

Numerical relativity simulations of compact binaries: Comparison of cell- and vertex-centered adaptive meshes

Boris Daszuta¹,¹ William Cook,¹ Peter Hammond²,² Jacob Fields^{2,3},^{2,3} Eduardo M. Gutiérrez^{2,3},^{2,3} Sebastiano Bernuzzi¹,¹ and David Radice^{2,3,4}

¹*Theoretisch-Physikalisches Institut, Friedrich-Schiller-Universität Jena, 07743, Jena, Germany*

²*Institute for Gravitation and the Cosmos, The Pennsylvania State University, University Park, Pennsylvania 16802, USA*

³*Department of Physics, The Pennsylvania State University, University Park, Pennsylvania 16802, USA*

⁴*Department of Astronomy and Astrophysics, The Pennsylvania State University, University Park, Pennsylvania 16802, USA*



(Received 24 June 2024; accepted 19 September 2025; published 5 November 2025)

Given the compact binary evolution problem of numerical relativity, in the finite-difference, block-based, adaptive mesh refinement context, choices must be made on how evolved fields are to be discretized. In GR-Athena++, the space-time solver was previously fixed to be vertex centered. Here, our recent extensions to a cell-centered treatment are described. Simplifications in the handling of variables during the treatment of general relativistic magneto-hydrodynamical (GRMHD) evolution are found. A novelty is that performance comparison for the two choices of grid sampling is made within a single code base. In the case of a binary black hole inspiral-merger problem, by evolving geometric fields on vertex centers, an average $\sim 20\%$ speed increase is observed, when compared against cell-centered sampling. The opposite occurs in the GRMHD setting. A binary neutron star inspiral-merger-collapse problem, representative of typical production simulations, is considered. We find that cell-centered sampling for the space-time solver improves performance, by a similar factor.

DOI: [10.1103/lxp5-c1g1](https://doi.org/10.1103/lxp5-c1g1)

I. INTRODUCTION

Fashioning an accurate description of astrophysically sourced, binary constituent encounters, under extreme conditions, is a pressing challenge for the numerical relativity (NR) community. Multimessenger detection efforts [1–3] benefit when supplemented by NR data. Observational efforts of next-generation gravitational-wave detectors [4–7] will also greatly benefit from higher fidelity reference numerical data that spans a larger range of parameter space (or, e.g., models [8] suitably informed by). To confront simulation challenges, sophisticated general relativistic, magneto-hydrodynamical (GRMHD) techniques (see, e.g., [9]), and microphysics descriptions [10,11] have been pursued. In tandem, the development of code infrastructure, and algorithms, with the aim of improved performance scaling on modern high-performance computing (HPC) resources, has also received recent, particular attention. Some examples of modern NR codes under active

development with these goals include: BAMPS [12,13], Dendro-GR [14], Einstein Toolkit [15] (GRaM-X [16]), GRChombo [17], NMesh [18], SpECTRE [19], and SPHINCS_BSSN [20]. A variety of numerical methods [21–23], algorithmic approaches [24–27], and NR formulations [28–35] are employed, to list but a few prominent examples. In spite of this, the actual time taken (HPC resource consumed) in the calculation of numerical evolutions, particularly for extreme scenarios that require highly resolved physical scales, can require weeks, or months [36]. Clearly, even minor performance optimizations, of a few percent, can have a significant impact on such HPC wait times, together with broader sustainability considerations [37].

In this work, we focus on finite-difference based evolution of the Z4c formulation [31,32] of NR, as coupled to the general relativistic, hydrodynamical (GRHD) formulation¹ of [38]. This is precisely the setting treated by our code GR-Athena++ [39,40] (see also [41]). In brief, this is our extension to the astrophysical (radiation), GRMHD code Athena++ [42–44], we have geared toward evolution with dynamical spacetime. Building upon the underlying octree, block-based adaptive-mesh-refinement (AMR) infrastructure, with the

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

¹For simplicity we do not investigate magnetic fields in this work.

addition of our Z4c-based spacetime solver both in vacuum [39] and for GRMHD with dynamical spacetime [40], has led to demonstrable scaling efficiencies in excess of 80% on (pre)exascale HPC infrastructure.

Specifically, here, we aim to address a seemingly innocuous question: should we select the geometric fields, that is, the dynamical fields of Z4c, to be discretized on vertex centers (VC) or cell centers? Our initial choice was based on considerations for binary black hole evolution, where transfer of data between differing levels of refinement is less computationally expensive for VC in contrast to cell centers. On the other hand, in consideration of GR (M)HD, typically hydrodynamical sampling is on cell centers (see, e.g., BAM [45], GRHydro [46], MHDuet [47], and WhiskyTHC [48]). This choice is in part due to greater amenability of the underlying hydrodynamical treatment as based on conservative, high-resolution-shock-capture (HRSC) methods (see, e.g., [49]). Athena++ also handles the hydrodynamical description on cell centers. Having previously fixed geometric fields to VC, discretized fields must consequently be transferred between grids, which incurs an overhead. *A priori* it is not necessarily clear as to how large this is. Our aim here is to investigate switching to a cell-centered treatment of the spacetime solver in GR-Athena++, thus providing an answer. Such a code-internal investigation does not appear to have been previously performed.² This gap we fill here. This paper proceeds as follows. In Sec. II we briefly explain the computational domain structure and the variable discretization strategy. Subsequently, we present a simplified refinement strategy, tailored to the binary merger problem, and independent of geometric field sampling selection. In Sec. III we demonstrate the robustness of the technique through numerically solving binary black hole (BBH) and binary neutron star (BNS) inspiral-merger problems, representative of typical production simulations. Section IV concludes.

II. METHOD

GR-Athena++ [39,40] builds upon Athena++ [42–44] and thus inherits overall features of the latter framework. In brief, problems are formulated over a target computational domain Ω (the Mesh), which is partitioned into a collection of subdomains Ω_i (MeshBlock objects). We have $\Omega = \sqcup_{i \in Z} \Omega_i$ where Z may be thought of as an ordered parametrization of multidimensional coordinates to a space-filling curve [26,27]. For a d -dimensional problem the overall extent of Ω together with suitable boundary data on $\partial\Omega$ must be provided. Following this, the underlying discretization is fixed through specification of Mesh and MeshBlock sampling

²Note that [50] described development of a flux conservative AMR treatment involving cell-centered sampling for the hydrodynamical sector, illustrating differences of a VC hydro AMR approach. We specifically focus on the effect of changes in the treatment of the spacetime sampling.

$N_M = (N_{M_1}, \dots, N_{M_d})$ and $N_B = (N_{B_1}, \dots, N_{B_d})$, respectively, where N_B must divide N_M componentwise.³ In the case that (local) refinement is desired, MeshBlock partitioning into 2^d new MeshBlock objects is recursively performed. This involves halving size along each axis at a fixed number of samples, which doubles the resolution at each step. The aforementioned is done while maintaining an (at most) 2:1 refinement ratio between neighboring $\{\Omega_i\}_{i \in Z}$. While the logical-rectilinear structure here described may then be further mapped to suitable curvilinear coordinates, for simplicity we focus on the Cartesian context.

A. Cell-centered extension

In GR-Athena++ [39,40] we extended sampling support to enable VC discretization of field variables together with infrastructure enabling high-order approximations to be carried out. In this work we focus on analogous extensions in description of geometric quantities on cell-centers (CX). A variable $\mathcal{V}(x)$ where $x \in [a, b]$ ($a < b$) is CX discretized⁴ to $\mathcal{V}_{i+1/2}$ over $x_{i+1/2} := a + (i + 1/2)\delta x$ where $\delta x := (b - a)/N_B$ and $i = 0, \dots, N_B - 1$. This yields N_B total samples. In practice the interval is extended by an additional $\mathcal{N}_g$ “ghost” nodes on either side.⁵ This extension facilitates evaluation of, e.g., derivative approximants, transfer of field data between neighboring MeshBlock objects that may differ in the level of refinement, and/or imposition of boundary conditions.

In order to populate ghost nodes, communication is required. In the case that $\mathcal{V}$ is discretized over a Mesh decomposed at a fixed, common level, under CX sampling, this entails that nearest-neighbor MeshBlock objects can be utilized to populate ghost layers through a straightforward copy. Suppose instead that subdomains differ in level l . In this situation, analogous to the VC case, each subdomain ${}^l\Omega_i$ (where prefix l indicates level) as described by a MeshBlock, features a fundamental and an additional coarse representation of the sampled variable. This latter is one level above at $l - 1$. Thus if N_B is the number of physical points on a MeshBlock extended by $2\mathcal{N}_g$ ghost nodes, then the complementary coarse representation features $N_B/2$ physical points extended by $2\mathcal{N}_{cg}$ coarse ghosts. The situation is illustrated (locally) in Fig. 1.

The general communication strategy and ordering of operations closely follows and extends that of Athena++. This minimizes the number of infrastructure changes required. Synchronization of data across levels may be understood

³For a uniform number of samples in each dimension we typically use a single scalar to represent all component values of the tuples.

⁴The grid sampling will also be referred to as CC. On CX fields are to be thought of as pointwise sampled, whereas on CC they carry a cell-averaged interpretation.

⁵In the case that $d = \dim(\Omega) > 1$ we utilize a tensor product grid of such extended one-dimensional discretizations.




FIG. 1. One-dimensional illustration of inter-MeshBlock communication given a refined local Mesh structure. Each MeshBlock features a fundamental representation of a sampled field depicted through square nodes, which is extended in each direction through ghost nodes (circles) that facilitate communication. Additionally a complementary coarse sampling is present, which is explicitly depicted for data on level l and suppressed on level $l-1$. Here $N_B = 12$, $\mathcal{N}_g = 2$, and $\mathcal{N}_{cg} = 3$. Prior to all communication data in filled square nodes is available; the task is thus to fill data at locations indicated by shaded circles and diamonds. Nodes are assigned MeshBlock-local labels for ease of discussion.

and implemented by considering the transfer operations: fine to coarse (restriction $\mathcal{R}$), and coarse to fine (prolongation $\mathcal{P}$). As an example consider synchronization of ghost-layer data between $^{l-1}\Omega_A$ and $^l\Omega_B$ as depicted in Fig. 1. *Initially* data on the fundamental grid of $^l\Omega_B$ over the nodes $(IS, IS+1, \dots, IE)$ is known. This is first restricted to a subset of the complementary coarse grid of $^l\Omega_B$ over the nodes $(CIME+1, \dots, CIME+\mathcal{N}_g)$; see green diamonds in the figure. Subsequently data at the nodes $(IE+1, \dots, IPE)$ of $^{l-1}\Omega_A$ may be filled through communication or a copy (green circles in figure). Conversely using data from $^{l-1}\Omega_A$ at the nodes $(IE-\mathcal{N}_{cg}+1, \dots, IE)$ allows for filling of $(CIMS, \dots, CIME)$ on $^l\Omega_B$ (purple circles) through communication or a copy. Finally, ghost layer nodes $(IMS, \dots, IS-1)$ of the fundamental grid over $^l\Omega_B$ may be filled (blue circles) through $\mathcal{P}$ based on data local to the MeshBlock. The procedure is analogous for the pair $^l\Omega_C$ and $^{l-1}\Omega_D$.

Comparing node positions in Fig. 1 across differing levels of refinement clearly shows that nodes do not align but are fully staggered.⁶ This is one aspect where CX differs from VC sampling [39]. The coarsening and refining of sampled field data via $\mathcal{R}$ and $\mathcal{P}$, respectively, is based on

⁶One could instead consider 3:1 refinement ratios to obviate this.

polynomial interpolation under CX and VC. Due to the interlevel node staggering occurring in CX, a nontrivial evaluation of interpolants for every target node is required. This is in contrast to the case of VC where a subset of operations reduces to injection (copy) [39]. The distinction heralds a performance implication as will be demonstrated in Secs. III A and III B.

We close this Subsection with a collection of technical points concerning implementation aspects.

- (i) The formal order of accuracy for the $\mathcal{R}$ and $\mathcal{P}$ operations on CX is controlled by the number of nodes selected. We fix this choice based on $\mathcal{N}_g$ selection to match the approximation order with that of the finite differencing.
- (ii) Given a target node, source data entering an interpolant describing CX $\mathcal{P}$ has a directional bias due to grid staggering. This can be understood by noting that the source data nodes are not symmetrically placed about the location of the target node, but offset. Consider at a desired formal order of accuracy the initial coarse, complementary restriction needed when communicating data to a finer neighboring MeshBlock. There may be insufficient nodes available to evaluate with symmetrically spaced source data. In this situation, we again bias the stencil selection.
- (iii) As is well known, floating point (FP) arithmetic is not necessarily associative: $(a+b)+c \neq a+(b+c)$ for FP numbers $\{a, b, c\}$. If care is not taken in how numerical expressions are ordered, errors can, e.g., lead to catastrophic cancellations, or result in spurious symmetry breaking during simulations [51]. A combination of Kahan-Babuska-Neumaier compensated summation [52], and explicit symmetrization of arithmetic operations involving averaging can partially mitigate this [51]. This can be optionally activated for interpolant evaluation. We have not, however, found it necessary for the simulations presented here.
- (iv) There are now two styles of cell-centered sampling where discretized field variables may be registered and stored: that inherited from Athena++, which we denote CC, and the newly introduced CX. The principal distinction is that for a Mesh featuring MeshBlock objects on distinct levels, the $\mathcal{R}$ and $\mathcal{P}$ operators are distinct: for CC there is minmod based slope limiting [53], whereas for CX we utilize the Barycentric Lagrange approach [54]. Special care has been taken to ensure that distinct fields may be suitably registered as either CC or CX with their respective logic available concurrently.
- (v) The VC double-restriction operation described in [39] we have found to not be strictly required. Its removal means that when utilizing vertex-centered sampling we now have $N_B \geq \max(4, \mathcal{N}_g)$. This allows for smaller MeshBlock size, which offers greater flexibility in tailoring resolution adapted regions.

B. Refinement strategy

In order to resolve dynamical features over widely disparate spatial scales we utilize adaptive mesh refinement AMR. Our purpose here is to fix notation, and to point to two simple, minimal criteria that will facilitate consistency tests of the new infrastructure. The aforementioned feature evolution of initial data describing a BBH (Sec. III A), and a binary neutron star (Sec. III B). For a comprehensive investigation of AMR criteria in GR-Athena++, and effect on waveform quality for BBH evolution, see [55].

The AMR conditions of GR-Athena++ are quite flexible. A user-specified function is defined for each `MeshBlock`; this function returns a flag that controls whether to (de)refine. Consider puncture-based evolution [45] of a BBH system. Here a black hole center is described through a puncture. The position of a puncture $\mathbf{x}_p(t)$ can be computed by augmenting the dynamical system evolved through an additional ODE based on gauge information [56]. In short, this is one variety of a so-called tracker. This allows us to construct a dynamically adjusted, octree-based, pseudo-box-in-box (pBIB) refined `Mesh` as detailed in [39] (or L_∞ -norm refinement in the language of [55]).

Let us suppose the `Mesh` is the Cartesian coordinatized, computational domain $\Omega := [-D, D]^3$. A parameter controlling the maximal level of refinement N_L is specified. A hierarchy of `MeshBlock` objects is then constructed, which increases in resolution as a puncture is approached (see, e.g., Fig. 6). On the finest resolution level the grid spacing will be $\delta x = 2D/(N_M 2^{N_L-1})$ [39]. As an alternative to pBIB, we can exploit information from the trackers and directly impose a desired region to be at fixed resolution (equivalently level). Denote a spherical region of radius R , centered at p as $\mathbb{S}_R(p)$. We center such a region on a puncture and impose a desired local level of refinement l . This we write as $\mathbb{S}_{l,R}(\mathbf{x}_p(t))$.

In a similar vein, consider a dynamical, user selected, control field χ featuring a local extrema χ^* . Another type of tracker is possible to construct based on following the time development of $\chi^*(t)$ through $\mathbb{S}_{l,R}(\chi^*)$. This allows for the region surrounding the extrema to be captured as its position changes in time. If multiple such regions are registered, occur at differing l , and happen to intersect, then we take the finest level as the target level. This is the principal condition that we employ for AMR during BNS evolution.

III. VALIDATION

Our first goal in this section is validating the extended cell-centered implementation in the context of vacuum binary black hole evolution. The space-time solver implemented in GR-Athena++ is based on the Z4c formulation [31,32]. Here the spacetime is constructed based on the evolution of the dynamical, geometric variables $\mathcal{Z} := (\chi, \tilde{\gamma}_{ij}, \hat{K}, \tilde{A}_{ij}, \Theta, \tilde{\Gamma}^i)$. The definition and meaning of

these variables is standard and may be found, together with implementation details, in [39]. Previously, motivated by AMR efficiency considerations in the vacuum context, we have chosen these to evolve over VC grid sampling. We compactly denote this through $\mathcal{Z}_{VC}$. This work details a relaxation of this restriction, extending functionality to allow for CX sampling, i.e., sampling and evolving directly as $\mathcal{Z}_{CX}$. We detail tests of this in Sec. III A.

Secondly, we investigate the BNS inspiral-merger problem. Our recently implemented GRMHD treatment in GR-Athena++ is based on the formulation of [38]. For full details and conventions, see [40]. Here we recall key ingredients of constructing the matter side of the evolution as based on selecting $\mathcal{Z}_{VC}$. For simplicity, we restrict our exposition to the GRHD sector. For the hydrodynamics, the (conserved) field variables are $\mathcal{Q}_A := (D, S_j, \tau)$, and satisfy a balance law. Recall that this is evolved based on conservative, finite-volume, HRSC techniques [49]. The natural choice of sampling that also leverages extant Athena++ infrastructure is cell centered, i.e., we take $\mathcal{Q}_{A,CC}$. We also work with a set of complementary, primitive variables $\mathcal{R}_A := (\rho, \tilde{u}^i, p)$. The relation $\mathcal{Q}_A(\mathcal{R})$ is nonlinearly implicit and is evaluated through numerical inversion [57]. In this work the conservative to primitive variable mapping is handled through the use of REPRIMAND [58]. This inversion requires geometric fields to be available on the same grid sampling as the matter. Hence an intergrid interpolation $\mathcal{Z}_{VC} \rightarrow \mathcal{Z}_{CC}$ is required.

In order to evolve the quantities $\mathcal{Q}_{A,CC}$ via HRSC we require assembly of fluxes at cell interfaces [that is, face-centers (FC)]. In this work, this is achieved based on the reconstruction of $\mathcal{R}_{A,CC}$ to FC with the fifth order method of [59] denoted WENO5Z. Flux assembly also requires a subset of $\mathcal{Z}_{FC}$ and hence another interpolation. The interface fluxes are combined according to the Local-Lax-Friedrichs scheme (see, e.g., [60]). Finally the sources appearing in the hydrodynamical balance law are assembled, which again requires $\mathcal{Z}_{CC}$. On the other hand, the gravitational matter density, momenta, and stress quantities that recouple the system to the spacetime solver, must be provided on VC. This entails the interpolation of hydrodynamical quantities to VC. It is clear from the above summary that by instead taking the evolved geometric variables as $\mathcal{Z}_{CX}$ from the start, fewer transfers of geometric data between grids would be required. We further validate our code functionality, for BNS, in Sec. III B.

In addition to ensuring consistency of the implementation, and that problems are solved correctly, a practical performance assessment is needed. We demonstrate that execution speed, and suitable choice of $\mathcal{Z}_{VC}$ ($\mathcal{Z}_{CX}$) is contingent on the class of physical problem simulated.

A. Binary black holes

We assess the performance of our code by focusing on the well-known calibration BBH problem of [45]. The Z4c

TABLE I. Properties of grid configurations utilized for vacuum tests. The total number of MeshBlock objects (#MB) at the initial time is indicated. In all cases the Mesh is the Cartesian coordinatized domain $\Omega := [-D, D]^3$, where $D = 1536M$ is selected. This entails that grid spacing is uniform along each spatial axis. In each case we select the maximal refinement through $N_L = 11$. The MeshBlock sampling is selected as $N_B = 16$. Resultant puncture resolutions δx_p are also provided.

N_M	$\delta x_p \times 10^{-2} [M]$	#MB
96	3.125	5144
128	2.344	6336
160	1.875	18640
192	1.563	22392

system is coupled to the moving puncture gauge using the $1 + \log$ -slicing condition [61] and the gamma driver shift [62,63], with parameters as described in [39]. Initial data is generated through the use of the `TWOPUNCTURES`⁷ library based on [64]. We model two initial nonspinning black holes of bare mass $m_{\pm} = 0.483M$, located on the x axis, with $x_{p,\pm}^1(t=0) = \pm 3.257M$, and initial momenta directed along the y axis, $p_{\pm}^2(t=0) = \mp 0.133M$. These parameter choices lead to an evolution featuring ~ 3 orbits to merger, which occurs at an evolution time $t \sim 170M$.

In these tests the Courant-Friedrich-Lewy (CFL) is fixed at $C_1 = 0.2$. This choice determines the global time step taken on each MeshBlock based on the finest grid-spacing (occurring in the vicinity of a puncture and denoted δx_p). Kreiss-Oliger dissipation is selected as $\sigma_D = 0.02$ throughout. The Z4c constraint damping parameters are taken as $(\kappa_1, \kappa_2) = (0.02, 0)$ (see [39]). Unless otherwise stated, in this Section we take $\mathcal{N}_g = \mathcal{N}_{cg}$, that is, the number of ghosts on the fundamental grid is set equal to the number on the complementary coarse representation (see Fig. 1). We select $\mathcal{N}_g = 4$, which induces a formal sixth order of spatial approximation. Time evolution is performed using the explicit fourth order RK4(4)2S [65].

The overall grid configurations based on pBIB refinement (see Sec. II B) are summarized in Table I. In principle the calibration BBH problem possesses a underlying bitant symmetry (i.e., reflection across the $z = 0$ plane to which the co-orbit and merger are confined). This allows for an approximate halving of the total computational resources required for a given problem. When imposing bitant symmetry, #MB as listed in Table I is roughly halved, and in order to keep the grid spacing fixed, we halve the Mesh sampling parameter along the z axis to $N_M/2$. As an initial consistency check, we select the geometric sampling as $\mathcal{Z}_{VC}$ and the $N_M = 96$ grid of Table I. We perform a comparable run with $\mathcal{Z}_{CX}$, for which we inspect the




FIG. 2. Tracker evolution during calibration BBH inspiral and merger. The one-parameter families $\mathbf{x}_{\pm}(t)$ describe where punctures are located. Consistent behavior for CX and VC choices of sampling is observed. Parameters: $N_M = 96$, and $N_B = 16$.

puncture tracker evolution in Fig. 2. Excellent agreement can be seen.

The BBH evolution considered here should result in an evolution with a strict symmetry about the $z = 0$ plane. During numerical evolution, however, this condition is not necessarily preserved. The z component of a puncture can potentially slowly drift out of the plane. The behavior is demonstrated in Fig. 3 for full 3D runs. In the case of $\mathcal{Z}_{CX}$, based on the structure of the level-to-level transfer operators, we do not observe, or expect, strict FP symmetry




FIG. 3. Evolution of the z component of the tracker $\mathbf{x}_-(t)$ during calibration BBH inspiral and merger. The CX and VC sampling choices are compared at fixed resolution. Parenthesis $(\mathcal{N}_g, \mathcal{N}_{cg})$ indicate ghost node selection. The “bit.ID” set evolved with the choice $\mathcal{Z}_{CX}$ has ID explicitly symmetrized. All runs of this figure are performed over full 3D grids, regardless of ID symmetrization. Parameters: $N_M = 96$, and $N_B = 16$. See text for further discussion.

⁷See <https://bitbucket.org/bernuzzi/twopunctures/>.

preservation. Instead, during the full 3D run we find a very minor drift developing, where the z component of the puncture reaches $\sim \mathcal{O}(10^{-10})$. We also verify in Fig. 3 that $\mathcal{N}_{cg}$ may indeed be independently selected, so as to increase the formal order of accuracy in the prolongation operation for $\mathcal{P}$. We have similarly verified this for $\mathcal{Z}_{CX}$. As this does not appear to substantively change simulation quality for the short evolution here, we proceed with runs that fix $\mathcal{N}_g = \mathcal{N}_{cg} = 4$.

One reason for this behavior is that the TwoPunctures library does not allow for explicit enforcement of an underlying symmetry about $z = 0$. Thus, when interpolating an initial data (ID) set to a Cartesian grid, the result violates the bitant property at the FP level. We have taken special care to allow for symmetrization of the prepared ID by rewriting the interpolation process. The purpose here is to provide a test on preservation of FP symmetry across various elements of the infrastructure, such as, e.g., operator approximants (finite-differencing kernels, interpolation, etc.).

In the case of $\mathcal{Z}_{VC}$, explicit bitant symmetrization of the ID does lead to an evolution with *exact* confinement of the punctures to the $z = 0$ plane. This is in contrast to runs in Fig. 3.

Having passed a tracker-based consistency check, we now check gravitational waveforms (GW). To do this, we construct the Weyl scalar Ψ_4 [45], and then interpolate onto geodesic spheres [66]. Subsequently, we project through numerical quadrature evaluation, onto spin-weighted spherical harmonics [67] ${}_s Y_{lm}$ of spin-weight $s = -2$. Our convention differs from the reference by a Condon-Shortley phase of $(-1)^m$. The above allows us to extract ψ_{lm} . The (ℓ, m) modes of the gravitational wave strain h satisfy $\psi_{\ell m} = \dot{h}_{\ell m}$. The strain is then given by the mode sum,

$$R(h_+ - ih_\times) = \sum_{\ell=2}^{\infty} \sum_{m=-\ell}^{\ell} h_{\ell m}(t) {}_{-2}Y_{\ell m}(\vartheta, \varphi). \quad (1)$$

Following the convention of the LIGO algorithms library [68] we set,

$$R h_{\ell m} = A_{\ell m} \exp(-i\phi_{\ell m}), \quad (2)$$

and the gravitational-wave instantaneous frequency is defined through,

$$\omega_{\ell m} = \frac{d}{dt} \phi_{\ell m}. \quad (3)$$

Note that the peak amplitude of the (2,2) mode of h is conventionally taken as the time of merger. We denote this u_{merger} .

The quantity ψ_{22} is illustrated for the calibration BBH in Fig. 4. Rather than just imposing ID bitant symmetrization,




FIG. 4. Consistent behavior between $\mathcal{Z}_{VC}$ and $\mathcal{Z}_{CX}$ for calibration BBH waveforms with and without bitant symmetry reduction. Black profiles depict the amplitude of the given mode based on calculations involving bitant symmetry reduction. Parameters: $N_M = 96$, and $N_B = 16$; extraction radius $R = 85M$.

symmetry reduction on the 3D grid is also imposed for a subset of runs. The waveforms continue to overlap consistently. This provides another important consistency test on implementation details.

Beyond verification of qualitative behavior, we next focus on the gravitational strain, impose bitant symmetry reduction, and assess convergence properties based on the Appendix. The result, for constructing h_{lm} , based on time-domain integration (see., e.g. [69]), is shown in Fig. 5. A sixth order compatible convergence trend is observed for the phase ϕ_{22} . This is present for both $\mathcal{Z}_{VC}$, and $\mathcal{Z}_{CX}$. We have also verified that behavior persists, consistently between grid sampling choices, for, e.g., the (4,4) mode.




FIG. 5. (2,2) mode of the gravitational strain (top panel) together with convergence properties in the associated phase for CX and VC sampling in the middle and lower panels, respectively. All quantities are taken as functions of the Schwarzschild tortoise coordinate u . We assess sixth order compatible convergence. Indicated in a common vertical black line across panels is $u_{\text{merger}} \simeq 158M$. We select calculations involving bitant symmetry reduction based on the coarse (c), medium (m), and fine (f) resolution triplet induced through selecting $N_M \in \{128, 160, 192\}$. Extraction performed at $R = 85M$.




FIG. 6. Mesh structure and values of $\tilde{\gamma}_{xx}$ under $\mathcal{Z}_{CX}$ at $t \simeq 110M$ in the orbital plane. Each quadrant depicts distinct resolutions (AMR choices) for the same physical calibration problem configuration. Counterclockwise, starting from the upper right quadrant we have pBIB AMR with $N_M \in \{96, 128, 192\}$. The lower right quadrant depicts evolution with the dynamical spherical AMR criterion, and $N_M = 96$ fixed. Each grey box (MeshBlock) is sampled with $N_B = 16$ points along each axis. For the cases involving the pBIB AMR criterion observe that reflection of outward propagating dynamical spherical features in $\tilde{\gamma}_{xx}$ is present in the vicinity of a change in refinement level. See text for further discussion.

As described in Sec. II B AMR conditions may be freely selected. In the approach to a puncture location, the conformal factor χ approaches a local minimum (see, e.g., [70]). Indeed, we have numerically verified that the trackers shown in Fig. 2 can also be constructed based on such local minimum. On account of the underlying dynamical (approximate) symmetries of the problem, together with their impact on conserved quantities [71], and GW extraction taking place on spheres, an approximately matched AMR condition is motivated. To this end, spherically matched Mesh refinement is induced by fixing levels based on the collection $\{\mathbb{S}_{11.3}(\chi_{\pm}^*), \mathbb{S}_{6,100}(\mathbf{0})\}$. That is, two spherical regions of radius $3M$, with $N_L = 11$ follow both punctures. A central fixed sphere of radius $100M$, and $N_L = 6$ is set to improve resolution in the wave-zone. This is compared against pBIB in Figs. 6 and 7.

Let us consider the behavior of $\tilde{\gamma}_{xx}$ in the vicinity of a change in refinement level as shown in Fig. 6. It is observed that small, spurious reflection occurs, from the underlying rectilinear grid structure. In contrast, by comparing the two $N_M = 96$ runs, we see that approximate spherically matched AMR partially mitigates these spurious reflections. The magnitude of the effect can also be seen to converge away when comparing $N_M = 192$. The change in




FIG. 7. Zoom of the strong-field region of Fig. 6. Trackers are overlaid where lighter color indicates increasing t . Observe that the finest resolution in both $N_M = 96$ runs is the same, however the region on the finest level is larger in the bottom right (dynamical spherical region) run. Note: color scalings are distinct between this figure and Fig. 6.

refinement strategy also has an effect on GW phasing, which we show in Fig. 8. At early times the amplitude and phasing error of the spherically matched AMR with $N_M = 96$ is comparable to that of a pBIB run at $N_M = 160$. This can be attributed to a combination of increased wave-zone resolution and better symmetry matching of the AMR




FIG. 8. Comparison of differences in the amplitude A of $(2,2)$ mode of the gravitational strain (top panel) and phase ϕ (bottom panel), for runs at different resolutions and refinement strategy, under $\mathcal{Z}_{CX}$. We consider pBIB AMR with (c, m, f) corresponding to $N_M \in \{96, 160, 192\}$. A single run with spherically matched AMR at $N_M = 96$ is denoted with $c(\dagger)$. Vertical black line across panels indicates u_{merger} . Extraction performed at $R = 100M$.

TABLE II. One hour timing test on SuperMUC-NG of 1r2z utilizing 16 nodes. The timing ratio r is defined as $r_t := (t_F(\text{CX}) - t_0)/(t_F(\text{VC}) - t_0)$. The number of MeshBlock objects is provided at t_0 . Bitant symmetry reduction has been performed. The AMR is based on the pBIB criterion.

N_M	$t_F(\text{CX})$	$t_F(\text{VC})$	r_t (2dp)	#MB	#MB/Core
96	26.48	35.17	0.71	2404	~3
128	18.01	21.98	0.77	3280	~4
160	10.41	10.80	0.93	8774	~11
192	8.49	9.03	0.87	11490	~15

(see also the L_2 versus L_∞ AMR discussion of [55]).

Another potential advantage of running the spherically matched AMR is a reduction in execution time. Running the spherically matched AMR at $N_M = 96$ in bitant results in #MB = 4532 initially, where δx_p is as in Table I. Compared to pBIB AMR under bitant symmetry reduction at $N_M = 160$, it results in a computational resource reduction of $\sim 3.4\times$ when coarser time step δt at fixed CFL is taken into account.

Finally, we provide a rough characterization of performance differences between $\mathcal{Z}_{\text{CX}}$ and $\mathcal{Z}_{\text{VC}}$ through a short timing test. To this aim we evolve the problem for the calibration BBH on the pBIB grid configurations of Table I. The initial checkpoint is taken at a time of $t_0 = 5M$, where the AMR condition has refined the Mesh. Timings are as summarized in Table II. We see that for pure vacuum, the spacetime solver has a distinct advantage when run with the choice $\mathcal{Z}_{\text{VC}}$ in terms of speed. Indeed, we observe an approximate average $\sim 20\%$ later final evolution time, across N_M , when compared to the $\mathcal{Z}_{\text{CX}}$ choice. This we attribute principally to a difference in complexity of the $\mathcal{R}/\mathcal{P}$ operations. A cautionary remark is, however, in order: the MeshBlock/core saturation has not been taken into account, which requires a certain threshold for efficiency [39,40]. With regard to performance, these timing results nonetheless strongly suggest that $\mathcal{Z}_{\text{VC}}$ is the better option for BBH.

B. Binary neutron stars

In the case of BNS evolution, in addition to the spacetime evolution, we need to address the evolution of the hydrodynamical variables. These variables satisfy a balance law of the form $\partial_t[\mathcal{Q}_A] + \partial_i \mathcal{F}_A^i[\mathcal{Q}] = \mathcal{S}_A[\mathcal{Q}]$ [38]. The source term $\mathcal{S}_A[\mathcal{Q}]$ does not involve derivatives of $\mathcal{Q}_A$, but does contain geometric fields and their derivatives. Here the finite-volume, HRSC methods we employ lead us to take the evolved $\mathcal{Q}_A$ on cell centers. In the $\partial_i \mathcal{F}_A^i[\mathcal{Q}]$ approximation step, primitive variables $\mathcal{R}_A[\mathcal{Q}]$ are assembled and reconstructed to FC [60]. This process requires geometric data on FC. Information about the evolved fluid variable state must be available on the grid used for the spacetime solver. If we select $\mathcal{Z}_{\text{VC}}$ the required

intergrid transfer operations can be summarized as: $\mathcal{Z}_{\text{VC} \rightarrow \text{CC}}$, $\mathcal{Z}_{\text{VC} \rightarrow \text{FC}}$, and finally $\mathcal{R}_{A, \text{CC} \rightarrow \text{VC}}$. On the other hand, selecting $\mathcal{Z}_{\text{CX}}$ means we only need $\mathcal{Z}_{\text{CX} \rightarrow \text{FC}}$.

The intergrid transfer operations summarized above are handled through the use of symmetric Lagrange interpolation. Specifically, during evaluation, N_I nearest neighbor nodes on either side of a target interpolation point are used along a salient axis. We have found it beneficial, for stability reasons, to fix $N_I = 1$ in a subset of operations. This we refer to as a hybrid mode. In this mode N_I is only allowed to vary freely for $\mathcal{Z}$ derivative assembly entering $\mathcal{S}_A[\mathcal{Q}]$, and primitives $\mathcal{R}_{A, \text{CC} \rightarrow \text{VC}}$ that recouple to the spacetime solver.

As initial consistency checks for geometry-matter field coupling and the above operations we consider static star tests in the Appendix. A more representative setting of BNS evolution problems is required. For this, we consider a quasicircular, equal-mass BNS inspiral-merger problem. Here irrotational, constraint-satisfying initial data is provided by Lorene [72]. Specifically, as a validation problem, we consider the G2_I14vs14_D4R33_45 km binary,⁸ which features baryon mass $M_b = 3.2500M_\odot$ and gravitational mass $M = 3.0297M_\odot$ at an initial separation of 45 km. The Arnowitt-Deser-Misner (ADM) mass of the binary is $M_{\text{ADM}} = 2.9984M_\odot$, the angular momentum $J_{\text{ADM}} = 8.8354M_\odot^2 \text{ G/c}$, and the initial orbital frequency is $f_0 \simeq 294 \text{ Hz}$. For these tests we utilize an ideal gas equation of state (see the Appendix) with $K = 123.64$. The evolution of this BNS system proceeds for three orbits prior to the formation of a massive remnant, thereafter leading to gravitational collapse.

The computational domain is again selected as $\Omega := [-D, D]^3$, where $D = 1536M_\odot \simeq 2268 \text{ km}$. The AMR is induced through tracking the local minimum of the lapse α , as positioned on each constituent of the BNS system. Resolution levels are fixed according to $\{\mathcal{S}_{8,15}(\alpha_\pm^*), \mathcal{S}_{7,30}(\alpha_\pm^*), \mathcal{S}_{6,110}(\mathbf{0})\}$. To perform a resolution study, we take a sequence of $N_M \in \{64, 96, 128\}$. This translates into a sequence of finest resolution levels of (554, 369, 277) m. The MeshBlock sampling is selected as $N_B = 16$. For the time evolution we make use of the SSPRK(3,3) method of [73], with a CFL of $\mathcal{C}_1 = 0.2$, and $\sigma_D = 0.5$. A tenuous, artificial atmosphere is added, where $\rho_{\text{atm}} = 10^{-16}$. We do not make use of thresholding [40]. For the RePrimAnd-based primitive recovery we set the tolerance to 10^{-10} . In this section no symmetry reductions are performed.

We select $\mathcal{Z}_{\text{CX}}$ and evolve the initial data with the aforementioned parameters. An additional run featuring $\mathcal{Z}_{\text{VC}}$ with $N_M = 64$ and all else fixed is also performed for internal comparison. In each case we find approximately three orbits prior to the formation of a massive remnant. A snapshot of the fluid rest mass-density at this point

⁸Dataset available at <https://lorene.obspm.fr/>.




FIG. 9. Mesh structure and values of fluid rest mass-density ρ [g cm^{-3}] shortly post-merger in the orbital plane for the BNS inspiral problem described in the text. Counter-clockwise, starting from the upper right quadrant we have evolution with CX sampling with $N_M \in \{64, 96, 128\}$. The lower right quadrant depicts evolution under VC sampling for geometric fields with $N_M = 64$. Each grey box (MeshBlock) is sampled with $N_B = 16$. The AMR criterion for each run has driven the distribution of refined region resolution towards an origin-centered approximate sphere. Trackers that follow local minima of $\alpha(t)$ are overlaid. Darker tracker color indicates increasing t .

($t \simeq 8.8$ ms) is shown in Fig. 9 comparing the four runs. A zoomed version is depicted in Fig. 10. We observe that the spiral arm structure is qualitatively similar for the two runs at $N_M = 64$, which provides an initial consistency check. The structure becomes sharper and more compact at higher N_M . The extrema tracking based AMR can also be seen to guide increased resolution toward regions that suitably follow the binary constituents.

Quantitative scalar monitors are shown in Fig. 11. We find that the time of collapse falls in the range $t \in [16.5, 20.2]$ ms for $\mathcal{Z}_{\text{CX}}$ selection. This depends on the selection of N_M . Importantly, we observe consistency to within 1% when comparing a verification run with $N_M = 64$ based on $\mathcal{Z}_{\text{VC}}$. Violation in baryon mass conservation grows to a plateau of $\sim \mathcal{O}(10^{-6})$ at merger for each run. This remains approximately constant until the point of gravitational collapse. Again we see that both choices of grid sampling for the geometric fields lead to consistent results.

Gravitational waveforms are the next item to verify. To construct the (2,2) mode of the strain, for the BNS system, we use fixed frequency integration (FFI) [74]. The FFI proceeds by transforming the signal to the frequency representation, applying a low-frequency cutoff f_0 , and




FIG. 10. Zoom of the strong-field region of Fig. 9. Observe that the finest resolution in both $N_M = 64$ runs is the same.

integrating. We select $f_0 = 0.002M^{-1}$. The result is shown in Fig. 12. Based on the $N_M = 128$ run, we find that the time of merger (retarded coordinate) is $u_{\text{merger}} \simeq 1583.2M$. This translates into $t_{\text{merger}} \simeq 1699.8M \simeq 7.9$ ms. Indeed, all




FIG. 11. Relative change of (global) maximum density ρ_c (upper panel), baryon mass conservation M_b (middle), and minimum of the lapse α (lower panel), for the BNS model problem. A resolution triplet induced by the choices $N_M \in \{64, 96, 128\}$ and selection of $\mathcal{Z}_{\text{CX}}$ is depicted, with a single run at $N_M = 64$ and $\mathcal{Z}_{\text{VC}}$ shown for reference. Collapse times, for each run, are signalled by: post-merger, global peak in the ρ_c monitor, blow up of the mass conservation monitor, and confirmed by reduction of lapse below ~ 0.2 . Vertical black line depicts merger time ($t \simeq 8.4$ ms), as based on the peak amplitude h_{22} associated with the $N_M = 128$ run (cf. Fig. 12). See text for further discussion. Note: legend common to all panels.




FIG. 12. Gravitational strain (2,2) mode extracted at $R = 100M$ for binary neutron star run of Sec. III B. (Top): real part of the strain with $\mathcal{Z}_{CX}$ for $N_M \in \{64, 96, 128\}$ is depicted. For $N_M = 64$ the amplitude is also shown. In the case of $N_M = 128$ we find $u_{\text{merger}} \simeq 1583.2M$ ($t_{\text{merger}} \simeq 1699.8M$). We observe consistent behavior at $N_M = 64$ for a comparable run based on VC sampling. (Middle): Instantaneous frequency is shown. (Top, middle): Vertical dashed lines indicate time of merger for the different data sets. (Bottom): Phase differences for the CX resolution triplet suggest convergence (to merger) as compatible with a second order trend.

CX runs fall within a band of ~ 0.2 ms. This indicates strong robustness of the simulations performed. The merger time for the $\mathcal{Z}_{VC}$ waveform is within 3% that of the CX at $N_M = 96$. Similarly, the morphology of the instantaneous frequency for CX when cross-compared against VC is seen to be alike. With the method described in the Appendix we also assess convergence. As can be seen in Fig. 12 we have a strong indication for a second order trend in the convergence of the h_{22} associated phase, as N_M is increased.

Overall, based on this discussion, when running with $\mathcal{Z}_{CX}$, good agreement is found with $\mathcal{Z}_{VC}$ here, and our prior results in [40]. This provides confidence in the correctness of our implementation, and the validity of the strategy outlined in Sec. II B.

To close this Section, as in the BBH case, we provide a rough characterization of performance differences between $\mathcal{Z}_{CX}$ and $\mathcal{Z}_{VC}$ through a short timing test. The initial checkpoint is taken at a time of $t_0 = 5M$ with timings summarized in Table III.

We see that for BNS, the space-time solver has a distinct advantage when run with the choice $\mathcal{Z}_{CX}$ in terms of speed. Indeed, we observe an approximate average $\sim 17\%$ later evolution time, across N_M , when compared to the $\mathcal{Z}_{VC}$ choice. This is in direct contrast to the BBH case (Table II). While it is the case that the level-to-level operators become more involved for CX, when contrasted against VC, the reduction in the number of intergrid transfer operations

TABLE III. One hour timing test on SuperMUC-NG of 1r2z utilizing 16 nodes. The timing ratio r is defined as $r_t := (t_F(\text{CX}) - t_0) / (t_F(\text{VC}) - t_0)$. The number of MeshBlock objects is provided at t_0 . No symmetry reduction has been performed. The AMR is based on spherical regions criteria.

N_M	$t_F(\text{CX})$	$t_F(\text{VC})$	r_t (2dp)	#MB	#MB/Core
64	170.45	151.02	1.13	3676	~ 5
96	70.84	61.11	1.17	6768	~ 9
128	34.78	28.36	1.27	12384	~ 16

appears to compensate. Indeed, it appears that the gap between CX and VC (r_t of Table III) increases with increasing N_M . This is also in contrast to the BBH case, where the gap appeared to decrease. Overall, we see that CX achieves later final evolved times in the BNS problem. Hence we conclude that $\mathcal{Z}_{CX}$ is the favored choice, for this class of problem.

IV. CONCLUSION

In this work we have proposed a new, CX variable treatment within the context of our space-time solver embedded within GR-Athena++. This has allowed us to solve the Z4c system in $\mathcal{Z}_{CX}$ mode, in a fashion consistent with our earlier work on vertex centers, i.e., $\mathcal{Z}_{VC}$ mode. Tests for both vacuum and GRHD problems demonstrated robust functionality for either grid sampling. The choice of $\mathcal{Z}_{CX}$ for GR(M)HD yields a substantial simplification in treatment of discretized variables.

To validate the new treatment, together with implementation details within the code, first consistency tests based on BBH evolution were performed. A step-by-step strategy was taken that featured, *inter alia*, a resolution study. This allowed us to assess a trend compatible with sixth order convergence in the error of the phase of the (2,2) mode of the gravitational wave strain, when working with evolution of $\mathcal{Z}_{CX}$. This is consistent with our prior work [39]. In order to provide insight as to which grid sampling strategy performs better, for BBH evolution, we considered a one hour timing test on SuperMUC-NG of the Leibniz-Rechenzentrum Munich. It was found that $\mathcal{Z}_{VC}$, on average, is $\sim 20\%$ faster. This we attribute to the differences in how the level-to-level refinement operators are evaluated for the two grid samplings. Motivated by the “symmetry-seeking” property of the L_2 AMR previously investigated in [55], we proposed a simplified, spherical-region based AMR strategy. This was demonstrated in the BBH validation problem to potentially partially mitigate some rectilinear grid imprinting.

Following this, a quasicircular, equal-mass, BNS inspiral, merger, and gravitational collapse problem was illustrated in the context of GRHD. We again utilized the (2,2) mode of the gravitational wave strain as a consistency diagnostic. A trend compatible with second order convergence was assessed in a

resolution study when working with $\mathcal{Z}_{\text{CX}}$. This trend is consistent with that observed in [40] when working with $\mathcal{Z}_{\text{VC}}$. With regard to performance, in contrast to the BBH case, it was found that $\mathcal{Z}_{\text{CX}}$ is, on average, $\sim 17\%$ faster for the runs considered. This we attribute to a reduction in the total number of intergrid interpolation operations required. For these tests, we also leveraged the simplified, spherical-region based AMR strategy. While we do not explicitly demonstrate GRMHD runs in this work, we have found another potential advantage of working with $\mathcal{Z}_{\text{CX}}$. First tests, involving the initial data of the BNS simulation presented here, allowed for evolution with magnetic fields through merger and collapse, *without* apparent horizon-based excision. This contrasts against the situation with $\mathcal{Z}_{\text{VC}}$ where we have found such excision techniques to be crucial.

The numerical simulations were performed on the national HPE Apollo Hawk at the High Performance Computing Center Stuttgart (HLRS). The authors acknowledge HLRS for funding this project by providing access to the supercomputer HPE Apollo Hawk under the grant numbers INTRHYGUE/44215 and MAGNETIST/44288. Simulations were also performed on SuperMUC_NG at the Leibniz-Rechenzentrum (LRZ) Munich. The authors acknowledge the Gauss Centre for Supercomputing e.V. ([75]) for funding this project by providing computing time on the GCS Supercomputer SuperMUC-NG at LRZ (allocations pn68wi and pn36jo). Simulations were also performed on TACC's Frontera (NSF LRAC allocation PHY23001) and on Perlmutter. This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. Postprocessing and development runs were performed on the ARA cluster at Friedrich Schiller University Jena. The ARA cluster is funded in part by DFG Grants No. INST 275/334-1 FUGG and No. INST 275/363-1 FUGG, and the ERC Starting Grant, Grant Agreement No. BinGraSp-714626.

ACKNOWLEDGMENTS

B. D. and S. B. acknowledge support by the EU Horizon under ERC Consolidator Grant No. InspiReM-101043372. P. H. acknowledges funding from the National Science Foundation under Grant No. PHY-2116686. E. G. and D. R. acknowledge funding from the National Science Foundation under Grant No. AST-2108467. E. G. acknowledges funding from an Institute for Gravitation and the Cosmos fellowship. J. F. and D. R. acknowledge the U.S. Department of Energy, Office of Science, Division of Nuclear Physics under Award No. DE-SC0021177. D. R. acknowledges support from NASA under Award No. 80NSSC21K1720 and from the Sloan Foundation. The authors are indebted to A. Celentano's PRISENCÓLINENSINÁINCIÚSOL.

APPENDIX: ASSESSING CONVERGENCE WITH A STATIC STAR

Evolving a stable, static star in the general relativistic setting provides for a simple sanity check on code correctness. To this end ID is prepared based on the assumption of a spherically symmetric, self-gravitating matter distribution at equilibrium satisfying the Tolman-Oppenheimer-Volkoff (TOV) equations (see, e.g., [76]). For ID preparation a polytropic equation of state (EOS) $p = K\rho^\Gamma$ is assumed with adiabatic index $\Gamma = 2$, and polytropic constant $K = 100$ as in [31]. The stable star model is taken to have a central density $\rho_c = 1.280 \times 10^{-3}$, gravitational mass $M = 1.400$, and circumferential radius $R = 9.586$. During evolution we make use of the more general EOS $p = (\Gamma - 1)\rho\epsilon$ with $\epsilon = K\rho^{\Gamma-1}/(\Gamma - 1)$. We perform an evolution of this data to a final time of $T = 1500$. This is done for both $\mathcal{Z}_{\text{VC}}$ and $\mathcal{Z}_{\text{CX}}$ choices, for a variety of intergrid interpolation orders (see Sec. III B). For these runs we have taken a CFL of $C_1 = 0.25$ and the Kreiss-Oliger dissipation $\sigma_D = 0.5$, and set artificial atmosphere to $\rho_{\text{atm}} = 10^{-18}$. The Z4c damping parameters are taken as




FIG. 13. Comparison of scalar quantities in stable star evolution (TOV test) on a fixed grid configuration involving CX and VC sampling where various varieties of intergrid interpolation of Sec. III B are utilized. The Mesh and MeshBlock samplings are $N_M = 64$ and $N_B = 16$, respectively. The computational domain is taken as $\Omega := [-1024, 1024]^3$ with a central, cubic region $\Omega_r := [-10, 10]^3$ refined with $N_L = 7$. The region Ω_r completely contains the star. As intergrid interpolation order is increased it can be seen (upper panel) that central density is better preserved, with oscillation magnitude of $\rho_c(t)$ reduced. Hamiltonian constraint violation (lower panel) is also reduced as interpolation order is increased. We find common mass conservation (middle panel) properties after $\sim t = 200$ regardless of interpolation order or geometric sampling. The vertical grey line indicates the first time of causal contact of $\partial\Omega$ with the origin.




FIG. 14. Convergence of scalar quantities in the TOV test for a fixed grid configuration involving CX and VC sampling. Overall grid structure is as in Fig. 13 where resolution triplets are constructed based selecting the level of Ω_r refinement through $N_L \in \{5, 6, 7\}$. In dashed lines we depict the rescaled coarse-medium difference according to Q_n of Eq. (A1). As additionally verified through least-squares fitting on $t \in [128, 950]$ we find: $n \sim 2$, i.e., $\sim 2^{\text{nd}}$ order convergence for ρ_c (upper panel); $n \sim 5$ for the Hamiltonian constraint (lower panel).

$(\kappa_1, \kappa_2) = (0.02, 0)$. Scalar diagnostic quantities are shown in Fig. 13.

To assess solution convergence we consider numerical evolution repeated at a triplet of coarse, medium, and fine resolutions $(\delta x_c, \delta x_m, \delta x_f)$ that satisfy $\delta x_c > \delta x_m > \delta x_f$.

The convergence rate of an approximation to the corresponding field data $\mathcal{F}$ may be investigated by comparing differences of solutions $\delta\mathcal{F}_{ab} := \mathcal{F}_a - \mathcal{F}_b$ at distinct resolutions. For δx^n one finds based on Taylor expansion that $\delta\mathcal{F}_{cm} \simeq \delta\mathcal{F}_{mf} \times Q_n(c, m, m, f)$ where we have introduced the so-called convergence factor,

$$Q_n(a, b, c, d) := \frac{\delta x_a^n - \delta x_b^n}{\delta x_c^n - \delta x_d^n}. \quad (\text{A1})$$

The $N_f = 2$ simulations featuring CX and VC of Fig. 13 show the smallest deviation from the initial value ρ_c during evolution, together with the smallest constraint violation. We construct a triplet of resolutions for convergence assessment by repeating the runs with $N_L \in \{5, 6\}$. Based on the refinement structure this entails $(\delta x_c, \delta x_m, \delta x_f) = (\delta x_c, \delta x_c/2, \delta x_c/4)$ on the finest level, covering the star, for each run. Behavior for the central density, together with the Hamiltonian constraint for our convergence assessment is shown in Fig. 14. As can be seen we observe $\sim 2^{\text{nd}}$ order, and $\sim 5^{\text{th}}$ order compatible trends for the aforementioned quantities, respectively. Based on the consistency between the qualitative features of the CX and VC scalar diagnostics depicted in Fig. 13, together with consistent order of convergence demonstrated in Fig. 14 we consider the TOV test to pass.

-
- [1] B. P. Abbott *et al.* (Virgo Collaboration and LIGO Scientific Collaboration), *Phys. Rev. Lett.* **119**, 161101 (2017).
 - [2] A. Goldstein *et al.*, *Astrophys. J.* **848**, L14 (2017).
 - [3] V. Savchenko *et al.*, *Astrophys. J.* **848**, L15 (2017).
 - [4] T. Akutsu *et al.*, *Prog. Theor. Exp. Phys.* **2021**, 05A102 (2021).
 - [5] P. Amaro-Seoane, [arXiv:1702.00786](#).
 - [6] M. Punturo, M. Abernathy, F. Acernese, B. Allen, N. Andersson *et al.*, *Classical Quantum Gravity* **27**, 194002 (2010).
 - [7] B. P. Abbott *et al.* (LIGO Scientific Collaboration), *Classical Quantum Gravity* **34**, 044001 (2017).
 - [8] A. Nagar *et al.*, *Phys. Rev. D* **98**, 104052 (2018).
 - [9] J. A. Font, *Living Rev. Relativity* **11**, 7 (2007).
 - [10] M. Shibata, K. Kiuchi, Y.-i. Sekiguchi, and Y. Suwa, *Prog. Theor. Phys.* **125**, 1255 (2011).
 - [11] D. Radice, S. Bernuzzi, A. Perego, and R. Haas, *Mon. Not. R. Astron. Soc.* **512**, 1499 (2022).
 - [12] D. Hilditch, A. Weyhausen, and B. Brügmann, *Phys. Rev. D* **93**, 063006 (2016).
 - [13] M. Bugner, T. Dietrich, S. Bernuzzi, A. Weyhausen, and B. Brügmann, *Phys. Rev. D* **94**, 084004 (2016).
 - [14] M. Fernando, D. Neilsen, H. Lim, E. Hirschmann, and H. Sundar, *SIAM J. Sci. Comput.* **41**, C97 (2019).
 - [15] F. Löffler *et al.*, *Classical Quantum Gravity* **29**, 115001 (2012).
 - [16] S. Shankar, P. Mösta, S. R. Brandt, R. Haas, E. Schnetter, and Y. de Graaf, *Classical Quantum Gravity* **40**, 205009 (2023).
 - [17] K. Clough, P. Figueras, H. Finkel, M. Kunesch, E. A. Lim, and S. Tunyasuvunakool, *Classical Quantum Gravity* **32**, 245011 (2015).
 - [18] W. Tichy, L. Ji, A. Adhikari, A. Rashti, and M. Pirog, *Classical Quantum Gravity* **40**, 025004 (2023).
 - [19] L. E. Kidder *et al.*, *J. Comput. Phys.* **335**, 84 (2017).
 - [20] S. Rosswog and P. Diener, *Classical Quantum Gravity* **38**, 115002 (2021).
 - [21] P. Grandclement and J. Novak, *Living Rev. Relativity* **12**, 1 (2009).
 - [22] G. Doulis, F. Atteneder, S. Bernuzzi, and B. Brügmann, *Phys. Rev. D* **106**, 024001 (2022).
 - [23] R. Alfieri, S. Bernuzzi, A. Perego, and D. Radice, *J. Low Power Electron. Appl.* **8**, 15 (2018).
 - [24] Q. F. Stout, D. L. De Zeeuw, T. I. Gombosi, C. P. T. Groth, H. G. Marshall, and K. G. Powell, in *Proceedings of the 1997 ACM/IEEE Conference on Supercomputing*, SC '97 (Association for Computing Machinery, New York, 1997), pp. 1–10.

- [25] M. J. Berger and J. Oliger, *J. Comput. Phys.* **53**, 484 (1984).
- [26] C. Burstedde, J. Holke, and T. Isaac, *Found. Comput. Math.* **19**, 843 (2019).
- [27] G. M. Morton, A computer oriented geodetic data base and a new technique in file sequencing, Tech. Rep. (International Business Machines Company, 1966).
- [28] T. Nakamura, K. Oohara, and Y. Kojima, *Prog. Theor. Phys. Suppl.* **90**, 1 (1987).
- [29] M. Shibata and T. Nakamura, *Phys. Rev. D* **52**, 5428 (1995).
- [30] T. W. Baumgarte and S. L. Shapiro, *Phys. Rev. D* **59**, 024007 (1999).
- [31] S. Bernuzzi and D. Hilditch, *Phys. Rev. D* **81**, 084003 (2010).
- [32] D. Hilditch, S. Bernuzzi, M. Thierfelder, Z. Cao, W. Tichy, and B. Bruegmann, *Phys. Rev. D* **88**, 084057 (2013).
- [33] H. Friedrich, *Commun. Math. Phys.* **100**, 525 (1985).
- [34] F. Pretorius, *Phys. Rev. Lett.* **95**, 121101 (2005).
- [35] L. Lindblom, M. A. Scheel, L. E. Kidder, R. Owen, and O. Rinne, *Classical Quantum Gravity* **23**, S447 (2006).
- [36] C. O. Lousto and J. Healy, *Phys. Rev. Lett.* **125**, 191102 (2020).
- [37] L. L. Lancelongue, H.-E. G. Aronson, A. Bateman, E. Birney, T. Caplan, M. Jukes, J. McEntyre, A. D. Morris, G. Reilly, and M. Inouye, *Nat. Comput. Sci.* **3**, 514 (2023).
- [38] F. Banyuls, J. A. Font, J. M. A. Ibanez, J. M. A. Marti, and J. A. Miralles, *Astrophys. J.* **476**, 221 (1997).
- [39] B. Daszuta, F. Zappa, W. Cook, D. Radice, S. Bernuzzi, and V. Morozova, *Astrophys. J. Suppl. Ser.* **257**, 25 (2021).
- [40] W. Cook, B. Daszuta, J. Fields, P. Hammond, S. Albanesi, F. Zappa, S. Bernuzzi, and D. Radice, *Astrophys. J. Suppl. Ser.* **277**, 3 (2025).
- [41] B. Daszuta and W. Cook, GR-Athena++: Magnetohydrodynamical evolution with dynamical space-time, in *New Frontiers in GRMHD Simulations* (Springer, Singapore, 2024).
- [42] C. J. White, J. M. Stone, and C. F. Gammie, *Astrophys. J. Suppl. Ser.* **225**, 22 (2016).
- [43] K. G. Felker and J. M. Stone, *J. Comput. Phys.* **375**, 1365 (2018).
- [44] J. M. Stone, K. Tomida, C. J. White, and K. G. Felker, *Astrophys. J. Suppl. Ser.* **249**, 4 (2020).
- [45] B. Brügmann, J. A. Gonzalez, M. Hannam, S. Husa, U. Sperhake, and W. Tichy, *Phys. Rev. D* **77**, 024027 (2008).
- [46] P. Mösta, B. C. Mundim, J. A. Faber, R. Haas, S. C. Noble, T. Bode, F. Löffler, C. D. Ott, C. Reisswig, and E. Schnetter, *Classical Quantum Gravity* **31**, 015005 (2014).
- [47] C. Palenzuela, S. Liebling, and B. Miñano, *Phys. Rev. D* **105**, 103020 (2022).
- [48] D. Radice, L. Rezzolla, and F. Galeazzi, *Classical Quantum Gravity* **31**, 075012 (2014).
- [49] M. Thierfelder, S. Bernuzzi, and B. Brügmann, *Phys. Rev. D* **84**, 044012 (2011).
- [50] C. Reisswig, R. Haas, C. D. Ott, E. Abdikamalov, P. Mösta, D. Pollney, and E. Schnetter, *Phys. Rev. D* **87**, 064023 (2013).
- [51] N. Fleischmann, S. Adami, and N. A. Adams, *Comput. Fluids* **189**, 94 (2019).
- [52] A. Neumaier, *Z. Angew. Math. Mech.* **54**, 39 (1974).
- [53] J. M. Stone, K. Tomida, C. J. White, and K. G. Felker, *Astrophys. J. Suppl. Ser.* **249**, 4 (2020).
- [54] J.-P. Berrut and L. N. Trefethen, *SIAM Rev.* **46**, 501 (2004).
- [55] A. Rashti, M. Bhattacharyya, D. Radice, B. Daszuta, W. Cook, and S. Bernuzzi, *Classical Quantum Gravity* **41**, 095001 (2024).
- [56] M. Campanelli, C. O. Lousto, P. Marronetti, and Y. Zlochower, *Phys. Rev. Lett.* **96**, 111101 (2006).
- [57] S. C. Noble, C. F. Gammie, J. C. McKinney, and L. Del Zanna, *Astrophys. J.* **641**, 626 (2006).
- [58] W. Kastaun, J. V. Kalinani, and R. Ciolfi, *Phys. Rev. D* **103**, 023018 (2021).
- [59] R. Borges, M. Carmona, B. Costa, and W. S. Don, *J. Comput. Phys.* **227**, 3191 (2008).
- [60] L. D. Zanna and N. Bucciantini, *Astron. Astrophys.* **390**, 1177 (2002).
- [61] C. Bona, J. Massó, J. Stela, and E. Seidel, in *The Seventh Marcel Grossmann Meeting: On Recent Developments in Theoretical and Experimental General Relativity, Gravitation, and Relativistic Field Theories*, edited by R. T. Jantzen, G. M. Keiser, and R. Ruffini (World Scientific, Singapore, 1996).
- [62] M. Alcubierre, B. Brügmann, P. Diener, M. Koppitz, D. Pollney, E. Seidel, and Ryoji Takahashi, *Phys. Rev. D* **67**, 084023 (2003).
- [63] J. R. van Meter, J. G. Baker, M. Koppitz, and D.-I. Choi, *Phys. Rev. D* **73**, 124011 (2006).
- [64] M. Ansorg, B. Brügmann, and W. Tichy, *Phys. Rev. D* **70**, 064011 (2004).
- [65] D. I. Ketcheson, *J. Comput. Phys.* **229**, 1763 (2010).
- [66] N. Wang and J.-L. Lee, *SIAM J. Sci. Comput.* **33**, 2536 (2011).
- [67] J. N. Goldberg, A. J. MacFarlane, E. T. Newman, F. Rohrlrich, and E. C. G. Sudarshan, *J. Math. Phys. (N.Y.)* **8**, 2155 (1967).
- [68] LIGO Scientific Collaboration, LIGO Algorithm Library—LALSuite, free software (GPL), (2018).
- [69] S. Albanesi, A. Rashti, F. Zappa, R. Gamba, W. Cook, B. Daszuta, S. Bernuzzi, A. Nagar, and D. Radice, *Phys. Rev. D* **111**, 024069 (2025).
- [70] M. Thierfelder, S. Bernuzzi, D. Hilditch, B. Brügmann, and L. Rezzolla, *Phys. Rev. D* **83**, 064022 (2011).
- [71] V. Mewes, Y. Zlochower, M. Campanelli, T. W. Baumgarte, Z. B. Etienne, F. G. Lopez Armengol, and F. Cipolletta, *Phys. Rev. D* **101**, 104007 (2020).
- [72] E.ourgoulhon, P. Grandclement, K. Taniguchi, J.-A. Marck, and S. Bonazzola, *Phys. Rev. D* **63**, 064029 (2001).
- [73] S. Gottlieb, D. I. Ketcheson, and C.-W. Shu, *J. Sci. Comput.* **38**, 251 (2009).
- [74] C. Reisswig and D. Pollney, *Classical Quantum Gravity* **28**, 195015 (2011).
- [75] www.gauss-centre.eu.
- [76] J. A. Font, M. A. Miller, W.-M. Suen, and M. Tobias, *Phys. Rev. D* **61**, 044011 (2000).