A.

For Setup C, which includes 8 levels of mesh refinement, we restrict the problem size to $480 \times 480 \times 480$ per level, and vary the computational load from 120^3 cells $\text{GPU}^{-1} \text{ level}^{-1}$ to 24^3 cells $\text{GPU}^{-1} \text{ level}^{-1}$. In this case, we find that the number of ZCs s^{-1} decreases at most by a factor of 2 as compared to the results of Setup B, and the parallel efficiency goes below 10% for a 1000 node run. In particular for the case without subcycling there is a known inefficiency in our code: Each level is independently distributed across all GPUs, which leads to much smaller per-GPU block sizes than if each GPU handled only blocks from a single level.

Overall, we find our results to be similar to the ones reported in [67].

5.4. Weak scaling

Weak scaling involves measuring the variation in the solution time with the number of GPUs, while maintaining a constant computational load per GPU. Ideally, for a perfect weak scaling, the solution time should remain constant, however, in practice, the parallel performance is affected by a number of factors stated earlier, such as inter-GPU and inter-node communication overheads, limited communication bandwidth, load imbalance, and synchronization overheads.

Our weak scaling results are illustrated in figures 12 and 13. For each of the Setups A and B, we consider two cases having a compute load of 128^3 cells GPU^{-1} and 240^3 cells GPU^{-1} , and vary the number of nodes (GPUs) from 8 (64) to 4096 (32 768). For Setup C instead, which involves 8 levels of AMR, we consider cases with 64^3 cells $\text{GPU}^{-1} \text{ level}^{-1}$ and 120^3 cells $\text{GPU}^{-1} \text{ level}^{-1}$. Summing up the compute load for all 8 levels effectively results in 128^3 cells GPU^{-1} and 240^3 cells GPU^{-1} for Setup C. Results for the case with higher compute load (240^3 cells GPU^{-1}) are presented in tables 4 and 5. Here too, the weak scaling efficiency η_w is normalized by the 8-node performance.

For 8 nodes with 240^3 cells GPU^{-1} , we find the number of $\text{ZCs s}^{-1} \text{ GPU}^{-1}$ to be about 0.16×10^8 and 0.072×10^8 for Setups A and B respectively. In general, Setup A runs about 2 to 2.5 times faster than Setup B. The weak scaling efficiency η_w remains as high as 85% and 91% upto 216 nodes (1728 GPUs) but decreases to 77% and 67% for 4096 nodes (32 768 GPUs) for Setups A and B respectively. As mentioned earlier, one reason for the decrease in performance is related to the communication overhead. The box-shaped grid structure used for the scaling test distributes the workload per GPU in cubical blocks, where the ghost cell communication scales with the number of blocks as $\sim n^3$. This, in turn, is limited by the communication bandwidth, as also pointed out in [67].

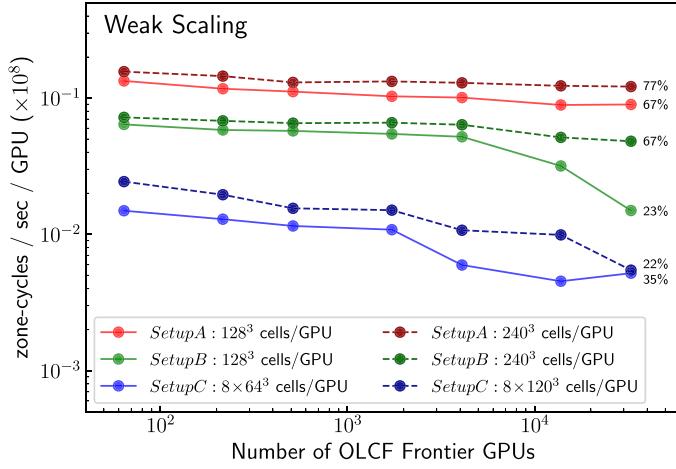


Figure 12. Results of weak scaling test performed for Setup A (static spacetime + uniform grid), Setup B (dynamical spacetime + uniform grid) and Setup C (dynamical spacetime + 8 levels AMR). Number of GPUs (nodes) is varied from 64 (8) to 32 768 (4096) along the x -axis, while keeping a fixed compute load of 128^3 cells GPU^{-1} and 240^3 cells GPU^{-1} for Setup A and B, and effectively for Setup C as well. Along y -axis, we show performance in terms of zone-cycles $\text{s}^{-1} \text{GPU}^{-1}$. In general, Setup A runs are about 2–2.5 times faster than Setup B, and about 6–22 times faster than Setup C. Overall, the runs with higher compute load perform better, for which, with 4096 GPUs, the weak scaling efficiency η_w remains at 83% and 88% for Setup A and B, but drops further to 77% and 67% respectively when using 32 768 GPUs. For Setup C instead, η_w decreases to 62% for the run with 1728 GPUs, and further down to 22% for 32 768 GPUs.

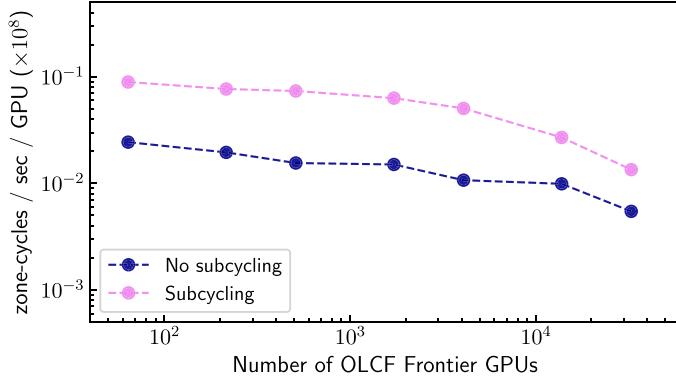


Figure 13. Comparison of weak scaling results for Setup C (dynamical spacetime + 8 levels AMR) without and with subcycling in time. An overall gain of factor 2.5–4.5 is obtained in performance when using subcycling, as expected since larger time steps are employed for coarser refinement levels, whereas runs without subcycling use the same time step size for all levels.

While the overheads could be minimized by adopting efficient schemes that would require less number of ghost zones, such as Discontinuous Galerkin methods [105] and compact FD [106], increasing the communication bandwidth instead requires improving the cluster infrastructure. Another strategy would be to improve the computation-communication overlap. As

Table 4. Weak scaling test results for Setup A (static spacetime + uniform grid) and Setup B (dynamical spacetime + uniform grid) with a fixed computational load of 240^3 cells GPU^{-1} . Here, the weak scaling efficiency η_w is normalized by the 8 node performance.

Nodes	GPUs	Grid	Setup A		Setup B	
			$\text{ZCs s}^{-1} \text{GPU}^{-1}$	η_w	$\text{ZCs s}^{-1} \text{GPU}^{-1}$	η_w
8	64	$960 \times 960 \times 960$	0.157×10^8	1.0	0.072×10^8	1.0
27	216	$1440 \times 1440 \times 1440$	0.145×10^8	0.92	0.068×10^8	0.94
64	512	$1920 \times 1920 \times 1920$	0.130×10^8	0.83	0.065×10^8	0.91
216	1728	$2880 \times 2880 \times 2880$	0.133×10^8	0.85	0.066×10^8	0.91
512	4096	$3840 \times 3840 \times 3840$	0.129×10^8	0.83	0.064×10^8	0.88
1728	13 824	$5760 \times 5760 \times 5760$	0.123×10^8	0.79	0.051×10^8	0.71
4096	32 768	$7680 \times 7680 \times 7680$	0.121×10^8	0.77	0.048×10^8	0.67

Table 5. Weak scaling test results for Setup C (dynamical spacetime + 8 levels AMR) with a fixed computational load of 120^3 cells $\text{GPU}^{-1} \text{level}^{-1}$ without and with subcycling in time. Here too, the weak scaling efficiency η_w is normalized by the 8 node performance.

Nodes	GPUs	Grid	Setup C		Setup C + subcycling	
			$\text{ZCs s}^{-1} \text{GPU}^{-1}$	η_w	$\text{ZCs s}^{-1} \text{GPU}^{-1}$	η_w
8	64	$8 \times (480 \times 480 \times 480)$	0.0244×10^8	1.0	0.0893×10^8	1.0
27	216	$8 \times (720 \times 720 \times 720)$	0.0195×10^8	0.8	0.0767×10^8	0.86
64	512	$8 \times (960 \times 960 \times 960)$	0.0155×10^8	0.64	0.0736×10^8	0.83
216	1728	$8 \times (1440 \times 1440 \times 1440)$	0.0150×10^8	0.62	0.0629×10^8	0.70
512	4096	$8 \times (1920 \times 1920 \times 1920)$	0.0107×10^8	0.44	0.0506×10^8	0.57
1728	13 824	$8 \times (2880 \times 2880 \times 2880)$	0.0099×10^8	0.41	0.0270×10^8	0.30
4096	32 768	$8 \times (3840 \times 3840 \times 3840)$	0.0055×10^8	0.22	0.0135×10^8	0.15

noted in [67], the current version of CarpetX does not launch the next GPU kernel until all the ghost-zones from the previous kernel are filled. The ability to launch subsequent kernels that may not immediately require ghost-zone values of the previous kernel while they are being filled, can increase the computation-communication overlap, a feature that is under development in CarpetX.

When considering AMR with 120^3 cells $\text{GPU}^{-1} \text{level}^{-1}$, we compare performance with and without the activation of subcycling in time. We notice that without subcycling, the 8-node performance goes down to $0.0244 \times 10^8 \text{ ZCs s}^{-1} \text{GPU}^{-1}$, and overall, we notice it to be about 6 to 22 times slower than Setup A (see figure 12). Whereas with subcycling, for 8 nodes, we get $0.0893 \times 10^8 \text{ ZCs s}^{-1} \text{GPU}^{-1}$, and in general, the performance is quantitatively similar to Setup B upto 512 nodes. Overall, we find that subcycling boosts the performance by a factor of about 2.5 to 4.5 when compared to the results without subcycling (see figure 13). This is expected since time steps are doubled as we move from finer to coarser grids when using subcycling, whereas a uniform time step is used for all levels for runs without subcycling. In terms of parallel efficiency, η_w reduces to 62% (70%) without (with) subcycling when using 216 nodes (1728 GPUs), and further decreases to 41% (30%) and 22% (15%) for 1728 nodes (13 824 GPUs) and 4096 nodes (32 768 GPUs) respectively. Here too, we find our results without subcycling to be consistent with the ones reported in [67].

The drop in $Z\text{Cs s}^{-1} \text{GPU}^{-1}$ as well as η_w can be attributed to reasons such as prolongation/restriction operations for interpolations to finer/coarser grids that demand additional computations and communication, as also highlighted in [67]. Also, the existing CarpetX infrastructure enforces each GPU to work on one or more blocks on each refinement level, resulting in serial calculations for the blocks assigned to each GPU. To make this handling efficient, one could let each GPU work on one or more blocks from any given level, thus enhancing weak scaling.

6. Conclusions and outlook

In this paper, we have presented a new three-dimensional GPU accelerated open-source numerical code *AsterX* [69] that solves the full set of GRMHD equations in Cartesian coordinates on a dynamical background. The code is based on CarpetX, the new driver for the Einstein Toolkit which, in turn, is built on top of AMReX, a software machinery designed for massively parallel, block-structured AMR applications. *AsterX* has been written from scratch in C++ but has adopted several algorithms from Spritz [57], and the RePrimAnd library [107], while also benefitting from other publicly available codes such as WhiskyTHC [108], GRHydro [48] and IllinoisGRMHD [52].

In *AsterX*, the employed GRMHD equations are based on the flux-conservative Valencia formulation, and the divergence-free character of magnetic fields is guaranteed by evolving the staggered vector potential. To accurately handle discontinuities, high resolution shock capturing schemes, enabled via reconstruction methods TVD MINMOD and PPM along with Lax-Friedrichs and HLLE flux solvers, have also been utilized. Our primary choice for primitive variable recovery is based on the 2D Noble scheme with 1D Palenzuela routine set as back up.

AsterX has been validated via a series of classical tests in special and general relativity. We demonstrated the code's ability to accurately reproduce exact solutions by simulating 1D MHD shock tube problems. Whereas, our convergence study illustrates the code to be at least second order accurate. A number of special relativistic MHD tests in 2D were also conducted, namely, cylindrical explosion, magnetic rotor, and loop advection, which proved to be in good agreement with literature. We also simulated the Kelvin–Helmholtz instability in 2D, to examine the handling of block-structured AMR, and confirmed that it performed reliably without issues. For our final test, we evolved a stable magnetized TOV configuration in a dynamical spacetime and successfully obtained peak oscillation frequency modes up to fourth harmonic.

Benchmarking the code on OLCF's Frontier cluster indicated that performance on a single GPU node is at least an order of magnitude faster than the respective CPU node. Strong scaling efficiency is maintained up to 80% for 1 GPU node when comparing with a single GPU throughput. Whereas, for weak scaling, parallel efficiencies are retained to 67%–77% for 4096 nodes (32 768 GPUs) relative to the 8 node results. Moreover, when employing AMR, use of subcycling provided a significant gain by a factor 2.5–4.5 in performance. Profiling the code also helped us to identify the bottlenecks which will be optimized in the near future.

At present, the code supports different analytical EOS, and is currently being extended to deal with tabulated microphysical EOSs. The primitive variable recovery algorithm of the RePrimAnd library [107] has also been ported and is undergoing testing, and we plan to add other routines of [109, 110]. In order to obtain higher order convergence, WENO-Z and MP5 reconstruction schemes have already been implemented and are being tested. Next, we intend to include support for advanced Riemann solver HLLD together with upwind CT-HLLD scheme [84, 111]. For modelling neutrino transport, we have also started looking into porting the moment-based M1 scheme [112] which would be coupled with the Monte–Carlo scheme,

to closely follow the approach of [113]. These developments will enable us to accurately model astrophysical systems such as core-collapse supernovae and BNS mergers, and will be presented in our next paper.

Meanwhile, we have already implemented a photon radiation leakage scheme in `AsterX`, which is currently being tested via black-hole accretion disk simulations and will be presented in another paper. This scheme will then be utilized in our production simulations to study gas dynamics in supermassive binary black-hole mergers, allowing us to reach unprecedented time and length scales via exascale computing.

Data availability statement

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

Acknowledgments

The authors would like to thank Michail Chabanov, Samuel Cupp, Zachariah Etienne, Vassilios Mewes, Phillip Moesta, Swapnil Shankar, Samuel Tootle, and Weiqun Zhang for insightful discussions. We are also grateful to Alexandru Dima, Cheng-Hsin Cheng, Michail Chabanov, and Jake Doherty for providing the `NewRadX` module. Moreover, we would like to thank the OLCF staff for their support during the Frontier Virtual Hackathon 2023.

We gratefully acknowledge the National Science Foundation (NSF) for financial support from Grant Nos. PHY-2110338, OAC-2004044/1550436/2004157/2004879, AST-2009330, OAC-1811228, OAC-2031744, OAC-2310548, OAC-2005572, OAC-OAC-2103680, and PHY1912632, as well as NASA for financial support from NASA TCAN Grant Nos. 80NSSC18K1488 and 80NSSC24K0100 to RIT. Research at Perimeter Institute is supported in part by the Government of Canada through the Department of Innovation, Science and Economic Development and by the Province of Ontario through the Ministry of Colleges and Universities. R Ciolfi and B Giacomazzo acknowledge support from the European Union under NextGenerationEU, PRIN 2022 Protocol No. 2022KX2Z3B.

This research used resources of the Oak Ridge Leadership Computing Facility provided by Director's Discretionary program on Frontier supercomputer (Grant No. AST182) at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725, and TACC's Frontera supercomputer allocations (Grant Nos. PHY20010 and AST-20021). Additional resources were provided by RIT's BlueSky and Green Pairie and Lagoon clusters acquired with NSF Grant Nos. PHY-2018420, PHY-0722703, PHY-1229173, and PHY1726215. Access to the Deep Bayou cluster was provided by Louisiana State University's HPC group.

ORCID iDs

- Jay V Kalinani  <https://orcid.org/0000-0002-2945-1142>
- Liwei Ji  <https://orcid.org/0000-0001-9008-0267>
- Lorenzo Ennoggi  <https://orcid.org/0000-0002-2771-5765>
- Federico G Lopez Armengol  <https://orcid.org/0000-0002-4882-5672>
- Lucas Timotheo Sanches  <https://orcid.org/0000-0001-6764-1812>
- Bing-Jyun Tsao  <https://orcid.org/0000-0003-4614-0378>

Steven R Brandt  <https://orcid.org/0000-0002-7979-2906>
 Manuela Campanelli  <https://orcid.org/0000-0002-8659-6591>
 Riccardo Ciolfi  <https://orcid.org/0000-0003-3140-8933>
 Bruno Giacomazzo  <https://orcid.org/0000-0002-6947-4023>
 Roland Haas  <https://orcid.org/0000-0003-1424-6178>
 Erik Schnetter  <https://orcid.org/0000-0002-4518-9017>
 Yosef Zlochower  <https://orcid.org/0000-0002-7541-6612>

References

- [1] Abbott B P *et al* 2017 GW170817: observation of gravitational waves from a binary neutron star inspiral *Phys. Rev. Lett.* **119** 161101
- [2] Abbott B P *et al* 2017 Gravitational waves and gamma-rays from a binary neutron star merger: GW170817 and GRB 170817A *Astrophys. J. Lett.* **848** L13
- [3] Pian E *et al* 2017 Spectroscopic identification of r-process nucleosynthesis in a double neutron-star merger *Nature* **551** 67–70
- [4] Abbott B P *et al* 2017 Multi-messenger observations of a binary neutron star merger *Astrophys. J. Lett.* **848** L12
- [5] Arcavi I *et al* 2017 Optical emission from a kilonova following a gravitational-wave-detected neutron-star merger *Nature* **551** 64–66
- [6] Kasen D, Metzger B, Barnes J, Quataert E and Ramirez-Ruiz E 2017 Origin of the heavy elements in binary neutron-star mergers from a gravitational-wave event *Nature* **551** 80–84
- [7] Mooley K P, Deller A T, Gottlieb O, Nakar E, Hallinan G, Bourke S, Frail D A, Horesh A, Corsi A and Hotokezaka K 2018 Superluminal motion of a relativistic jet in the neutron-star merger GW170817 *Nature* **561** 355–9
- [8] Lazzati D, Perna R, Morsany B J, Lopez-Camara D, Cantiello M, Ciolfi R, Giacomazzo B and Workman J C 2018 Late time afterglow observations reveal a collimated relativistic jet in the ejecta of the binary neutron star merger GW170817 *Phys. Rev. Lett.* **120** 241103
- [9] Ghirlanda G *et al* 2019 Compact radio emission indicates a structured jet was produced by a binary neutron star merger *Science* **363** 968–71
- [10] Margalit B and Metzger B D 2017 Constraining the maximum mass of neutron stars from multi-messenger observations of GW170817 *Astrophys. J. Lett.* **850** L19
- [11] Abbott B P *et al* 2018 GW170817: measurements of neutron star radii and equation of state *Phys. Rev. Lett.* **121** 161101
- [12] Raithel C A 2019 Constraints on the neutron star equation of state from GW170817 *Eur. Phys. J. A* **55** 80
- [13] Nicholl M, Margalit B, Schmidt P, Smith G P, Ridley E J and Nuttall J 2021 Tight multimessenger constraints on the neutron star equation of state from GW170817 and a forward model for kilonova light-curve synthesis *Mon. Not. R. Astron. Soc.* **505** 3016–32
- [14] Abbott B P *et al* 2020 GW190425: observation of a compact binary coalescence with total mass $3.4 m_{\odot}$ *Astrophys. J.* **892** L3
- [15] Abbott R *et al* 2021 Observation of gravitational waves from two neutron star–black hole coalescences *Astrophys. J. Lett.* **915** L5
- [16] Abac A G *et al* (LIGO Scientific Collaboration, Virgo Collaboration, KAGRA Collaboration) 2024 Observation of gravitational waves from the coalescence of a 2.5–4.5m compact object and a neutron star (arXiv:2404.04248)
- [17] Punturo M *et al* 2010 The third generation of gravitational wave observatories and their science reach *Class. Quantum Grav.* **27** 084007
- [18] Ballmer S W *et al* 2022 Snowmass2021 cosmic frontier white paper: future gravitational-wave detector facilities *Snowmass 2021*
- [19] Agazie G *et al* 2023 The nanograv 15 yr data set: constraints on supermassive black hole binaries from the gravitational-wave background *Astrophys. J. Lett.* **952** L37
- [20] Amaro-Seoane P *et al* 2017 Laser interferometer space antenna (arXiv:1702.00786)
- [21] Amaro-Seoane P *et al* 2023 Astrophysics with the laser interferometer space antenna *Living Rev. Relativ.* **26** 2

[22] Bogdanović T, Miller M C and Blecha L 2022 Electromagnetic counterparts to massive black-hole mergers *Living Rev. Relativ.* **25** 3

[23] Ruiz M, Lang R N, Paschalidis V and Shapiro S L 2016 Binary neutron star mergers: a jet engine for short gamma-ray bursts *Astrophys. J. Lett.* **824** L6

[24] Ciolfi R, Kastaun W, Kalinani J V and Giacomazzo B 2019 First 100 ms of a long-lived magnetized neutron star formed in a binary neutron star merger *Phys. Rev. D* **100** 023005

[25] Ciolfi R 2020 Collimated outflows from long-lived binary neutron star merger remnants *Mon. Not. R. Astron. Soc. Lett.* **495** L66–L70

[26] Mösta P, Radice D, Haas R, Schnetter E and Bernuzzi S 2020 A magnetar engine for short GRBs and kilonovae *Astrophys. J. Lett.* **901** L37

[27] Sun L, Ruiz M, Shapiro S L and Tsokaros A 2022 Jet launching from binary neutron star mergers: incorporating neutrino transport and magnetic fields *Phys. Rev. D* **105** 104028

[28] Combi L and Siegel D M 2023 Jets from neutron-star merger remnants and massive blue kilonovae *Phys. Rev. Lett.* **131** 231402

[29] Most E R 2023 Impact of a mean field dynamo on neutron star mergers leading to magnetar remnants *Phys. Rev. D* **108** 123012

[30] Kiuchi K, Reboul-Salze A, Shibata M and Sekiguchi Y 2024 A large-scale magnetic field produced by a solar-like dynamo in binary neutron star mergers *Nat. Astron.* **8** 298–307

[31] Fernández R, Tchekhovskoy A, Quataert E, Foucart F and Kasen D 2019 Long-term grmhd simulations of neutron star merger accretion discs: implications for electromagnetic counterparts *Mon. Not. R. Astron. Soc.* **482** 3373–93

[32] Ciolfi R and Vijay Kalinani J 2020 Magnetically driven baryon winds from binary neutron star merger remnants and the blue kilonova of 2017 August *Astrophys. J. Lett.* **900** L35

[33] Curtis S, Mösta P, Wu Z, Radice D, Roberts L, Riciglio G and Perego A 2023 *r*-process nucleosynthesis and kilonovae from hypermassive neutron star post-merger remnants *Mon. Not. R. Astron. Soc.* **518** 5313–22

[34] Combi L and Siegel D M 2023 GRMHD simulations of neutron-star mergers with weak interactions: *r*-process nucleosynthesis and electromagnetic signatures of dynamical ejecta *Astrophys. J.* **944** 28

[35] Noble S C, Mundim B C, Nakano H, Krolik J H, Campanelli M, Zlochower Y and Yunes N 2012 Circumbinary MHD accretion into inspiraling binary black holes *Astrophys. J.* **755** 51

[36] Lopez Armengol F G, Combi L, Campanelli M, Noble S C, Krolik J H, Bowen D B, Avara M J, Mewes V and Nakano H 2021 Circumbinary disk accretion into spinning black hole binaries *Astrophys. J.* **913** 16

[37] Noble S C, Krolik J H, Campanelli M, Zlochower Y, Mundim B C, Nakano H and Zilhão M 2021 Mass-ratio and magnetic flux dependence of modulated accretion from circumbinary disks *Astrophys. J.* **922** 175

[38] Bowen D B, Campanelli M, Krolik J H, Mewes V and Noble S C 2017 Relativistic dynamics and mass exchange in binary black hole mini-disks *Astrophys. J.* **838** 42

[39] Bowen D B, Mewes V, Campanelli M, Noble S C, Krolik J H and Zilhão M 2018 Quasi-periodic behavior of mini-disks in binary black holes approaching merger *Astrophys. J. Lett.* **853** L17

[40] Bowen D B, Mewes V, Noble S C, Avara M, Campanelli M and Krolik J H 2019 Quasi-periodicity of supermassive binary black hole accretion approaching merger *Astrophys. J.* **879** 76

[41] Combi L, Lopez Armengol F G, Campanelli M, Ireland B, Noble S C, Nakano H and Bowen D 2021 Superposed metric for spinning black hole binaries approaching merger *Phys. Rev. D* **104** 044041

[42] Avara M J, Krolik J H, Campanelli M, Noble S C, Bowen D and Ryu T 2023 Accretion onto a supermassive black hole binary before merger (arXiv:2305.18538)

[43] Gutiérrez E M, Combi L, Noble S C, Campanelli M, Krolik J H, López Armengol F G and García F 2022 Electromagnetic signatures from supermassive binary black holes approaching merger *Astrophys. J.* **928** 137

[44] Gutiérrez E M, Combi L, Romero G E and Campanelli M 2024 Non-thermal radiation from dual jet interactions in supermassive black hole binaries (arXiv:2301.04280)

[45] Ennoggi L *et al* 2024 Accretion onto merging supermassive black hole binaries (in preparation)

[46] Giacomazzo B and Rezzolla L 2007 WhiskyMHD: a new numerical code for general relativistic magnetohydrodynamics *Class. Quantum Grav.* **24** 235

[47] Giacomazzo B, Rezzolla L and Baiotti L 2011 Accurate evolutions of inspiralling and magnetized neutron stars: equal-mass binaries *Phys. Rev. D* **83** 044014

[48] Mösta P *et al* 2014 GRHydro: a new open-source general-relativistic magnetohydrodynamics code for the Einstein toolkit *Class. Quantum Grav.* **31** 015005

[49] Gammie C F, McKinney J C and Tóth G 2003 Harm: a numerical scheme for general relativistic magnetohydrodynamics *Astrophys. J.* **589** 444

[50] Noble S C, Gammie C F, McKinney J C and Del Zanna L 2006 Primitive variable solvers for conservative general relativistic magnetohydrodynamics *Astrophys. J.* **641** 626

[51] Murguia-Berthier A 2021 *et al* HARM3D+NUC: a new method for simulating the post-merger phase of binary neutron star mergers with GRMHD, tabulated EOS and neutrino leakage (arXiv:2106.05356)

[52] Etienne Z B, Paschalidis V, Haas R, Mösta P and Shapiro S L 2015 IllinoisGRMHD: an open-source, user-friendly GRMHD code for dynamical spacetimes *Class. Quantum Grav.* **32** 175009

[53] Werneck L *et al* 2023 Addition of tabulated equation of state and neutrino leakage support to Illinois GRMHD *Phys. Rev. D* **107** 044037

[54] Palenzuela C, Miñano B, Arbona A, Bona-Casas C, Bona C and Massó J 2021 Simflowny 3: an upgraded platform for scientific modeling and simulation *Comput. Phys. Commun.* **259** 107675

[55] Palenzuela C, Liebling S and Miñano B 2022 Large eddy simulations of magnetized mergers of neutron stars with neutrinos *Phys. Rev. D* **105** 103020

[56] Izquierdo M R, Bezares M, Liebling S and Palenzuela C 2024 Large eddy simulations of magnetized mergers of black holes and neutron stars *Phys. Rev. D* **110** 083017

[57] Cipolletta F, Kalinani J, Giacomazzo B and Ciolfi R 2020 Spritz: a new fully general-relativistic magnetohydrodynamic code *Class. Quantum Grav.* **37** 135010

[58] Cipolletta F, Vijay Kalinani J, Giangrandi E, Giacomazzo B, Ciolfi R, Sala L and Giudici B 2021 Spritz: general relativistic magnetohydrodynamics with neutrinos *Class. Quantum Grav.* **38** 085021

[59] Kalinani J V, Ciolfi R, Kastaun W, Giacomazzo B, Cipolletta F and Ennoggi L 2022 Implementing a new recovery scheme for primitive variables in the general relativistic magnetohydrodynamic code spritz *Phys. Rev. D* **105** 103031

[60] Cook W, Daszuta B, Fields J, Hammond P, Albanesi S, Zappa F, Bernuzzi S and Radice D 2023 GR-Athena++: general-relativistic magnetohydrodynamics simulations of neutron star spacetimes (arXiv:2311.04989)

[61] Grete P *et al* 2023 Parthenon—a performance portable block-structured adaptive mesh refinement framework *Int. J. High Perform. Comput. Appl.* **37** 465–86

[62] Carter Edwards H, Trott C R and Sunderland D 2014 Kokkos: enabling manycore performance portability through polymorphic memory access patterns *J. Parallel Distrib. Comput.* **74** 3202–16

[63] Zhang W *et al* 2019 AMReX: a framework for block-structured adaptive mesh refinement *J. Open Source Softw.* **4** 1370

[64] Zhang W, Myers A, Gott K, Almgren A and Bell J 2021 AMReX: block-structured adaptive mesh refinement for multiphysics applications *Int. J. High Perform. Comput. Appl.* **35** 508–26

[65] Grete P, Glines F W and O’Shea B W 2020 K-athena: a performance portable structured grid finite volume magnetohydrodynamics code *IEEE Trans. Parallel Distrib. Syst.* **32** 85–97

[66] Liska M, Hesp C, Tchekhovskoy A, Ingram A, van der Klis M, Markoff S B and Van Moer M 2021 Disc tearing and bardeen-petterson alignment in GRMHD simulations of highly tilted thin accretion discs *Mon. Not. R. Astron. Soc.* **507** 983–90

[67] Shankar S, Mösta P, Brandt S R, Haas R, Schnetter E and de Graaf Y 2023 GRaM-X: a new GPU-accelerated dynamical spacetime GRMHD code for exascale computing with the Einstein Toolkit *Class. Quantum Grav.* **40** 205009

[68] Cupp S *et al* 2023 The Einstein Toolkit: ET_2023_11 (<https://doi.org/10.5281/zenodo.1038040>)

[69] Kalinani J V *et al* 2024 The AsterX code (v1.0) (Zenodo) (<https://doi.org/10.5281/zenodo.14502139>)

[70] Bernuzzi S and Hilditch D 2010 Constraint violation in free evolution schemes: comparing the BSSNOK formulation with a conformal decomposition of the Z4 formulation *Phys. Rev. D* **81** 084003

[71] Hilditch D, Bernuzzi S, Thierfelder M, Cao Z, Tichy W and Brügmann B 2013 Compact binary evolutions with the Z4C formulation *Phys. Rev. D* **88** 084057

[72] Etienne Z B, Paschalidis V, Liu Y T and Shapiro S L 2012 Relativistic magnetohydrodynamics in dynamical spacetimes: improved electromagnetic gauge condition for adaptive mesh refinement grids *Phys. Rev. D* **85** 024013

[73] Etienne Z B, Liu Y T and Shapiro S L 2010 Relativistic magnetohydrodynamics in dynamical spacetimes: a new adaptive mesh refinement implementation *Phys. Rev. D* **82** 084031

[74] Farris B D, Gold R, Paschalidis V, Etienne Z B and Shapiro S L 2012 Binary black-hole mergers in magnetized disks: simulations in full general relativity *Phys. Rev. Lett.* **109** 221102

[75] Roe P L 1986 Characteristic-based schemes for the Euler equations *Ann. Rev. Fluid Mech.* **18** 337

[76] Colella P and Woodward P R 1984 The piecewise parabolic method (PPM) for gas dynamical simulations *J. Comput. Phys.* **54** 174

[77] McCorquodale P and Colella P 2011 A high-order finite-volume method for conservation laws on locally refined grids *Commun. Appl. Math. Comput. Sci.* **6** 1

[78] Borges R, Carmona M, Costa B and Sun Don W 2008 An improved weighted essentially non-oscillatory scheme for hyperbolic conservation laws *J. Comput. Phys.* **227** 3191–211

[79] Suresh A and Huynh H T 1997 Accurate monotonicity-preserving schemes with Runge–Kutta time stepping *J. Comput. Phys.* **136** 83–99

[80] Toro E F 2013 *Riemann Solvers and Numerical Methods for Fluid Dynamics: A Practical Introduction* (Springer)

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

[82] Del Zanna L, Bucciantini N and Londrillo P 2003 An efficient shock-capturing central-type scheme for multidimensional relativistic flows-II. Magnetohydrodynamics *Astron. Astrophys.* **400** 397

[83] Balsara D S and Spicer D S 1999 A staggered mesh algorithm using high order Godunov fluxes to ensure solenoidal magnetic fields in magnetohydrodynamic simulations *J. Comput. Phys.* **149** 270

[84] Mignone A and Del Zanna L 2021 Systematic construction of upwind constrained transport schemes for mhd *J. Comput. Phys.* **424** 109748

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

[86] Siegel D M and Metzger B D 2018 Three-dimensional GRMHD simulations of neutrino-cooled accretion disks from neutron star mergers *Astrophys. J.* **858** 52

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

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

[89] Berger M and Rigoutsos I 1991 An algorithm for point clustering and grid generation *IEEE Trans. Syst. Man Cybern.* **21** 1278–86

[90] Mongwane B 2015 Toward a consistent framework for high order mesh refinement schemes in numerical relativity *Gen. Relativ. Gravit.* **47** 60

[91] Guillaume Q et al 2019 NSIMD: agenium scale vectorization library for CPUs and GPUs (available at: <https://github.com/agenium-scale/nsimd>)

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

[93] Giacomazzo B and Rezzolla L 2006 The exact solution of the riemann problem in relativistic magnetohydrodynamics *J. Fluid Mech.* **562** 223

[94] Tóth G 2000 The $\nabla \cdot b = 0$ constraint in shock-capturing magnetohydrodynamics codes *J. Comput. Phys.* **161** 605

[95] Beckwith K and Stone J M 2011 A second-order Godunov method for multi-dimensional relativistic magnetohydrodynamics *Astrophys. J. Suppl. Ser.* **193** 6

[96] Gardiner T A and Stone J M 2005 An unsplit Godunov method for ideal MHD via constrained transport *J. Comput. Phys.* **205** 509

[97] Stone J M, Gardiner T A, Teuben P, Hawley J F and Simon J B 2008 Athena: a new code for astrophysical MHD *Astrophys. J. Suppl. Ser.* **178** 137

[98] Springel V 2010 Smoothed particle hydrodynamics in astrophysics *Annu. Rev. Astron. Astrophys.* **48** 391–430

[99] Oppenheimer J R and Volkoff G M 1939 On massive neutron cores *Phys. Rev.* **55** 374–81

[100] Tolman R C 1939 Static solutions of Einstein's field equations for spheres of fluid *Phys. Rev.* **55** 364

[101] Ciolfi R and Rezzolla L 2013 Twisted-torus configurations with large toroidal magnetic fields in relativistic stars *Mon. Not. R. Astron. Soc. Lett.* **435** L43–L47

[102] Suvorov A G and Glampedakis K 2023 Magnetic equilibria of relativistic axisymmetric stars: the impact of flow constants *Phys. Rev. D* **108** 084006

[103] Font J A, Goodale T, Iyer S, Miller M, Rezzolla L, Seidel E, Stergioulas N, Suen W and Tobias M 2002 Three-dimensional numerical general relativistic hydrodynamics. II. Long-term dynamics of single relativistic stars *Phys. Rev. D* **65** 084024

[104] Adhianto L, Banerjee S, Fagan M, Krentel M, Marin G, Mellor-Crummey J and Tallent N R 2010 HPCTOOLKIT: tools for performance analysis of optimized parallel programs *Concurr. Comput. Pract. Exp.* **22** 685–701

[105] Miller J M and Schnetter E 2016 An operator-based local discontinuous Galerkin method compatible with the BSSN formulation of the Einstein equations *Class. Quantum Grav.* **34** 015003

[106] Daszuta B 2024 Spectrally-tuned compact finite-difference schemes with domain decomposition and applications to numerical relativity *J. Comput. Phys.* **508** 112958

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

[108] Radice D, Rezzolla L and Galeazzi F 2014 High-order fully general-relativistic hydrodynamics: new approaches and tests *Class. Quantum Grav.* **31** 075012

[109] Cerdá-Durán P, Font J A and Dimmelmeier H 2007 General relativistic simulations of passive-magneto-rotational core collapse with microphysics *Astron. Astrophys.* **474** 169–91

[110] Newman W I and Hamlin N D 2014 Primitive variable determination in conservative relativistic magnetohydrodynamic simulations *SIAM J. Sci. Comput.* **36** B661–83

[111] Kiuchi K, Held L E, Sekiguchi Y and Shibata M 2022 Implementation of advanced riemann solvers in a neutrino-radiation magnetohydrodynamics code in numerical relativity and its application to a binary neutron star merger *Phys. Rev. D* **106** 124041

[112] Radice D, Bernuzzi S, Perego A and Haas R 2022 A new moment-based general-relativistic neutrino-radiation transport code: methods and first applications to neutron star mergers *Mon. Not. R. Astron. Soc.* **512** 1499–521

[113] Izquierdo M R, Abalos F and Palenzuela C 2024 Guided moments formalism: a new efficient full-neutrino treatment for astrophysical simulations *Phys. Rev. D* **109** 043044