



CHORUS

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

Three-dimensional general-relativistic hydrodynamic simulations of binary neutron star coalescence and stellar collapse with multipatch grids

C. Reisswig, R. Haas, C. D. Ott, E. Abdikamalov, P. Mösta, D. Pollney, and E. Schnetter

Phys. Rev. D **87**, 064023 — Published 18 March 2013

DOI: [10.1103/PhysRevD.87.064023](https://doi.org/10.1103/PhysRevD.87.064023)

Three-Dimensional General-Relativistic Hydrodynamic Simulations of Binary Neutron Star Coalescence and Stellar Collapse with Multipatch Grids

C. Reisswig,^{1,*} R. Haas,¹ C. D. Ott,^{1,2,†} E. Abdikamalov,¹ P. Mösta,¹ D. Pollney,³ and E. Schnetter^{4,5,6}

¹*TAPIR, MC 350-17, California Institute of Technology,
1200 E California Blvd., Pasadena, CA 91125, USA*

²*Kavli Institute for the Physics and Mathematics of the Universe (Kavli IPMU), The University of Tokyo, Kashiwa, Japan*

³*Department of Mathematics, Rhodes University, Grahamstown, 6139 South Africa*

⁴*Perimeter Institute for Theoretical Physics, 31 Caroline St. N., Waterloo, ON N2L 2Y5, Canada*

⁵*Department of Physics, University of Guelph, 50 Stone Road East, Guelph, ON N1G 2W1, Canada*

⁶*Center for Computation & Technology, 216 Johnston Hall,
Louisiana State University, Baton Rouge, LA 70803, USA*

(Dated: February 7, 2013)

We present a new three-dimensional general-relativistic hydrodynamic evolution scheme coupled to dynamical spacetime evolutions which is capable of efficiently simulating stellar collapse, isolated neutron stars, black hole formation, and binary neutron star coalescence. We make use of a set of adapted curvi-linear grids (multipatches) coupled with flux-conservative cell-centered adaptive mesh refinement. This allows us to significantly enlarge our computational domains while still maintaining high resolution in the gravitational-wave extraction zone, the exterior layers of a star, or the region of mass ejection in merging neutron stars. The fluid is evolved with a high-resolution shock capturing finite volume scheme, while the spacetime geometry is evolved using fourth-order finite differences. We employ a multi-rate Runge-Kutta time integration scheme for efficiency, evolving the fluid with second-order and the spacetime geometry with fourth-order integration, respectively. We validate our code by a number of benchmark problems: a rotating stellar collapse model, an excited neutron star, neutron star collapse to a black hole, and binary neutron star coalescence. The test problems, especially the latter, greatly benefit from higher resolution in the gravitational-wave extraction zone, causally disconnected outer boundaries, and application of Cauchy-characteristic gravitational-wave extraction. We show that we are able to extract convergent gravitational-wave modes up to $(\ell, m) = (6, 6)$. This study paves the way for more realistic and detailed studies of compact objects and stellar collapse in full three dimensions and in large computational domains. The multipatch infrastructure and the improvements to mesh refinement and hydrodynamics codes discussed in this paper will be made available as part of the open-source Einstein Toolkit.

PACS numbers: 04.25.D-, 04.30.Db, 97.60.Bw, 02.70.Bf

I. INTRODUCTION

Some of the most interesting relativistic astrophysical phenomena such as stellar collapse, black hole formation, or binary neutron star coalescence, require numerical simulations on large computational domains, involve many different length scales, and are intrinsically three-dimensional (3D). Due to their extreme nature in terms of fluid densities and velocities, an accurate treatment of general-relativistic (GR) gravity is required. Depending on the problem, magnetic field evolution and neutrino interactions may also be required. Thus, numerical computations in relativistic astrophysics are truly multi-physics, and as such, are especially demanding in terms of computational modeling technology and resources.

Current state of the art 3D GR hydrodynamic simulations in the context of stellar collapse [1–5] or binary neutron star coalescence [6–13] (see [14] for a recent review) are based on Cartesian grids with adaptive mesh refinement (AMR). As the domain is enlarged or the resolution

increased, such grids pose a serious bottleneck in terms of the computational power that is required, even with AMR. Since Cartesian grids scale as N^3 in terms of the number N of grid points along one spatial direction in 3D, available computational resources are rapidly exhausted when additional points in each coordinate direction are added. The symmetry of the computational problem, however, is essentially spherical, at least at some distance from the central region of the simulation. Thus, Cartesian grids are wasteful with respect to angular resolution when the problem becomes symmetrically spherical.

For instance, stellar collapse proceeds in approximately spherical or axisymmetric terms (e.g. [15–18]). At later times, various hydrodynamic instabilities (e.g. convection and instabilities of the shock) break this symmetry. The global features, however, remain approximately spherical or axisymmetric.

In the case of coalescing binary neutron stars, the central region containing the two neutron stars is not of spherical symmetry. At larger distances and in the gravitational-wave (GW) zone, however, the problem becomes spherical. The gravitational-wave extraction zone must generally be located at large radii in order to limit near-zone effects in the extracted wave. But

* Einstein Fellow; reisswig@tapir.caltech.edu

† Alfred P. Sloan Research Fellow

even with more sophisticated techniques such as Cauchy-characteristic extraction [19–24] that allow us to extract gauge-invariant GWs at future null infinity \mathcal{I}^+ , it is necessary to enlarge the domain sufficiently so that constraint-violating modes generated at the outer boundary are causally disconnected from the interior evolution and the wave-extraction zone. These constraint-violating modes are generated due to the lack of constraint-preserving outer boundary conditions for the Einstein equations (see [25] for a recent review) for certain types of evolution systems (including the common BSSN system), and travel at the speed of light [26, 27] to the interior of the domain. Without these systematic errors, the evolution and wave extraction would generally be more accurate. Furthermore, in case mass is ejected during and after merger, enlarging refinement levels to track the evolution of the ejected material becomes very expensive.

It therefore seems natural to apply spherical grids to maintain high resolution also in the outer regions of the domain. The computational effort when using spherical grids scales linearly with the number of radial points N , assuming constant angular resolution. Thus, spherical grids can give a tremendous performance improvement when the domain is enlarged or the (radial) resolution increased.

Spherical grids have been widely used for many astrophysical problems, including stellar collapse (e.g., [28–32]), core-collapse supernovae (e.g., [33–35]), oscillations of neutron stars (e.g., [36, 37]), neutron star magnetosphere (e.g., [38]), accretion onto black holes [39], and simulations of accretion disks (e.g., [40–42]). Unfortunately, the standard spherical-polar coordinate system imposes a serious difficulty along the axis and the poles, where special care must be taken to regularize the fields (e.g., [43]). Recent approaches make use of implicit Runge-Kutta methods that yield stable evolutions in spherical coordinates without regularizations [44–46]. But even with these fixes, the angular and radial distribution of grid points is non-optimal in the sense that they cluster at the poles and at the coordinate origin. In addition, spherical grids are less suited in regions where the underlying symmetry is non-spherical, e.g., in the vicinity of a binary neutron star system, or the highly turbulent and convective region behind the accretion shock in a core-collapse supernova.

In order to handle multiple regions of different symmetry within the same simulation, *multipatch* (sometimes also called *multiblock*) schemes have been developed for a wide range of physics and engineering applications. The idea is to cover the simulation domain with multiple curvi-linear coordinate “patches”. Each patch is locally uniform. Diffeomorphic mappings from local to the global coordinates enable to represent a wide range of grid shapes in different regions of the simulation. One such example is given in Fig. 1. In this setup, a central Cartesian patch is surrounded by six “inflated cube” spherical grid patches. This is a natural configuration for our purposes. The aspherical region of a collapsing star

or a merging binary is best modeled by a central Cartesian patch, capable of AMR. The gravitational-wave zone and/or the outer layers of a star are best modeled by the more efficient spherical grids. This allows us to employ large domains at high resolution with modest computational cost. Notably, the outer boundary can be causally disconnected from the interior evolution and the gravitational-wave extraction zone.

Within the context of numerical relativity and relativistic astrophysics, multipatch schemes have already been successfully applied in a range of different problems ranging from simulations of accretion disks [47, 48], horizon finding [49], wave extraction [50], single black holes [51, 52], orbiting black holes [53], relativistic fluid evolutions on fixed backgrounds [54], elliptic and initial data solvers [55–59], to characteristic evolutions of Einstein’s equations [20, 60, 61]. Multidomain *spectral* methods have been successfully applied to vacuum binary black hole evolutions yielding high accuracy and efficiency [62–67] using a dual-coordinate frame method [68]. The same multidomain spectral code `SpEC`, coupled to a finite volume fluid solver, has also been used to simulate neutron star black hole mergers [69–72]. Neither of the works above, however, make use of AMR for the fluid fields, and thus are limited in the respective range of astrophysical applications. In particular, efficient simulations of stellar collapse and black hole formation require AMR in the central region of the collapsing star. Also, the near-field region in simulations of binary neutron star coalescence substantially benefit from AMR, in particular when material is ejected in the post-merger phase.

In the context of vacuum binary black hole merger simulations, multipatch schemes combined with AMR have been successfully applied [73–79]. We base our code on the `Llama` infrastructure developed in [73], which makes use of the `Cactus` computational toolkit [80] and the `Carpets` AMR driver [81, 82]. We extend the original pure vacuum scheme to include full matter dynamics using the publicly available GR hydrodynamics code `GRHydro`, which is part of the `EinsteinToolkit` [83]. We thus present the first successful multipatch scheme capable of AMR that can stably evolve fluid dynamics coupled to fully GR spacetime dynamics.

In addition, we make a number of improvements: (i) We extend the AMR driver `Carpets` to support cell-centered mesh refinement, which allows us to apply *refluxing*, a technique to maintain conservation of mass, energy and momentum fluxes across mesh refinement boundaries [84] (see [85] for a recent application to GR hydrodynamics). This greatly improves conservation of mass in our simulations of stellar collapse, especially in the postbounce evolution. (ii) We apply enhanced PPM (piecewise parabolic method) reconstruction [86, 87], which significantly improves the numerical accuracy and the behavior of the constraints. (iii) To improve the execution speed of the simulations, we apply multirate Runge-Kutta (RK) time integration (e.g. [88, 89]) in which the spacetime is evolved with a standard fourth-

order RK method, whereas the fluid is evolved with a second order RK scheme without significant loss of accuracy. This reduces the number of intermediate steps in the fluid evolution, which dominates in terms of processor cycles compared to spacetime evolution, in particular when using a microphysical equation of state.

We apply the new code to a number of benchmark problems, including the evolution of a single isolated and perturbed neutron star, the collapse of a rotating stellar core, the collapse of a neutron star to a black hole, and the merger of a binary neutron star system. We investigate the accuracy and convergence of each test problem. This is an important code verification towards our program to carry out fully 3D simulations of core-collapse supernovae (see [90] for a recent application of our scheme) and black hole formation in the context of the collapsar scenario for long gamma-ray bursts. The new multipatch scheme allows us to significantly enlarge the computational domain by maintaining a fixed angular resolution. This is useful in many ways: (i) we are able to causally disconnect the outer boundary from the interior evolution and the gravitational-wave extraction zone, thus avoiding systematic errors from the approximate and non-constraint preserving artificial outer boundary condition, (ii) we have a larger wave-extraction zone with higher overall resolution, thus making it possible to extract higher-order than the dominant GW modes, (iii) in binary neutron star mergers, ejected material can be tracked out to large radii with relatively high resolution, (iv) the number of mesh refinement levels can be decreased, leading to better parallel scaling. As a result, our multipatch scheme can efficiently evolve models of stellar collapse in full 3D (see also [90]), and is capable of more accurate gravitational-wave extraction in models of binary neutron star mergers. In the latter test problem, we extract convergent gravitational-wave modes up to $(\ell, m) = (6, 6)$.

This paper is organized as follows. In Sec. II A and II B, we first review the underlying hydrodynamic and spacetime evolution systems and how we solve them numerically. Subsequently, in Sec. II C, we present our approach to multipatches and their numerical implementation. We also discuss our implementation of cell-centered AMR (Sec. II D), and describe multirate RK time integration (Sec. II E). Finally, in Sec. III, we present detailed tests of isolated perturbed and unperturbed neutron stars, collapsing stellar cores, neutron star collapse to a black hole, and merging binary neutron stars. We conclude and summarize our findings in Sec. IV. In an appendix, we presents basic tests with shock tubes (Appendix A), we review the enhanced PPM scheme as developed in [86, 87] (Appendix B), discuss our treatment of the artificial low-density atmosphere (Appendix C), present an optimized ghost-zone update scheme to improve the parallel scaling (Appendix D), describe our volume integration scheme for overlapping grids (Appendix E), and investigate the influence of boundary effects on binary neutron star merger dynamics and wave

extraction (Appendix F).

II. METHODS

A. General-Relativistic Hydrodynamics

We base our code on the open-source GR hydrodynamics code `GRHydro` that is part of the `EinsteinToolkit` [91] and is described in [21, 83, 92].

We introduce *primitive* variables in the form of the fluid density ρ , the fluid's specific internal energy ϵ , and the fluid 3-velocity as seen by Eulerian observers at rest in the current spatial 3-hypersurface [93],

$$v^i = \frac{u^i}{W} + \frac{\beta^i}{\alpha}, \quad (1)$$

where u^i is the fluid 4-velocity, $W = (1 - v^i v_i)^{-1/2}$ is the Lorentz factor, and α and β^i are lapse and shift, respectively (to be introduced in Sec. II B). In terms of the 3-velocity, the contravariant 4-velocity is then given by

$$u^0 = \frac{W}{\alpha}, \quad u^i = W \left(v^i - \frac{\beta^i}{\alpha} \right), \quad (2)$$

and the covariant 4-velocity is

$$u_0 = W(v^i \beta_i - \alpha), \quad u_i = W v_i. \quad (3)$$

The evolution equations are written in the Valencia form of GR hydrodynamics [94, 95] as a first-order hyperbolic flux-conservative evolution system for the *conserved* variables D , S^i , and τ which are defined in terms of the primitive variables ρ , ϵ , v^i ,

$$\begin{aligned} D &= \sqrt{\gamma} \rho W, \\ S^i &= \sqrt{\gamma} \rho h W^2 v^i, \\ \tau &= \sqrt{\gamma} (\rho h W^2 - P) - D, \end{aligned} \quad (4)$$

where γ is the determinant of the 3-metric γ_{ij} (see Sec. II B), and the quantities P , and $h = 1 + \epsilon + P/\rho$ denote pressure, and specific enthalpy, respectively. The evolution system then becomes

$$\frac{\partial \mathbf{U}}{\partial t} + \frac{\partial \mathbf{F}^i}{\partial x^i} = \mathbf{S}, \quad (5)$$

with

$$\begin{aligned} \mathbf{U} &= [D, S_j, \tau], \\ \mathbf{F}^i &= \alpha [D \tilde{v}^i, S_j \tilde{v}^i + \delta_j^i P, \tau \tilde{v}^i + P v^i], \\ \mathbf{S} &= \alpha \left[0, T^{\mu\nu} \left(\frac{\partial g_{\nu j}}{\partial x^\mu} - \Gamma_{\mu\nu}^\lambda g_{\lambda j} \right), \right. \\ &\quad \left. \alpha \left(T^{\mu 0} \frac{\partial \ln \alpha}{\partial x^\mu} - T^{\mu\nu} \Gamma_{\mu\nu}^0 \right) \right]. \end{aligned} \quad (6)$$

Here, $\tilde{v}^i = v^i - \beta^i/\alpha$, $\Gamma_{\mu\nu}^\lambda$ are the 4-Christoffel symbols, and $T^{\mu\nu}$ is the stress-energy tensor. The pressure $P = P(\rho, \epsilon, \{X_i\})$ is obtained via our equation of state module, which is capable of handling a set of different equations of state, including microphysical finite-temperature variants. The $\{X_i\}$ are additional compositional variables of the matter such as the electron fraction Y_e , which are used for microphysical equations of state. In the present work, however, we resort to simple (piecewise) polytropic and ideal gas (Γ -law) equations of state.

The above evolution equations are spatially discretized by means of a high-resolution shock-capturing (HRSC) scheme using a second-order accurate finite-volume algorithm. The equations are kept in semi-discrete form and first-order (in space) Riemann problems are solved at cell interfaces with the approximate HLL solver¹ [97].

The states at cell interfaces are reconstructed using a new and improved variant of the piecewise parabolic method (PPM) [86, 87, 98]. As noted in [86, 87], the original PPM (oPPM) scheme [98] has the side-effect of flattening local smooth extrema which are physical, thus limiting the accuracy. In the present context of simulating compact objects, one naturally has extrema at the stellar center(s) where the matter density is largest. We find that the oPPM scheme reduces the accuracy there, which then strongly affects the overall accuracy of our simulations (see Sec. III, and also Fig. 25). Ref. [86], further refined by Ref. [87], suggests modifications to the original limiter which can distinguish between smooth maxima that are part of the solution, and artificial maxima that may be introduced at shocks and other discontinuities. While smooth maxima need to be retained as part of the solution, artificial maxima must be avoided to suppress Gibbs phenomenon at shocks and other discontinuities. We summarize the procedure for “enhanced” PPM (ePPM) reconstruction in Appendix B.

We note that under certain conditions, the requirement that the modulus of the reconstructed primitive velocity must stay below the speed of light c may be violated. This can happen, since the primitive velocity is a bounded function (bounded by the requirement $v_i v^i \leq c^2$), and the ePPM reconstruction scheme does not enforce this constraint close to any occurring extrema. Thus, the ePPM scheme may reconstruct velocity components that result in a velocity modulus equal to or slightly larger than the speed of light near extrema. To avoid this problem, we reconstruct Wv^i , i.e. the Lorentz factor W times the primitive velocity v^i . The quantity Wv^i is unbounded and thus does not require special treatment near extrema.

The time integration and coupling with curvature (Sec. II B) are carried out with the Method of Lines [99]

(see Sec. II E).

After each evolution step, we compute the primitive quantities from the evolved conserved quantities. Since the primitive quantities are implicit functions of the conserved ones, it is necessary to use a numerical root finding algorithm. As described in, e.g. [83], this is done via a Newton-Raphson scheme.

In some rare situations, the initial guesses for the root finding procedure are not well-posed, and cause the Newton-Raphson scheme to fail to converge. In particular, we find this behavior at the surface of a neutron star, when the latter is threaded by an AMR boundary and refluxing is active. In this case, we resort to a simple bisection algorithm which converges more slowly, but is more robust.

In regions of the computational domain, where we have physical vacuum, we employ an artificial low density “atmosphere” (see Appendix C). In order to reduce the influence of the artificial atmosphere on the curvature evolution, we exponentially damp the stress-energy tensor $T_{\mu\nu}$ to zero outside a given radius. More specifically, we introduce the radius dependent stress-energy damping $T_{\mu\nu} \rightarrow \lambda(r)T_{\mu\nu}$ with the damping factor

$$\lambda(r) = \begin{cases} 1 & \text{for } r \leq R_0, \\ \frac{1}{2} \left(1 - \tanh \left(\frac{8r - 4(R_1 + R_0)}{R_1 - R_0} \right) \right) & \text{otherwise,} \\ 0 & \text{for } r \geq R_1, \end{cases} \quad (7)$$

where the damping is applied between the two radii $R_0 < R_1$.

At outer boundaries, we apply a copy-from-neighbor (flat) boundary condition for the evolved fluid quantities.

Finally, in order to be compatible with multipatch discretization, we need to introduce additional coordinate transformations as described in Sec. II C 3 below.

B. Curvature Evolution

The spacetime evolution is performed by a variant of the BSSN evolution system [100–103] and is implemented in the CTGamma curvature evolution code [73], which was developed for arbitrary coordinate systems mapping the spatial domain.

The standard BSSN system is derived from a 3+1 split of spacetime resulting in a foliation in terms of spatial hypersurfaces along a timelike vector field. It introduces the following set of evolved variables

$$\phi, \quad \tilde{\gamma}_{ab}, \quad K, \quad \tilde{A}_{ab}, \quad \tilde{\Gamma}^a, \quad (8)$$

¹ More sophisticated Riemann solvers are available within GRHydro. In our experience, however, HLL is a robust and fast choice. We find that more sophisticated solvers do not lead to a significant accuracy improvement (see also, e.g., [96]).

which are solved according to

$$\partial_t \phi = -\frac{1}{6} \alpha K + \frac{1}{6} \partial_i \beta^i, \quad (9a)$$

$$\begin{aligned} \partial_t \tilde{\gamma}_{ab} = & -2\alpha \tilde{A}_{ab} + \beta^i \partial_i \tilde{\gamma}_{ab} + 2\tilde{\gamma}_{i(a} \partial_b) \beta^i \\ & - \frac{2}{3} \tilde{\gamma}_{ab} \partial_i \beta^i, \end{aligned} \quad (9b)$$

$$\begin{aligned} \partial_t K = & -D_i D^i \alpha + \alpha (A_{ij} A^{ij} + \frac{1}{3} K^2) + \beta^i \partial_i K \\ & + 4\pi \alpha (\rho_{\text{ADM}} + S), \end{aligned} \quad (9c)$$

$$\begin{aligned} \partial_t \tilde{A}_{ab} = & e^{-4\phi} (-D_a D_b \alpha + \alpha R_{ab})^{\text{TF}} + \beta^i \partial_i \tilde{A}_{ab} \\ & + 2\tilde{A}_{i(a} \partial_b) \beta^i - \frac{2}{3} A_{ab} \partial_i \beta^i \\ & - 8\pi e^{-4\phi} \alpha (S_{ab})^{\text{TF}}, \end{aligned} \quad (9d)$$

$$\begin{aligned} \partial_t \tilde{\Gamma}^a = & \tilde{\gamma}^{ij} \partial_i \beta_j \beta^a + \frac{1}{3} \tilde{\gamma}^{ai} \partial_i \partial_j \beta^j + \beta^j \partial_j \tilde{\Gamma}^a \\ & - \tilde{\Gamma}^i \partial_i \beta^a + \frac{2}{3} \tilde{\Gamma}^a \partial_i \beta^i - 2\tilde{A}^{ai} \partial_i \alpha \\ & + 2\alpha (\tilde{\Gamma}_{ij}^a \tilde{A}^{ij} + 6\tilde{A}^{ai} \partial_i \phi - \frac{2}{3} \tilde{\gamma}^{ai} \partial_i K) \\ & - 16\pi \alpha \tilde{\gamma}^{ai} S_i, \end{aligned} \quad (9e)$$

where D_a is the covariant derivative determined by the conformal 3-metric $\tilde{\gamma}_{ab}$, and ‘‘TF’’ indicates that the trace-free part of the bracketed term is used.

Above, we show the ‘‘ ϕ ’’-variant of the BSSN system. Our curvature evolution code also provides the ‘‘ χ ’’- and ‘‘ W ’’-variants of the evolution system (see [73] for details). Here, we employ the ϕ -variant.

The stress-energy tensor $T_{\mu\nu}$ is incorporated via the projections

$$\rho_{\text{ADM}} := \frac{1}{\alpha^2} (T_{00} - 2\beta^i T_{0i} + \beta^i \beta^j T^{ij}), \quad (10)$$

$$S := \tilde{\gamma}^{ij} T_{ij}, \quad (11)$$

$$S_a := -\frac{1}{\alpha} (T_{0a} - \beta^j T_{aj}), \quad (12)$$

$$(S_{ab})^{\text{TF}} := \left(T_{ab} - \frac{1}{3} e^{4\phi} S \tilde{\gamma}_{ab} \right). \quad (13)$$

After each evolution step, the evolved curvature variables (8) are transformed (via an algebraic relation) to the standard ADM variables $\{g_{ij}, K_{ij}\}$ (e.g., [104]), where g_{ij} is the (physical) 3-metric, and K_{ij} the extrinsic curvature. The ADM variables are used to couple the curvature evolution to the hydrodynamic evolution scheme, i.e., our hydrodynamic scheme uses the physical 3-metric g_{ij} rather than the evolved conformal 3-metric $\tilde{\gamma}_{ab}$ above.

The *lapse* gauge scalar α is evolved using the $1 + \log$ condition [105],

$$\partial_t \alpha - \beta^i \partial_i \alpha = -2\alpha K, \quad (14)$$

while the *shift* gauge vector β^a is evolved using the hy-

perbolic $\tilde{\Gamma}$ -driver equation [106],

$$\partial_t \beta^a - \beta^i \partial_i \beta^a = \frac{3}{4} B^a, \quad (15a)$$

$$\partial_t B^a - \beta^j \partial_j B^i = \partial_t \tilde{\Gamma}^a - \beta^i \partial_i \tilde{\Gamma}^a - q(r) \eta B^a, \quad (15b)$$

where η is a parameter which acts as a (mass dependent) damping coefficient. To avoid certain stability issues with the gauge arising in the far-field regime [107], the damping coefficient is allowed to spatially change, either by some dynamic evolution [108], or by a fixed prescription. We use the simple prescription for a radial fall-off of $q(r)\eta$ with $0 \leq q(r) \leq 1$ (see [107] for details). If not stated otherwise, we use a fall-off radius of $R = 250 M_\odot$.

The 3+1 decomposition of the Einstein equations also results in a set of constraint equations. The Hamiltonian constraint equation reads

$$H \equiv R^{(3)} + K^2 - K_{ij} K^{ij} - 16\pi \rho_{\text{ADM}} = 0, \quad (16)$$

where $R^{(3)}$ denotes the 3-Ricci scalar, and the momentum constraint equations read

$$M^a \equiv D_i (K^{ai} - \gamma^{ai} K) - 8\pi S^a = 0. \quad (17)$$

We do not actively enforce the constraints during evolution, but rather check how well our numerically obtained metric quantities satisfy the constraints over the course of the evolution. Thus, this offers a valuable accuracy monitor for the curvature evolution.

The spacetime equations are discretized using fourth-order finite difference operators [109]. The finite difference stencils are centered. An exception are the advection terms of the form $\beta^i \partial_i$, which use operators that are upwinded by one stencil point towards the local direction of the shift vector β^i [73].

Consistent with the order of accuracy of spatial finite difference derivatives, we also apply Kreiss-Oliger dissipation [109] which is of one order higher than the spatial discretization order. In the case of fourth-order differencing, we thus apply fifth-order dissipation operators. Dissipation is added to the right-hand-sides (RHS) of the curvature evolution quantities at any time integration substep. The strength of the dissipation can be controlled by a parameter $\epsilon_{\text{diss}} \in [0, 1]$. Unless otherwise specified, we use $\epsilon_{\text{diss}} = 0.1$ throughout this work.

At outer boundaries, we impose a simple approximate radiative boundary condition as described in [73]. Since data from this condition are not strictly constraint satisfying, constraint violating modes are generated at the boundary, and travel with the speed of light [26, 27] to the interior of the domain where they introduce a systematic error in the curvature evolution.

C. Multipatches

We build our code on the `Llama` infrastructure described in detail in [73]. This infrastructure implements

multipatches via an arbitrary number of curvi-linear overlapping grid patches using fourth-order Lagrange and second-order essentially non-oscillatory (ENO) interpolation for exchanging data in inter-patch ghost zones between neighboring patches. In [73], only the pure vacuum problem was considered. Here, we extend the multipatch evolution scheme to include matter.

1. Patch Systems

A useful patch system is shown in Fig. 1: the central Cartesian patch is surrounded by six spherical inflated-cube patches. The *nominal*² grids of the spherical patches have inner radius R_S , outer radius R_B , radial spacing ΔR_1 , which is allowed to stretch to ΔR_2 within some finite region, and angular resolution $(\Delta\rho, \Delta\sigma)$ per angular direction (ρ, σ) . Note that the angles (ρ, σ) used to define the local coordinates of each inflated-cube patch do not coincide with standard spherical-polar coordinates (see below). The central patch contains a hierarchy of refined regions, allowing to place resolution where necessary. This patch system is particularly useful in problems with spherical symmetry at some radius from the central source.

Each grid patch defines *local* uniform coordinates (u, v, w) related to the *global* Cartesian (x, y, z) coordinate space by a diffeomorphic relation. For the central Cartesian patch depicted in Fig. 1, this relation is trivially given by the identity function. The inflated-cube coordinates, however, are defined by non-trivial coordinate functions. For each angular patch, we define local angular coordinates (ρ, σ) that range over $(-\pi/4, +\pi/4) \times (-\pi/4, +\pi/4)$ and can be related to global angular coordinates (μ, ν, ϕ) (see Fig. 1) which are given by

$$\mu \equiv \text{rotation angle about the x-axis} = \arctan(y/z), \quad (18a)$$

$$\nu \equiv \text{rotation angle about the y-axis} = \arctan(x/z), \quad (18b)$$

$$\phi \equiv \text{rotation angle about the z-axis} = \arctan(y/x). \quad (18c)$$

For each angular patch, we have two unique angles (ρ, σ) out of the three global angles (μ, ν, ϕ) that parametrize the local coordinates. For instance, for the patch normal to the positive x -direction, we select

$$\rho \equiv \nu = \arctan(z/x), \quad (19a)$$

$$\sigma \equiv \phi = \arctan(y/x), \quad (19b)$$

$$R = f(r), \quad (19c)$$

where $r = \sqrt{x^2 + y^2 + z^2}$. Similarly, the coordinates of the patch along the positive y and z axes are parametrized by $(\rho, \sigma) \equiv (\mu, \phi)$ and $(\rho, \sigma) \equiv (\mu, \nu)$, respectively. The remaining three patches along the negative axes are related in a similar way.

In the radial coordinate direction, we apply radial stretching with an appropriate stretching function $R = f(r)$. In the stretching region, the physical coordinate radius is stretched, corresponding to a smooth decrease in radial resolution from spacing ΔR_1 to spacing ΔR_2 . Outside the stretching region, we keep the radial spacing constant. Details can be found in [73].

2. Spacetime Evolution Scheme

Here, and as described in [73], the spacetime evolution is solved in the *global* Cartesian (x, y, z) tensor basis, where the grid patches are generally distorted, i.e., they are *not* uniform. Derivatives are approximated via finite differences in the local coordinate system (u, v, w) of each grid patch, where, as required by our finite difference scheme, the grid patches are uniform. In order to transform to the global tensor basis, Jacobian transformations of the form $J^i_j = \partial u^i / \partial x^j$ are applied to the first and second derivatives at each point,

$$\frac{\partial}{\partial x_i} = \left(\frac{\partial u_j}{\partial x_i} \right) \frac{\partial}{\partial u_j}, \quad (20a)$$

$$\frac{\partial^2}{\partial x_i \partial x_j} = \left(\frac{\partial^2 u_k}{\partial x_i \partial x_j} \right) \frac{\partial^2}{\partial u_k^2} + \left(\frac{\partial u_k}{\partial x_i} \frac{\partial u_l}{\partial x_j} \right) \frac{\partial^2}{\partial u_k \partial u_l}, \quad (20b)$$

thus obtaining the derivatives in the global (x, y, z) coordinate space. The Jacobians are precomputed at each grid point. The main advantage of solving the equations in the global (x, y, z) basis is simplicity. There is no need for inter-patch coordinate basis transformations. Perhaps more importantly, the existing code infrastructure, and especially analysis tools, do not need to be changed, since the assumption of a global Cartesian tensor basis is still maintained.

3. Hydrodynamic Evolution Scheme

Finite volume schemes work well on general unstructured meshes. The original implementation of the hydrodynamic evolution code GRHydro, however, assumes uniform coordinates. Without a major rewrite of the code, we can keep our original scheme by solving the Riemann problem in the *local* frame, where the coordinates are uniform. This requires no changes to the core of the scheme. Any computation simply carries over to the local coordinate basis. Effectively, this means that the primitive and conserved quantities are thus represented in the local coordinate basis.

² We define the *nominal* grid as the unique set of points covering the entire computational domain, i.e. the nominal grid of a single patch excludes ghost points (and additional overlap points; see further below) that are shared with a neighboring patch.

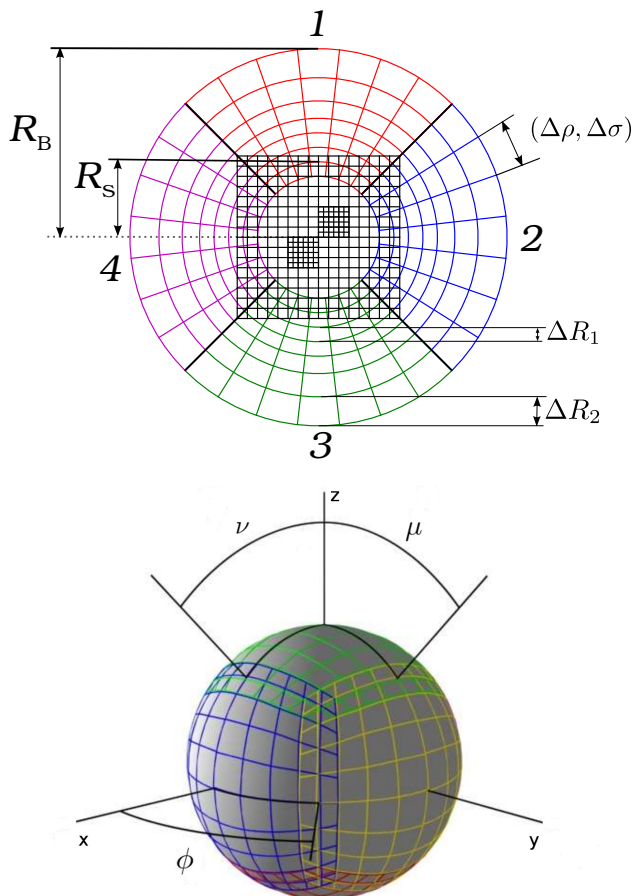


FIG. 1. Depiction of a typical patch system used in our simulations. The upper figure schematically shows a $z = 0$ slice of the employed grids: a central Cartesian grid (patch 0) is surrounded by spherical *inflated-cube* grid patches (patches 1 – 4 out of a total of six spherical patches). The central grid is capable of AMR allowing to refine the resolution at the central region of, e.g., a star where the density and curvature gradients become large. R_B and R_S denote the radii of the outer computational boundary and of the boundary between spherical and Cartesian grids, respectively. The spherical grid has a fixed angular resolution denoted by $(\Delta\rho, \Delta\sigma)$, while the radial resolution is allowed to stretch from radial resolution ΔR_1 to ΔR_2 . The lower figure shows a radial $R = \text{const.}$ shell of the outer spherical grid, comprised of six inflated-cube grid patches. Angular points can be uniquely determined by two out of three angular coordinates (μ, ν, ϕ) (18). Interpolation at patch boundaries reduces to 1D interpolation. Points are almost uniformly distributed across the sphere.

Special attention is required when coupling the hydrodynamics solver to the metric solver (Sec. II C 2). The metric solver explicitly computes the metric components in the *global* frame and is thus generally incompatible with the hydrodynamic quantities defined in the *local* frame. We therefore introduce the additional step of transforming the metric components to the *local* basis before each hydrodynamic RHS step. Correspondingly, after each hydrodynamic step, we need to compute the

TABLE I. Required quantities for the hydrodynamic evolution scheme and their coordinate bases. A tilde denotes quantities that need to be obtained by applying a Jacobian transformation. The last four quantities are only required for microphysical equations of state.

Quantity	Type	global	local
metric	tensor	g_{ij}	\tilde{g}_{ij}
extrinsic curvature	tensor	K_{ij}	\tilde{K}_{ij}
shift	vector	β^i	$\tilde{\beta}^i$
lapse	scalar	α	α
prim. density	scalar	ρ	ρ
specific internal energy	scalar	ϵ	ϵ
prim. velocity	vector	\tilde{v}^i	v^i
cons. density	densitized scalar	-	D
cons. internal energy	densitized scalar	-	τ
momentum	densitized vector	-	S_i
stress-energy tensor	tensor	$T^{\mu\nu}$	-
Lorentz factor	scalar	W	-
pressure	scalar	P	-
prim. electron fraction	scalar	Y_e	Y_e
cons. electron fraction	densitized scalar	-	Y_e^{con}
temperature	scalar	T	T
entropy	scalar	s	s

stress-energy tensor $T^{\mu\nu}$ in the *global* basis as required by the metric solver.

Since the various analysis tools explicitly assume a global coordinate frame for the primitive variables, we introduce a separate set of *global* primitive variables. Effectively, this only requires extra memory for the primitive 3-velocity $\{\tilde{v}^i\}$, since the primitive density ρ and the specific internal energy ϵ are scalars. In case of microphysical equation of states, no extra memory is required for the temperature T , the entropy s , and the primitive electron fraction Y_e , since they are scalars as well. Once the primitive quantities are known in the global frame, the stress-energy tensor can be directly computed in the global frame.

For clarity, we list the various quantities in their corresponding available coordinate basis in Table I.

4. Inter-Patch Interpolation and Coordinate Transformation

Data in the ghost zones of a given grid patch are exchanged via high-order Lagrange polynomial interpolation for those quantities that are smooth (such as the curvature evolution variables), and optionally second-order essentially non-oscillatory (ENO) interpolation [110] for those variables that may contain discontinuities (such as the hydrodynamic evolution variables). The scheme is depicted in Fig. 2. Ghost points (indicated by empty boxes) on some patch p must be interpolated from points from a neighboring overlapping patch q . The inter-patch boundary is indicated by a vertical line. For Lagrange interpolation, in order to maintain maximal accuracy, we




FIG. 2. Depiction of the second-order ENO inter-patch interpolation scheme used for the fluid variables between two overlapping patches p and q . The inter-patch boundary is indicated by the vertical line. Each interpolated point in the ghost zones (empty boxes) is obtained from an interpolation polynomial whose stencil is selected based on the local smoothness of the interpolated quantity. There are three possible choices: left (L) stencil using blue and green points, right (R) stencil using green and red points, and first-order (f) stencil using only green points. Since none of the stencil points on p are allowed to be inter-patch boundary points of p , we need to introduce a certain number of *additional* overlap points (filled boxes) to ensure that this is the case.

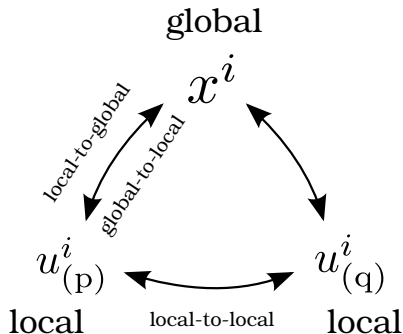


FIG. 3. Coordinate systems and their transformations. Local coordinates $u^i_{(p)}$ and $u^i_{(q)}$ of patches p and q , respectively, are related via “local-to-local” transformations. “Local-to-local” transforms are necessary for fluid variable inter-patch interpolation. The global Cartesian coordinates x^i are used to represent the curvature variables and to carry out any analysis on the curvature or fluid variables, such as gravitational-wave extraction, or fluid density oscillation mode analysis. Therefore, “global-to-local” and “local-to-global” transforms are necessary.

center the interpolation stencils around the interpolation point. The ENO operator, on the other hand, is allowed to use second-order off-centered Lagrange interpolation stencils according to the local smoothness of the interpolated fields [110]. In addition, we check if the interpolant introduces a local maximum and switch to first order in that case. In order to speed up the computation, we precompute and store all possible stencil configurations for each inter-patch ghost point.

To yield a consistent boundary treatment, we have to ensure that an interpolation stencil does not contain any ghost points from the source patch. For this to be the case, we need to introduce *additional overlap* points (indicated by colored boxes in Fig. 2) that lead to an over-

lap of the evolved region. Effectively, this means that the equations are solved twice in the additional overlap region, which introduces a small computational overhead.

We note that quantities which are defined in the global Cartesian tensor basis such as the curvature evolution variables (8) do not need to be transformed between local coordinates patches. In our present hydrodynamics scheme, however, the evolved conserved variables are defined in local coordinates. Hence, for inter-patch ghost zone interpolation, they must be transformed between local coordinate systems. Let us denote the local coordinates of source patch p as $u^i_{(p)}$, and the local coordinates of target patch q as coordinates $u^i_{(q)}$. In the Valencia formulation of the hydrodynamic evolution equations (4), the conserved density D is a densitized scalar: it transforms as a scalar *tensor density* of tensor weight $+1$. Accordingly, the “local-to-local” transformation between local coordinates $u_{(p)}$ of patch p and local coordinates $u_{(q)}$ of patch q is given by

$$D_{(q)} = \left| \det \frac{\partial u^i_{(p)}}{\partial u^j_{(q)}} \right| D_{(p)}. \quad (21)$$

Hence, after having obtained the interpolated value of D in the “old” basis defined by the local coordinates of patch p , we need to represent it in the “new” basis defined by the local coordinates of patch q according to transformation (21), before we assign its transformed value to one of the ghost points of q . Similarly, we also need to transform the conserved 3-momentum S^i , which transforms as a densitized contravariant vector of tensor weight $+1$ according to

$$S^j_{(q)} = \left| \det \frac{\partial u^k_{(p)}}{\partial u^l_{(q)}} \right| \frac{\partial u^j_{(q)}}{\partial u^i_{(p)}} S^i_{(p)}. \quad (22)$$

The various coordinate transformations that are required in our code are depicted in Fig. 3.

D. Cell-centered AMR and Refluxing

We introduce cell-centered AMR in combination with a *refluxing* scheme at refinement level boundaries to ensure conservation of rest mass and – in the absence of GR effects – also momentum and energy of the fluid [84, 111]. Because gravity leads to sources and sinks for fluid momentum and energy, these quantities are generally not conserved in curved spacetimes. This is reflected in the source terms of the fluid conservation laws (5), which are zero only in flat space. The numerical fluxes in our finite volume scheme between grid cells, however, must be conserved. Since we employ subcycling in time where finer grids take multiple small time steps for each coarse grid time step [81], the conservation properties of our finite volume approach do not hold at mesh refinement boundaries without refluxing.




FIG. 4. Vertex-centered AMR (upper figure) versus cell-centered AMR (lower figure). In cell-centered AMR, two fine grid cell faces always coincide with a coarse grid cell face (red line). Thus, it becomes possible to sum up the two fine grid fluxes computed on cell faces to become one coarse grid cell flux. Cell-centered quantities, however, always need to be interpolated in the prolongation and restriction operation. In the vertex-centered case, every second grid point coincides with one coarse grid point. Thus, interpolation is not necessary for every point, and restriction becomes exact.

In cell-centered AMR schemes, coarse cells are subdivided into multiple smaller cells, ensuring that coarse grid and fine grid cell faces align (see red line in the lower part of Fig. 4). In contrast, the cell centers do not align. This is different from vertex-centered AMR schemes, where one aligns coarse and fine grid cell centers but not their faces (red line in the upper part of Fig. 4).

One may argue that vertex-centered schemes are more natural for wave-type equations such as the Einstein equations, which is why vertex-centered refinement was originally implemented in the Carpet AMR driver. However, refluxing requires cell-centered refinement, and this comes with a certain added complexity that we describe below.

a. Prolongation. Prolongation is the interpolation from coarse to fine-grid cells. In a vertex-centered scheme (and when assuming a refinement factor of two), every second fine-grid point is aligned with a coarse-grid point, and prolongation there corresponds to a copy. In between coarse-grid points, one needs to interpolate. Curvature quantities are interpolated via a fifth-order Lagrange polynomial. Hydrodynamics quantities are interpolated via a second-order ENO interpolator [110] (also see Sec. II C 3) to avoid oscillations near discontinuities.

In a cell-centered scheme, *every* fine-grid cell requires interpolation. We interpolate curvature quantities via a fourth-order Lagrange polynomial, and interpolate hydrodynamics quantities via a second-order ENO interpolator.

b. Restriction. Restriction transfers fine-grid information to the next coarser grid via a restriction operator \mathcal{R} , after both have been evolved in time, and are aligned in time again. Different discretization errors will have led to slightly different results, and one overwrites the

coarse-grid results by respective fine-grid results. For a vertex-centered scheme, the restriction operator \mathcal{R} is trivial, since each coarse-grid point is aligned with a fine-grid point, and hence the variable on the fine-grid point can simply be copied.

For cell-centered schemes, the restriction operator \mathcal{R} is non-trivial: it involves interpolation. We interpolate curvature quantities via a third-order Lagrange polynomial. Hydrodynamics quantities are averaged, corresponding to linear interpolation. This is a conservative operation, so that e.g. the mass in a coarse-grid cell is the sum of the masses in all contained fine-grid cells.

The distinction between curvature and hydrodynamics quantities is crucial to achieving high accuracy. If one does not use higher-order operations for the curvature quantities, then the accuracy of the overall simulation is significantly reduced. On the other hand, one needs to employ a conservative interpolation scheme for the hydrodynamics quantities, but can accept a lower order of accuracy there. For restricting curvature quantities, we therefore use third-order polynomial interpolation.

c. Refluxing. Refluxing is an algorithm to ensure flux conservation across mesh refinement boundaries [85, 111]. Since coarse and fine grids are evolved in time independently, it is not guaranteed that the fluxes leaving the fine grid are identical to those entering an abutting coarser grid. Fig. 5 illustrates the situation: two fine-grid steps of stepsize Δt^{l+1} are taken for one coarse grid step of stepsize Δt^l . After coarse and fine grids have been *independently* evolved, the fine-grid fluid state is restricted on the coarse grid when both grids are aligned in time again, as described above. Now, the fluid state of a coarse grid cell at the boundary (blue cell in Fig. 5) was updated from numerical fluxes between the *original*, i.e., non-restricted neighboring coarse-grid cells (flux through red line in Fig. 5). The restricted fine-grid fluid state, however, was computed from numerical fluxes between neighboring fine-grid cells (fluxes through green line in Fig. 5), which is not guaranteed to coincide with the original, non-restricted coarse-grid flux. The true numerical flux through the coarse/fine grid boundary of a coarse cell, however, is given by the sum of all fine grid fluxes through that boundary. The fluid state of a coarse-grid cell at a mesh-refinement boundary therefore needs to be corrected by the difference between the sum of the underlying fine-grid fluxes and the original, non-restricted coarse-grid flux. Only then it is guaranteed that the fluid state of a coarse-grid cell at a mesh-refinement boundary is a result of a conservative operation.

Refluxing integrates the coarse-grid and fine-grid fluxes across these faces, and then adjusts the coarse grid cell just outside the refined region according to the flux difference.

We outline the generic refluxing algorithm for a conserved quantity f in the steps below.

1. We start with a fine grid level $l + 1$ and a coarse grid level l which are momentarily aligned in time, i.e. $t_i^l = t_{2j}^{l+1}$, where i denotes the i -th step on the




FIG. 5. AMR time evolution, showing fluxes across cell faces, for both coarse (upper row) and fine cells (lower row). Time moves upwards; the fine grid (lower panel) takes multiple steps for each coarse grid step (upper panel). In Berger-Oliger AMR, the coarse and fine levels are evolved independently, and the sum of the fine grid fluxes crossing the green faces are not guaranteed to be equal to the coarse grid flux crossing the red face. At the end of a time step, the neighboring bold-faced coarse and fine cells may be in an inconsistent state, requiring *refluxing* to add a correction to the light blue coarse grid cell.

coarse level, and j denotes the j -th step on the fine grid. Due to subcycling in time, for any coarse-grid time step, there are twice as many fine-grid time steps, i.e. $i = 2j$.

2. At the refinement boundary (red line of Fig. 4, or red and green lines in Fig. 5), we store integrated coarse and fine grid *flux registers*³ I^l and I^{l+1} for some conserved quantity f . Due to the 2:1 mesh refinement, there are four integrated fine grid flux registers for every integrated coarse grid flux register. (Only two are visible in Fig. 4). At $t_i^l = t_{2j}^{l+1}$, all registers are zero.
3. Each refinement level is independently integrated forward in time until the two refinement levels are aligned in time again, i.e., until we have $t_{i+1}^l = t_{2j+2}^{l+1}$. During each integration step, the hydrodynamic evolution scheme computes fluxes F for a quantity f located at all cell interfaces. At the refinement boundary, we use the computed fine grid fluxes F^{l+1} on the fine grid cell interfaces, and

coarse grid fluxes F^l on the coarse grid cell interfaces to integrate coarse and fine grid flux registers forward in time, i.e., we independently integrate

$$\partial_t I^{l+1} = F^{l+1}, \quad \partial_t I^l = F^l, \quad (23)$$

at the refinement boundary.

4. After restriction, when $t_{i+1}^l = t_{2j+2}^{l+1}$, we use I^{l+1} and I^l to compute a correction for conserved quantity f . The correction is obtained as follows.

- (a) The integrated fine grid flux register I^{l+1} is restricted to the coarse grid via

$$I_{\text{fine}}^l = \mathcal{R}_{\mathcal{I}} I^{l+1}, \quad (24)$$

where $\mathcal{R}_{\mathcal{I}}$ denotes the *cell interface* restriction operator. Note that since the flux registers are stored on cell *faces*, this operator is different from the restriction operator \mathcal{R} used for the fluid state vector.

- (b) A correction C_f^l for conserved quantity f on coarse grid level l is now obtained via

$$C_f^l = (I_{\text{fine}}^l - I^l) / \Delta^l x \quad (25)$$

where $\Delta^l x$ denotes the grid spacing of refinement level l .

5. The correction C_f^l is added to the coarse grid cell on level l next to the refinement boundary (blue cell in Fig. 5), i.e.

$$f_{\text{corrected}}^l = f^l + C_f^l. \quad (26)$$

This completes the refluxing operation. We repeat the steps 1-5 until the evolution is complete.

The steps above are performed for any of the evolved conserved quantities D , S^i , τ , and Y_e^{con} .

We note that the state thus obtained in the corrected coarse grid cells may be thermodynamically inconsistent in the sense that a solution for the primitive variables can not be found (even with bisection as mentioned in Sec. II A). We have found this behavior at the very steep contact discontinuity near the surface of a neutron star, which is expected with our atmosphere treatment, as we discuss in Appendix C. In that case, we may not apply refluxing when cells at atmosphere values are involved⁴.

E. Time Integration and Multirate Runge-Kutta Schemes

We carry out time integration using the Method of Lines (MoL) [99]. MoL is based on a separate treatment

³ A flux register is a grid variable which stores the integrated flux on cell interfaces which are on a mesh-refinement boundary.

⁴ Note that the atmosphere treatment is not conservative anyway.

of the spatial derivatives (the right-hand sides), and the time derivatives. This allows one to employ integration methods for ordinary differential equations (ODE) such as Runge-Kutta (RK) schemes for the time integration.

We evolve the spacetime and hydrodynamic sector of our evolution system simultaneously using full matter-spacetime coupling. The coupling between the two sectors is achieved via source terms. The spacetime evolution is sourced by the stress-energy tensor computed by the hydrodynamic sector. Vice versa, the hydrodynamic part contains additional source terms which are a result of the coupling to a curved spacetime metric. Written in simplified form, our system is given by

$$\partial_t \mathbf{g} = \mathbf{F}(\mathbf{g}, \mathbf{q}), \quad (27)$$

$$\partial_t \mathbf{q} = \mathbf{G}(\mathbf{g}, \mathbf{q}), \quad (28)$$

where \mathbf{g} denotes curvature evolution quantities, \mathbf{q} denotes fluid evolution quantities, and \mathbf{F} and \mathbf{G} denote the RHS functions.

Traditionally, spacetime metric and hydrodynamic variables are evolved simultaneously using the same time integration scheme. A standard choice in our case is the classical fourth-order Runge-Kutta (RK4) method⁵. The timestep is chosen such that the Courant-Friedrich-Lewy (CFL) factor, defined as $C = \Delta t / \Delta x$, becomes $C = 0.4$. The CFL factor is limited by the stability region of the numerical scheme, which in turn is limited by the speed of light.

We observe two important points in our simulations. First, the error in our numerical evolution is in most cases not dominated by the time integration (see Sec. III). The choice of Δt is not guided by accuracy requirements, but rather by the restrictions imposed by the CFL condition. This is unfortunate since a larger timestep would speed up our simulation with only small negative impact on the accuracy. Second, we find that the CFL factor is largely determined by the spacetime evolution. In the Cowling approximation, i.e. when the spacetime sector is not evolved and held fixed at its initial setup, we typically can use more than twice as large CFL factors (up to $C \approx 1$) without encountering any numerical instabilities.

Since our timestep is fixed, rather than enlarging the timestep Δt (and hence C), we switch to the classical second-order Runge-Kutta (RK2) method instead. This scheme has a smaller stability region by roughly a factor of two compared to RK4. Due to the less restrictive CFL factor for the fluid evolution compared to the curvature evolution, however, we can still use the same timestep as for the curvature evolution with RK4. The advantage of the RK2 schemes is that they require half as many RHS evaluations compared to RK4. The accuracy of RK2,

TABLE II. Butcher tableau for an explicit multirate RK4/RK2 scheme. The right table (separated by the double vertical line) shows the coefficients b_i (bottom line), c_i (first vertical column), and a_{ij} for the classical RK4 scheme. The left table shows the corresponding RK2 coefficients evaluated at timesteps that coincide with RK4 timesteps.

0					0				
0	0				1/2	1/2			
0	0	0			1/2	0	1/2		
1	1	0	0		1	0	0	1/2	
	1/2	0	0	1/2		1/3	1/6	1/6	1/3

however, is typically much lower than that of an RK4 integration. In practice, we find that the reduction in accuracy is not a severe limitation for most cases (see Sec. III).

We therefore apply the RK2 integrator for the hydrodynamic sector, while maintaining the RK4 integrator for the spacetime part.

A scheme for coupling different parts of a system of equations with different RK integrators is given by *multirate* RK schemes (e.g. [88, 89]). Here, we make the simple Ansatz of performing one RK2 intermediate RHS evaluation for two RK4 intermediate RHS evaluations. That is, the additional RK4 intermediate RHS evaluations simply use the results from the last intermediate RK2 step.

To be more explicit, given the equation

$$\partial_t y = f(t, y), \quad (29)$$

where f corresponds to the RHS, we write a generic RK scheme according to

$$y_{n+1} = y_n + \Delta t \sum_{i=1}^s b_i k_i, \quad (30)$$

$$k_i = f(t_n + c_i \Delta t, y_n + \Delta t \sum_{j=1}^s a_{ij} k_j). \quad (31)$$

The coefficients b_i , c_i , and a_{ij} can be written in the standard Butcher notation (see, e.g. [112]).

In our multirate scheme, we use two different sets of coefficients. The coefficients for the RK2 scheme are arranged such that RHS evaluations coincide with RK4 RHS evaluations. We list the corresponding multirate Butcher tableau in Table II.

F. Gravitational Wave Extraction

GWs are extracted in the *wave-extraction zone* of our simulation. We define the wave-extraction zone as the region on the computational grid which is at sufficient distance from the gravitating source to avoid near-zone effects, and at the same time offers sufficient resolution to resolve the waves. Beyond the wave-extraction zone,

⁵ The classical RK4 does not have the total variation diminishing (TVD) property. Strictly, this property is necessary to avoid artificial oscillations at shocks. In practice, we find that this is not an issue.

we typically use radial stretching to gradually decrease the radial resolution up to a certain radius (e.g., Fig. 1).

We use the techniques described in detail in [21]. Among those are (i) the standard slow-motion weak-field quadrupole formalism (see e.g., [15, 16, 29, 113–115]) which is purely based on the quadrupolar matter distribution and does not take into account any curvature effects, (ii) Regge-Wheeler-Zerilli-Moncrief (RWZM) extraction based on gauge-invariant spherical perturbations about a fixed Schwarzschild background (see [116] for a review), (iii) Newman-Penrose extraction based on complex spin-weighted components of the Weyl tensor [73, 117, 118], and (iv) Cauchy-characteristic extraction (CCE) [19–24] making use of nonlinear nullcone evolutions of the Einstein equations out to future null infinity \mathcal{J}^+ (see [119] for a new high-order algorithm). Extraction technique (iv) is the only one capable of determining the gravitational radiation content unambiguously and without finite-radius and gauge errors [21–24].

The curvature-based techniques (ii)-(iv) require one or two integrations in time in order to compute the strain, which may lead to strong non-linear and unphysical artificial drifts. This can be overcome by the *fixed frequency integration* (FFI) technique presented in [120]. FFI requires the choice of a cut-off frequency f_0 , which ideally must be below the physical frequency components contained in the signal. For instance, for a typical binary neutron star inspiral signal, $f_0^m < m\Omega_{\text{orbital}}/2\pi$, where Ω_{orbital} is the initial orbital frequency, and m is the associated harmonic m -mode number.

The energy and angular momentum that is lost due to the emission of GWs can be computed in terms of spin-weighted spherical harmonic coefficients of Ψ_4 as derived in [121, 122]. We use the expressions for the radiated energy flux dE_{rad}/dt and angular momentum flux dJ_{rad}/dt in terms of the Weyl scalar Ψ_4 from [122]. In the expressions for dE_{rad}/dt and dJ_{rad}/dt , we evaluate the appearing time integrals of the harmonic modes using FFI with $f_0^m = mf_0$ for each given m -mode. In order to obtain the total radiated energy E_{rad} and angular momentum J_{rad} from their fluxes, respectively, we time integrate in the time domain⁶.

1. Numerical Setup

We report the numerical settings employed for the various wave extraction techniques that are used in this work. Since we are not interested in the numerical convergence properties of the wave extraction methods themselves (this has been analyzed elsewhere, e.g. [21–23, 50, 73, 74]), we stick to fixed settings for all test cases and numerical resolutions considered in Sec. III. Guided

by previous work [21, 23], we find that the numerical error in the wave extraction is negligible provided appropriate settings.

The most involved GW extraction technique is CCE. In that method, we solve the Einstein equations along null hypersurfaces between a worldtube Γ and future null infinity \mathcal{J}^+ . The worldtube Γ is typically located at some radius R_Γ in the wave-extraction zone, and is simulation dependent [23] (and references therein). Specific to the present work, the CCE grid consists of $N_r = 301$ points along the radial direction. Each radial shell is discretized by two stereographic patches comprised of $N_{\text{ang}} = 81$ points per direction per patch. At the inner-boundary worldtube Γ , we use up to $\ell_{\text{max}} = 8$ harmonic modes for the decomposed Cauchy metric data. The metric data is decomposed on spheres with $N_\theta = 120$ and $N_\phi = 240$ points in θ and ϕ direction, respectively. The compactification parameter⁷ r_{wt} is set to the particular extraction radius for a given simulation, e.g. $r_{\text{wt}} = 100M_\odot$. In all cases, the innermost radial compactified coordinate point is given by $x_{\text{in}} = 0.49$. Together with an appropriate setting of r_{wt} , this ensures that the worldtube Γ is located close to the first few radial points on the characteristic grid. The timestep and extraction radius must be picked on a case by case basis. The wave-extraction zone is always located on the spherical “inflated-cube” grids. For the stellar collapse model A3B3G3 (Sec. III B), the wave-extraction zone is located between radii $1000 M_\odot < R_\Gamma < 2500 M_\odot$. For all remaining tests, the wave-extraction zone is located at $100 M_\odot < R_\Gamma < 250 M_\odot$. The wave-extraction output frequency is dictated by the timestep of the spherical “inflated-cube” grids.

The remaining wave-extraction techniques are much simpler and only require single spheres at some finite radius R .

To project metric data from the 3D grid onto spheres, we use fourth-order Lagrange interpolation.

G. Horizon Finding and Hydrodynamic Excision at the Puncture

To track the appearance and shape of an apparent horizon, we use `AHFinderDirect` [49, 123] which is part of the `EinsteinToolkit` [83]. As soon as an apparent horizon is found during an evolution, we excise the fluid variables within a fraction of the radius of the apparent horizon and set them to their corresponding atmosphere values. We get stable evolutions when excising about 85% of the interior of the apparent horizon volume.

In order to compute angular momentum J_{AH} and mass M_{AH} of a black hole, we use the isolated / dynamical horizon framework provided by `QuasiLocalMeasures` [124],

⁶ FFI cannot be applied since the radiated fluxes are non-oscillatory.

⁷ See [23] for a description of CCE relevant parameters.

TABLE III. Initial parameters and properties of the (perturbed) TOV star used to construct the initial data. The density perturbation is only applied in the perturbed TOV test case. Units are in $c = G = M_\odot = 1$, unless otherwise specified.

Polytropic scale	K	100
Polytropic index	Γ	2
Central rest-mass density [M_\odot^{-2}]	ρ_c	1.28×10^{-3}
ADM mass [M_\odot]	M_{ADM}	1.4002
Baryonic mass [M_\odot]	M_B	1.5062
Equatorial radius [M_\odot] ([km])	R_e	9.586 (14.16)
Density pert. mode	ℓ	2
Density pert. amplitude	λ	0.01
Monopole fundamental mode [kHz]	F	1.458
First overtone [kHz]	H_1	3.971
Quadrupole fundamental mode [kHz]	2f	1.586
First overtone [kHz]	2p_1	3.726

which is part of the `EinsteinToolkit`. This framework defines mass and angular momentum in terms of particular closed 2-surfaces, such as the apparent horizon.

The spherical surface defining the apparent horizon shape uses $N_\theta = 41$ points along the θ -direction and $N_\phi = 80$ points along the ϕ -direction.

III. RESULTS

We revisit a number of “benchmark“ problems commonly found in the literature: an isolated perturbed and unperturbed neutron star, a rotating core collapse model, a collapsing neutron star to a black hole, and a binary neutron star coalescence. Basic code tests such as shock tubes can be found in the Appendix. We describe our analysis in more detail in corresponding sections below.

A. Isolated Neutron Star

We investigate convergence and accuracy of an isolated unperturbed neutron star and an isolated perturbed neutron star using full GR matter-spacetime coupling in three spatial dimensions. The neutron stars are given by the solution of the Tolman-Oppenheimer-Volkoff (TOV) equations [125, 126].

This test aims at showing the correctness of our cell-centered AMR scheme, and enhanced PPM reconstruction.

1. Initial Conditions and Equation of State

We use a polytropic equation of state $P = K\rho^\Gamma$ with scale $K = 100$ and index $\Gamma = 2$ in the initial data construction. Although this choice does not represent a realistic choice for real neutron stars, these parameters have




FIG. 6. Unperturbed TOV star: normalized central density $\rho_c(t)/\rho_c(t=0) - 1$ on the three resolutions r_0 , r_1 , and r_2 (top panel), difference in normalized central density between low and medium resolutions, and medium and high resolution (center panel), and the L_2 -norm of the Hamiltonian constraint $\|H\|_2$ on all three resolutions (bottom panel). As the resolution is increased, the amplitude of the central density oscillations, the offset, and the slope decrease as expected. The differences in resolutions of the central density are scaled for second-order convergence. The L_2 -norms of the Hamiltonian constraint $\|H\|_2$ are scaled for first-order convergence. The resolution study is performed using cell-centered AMR and ePPM.




FIG. 7. Unperturbed TOV star: the L_2 -norm of the Hamiltonian constraint $\|H\|_2$ (upper panel), and conservation of total baryonic mass M_B (lower panel) for different numerical setups. We compare vertex-centered (vc) with cell-centered (cc) AMR using oPPM and/or ePPM. In addition, we also show a simulation with “ePPM, cc” using multirate time integration. $\|H\|_2$ is strongly affected by the choice of numerical scheme, while M_B is essentially unaffected. The setup “ePPM, cc” performs best, while “oPPM, cc” performs worst. The standard scheme “vc, oPPM” used in other codes (e.g. [83, 92]) is slightly worse than the new scheme “ePPM, cc”. Multirate time integration leads to nearly identical results.

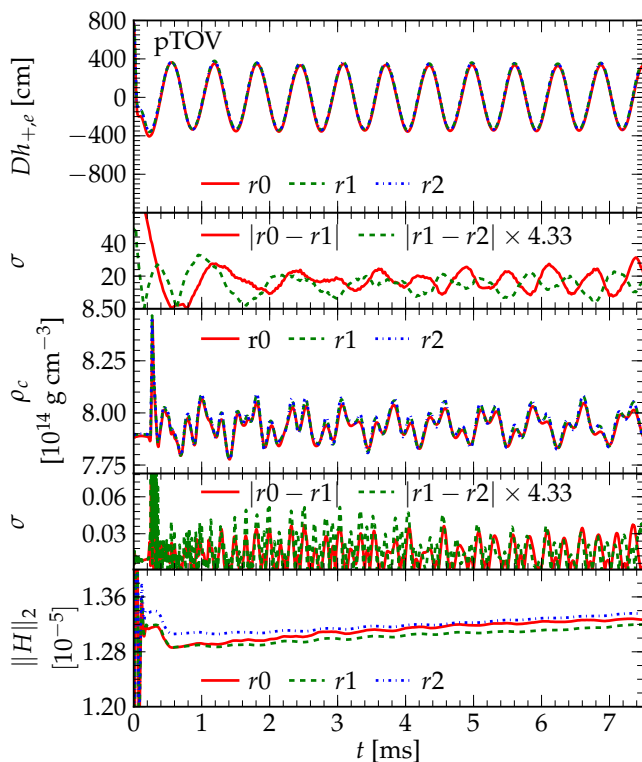


FIG. 8. Perturbed TOV star: the top panel shows the “+” polarization of the GW strain $Dh_{+,e}$ as emitted in the equatorial plane and rescaled by distance D for the three resolutions r_0 , r_1 , and r_2 . The waveforms are computed with CCE. In the panel below, we show the differences in GW strain between r_0 and r_1 , and r_1 and r_2 , where the latter is rescaled for second-order convergence. In the third panel from the top, we show the absolute central density evolution $\rho_c(t)$ for the three resolutions. Below, we show the differences in central density scaled for second-order convergence. In the bottom panel, we show the L_2 -norms of the Hamiltonian constraint $\|H\|_2$. Since the initial data for the perturbed case are not constraint satisfying, the constraints do not exhibit clean convergence. The convergence study is performed using cell-centered AMR and ePPM.

been used in previous work (e.g. [36, 127]), and can be used as code verification. During evolution, we use an ideal fluid Γ -law equation of state with $\Gamma = 2$. The key parameters are given in Table III. The initial data are generated via Hachisu’s self-consistent field method [128, 129] which requires as input the central density ρ_c of the star, and a polar-to-equatorial axes ratio between 0 and 1 to define rotation. In the present case, we set $\rho_c = 1.28 \times 10^{-3} M_\odot^{-2}$ and use an axes ratio of 1 (no rotation). In the case of the perturbed TOV star, we perturb the star by a spherical harmonic $(\ell, m) = (2, 0)$ density perturbation of amplitude $\lambda = 0.01$.




FIG. 9. Perturbed TOV star: the impact of different numerical settings on the L_2 -norm of the Hamiltonian constraints $\|H\|_2$ (upper panel), and on the conservation of baryonic mass M_B (lower panel). The setup using vertex-centered (vc) AMR and oPPM (blue dashed curve) leads to larger constraint violations than the setup using cell-centered (cc) AMR and ePPM. Multirate time integration does not change the accuracy of the results. In all cases, M_B is nearly equally well conserved.




FIG. 10. Perturbed TOV star: Power spectrum of ρ_{00} , ρ_{20} and $h_{+,e}$ (individually scaled for better visibility), and the first few fundamental neutron star oscillation modes (vertical lines) computed in [36].

2. Numerical Setup

The grid is similar to the one depicted in Fig. 1, except that here, we have just one refinement region. The fine grid spacing is $\Delta x = 0.2 M_\odot$ for the low resolution (r_0), $\Delta x = 0.125 M_\odot$ for the medium resolution (r_1), and $\Delta x = 0.1 M_\odot$ for the high resolution simulation (r_2). The fine grid extends to $R = 11 M_\odot$ and encompasses the entire star. The inter-patch boundary between central Cartesian patch and outer spherical grid is located at $R_S = 65 M_\odot$. We use 15, 24 and 30 cells per angular direction per spherical patch for the low, medium and high resolutions, respectively. The radial resolution is chosen based on the Cartesian

coarse grid resolution $\Delta r = 1.6 M_\odot$, $\Delta r = 1.0 M_\odot$, and $\Delta r = 0.8 M_\odot$, for low, medium, and high resolutions, respectively. We use radial stretching outside the wave extraction zone to efficiently extend the computational domain so that the outer boundary is causally disconnected from wave-extraction zone and interior evolution. Accordingly, we stretch the radial resolution to $\Delta r = 6.4 M_\odot$, $\Delta r = 4.0 M_\odot$, and $\Delta r = 3.2 M_\odot$ for low, medium, and high resolution simulations, respectively, in the region between radii $R_1 = 100 M_\odot$ and $R_2 = 800 M_\odot$. The outer boundary is located at $R_B = 3500 M_\odot$.

3. Discussion

a. Unperturbed TOV star. We first consider a single isolated non-rotating, unperturbed TOV star with parameters reported in Table III. In the top panel of Fig. 6, we show the normalized central density evolution $\rho_c(t)/\rho_c(t=0)$ as a function of time on the three resolutions $r0$, $r1$, and $r2$, using our new cell-centered AMR and enhanced PPM scheme. In an ideal setting, the central density evolution should be constant as a function of time since the TOV solution represents a static fluid configuration. Numerical errors induced by interpolation from the initial data solver grid onto the evolution grid, however, lead to an artificial excitation of the star, and, hence, to non-trivial central density oscillations, which must converge to zero as the resolution is increased. Due to the interpolation of the fluid initial data onto the evolution grid, we observe a large initial spike and an overall offset in the density oscillations. We additionally see an overall non-zero slope in the central density evolution caused by numerical errors during evolution. As the resolution is increased, we consistently observe that the amplitudes of the oscillations decrease, the offset becomes smaller, and the overall slope is reduced. In the center panel, we show the difference in normalized central density $\rho_c(t)/\rho_c(t=0)$ between resolutions $r0$ and $r1$, and $r1$ and $r2$. We perform a three-level convergence test by computing the ratio of the differences in a given quantity F between the three resolutions,

$$C = \frac{|F^{\text{medium}} - F^{\text{low}}|}{|F^{\text{high}} - F^{\text{medium}}|}. \quad (32)$$

The ratio C defines the *measured* convergence rate of the solution (e.g. [104]). Given three resolutions with spacing Δx_{low} , Δx_{medium} , and Δx_{high} , the *theoretical* convergence rate for a particular order of convergence p can be computed via

$$C = \frac{|\Delta x_{\text{medium}}^p - \Delta x_{\text{low}}^p|}{|\Delta x_{\text{high}}^p - \Delta x_{\text{medium}}^p|}. \quad (33)$$

Given our numerical resolutions, according to (33), we expect that the difference between medium and high resolution, $r1$ and $r2$, decreases by a factor of $C = 4.33$ for

second-order convergence compared with the difference between medium and low resolution, $r1$ and $r0$.

In the bottom panel of Fig. 6, we show the time evolutions of the L_2 -norm of the Hamiltonian constraint $\|H(t)\|_2$ (16) for the three resolutions $r0$, $r1$, and $r2$. As the resolution is increased, the error drops consistent with first-order convergence, since the rescaled medium and high resolution curves are on top of each other. We note that while the fluid body itself is smooth, the surface of the star is non-smooth, hence inducing a dominant first-order error (compare Fig. 25).

In the top panel of Fig. 7, we show the L_2 -norm of the Hamiltonian constraint $\|H\|_2$ of a static TOV star using vertex-centered (vc) AMR, and cell-centered (cc) AMR. Both AMR setups are run with the oPPM and ePPM reconstruction method. In addition, we also perform a simulation using cell-centered AMR and ePPM reconstruction with multirate time integration. We observe that the setup "ePPM, cc" exhibits the lowest constraint violations. The setup "ePPM, cc, multirate" is right on top of the red curve, hence indicating comparable accuracy. The setup "vc, oPPM", which is the setup used in previous work (e.g. [1, 2, 21, 83, 92]) yields slightly less accurate evolution. Finally, the setup "cc, oPPM" yields significantly reduced accuracy compared to all other setups. This is mainly due to the oPPM scheme, which is known to reduce the order of accuracy at smooth maxima to first order (see Appendix B, Fig. 25). This effect is not seen in the vertex-centered setup "vc, oPPM", since the central density is exactly located on a grid point.

In the bottom panel of Fig. 7, we show conservation of mass for the considered numerical setups. In all cases, the total mass loss is on the order of 10^{-7} over the course of the evolution. Since the AMR boundaries are all located in the vacuum region outside the star, refluxing at AMR boundaries is not relevant. The mass loss is entirely due to interaction with the artificial low-density atmosphere in the vacuum region (see also Appendix C).

b. Perturbed TOV star. As a second test, we apply an initial $(\ell, m) = (2, 0)$ density perturbation with amplitude $\lambda = 0.01$ onto the same TOV star considered above. A more complete study of this configuration including variations on perturbation parameters has been performed in [36, 127]. Numerical grids and setups are identical to those of the static TOV star, and we perform the same analysis as above. In addition, we also analyze the non-trivial $(\ell, m) = (2, 0)$ mode of the GW signal that is induced by fundamental mode oscillations. In the upper panel of Fig. 8, we plot the "+" polarization of the GW signal $Dh_{+,e}$ as emitted in the equatorial plane from the three resolutions $r0$, $r1$, and $r2$. Since only the $(\ell, m) = (2, 0)$ mode is excited, the entire wave signal can be written as

$$Dh_{+,e} = Dh_+^{20} {}_{-2}Y_{20}(\theta = \frac{\pi}{2}, \phi = 0). \quad (34)$$

Here, D is the distance from the source. We compute h_+^{20} with CCE and use an FFI cut-off frequency of $f_0 = 812$

Hz (see Sec. IIF). We also show the differences of the GW strain between low and medium, and medium and high resolutions, where the latter is scaled for second-order convergence. In addition, we show the central density evolution $\rho_c(t)$ for the three resolutions which converge. Similarly to the above, we plot the differences between low and medium, and medium and high resolutions scaled for second-order convergence. We also show the L_2 -norm of the Hamiltonian constraints $\|H\|_2$ of the three resolutions. Since the initial data solver does not take into account the effects of the perturbation onto the initial spacetime metric, the constraints do not converge initially, and only slowly converge at later times. In the present plot, we have not used any rescaling. We note, however, that the slopes of the medium and high resolutions are slightly smaller than for the low resolution case.

When comparing the strain $Dh_{+,e}^{\text{CCE}}$ as computed with CCE to the strain $Dh_{+,e}^{\text{Q}}$ as computed from the RWZM formalism, we generally find that the strain computed via the RWZM formalism is prone to numerical noise. In addition, we find that the finite-radius error and gauge error inherent in the waveform obtained from RWZM master functions at radii $R = 100 M_\odot$ and $R = 250 M_\odot$ is on the order of 10%. A similar behavior applies to the strain $Dh_{+,e}^{\text{NP}}$ as extracted via the NP formalism at a finite radius.

Finally, we also check that the correct fundamental oscillation modes are excited. In Fig. 10, we compare the frequency spectrum of the density ρ and the strain $Dh_{+,e}$ to the eigenmodes found in [36]. In order to compute the spectrum of ρ , we first project ρ from the 3D grid onto spherical shells inside the star, and then decompose in terms of spherical harmonics. The vertical lines in Fig. 10 correspond to the fundamental monopole mode F and its first overtone H_1 , and the fundamental quadrupole mode 2f and its first overtone 2p_1 . As expected, the spectrum of the strain $Dh_{+,e}$ and the $(\ell, m) = (2, 0)$ mode of the density ρ_{20} both peak at the correct quadrupole eigenmode frequencies. Likewise, the spectrum of the $(\ell, m) = (0, 0)$ density mode correctly peaks at the monopole eigenmode frequencies.

B. Rotating Stellar Collapse

We investigate convergence and accuracy of the benchmark rotating stellar collapse model A3B3G3, which has been previously considered in the literature [28, 29]. This tests the ability of the code to simulate the collapse of a rapidly differentially spinning iron core in full 3D with causally disconnected outer boundaries, albeit with simplified microphysics. We show that due to larger wave extraction radii, the waveforms extracted via curvature-based methods such as CCE are more accurate than what has been computed before [21].




FIG. 11. Stellar collapse: The GW strain $Dh_{+,e}$ extracted via the quadrupole formula (upper panel), the central density ρ_c (third panel from the top), and the L_2 -norm of Hamiltonian constraint $\|H\|_2$, all on the three resolutions r_0 , r_1 , and r_2 . The panels directly below the top panel and the third panel from the top show the difference in strain and central density between low and medium resolutions, and medium and high resolutions. The differences are scaled for second-order convergence. The L_2 -norm of Hamiltonian constraint is scaled for first-order convergence. Before core bounce, the constraint exhibits second order convergence. After shock formation, the convergence rate is reduced to first order. The convergence study is performed using cell-centered AMR and ePPM.

1. Initial Data and Equation of State

For the purpose of this test, we employ a *hybrid* equation of state [29, 30, 130] that combines a 2-piece piecewise polytropic pressure P_{P} with a thermal component P_{th} , i.e., $P = P_{\text{P}} + P_{\text{th}}$. To model the stiffening of the equation of state at nuclear density $\rho_{\text{nuc}} \cong 2 \times 10^{14} \text{ g cm}^{-3}$, we assume that the polytropic index Γ jumps from Γ_1 below nuclear density to Γ_2 above. The equation of state parameters are given in Table IV.

The initial data are constructed from $n = 3$ ($\Gamma_{1,\text{ini}} = \Gamma_1 = 4/3$) polytropes in rotational equilibrium generated via Hachisu's self-consistent field method [128, 129] which not only provides fluid, but also spacetime curvature initial data. While being set up as marginally stable polytropes with $\Gamma_{1,\text{ini}} = 4/3$, during evolution, the initial sub-nuclear polytropic index Γ_1 is reduced to




FIG. 12. Comparison of vertex-centered (vc) AMR with oPPM versus cell-centered (cc) AMR with ePPM for stellar collapse model A3B3G3. We show the central density ρ_c (upper panel), the L_2 -norm of the Hamiltonian constraint $\|H\|_2$ (middle panel), and conservation of total baryonic mass M_B (bottom panel). Due to refluxing in the cell-centered case, the mass is almost perfectly conserved, while in the vertex-centered case, the mass is rapidly growing (bottom panel). Due to ePPM, the constraints in the cell-centered case exhibit almost no growth after core bounce, while in the vertex-centered case with oPPM the constraints are clearly growing (lower panel). The results are not changed when multirate time integration is used. The comparison is done using baseline resolution $r1$.

TABLE IV. Initial parameters and properties of the rotating stellar collapse model A3B3G3. Units are in $c = G = M_\odot = 1$, unless otherwise specified.

Polytropic scale	K	0.4640517
Initial polytropic index	$\Gamma_{1,\text{ini}}$	1.3
Evolved polytropic index 1	Γ_1	1.31
Evolved polytropic index 2	Γ_2	2.5
Thermal polytropic index	Γ_{th}	1.5
Central rest-mass density [M_\odot^{-2}]	ρ_c	1.6193×10^{-8}
Axes ratio		0.93
Degree of differential rotation [km]	A	500
Rotational / binding energy [%]	$T/ W $	0.9
Equatorial radius [M_\odot]	R_e	1.0661×10^3
Baryonic mass [M_\odot]	M_B	1.4596
ADM mass [M_\odot]	M_{ADM}	1.4596
ADM ang. mom. [M_\odot^2]	J_{ADM}	2.4316
Spin	a	1.1413

$\Gamma_1 < \Gamma_{1,\text{ini}}$ to accelerate collapse. Following previous studies [16, 28, 29], we use $\Gamma_2 = 2.5$ in the super-nuclear regime.

In the present test, we revisit model A3B3G3 from [28, 29]. This configuration uses $\Gamma_1 = 1.31$. It is strongly differentially rotating, with its initial central angular velocity dropping by a factor of two over $A = 500$ km. This,




FIG. 13. The GW strain extracted $Dh_{+,e}$ from the rotating stellar collapse model A3B3G3 (upper panel). We show the strain extracted via CCE from different worldtube locations $R_\Gamma = 1000 M_\odot$, $R_\Gamma = 1500 M_\odot$, and $R_\Gamma = 2500 M_\odot$, as well as the strain computed via the quadrupole formula. Larger CCE worldtube radii permit lower FFI cut-off frequencies without introducing unphysical drifts in the GW strain. All waveforms extracted via CCE are in good agreement to within a few percent with the waveform computed via the quadrupole formula. The lower panel shows the differences in strain amplitude of the inner extraction radii to the outermost extraction radius. The differences converge as the extraction radius is increased. The comparison is done using baseline resolution $r1$.

in combination with $T/|W| = 0.9\%$, leads to rapid rotation in the inner core, resulting in a very strong GW signal at core bounce and dynamics that are significantly affected by centrifugal effects. It produces a “Type-F” GW signal with a centrifugally-widened broad peak at core bounce [28, 29].

2. Numerical Setup

We use five refinement levels located at the center of the domain. The refinement boxes of each level have a half-width of $R_{r1} = [192M_\odot, 144M_\odot, 98M_\odot, 40M_\odot, 12M_\odot]$, respectively. The coarsest level is comprised of cubed-sphere multipatch grids (Fig. 1). The inner radius of the spherical grids is $R_S = 384M_\odot$, and the outer boundary is $R_B = 16000M_\odot$. Initially, only the coarsest level is active. Additional levels are progressively added as the central density increases during collapse. The initial stellar radius of model A3B3G3 is $R_e = 1066.1M_\odot = 1574.84$ km in the equatorial plane. Thus, the inter-patch boundaries thread the star in this particular setup. The finest refinement level is picked such that the protoneutron star is fully contained

on that level. The GW extraction zone extends to a radius of $R = 2500M_\odot$. Beyond that radius, we apply radial stretching up to a radius $R = 6000M_\odot$. In this stretching region, the radial grid spacing is increased by a factor of 16, and the resolution becomes too coarse for reliable wave extraction.

For our baseline resolution (denoted by $r1$), we pick a radial grid spacing of $\Delta r = 8.0M_\odot$ on the non-stretched spherical inflated-cube grids, and a Cartesian resolution of $\Delta x = 8.0M_\odot$ on the central Cartesian patch. Given our five refinement levels above, this results in a resolution of $0.25M_\odot = 369.3\text{m}$ for the protoneutron star. The angular resolution of the cubed-sphere grids is set to $N_{\text{ang}} = 30$ cells per patch and direction. This makes a total of $N_{\text{ang,total}} = 120$ points across the equatorial plane.

In addition to our baseline resolution $r1$, we also consider a low resolution run $r0$, and a high resolution run $r2$ to check for convergence. Resolution $r0$ uses $\Delta r = \Delta x = 9.6M_\odot$ and $N_{\text{ang}} = 24$ (20% lower), and resolution $r2$ uses $\Delta r = \Delta x = 6.4M_\odot$ and $N_{\text{ang}} = 36$ (20% higher).

In all considered cases, we set the damping coefficient of the Γ -driver gauge condition to $\eta = 1/2$. Dissipation is set to $\epsilon_{\text{diss}} = 0.1$ on the fine levels, and $\epsilon_{\text{diss}} = 0.01$ on the multipatch grid. The atmosphere level is set to be 10^{-10} of the central density, and we damp the stress-energy tensor in the atmosphere using (7) with $R_0 = 1300M_\odot$ and $R_1 = 1400M_\odot$.

3. Discussion

In Fig. 11, we show convergence of the plus polarization of the GW strain $Dh_{+,e}$ measured in the equatorial plane, the central density ρ_c , and the L_2 -norm of the Hamiltonian constraint $\|H\|_2$. The GW strain is computed using the quadrupole formula, though a similar analysis and result apply to all extraction methods. All three quantities are shown for the three resolutions $r0$, $r1$, and $r2$, using multipatches, cell-centered AMR, refluxing, and enhanced PPM (see Sec. IIIB2). We align the results from all three resolutions at the time when the central density ρ_c reaches its maximum at core bounce. We observe first order convergence in $\|H\|_2$ after core bounce. In the prebounce phase, $\|H\|_2$ exhibits second-order convergence. This behavior is expected since the numerical scheme reduces to first order at the shock front after bounce where the error are greatest.

In Fig. 11, we also show the absolute difference of the GW strain $Dh_{+,e}$ and the central density ρ_c between low ($r0$) and medium ($r1$) resolutions, and medium and high ($r2$) resolutions. The convergence behavior of the two quantities is less clean than what can be observed for the Hamiltonian constraint due to their oscillatory nature. The convergence is between the expected first and second-order accuracy.

In Fig. 12, we compare vertex-centered AMR with

original PPM reconstruction versus cell-centered AMR with refluxing and enhanced PPM. In addition, we show the behavior of the latter case when multirate RK time evolution is applied. As is clear from the bottom two panels, the cell-centered scheme with refluxing and enhanced PPM (“cc, ePPM”) outperforms the vertex-centered scheme with original PPM (“vc, oPPM”). While in the cell-centered case, $\|H\|_2$ essentially remains constant after core bounce, it clearly grows in the vertex-centered case. Even worse, the vertex-centered case exhibits a rapid growth in total baryonic mass after core bounce. The evolution with multirate RK performs equally well as the “cc, ePPM” setup, which uses standard RK4 time integration. The multirate setup offers a speed up of $\sim 20\%$ for the current test problem. The speed-up can be significantly larger when full microphysics and neutrino transport is employed (e.g. [90]).

In Fig. 13, we revisit our study of extracting gravitational radiation using curvature-based methods [21]. In [21], we found a radial dependence of the accuracy of the curvature-based extraction methods. This study made use of purely Cartesian simulation domains, and was thus limited in terms of possible domain sizes and extraction radii. The maximum extraction radius was limited to $R = 1000M_\odot$. This is still fairly close and means that the waveforms are extracted well inside the star. Our curvature-based extraction methods, however, assume vacuum, i.e. a vanishing stress-energy tensor at the extraction location. In [21], we thus conjectured that increased extraction radii that are located outside the star would further improve the accuracy of the extracted waveforms. Given our new multipatch setup, we can confirm this conjecture. We have placed three extraction radii at $R = [1000M_\odot, 1500M_\odot, 2500M_\odot]$ in a region with constant radial spacing $\Delta r = 8.0M_\odot$ where the radial direction is not yet stretched. The upper panel of Fig. 13 shows the “+” polarization of the GW strain $Dh_{+,e}$ measured in the equatorial plain extracted via CCE. As a comparison, in the same panel, we also show $Dh_{+,e}$ computed via the quadrupole formula. We apply FFI to compute the strain Dh from Ψ_4 extracted with CCE (see Sec. IIF). In [21], we conjectured that the low cut-off frequency that must be picked for FFI can be reduced as the extraction radius is increased. Here, we confirm that this is indeed the case. While extraction radius $R = 1000M_\odot$ requires a low cut-off frequency $f_0 = 100\text{Hz}$ which is well inside the LIGO sensitivity band, we find that at radius $R = 1500M_\odot$ we can get away with $f_0 = 60\text{Hz}$. At radius $R = 2500M_\odot$, we can further reduce this to $f_0 = 30\text{Hz}$ without introducing artificial non-linear drifts in the strain. In the bottom panel of Fig. 13, we show the difference in GW amplitude of the waveforms computed from the inner extraction radii to the waveform computed from the outer most extraction radius. We confirm that as the extraction radius is increased, the differences further decrease similar to what has been found in [21].

The waveform computed via the quadrupole formula

does not suffer from amplification of low frequency errors [21]. We observe that the waveforms extracted via CCE at larger radius and decreased f_0 more closely resemble the monotonically rising signal in the prebounce phase that the waveform computed via the quadrupole formula exhibits. Overall, in accordance with [21], we still measure the same deviations between GW amplitudes computed from CCE and the quadrupole formula to within a few percent at core bounce. This is not surprising, since the error in CCE due to different worldtube extraction locations is much smaller than the observed deviation from the waveform extracted via the quadrupole formula.

Finally, we note that we have also computed the GW strain via the RWZM formalism (not shown). In our previous more detailed study on GW extraction in the context of rotating stellar collapse [21], we found that the RWZM formalism leads to waveforms which are contaminated by high frequency noise. Unfortunately, in the current study, which allows us to use larger extraction radii than $R = 1000M_\odot$, we find that the systematic high-frequency noise inherent in the RWZM waveforms is not reduced, but instead even increases with increased extraction radius. As already conjectured in [21], this is most likely due to the perturbative manner the waves are extracted from the spacetime in the RWZM formalism. In this formalism, the spherical background geometry is projected out, which can result in very small values for the aspherical perturbation coefficients that are prone to numerical noise and cancellation effects. At larger radii, the aspherical perturbations are even smaller since they fall off as $1/r$, and thus are harder to capture accurately. The RWZM approach may therefore be less suited for the extraction of the generally weak GW signals emitted in core collapse.

C. Neutron Star Collapse

Three-dimensional collapse of an isolated neutron star to a black hole is a valuable test of accuracy and convergence of our code for black hole formation in massive stars. We consider the uniformly rapidly rotating model *D4* previously studied in [5, 92] as a benchmark problem. Apart from showing convergence and consistency with previous results, we improve the simulations by causally disconnecting the outer boundary from the interior evolution and the wave-extraction zone. We show that cell-centered AMR with refluxing leads to better conservation of mass than vertex-centered AMR. We also employ CCE for GW extraction.

1. Initial Data and Equation of State

The initial condition is given by a stable relativistic polytrope. Specifically, we use a polytrope $P = K\rho^\Gamma$ with $\Gamma = 2$ and $K_{\text{ini}} = 100$ in the initial data construction. The initial data are generated via Hachisu’s self-




FIG. 14. Rotating neutron star collapse: convergence analysis of the “+” polarization of the GW strain $Dh_{+,e}$ as emitted in the equatorial plane and extracted via CCE (top two panels), central density ρ_c evolution (next two panels), and L_2 -norm of the Hamiltonian constraint $\|H\|_2$ (bottom panel). The differences in $Dh_{+,e}$ and ρ_c between medium and high resolution are scaled for second-order convergence. At $t - t_{\text{BH}} = 0$, the density drops to zero due to hydrodynamic excision within the horizon. The L_2 -norm of the Hamiltonian constraint (bottom panel) does not converge initially due to numerical artifacts from the initial data solver, however, later converges at second-order during black hole formation $t - t_{\text{BH}} \simeq 0$ and black hole ring-down $t - t_{\text{BH}} > 0$. The convergence study is performed using cell-centered AMR with ePPM.

consistent field method [128, 129]. The central density is set to $\rho_c = 3.116 \times 10^{-3} M_\odot^{-2} = 1.924 \times 10^{15} \text{ g cm}^{-3}$. We use an axes ratio of 0.65, which results in $\beta = T/|W| = 7.6796 \times 10^{-2}$ corresponding to a dimensionless spin of $a = J/M^2 = 0.54354$. In order to induce the gravitational collapse, we introduce an artificial pressure depletion of 2% by setting $K = 98$ at the onset of the evolution. During evolution, we use an ideal fluid Γ -law equation of state with $\Gamma = 2$. The initial parameters and properties of the test case are summarized in Table V.

2. Numerical Setup

The GW extraction is carried out on the cubed-sphere grid setup shown in Fig. 1. We pick the radius of the outer boundary such that the wave-extraction zone and

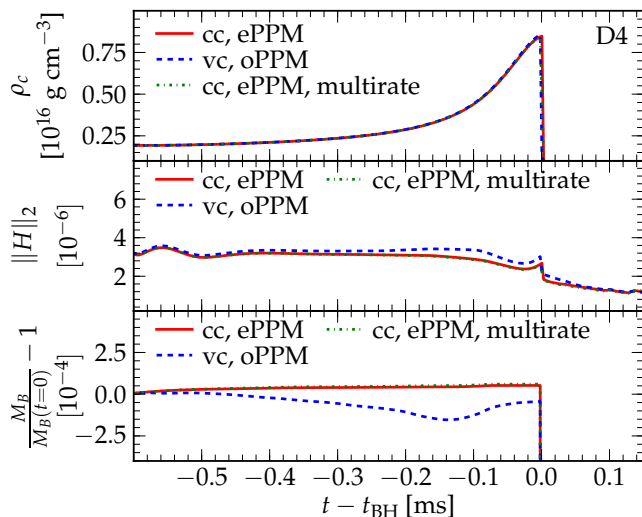


FIG. 15. Rotating neutron star collapse: we compare vertex-centered (vc) AMR and oPPM reconstruction with cell-centered (cc) AMR and ePPM reconstruction. The latter setup is also shown using multirate RK time integration. The top panel compares the central density evolution profile $\rho_c(t)$. The center panel compares the evolution of the L_2 -norm of the Hamiltonian constraint $\|H\|_2$. The bottom panel compares the conservation of baryonic mass M_B . The setup “vc, oPPM” produces slightly larger violations in the Hamiltonian constraints, especially in the late collapse phase shortly before the black hole forms. Due to refluxing, the cell-centered case exhibits much better conservation of baryonic mass. Multirate RK time integration does not lead to different results. The comparison is done using baseline resolution $r1$.

the interior evolution are causally disconnected from the outer boundary, which we set to $R_B = 800M_\odot$.

For our baseline grid setup $r1$, we make use of a radial and Cartesian resolution of $\Delta r = \Delta x = 1.28M_\odot$ and $N_{\text{ang}} = 25$ cells per patch and per angular direction. The boundary between central Cartesian and cubed-sphere grids is located at $R_S = 65M_\odot$. The radial coordinate spacing is increased from Δr to $2\Delta r$ in the region between $R = 250M_\odot$ and $R = 600M_\odot$.

We employ five additional levels of AMR with half-widths $R_{r1} = [30M_\odot, 18M_\odot, 11M_\odot, 5M_\odot, 3M_\odot]$ located at the center of the Cartesian domain. With an initial radius of $R_{\text{NS}} \approx 10M_\odot$ along the equatorial plane, this means that the finest two levels thread through the neutron star. These two levels are required to resolve the black hole formed in the collapse. For our baseline resolution $r1$, we therefore have a grid spacing of $\Delta x = 0.16M_\odot = 0.24\text{km}$ on the third finest level encompassing the entire neutron star, and a resolution of $\Delta x = 0.04M_\odot = 0.06\text{km}$ on the finest level containing the black hole.

In addition to $r1$, we also use a low resolution $r0$ with a coarse grid spacing of $\Delta r = \Delta x = 1.6M_\odot$ and $N_{\text{ang}} = 20$ cells per patch and per angular direction, and a high resolution setup $r2$ with a coarse grid spacing of $\Delta r =$




FIG. 16. Rotating neutron star collapse: we show the total ADM mass M_{ADM} (top panel; red dashed line) and the mass of the apparent horizon M_{AH} plus energy radiated in GWs E_{rad} as a function of time (blue straight line). The total ADM angular momentum J_{ADM} of the spacetime (red dashed line), and the angular momentum J_{AH} as measured on the apparent horizon (blue straight line) is shown in the bottom panel. The inset plots show a close-up of the time evolution of $M_{\text{AH}} + E_{\text{rad}}$ and J_{AH} . As all matter becomes trapped in the event horizon, both, $M_{\text{AH}} + E_{\text{rad}}$ and J_{AH} , quickly asymptote to the conserved ADM values of the spacetime. Due to systematic (atmosphere) and numerical errors, the asymptoted values do not agree with the initial ADM values. Note that the mass radiated in GWs is negligible compared to the total mass of the black hole and thus barely contributes to $M_{\text{AH}} + E_{\text{rad}}$. No angular momentum is radiated in GWs. The results are shown for resolution $r2$ using cell-centered AMR with ePPM.

$\Delta x = 1.024M_\odot$ and $N_{\text{ang}} = 31$ cells per patch and per angular direction.

We set the damping coefficient of the Γ -driver gauge condition to $\eta = 1/2$, and exponentially damp η to zero starting from radius $R_\eta = 65M_\odot$.

The artificial low-density atmosphere is 10^{-8} of initial central density. We also perform a simulation with an atmosphere density 10^{-10} of the central density, however, we find only negligible differences in the accuracy of our results.

3. Discussion

Following initial pressure depletion, the uniformly rotating polytrope collapses. During collapse, the central density ρ_c increases until time $t - t_{\text{BH}} = 0$, the time when an apparent horizon, and thus a black hole forms. After formation of the horizon, the matter inside the horizon is excised from the grid, and the remaining exterior matter is rapidly dragged into the nascent black hole, leaving behind the artificial low-density atmosphere. Upon for-




FIG. 17. Rotating neutron star collapse: power spectral density of “+” polarization of GW strain $Dh_{+,e}$ as emitted in the equatorial plane and extracted via CCE. The blue straight line is the spectrum of the entire waveform, while the green dashed line is the spectrum of the ring-down signal. The red vertical line denotes the $(\ell, m) = (2, 0)$ prograde fundamental ($N = 0$) quasi-normal mode frequency $f_{\text{QNM}} = 6.68$ kHz of a spinning black hole of mass $M = 1.8602 M_{\odot}$ and dimensionless spin $a = 0.5435$ as computed in [131]. Mass and spin of the nascent black hole are determined on its apparent horizon using the isolated horizon framework. The analysis is done using baseline resolution $r1$ with cell-centered AMR with ePPM.

mation, the black hole is highly excited and radiates GWs until it settles to a Kerr state. This produces a characteristic ring-down GW signal with a particular quasi-normal mode frequency which depends only on mass and spin of the black hole.

In Fig. 14, we show the emitted GW signal $Dh_{+,e}$, and the evolution of the central density ρ_c for the three resolutions $r0$, $r1$ and $r2$, respectively. The simulations are performed using cell-centered AMR, refluxing, and ePPM reconstruction. The GW signal is extracted using CCE and we use FFI with a cut-off frequency of $f_0 = 1$ kHz to obtain $Dh_{+,e}$. We note that the only significant non-zero signal is contained in the $(\ell, m) = (2, 0)$ wave mode⁸ and we use (34) to get $Dh_{+,e}$. When comparing the waveform obtained from CCE to the one obtained from RWZM (not shown), we notice that the waveforms from RWZM are more susceptible to numerical noise and contain spurious high-frequency oscillations. This is consistent with our findings in [21] (see also Sec. III B). The waveforms extracted via RWZM are similar to those obtained in [5, 92], which also use RWZM extraction. We

⁸ Earlier studies [5, 92] also found an $(\ell, m) = (4, 0)$ wave mode. In our case, this mode is three orders of magnitudes smaller than the $(\ell, m) = (2, 0)$ mode amplitude and comparable to the level of numerical noise. Since the earlier study did not use causally disconnected outer boundaries, did not compute the waveform at future null infinity \mathcal{I}^+ , and had less resolution in the wave-extraction zone, we argue that a $(\ell, m) = (4, 0)$ could have been excited because of numerical artifacts and systematic errors.

TABLE V. Initial parameters and properties of the collapsing neutron star. ADM mass M_{ADM} and angular momentum J_{ADM} are computed from the initial data solver at spatial infinity i^0 . The radiated energy E_{rad} and angular momentum J_{rad} are computed from waves extracted via the method of CCE including modes up to $\ell = 6$. The apparent horizon mass M_{AH} and angular momentum J_{ADM} are computed on the apparent horizon surface after the black hole has settled to an approximate Kerr state. The data are reported for high resolution simulation $r2$. The value in brackets denotes the numerical error in the last reported digit. Units are in $c = G = M_{\odot} = 1$, unless otherwise specified.

Initial polytropic scale	K_{ini}	100
Evolved polytropic scale	K	98
Polytropic index	Γ	2
Central rest-mass density [M_{\odot}^{-2}]	ρ_c	3.116×10^{-3}
Axes ratio		0.65
Rotational / binding energy [%]	$T/ W $	7.68
Equatorial radius [M_{\odot}]	R_e	9.6522
Baryonic mass [M_{\odot}]	M_B	2.0443
ADM mass [M_{\odot}]	M_{ADM}	1.8605
ADM ang. mom. [M_{\odot}^2]	J_{ADM}	1.8814
Spin	a	0.5435
Rad. energy [M_{\odot}]	E_{rad}	$8.14(3) \times 10^{-7}$
Rad. ang. mom. [M_{\odot}^2]	J_{rad}	$0(1) \times 10^{-10}$
AH mass [M_{\odot}]	M_{AH}	1.8602(3)
AH ang. mom. [M_{\odot}^2]	J_{AH}	1.874(7)

thus believe that the results of [5, 92] also suffer from the same spurious high-frequency noise.

We align all quantities at the coordinate time when an apparent horizon appears ($t - t_{\text{BH}} = 0$). By computing the differences in low and medium, and medium and high resolutions, we get an estimate for the convergence of our simulations. In panels below the emitted GW signal $Dh_{+,e}$, and central density evolution ρ_c of Fig. 14, respectively, we show the differences in GW signal and central density using the three different resolutions. The differences between medium and high resolutions are scaled for second-order convergence. At black hole formation, the GW signal and central density exhibit clear second-order convergence. During collapse, while the central density shows second-order convergence, the convergence of the GW signal is somewhat obscured due to the oscillatory nature of the latter, especially when the signal is not perfectly in phase. In the lower panel of Fig. 14, we show the L_2 -norms of the Hamiltonian constraint $\|H\|_2$ for the three resolutions. Since the artificial initial pressure depletion is not constraint satisfying, the constraints do not converge initially. For this reason, we do not introduce any rescaling for convergence. However, the slopes for higher resolutions are smaller, resulting in somewhat smaller constraint violations at later times. At the time when an apparent horizon appears, and during ring-down, the constraints exhibit second-order convergence.

In Fig. 15, we compare performance of cell-centered AMR with ePPM, vertex-centered AMR with oPPM,

and cell-centered AMR with ePPM and multirate RK time integration using baseline resolution $r1$. The vertex-centered case with oPPM exhibits slightly larger constraint violations than the cell-centered setup using ePPM. Before the horizon forms, baryonic mass should be exactly conserved. In practice, this is not the case, even in the cell-centered case with refluxing. One reason for non-conservation is the artificial low-density atmosphere (see Appendix C). Another reason is the buffer-zone prolongation in regions that thread the surface of the star. Here, prolongation involving cells in the atmosphere can amplify mass non-conservation. We note, however, that the cell-centered case with refluxing performs better than the vertex-centered case. The simulation using multirate time integration performs equally well compared to the same simulation using standard RK4 time integration.

In Fig. 16, we show the mass and spin evolution of the apparent horizon. After $t - t_{\text{BH}} = 0$, horizon mass and spin are quickly growing until they asymptote towards the ADM mass and angular momentum of the spacetime, respectively. For a given spacetime, ADM mass and angular momentum are always constant. Both quantities are calculated in the initial data solver and evaluated at spatial infinity. Since all matter falls into the horizon, the black hole mass plus the radiated energy must be equal to the ADM mass. The same applies to the angular momentum. In the present case, we have $M_{\text{ADM}} = 1.8605 M_{\odot}$. The black hole settles to a horizon mass of $M_{\text{AH}} = 1.8602 M_{\odot}$. Thus, the difference is 0.016%. Similarly, the angular momentum initially is $J_{\text{ADM}} = 1.8814 M_{\odot}^2$, and the black hole settles to $J_{\text{AH}} = 1.874 M_{\odot}^2$. This makes a difference of 0.39%. The radiated energy is $E_{\text{rad}} = 8.14 \times 10^{-7} M_{\odot}$ and hence is tiny compared to the rest mass of the system. This value agrees to the estimate given in [5, 92]. Since the only significant non-zero GW mode is the $(\ell, m) = (2, 0)$ mode, no angular momentum is radiated. We find that by decreasing the atmosphere level and increasing the resolution, the differences in horizon mass and angular momentum compared to the initial ADM values are decreased. Hence, the error in mass and angular momentum conservation is due to systematic (atmosphere) and numerical error.

In Fig. 17, we investigate the power spectrum of the emitted GW signal $D\hat{h}_{+,e}$. The blue straight curve is the power spectrum of the entire signal which peaks at $f_{\text{peak}} = 5.06$ kHz. The green dashed curve is produced by first applying a time-domain window function around the black hole ring-down part of the waveform before taking the Fourier transform. Thus, the green dashed curve is the power spectrum of the black hole ring-down part of the waveform. This curve peaks at $f_{\text{peak, ring-down}} = 6.47$ kHz. We can compare this frequency with the theoretically obtained quasi-normal (QNM) ring-down frequency for a perturbed black hole in vacuum. For the black hole mass $M_{\text{AH}} = 1.8602 M_{\odot}$ and dimensionless spin $a = J_{\text{AH}}/M_{\text{AH}}^2 = 0.5414$, the $(\ell, m) = (2, 0)$ pro-




FIG. 18. Binary neutron stars: Convergence study of the $(\ell, m) = (2, 2)$ mode of the GW strain Dh , and the L_2 -norm of the Hamiltonian constraint $\|H\|_2$. The top panel shows the “+” polarization of the $(\ell, m) = (2, 2)$ mode for all three resolutions. The panel below shows the GW phase ϕ of the $(\ell, m) = (2, 2)$ mode. The third panel from the top shows the difference in phase ϕ , scaled for second-order convergence. The vertical dashed line indicates appearance of an apparent horizon in the high-resolution simulation. The bottom panel shows the L_2 -norm of the Hamiltonian constraint scaled for first-order convergence. The simulations were performed using cell-centered AMR, refluxing, and ePPM reconstruction.

grade fundamental ($N = 0$) quasi-normal frequency is $f_{\text{QNM}} = 6.68$ kHz [131]. Thus, the relative difference is $\simeq 3.3\%$. This is consistent with [5] who find “good agreement” (unfortunately they do not provide numbers). Note that we do not expect the two values to exactly coincide. The theoretical QNM frequency is strictly only valid for perturbed Kerr black holes in vacuum. Since matter is crossing the horizon initially, the ring-down signal will naturally be affected by black hole growth and spin-up.

D. Binary Neutron Stars

We investigate accuracy and convergence of the inspiral and coalescence of a binary neutron star (BNS) system. Previous studies in full general relativity were restricted by the employed purely Cartesian grids (e.g. [6–11, 132–134], also see [14] for a recent review), and thus the accuracy of the GW extraction was limited.

For the first time in the context of binary neutron star mergers, we use CCE for GW extraction at future null infinity \mathcal{J}^+ (see Sec. II F). This removes finite radius




FIG. 19. Binary neutron stars: comparison between cell-centered (cc) AMR with ePPM and vertex-centered (vc) AMR with oPPM. The top panel shows the $(\ell, m) = (2, 2)$ mode of the “+” polarization of the GW strain Dh . The center panel shows the L_2 -norm of the Hamiltonian constraint $\|H\|_2$. The bottom panel shows conservation of baryonic mass M_B . The vertical dashed line indicates the appearance of an apparent horizon in the baseline resolution simulation. The simulations were performed using resolution $r1$, though for the conservation of mass, we also show the high resolution ($r2$) result. The error in mass conservation converges with better than second-order as the resolution is increased up to the point when a new refinement level is switched on at $t \simeq 7.5$ ms.

and gauge errors and, combined with our multipatch grid, allows us to extract the higher than leading order modes.

Finally, we also compare vertex centered AMR with oPPM with cell-centered AMR with refluxing and ePPM.

1. Initial Conditions and Equation of State

The particular system we evolve is the initial data set `G2_I12vs12_D5R33_60km` produced by the LORENE code [57, 135]. This system, with the same parameters as described below, has also been considered in [136, 137].

The system consists of two neutron stars initially described by a polytropic equation of state $P = K\rho^\Gamma$ with $K = 123.6$ and $\Gamma = 2$ with an initial coordinate separation of 45 km. We evolve the system using a Γ -law equation of state of the form

$$P = (\Gamma - 1)\rho\epsilon. \quad (35)$$

These parameters yield neutron stars of individual baryonic mass $M_B = 1.78 M_\odot$ and ADM-mass in isolation of $M_{NS} = 1.57 M_\odot$. The total ADM mass of the system is $M_{ADM} = 3.2515 M_\odot$, and the total ADM angular momentum is $J_{ADM} = 10.1315 M_\odot^2$. The initial orbital angular frequency of the binary is $\Omega_{ini} = 302$ Hz. The initial parameters and properties are listed in Table VI.




FIG. 20. Binary neutron stars: GW modes $(\ell, m) = (3, 2), (4, 4), (6, 6), (8, 8)$ of “+” polarization of the strain Dh unambiguously extracted via CCE. The waveforms are shown for high resolution simulation $r2$. The vertical line indicates the time of appearance of an apparent horizon. Following appearance of an apparent horizon, a black hole ring-down signal is visible.

2. Numerical Setup

The numerical setup consists of the six spherical *inflated-cube* grids that surround the central Cartesian cube. The inner spherical radius of the inflated cube grids is located at a coordinate radius of $R_S = 75.84 M_\odot$ and the outer (spherical) boundary is located at a radius of $R_B = 2800 M_\odot$. The radial resolution at the inner spherical inter-patch boundary matches the coarse-grid Cartesian resolution of the central cube and is $\Delta x = 1.5 M_\odot = 2.22$ km, $\Delta x = 1.2 M_\odot = 1.77$ km and $\Delta x = 0.96 M_\odot = 1.42$ km for the low, medium and high resolution runs, respectively. In the region $250 M_\odot < r < 800 M_\odot$ we smoothly transition to a coarser resolution of $6.0 M_\odot$, $4.8 M_\odot$ and $3.84 M_\odot$ for low ($r0$), medium ($r1$) and high resolution ($r2$), respectively. The angular resolution is constant along radial distances and we use 21, 25 and 31 angular grid points per angular direction and spherical patch for the three resolutions. We use 4 initial levels of mesh refinement in the inner Cartesian cube to resolve the neutron stars. We surround each neutron star with a set of nested, refined cubes of half-width $13 M_\odot$, $17.875 M_\odot$ and $26.125 M_\odot$, where the finest level




FIG. 21. Binary neutron stars: Phase convergence of the $(\ell, m) = (6, 6)$ mode of the GW strain Dh . The top panel shows the “+” polarization component Dh_+^{66} , and the panel below shows the phase ϕ , for low $r0$, medium $r1$, and high $r2$ resolutions. The bottom panel shows the phase differences between low and medium, and medium and high resolutions, scaled for second-order convergence. Convergence is maintained throughout inspiral and merger. In the ring-down phase, the coarsest resolution $r0$ is insufficient to accurately resolve this mode, and the results cease to converge properly. The vertical line indicates appearance of an apparent horizon in the high resolution simulation.

completely covers the neutron star. All refined cubes surrounding the stars are contained in the common, coarse cube of half-width R_S . In each refined level the resolution is twice that of the previous level. On the finest level, the neutron stars are covered with a resolution of $\Delta x = 0.1875 M_\odot = 0.278$ km, $\Delta x = 0.15 M_\odot = 0.222$ km and $\Delta x = 0.12 M_\odot = 0.176$ km for the three resolutions $r0$, $r1$ and $r2$, respectively.

When the two neutron stars are about to come into contact, we remove the nested set of cubes surrounding each individual star and surround the binary with a common set of nested cubes of half-width R_S , $30 M_\odot$, $15 M_\odot$ and $7.5 M_\odot$ ensuring uniform resolution in the central region. Once the lapse function drops to values that indicate that an apparent horizon is about⁹ to form, we switch on a final level of radius $3.5 M_\odot$ and resolution $9.38 \times 10^{-2} M_\odot$, $7.5 \times 10^{-2} M_\odot$ and $6.00 \times 10^{-2} M_\odot$ for the low, medium and high resolution runs respectively. This level allows us to handle the steep metric gradients developing inside of the newly formed apparent horizon.

During inspiral, we track the center of mass of each

⁹ This is a consequence of the $1 + \log$ slicing condition (14) which locally slows down time evolution (i.e. $\alpha < 1$) in regions of strong curvature. A closed surface of lapse of $\alpha \lesssim 0.3$ has been found to approximately resemble the apparent horizon shape.




FIG. 22. Binary neutron stars: Amplitude convergence of the $(\ell, m) = (6, 6)$ mode of the GW strain $Dh = Dh(\phi)$ as a function of phase ϕ . The top panel shows the “+” polarization component $Dh_+^{66}(\phi)$, and the panel below shows the amplitude $A(\phi)$, for low ($r0$), medium ($r1$), and high ($r2$) resolutions. The bottom panel shows the amplitude differences between low and medium, and medium and high resolutions, scaled for second-order convergence. Convergence is maintained throughout inspiral and merger. In the ring-down phase, however, the coarsest resolution $r0$ is insufficient to accurately resolve this mode, and the results cease to properly converge. The vertical line indicates appearance of an apparent horizon in the high resolution simulation.

neutron star to keep the two fluid bodies close to the center of their refined regions. We compute the center of mass of an individual neutron star by integrating over the conserved density within a radius $R = 4.0 M_\odot$ of the densest point on the grid. This method produces smoother tracks than directly using the location of the densest point, and helps reducing the jitter in the mesh refinement boxes observed otherwise.

We set the damping coefficient of the Γ -driver gauge condition to $\eta = 1$.

We set the dissipation strength to $\epsilon_{\text{diss}} = 0.1$ everywhere on the grid. The artificial low-density atmosphere is 10^8 times lower than the initial central density.

3. Discussion

While the two neutron stars orbit each other, they lose energy due to gravitational radiation, inspiral, and finally merge. The nascent hypermassive neutron star remnant has a mass which is well above the maximum mass of neutron stars. It forms a black hole on a dynamical timescale. The black hole is initially highly excited, and relaxes to a Kerr state by emitting gravitational ring-down radiation.

In Fig. 18, we show convergence of the dominant

TABLE VI. Parameters of the binary neutron star system. ADM mass M_{ADM} and angular momentum J_{ADM} are computed by the initial data solver at spatial infinity i^0 . The radiated energy E_{rad} and angular momentum J_{rad} are computed from waves extracted via CCE including modes up to $\ell = 6$. The apparent horizon mass M_{AH} and angular momentum J_{ADM} are computed after the black hole has settled to an approximate Kerr state. Gravitational disk mass M_{disk} and angular momentum J_{disk} are calculated from energy and angular momentum conservation. The data are reported for simulation $r2$. The value in brackets denotes the numerical error in the last reported digit. Units are in $c = G = M_{\odot} = 1$, unless otherwise specified.

Lorene initial data set	G2_I12vs12_D5R33_60km
Initial separation [km]	d 45
Polytropic scale	K 123.6
Polytropic index	Γ 2
Initial orbital frequency [Hz]	Ω_{ini} 302
ADM mass [M_{\odot}]	M_{ADM} 3.2515
ADM ang. mom. [M_{\odot}^2]	J_{ADM} 10.1315
Rad. energy [M_{\odot}] (%)	E_{rad} $2.51(5) \times 10^{-2}$ (0.77%)
Rad. ang. mom. [M_{\odot}^2] (%)	J_{rad} 1.206(9) (11.9%)
AH mass [M_{\odot}]	M_{AH} 3.2249(3)
AH ang. mom. [M_{\odot}^2]	J_{AH} 8.75(2)
AH spin	a 0.841(2)
Grav. mass disk [M_{\odot}]	M_{disk} $1.4(4) \times 10^{-3}$
Bary. mass disk [M_{\odot}]	$M_{B,\text{disk}}$ $1.3(2) \times 10^{-3}$
Ang. mom. disk [M_{\odot}^2]	J_{disk} 0.16(4)

$(\ell, m) = (2, 2)$ mode of the GW strain Dh , the L_2 -norm of the Hamiltonian constraint $\|H\|_2$. The upper panel shows the “+” polarization of the $(\ell, m) = (2, 2)$ mode of the GW strain for the resolutions $r0$, $r1$, and $r2$. The waveform is extracted via CCE. To obtain Dh , we use a cut-off parameter $f_0 = 507$ Hz, which is below the initial instantaneous $(\ell, m) = (2, 2)$ mode frequency f_{ini}^{22} determined from the initial orbital frequency by $f_{\text{ini}}^{22} = 2\Omega_{\text{ini}}$. To assess the phase convergence, we plot the differences in phase between low $r0$ and medium $r1$ resolution, and medium and high $r2$ resolution, scaled for second-order convergence. We also plot the L_2 -norm of the Hamiltonian constraint $\|H\|_2$ scaled for first-order convergence. Similar to the isolated neutron star tests in Sec. III A, the dominant constraint error is generated at the contact discontinuity at the neutron star surface, where our scheme locally reduces to first-order accuracy.

In Fig. 19, we compare cell-centered (cc) AMR and ePPM reconstruction with vertex-centered (vc) AMR and oPPM. The $(\ell, m) = (2, 2)$ mode of the GW strain Dh and the L_2 -norm of the Hamiltonian constraint $\|H\|_2$ do not show any significant differences between the two numerical setups at this point. After black hole and disk formation, the vertex-centered scheme exhibits a slightly larger slope in constraint growth. In the bottom panel, we show conservation of total baryonic mass M_B . During early inspiral, both setups conserve mass to a high degree, only affected by small errors due to our artificial atmosphere (see Appendix C). Note that both neu-

tron stars are completely contained on their finest grids. Thus, there are no refinement boundaries directly influencing the evolution of the two fluid bodies. As the inspiral progresses, we find that mass conservation is violated in the cell-centered case to a higher degree than in the vertex-centered case (though the error converges as the resolution is increased). This appears to be an artifact of buffer zone prolongation close to the neutron star surface in combination with low density matter slightly above and at atmosphere values. Due to numerical errors, small amounts of mass are leaking out of the neutron star during inspiral and interact with the atmosphere. As this low density matter reaches the buffer zones, numerical errors due to prolongation, which are by construction larger in the cell-centered case, tend to amplify the negative effects of the atmosphere treatment. In experiments with isolated neutrons stars, however, we find that when the refinement boundaries are sufficiently far removed, and/or the atmosphere level is further decreased, mass can be conserved to a higher degree.

We also compare the simulations to a setup using multirate RK time integration and cell-centered AMR with ePPM. Unfortunately, due to the large fluid bulk velocities in the inspiral phase, the orbital phase accuracy is significantly affected by the lower order fluid time integration. Thus, we do not recommend application of multirate RK schemes in the context of binary neutron star mergers, especially when orbital phase accuracy is paramount. The problem may be ameliorated by the use of co-rotating coordinates (see, e.g., [71]).

In order to demonstrate the potential of the multipatch scheme for more accurate wave extraction, we show in Fig. 20 some of the higher harmonic GW modes that are emitted during inspiral, merger, and ring-down. We show (from top to bottom) the $(\ell, m) = (3, 2)$, $(\ell, m) = (4, 4)$, $(\ell, m) = (6, 6)$, and $(\ell, m) = (8, 8)$ modes of “+” polarization of the strain Dh . The modes are extracted from a simulation using resolution $r2$, cell-centered AMR, and ePPM. All modes up to $(\ell, m) = (4, 4)$ show a clean inspiral, merger and ring-down signal, and converge with resolution (see below). For higher modes, our lowest resolution run $r0$ is insufficient to also allow for clean convergence of the corresponding ring-down signals. Accordingly, those should be taken with a grain of salt. As an example, in Figs. 21 and 22, we show convergence of phase and amplitude of the $(\ell, m) = (6, 6)$ mode of the GW strain, respectively. Fig. 21 shows the GW amplitude A reparametrized in terms of the gravitational phase ϕ to disentangle phase from amplitude. Both figures indicate that second-order convergence is maintained during inspiral up to merger. The ring-down part, however, does not exhibit clean second-order convergence. In that case, the coarse resolution becomes insufficient, and the result ceases to converge properly. We note that for the highest extracted mode, $(\ell, m) = (8, 8)$, the coarsest resolution is insufficient to allow for clean convergence also in the inspiral phase.

We compute the radiated energy E_{rad} , radiated angu-

lar momentum J_{rad} , the horizon mass M_{AH} , and horizon angular momentum J_{AH} . For the computation of the radiated quantities, we include modes $\ell \leq 6$ as extracted via CCE. After the black hole has formed and settled to an approximate Kerr state, some amount of material is located in an accretion disk surrounding the black hole. Hence, we do not expect that horizon mass and radiated energy balance with the total ADM mass at this time. Rather, the difference denotes the gravitational mass of the accretion disk that has formed. Likewise, the same is true for the balance of angular momentum. Given the horizon mass, the spacetime’s total ADM mass, and the radiated energy, we estimate the gravitational mass of the accretion disk to be $M_{\text{disk}} = M_{\text{ADM}} - M_{\text{AH}} - E_{\text{rad}} = (1.4 \pm 0.4) \times 10^{-3} M_{\odot}$. The disk’s baryonic mass is $M_{B,\text{disk}} = (1.3 \pm 0.2) \times 10^{-3} M_{\odot}$, which we compute by integrating over all material outside of the apparent horizon and within a radius $R < 40 M_{\odot}$. Both, baryonic and gravitational mass agree within their error bars. We note that the mass of the disk, though clearly visible in density contour plots of our simulation (not shown), is tiny and thus not much above the numerical error. Given the horizon angular momentum, the spacetime’s total ADM angular momentum, and the radiated angular momentum, we estimate the disk’s angular momentum to be $J_{\text{disk}} = J_{\text{ADM}} - J_{\text{AH}} - J_{\text{rad}} = 0.16 \pm 0.04 M_{\odot}^2$. For convenience, we list spacetime, black hole, disk, and radiated mass (and angular momentum) in Table VI. All error bars are estimated using medium and high resolution results. The results for mass and spin of the black hole agree with the values that were found in [136].

In our binary neutron star merger problem, we also investigate the error inherent to finite-radius GW extraction. We compare Ψ_4 as extracted via the NP formalism at a finite radius with Ψ_4 as extracted via CCE at future null infinity \mathcal{J}^+ . We align two given waveforms in the early inspiral phase by minimizing their phase difference over an interval $t \in [2.5 \text{ ms}, 3.5 \text{ ms}]$ using the method described in [138]. For the $(\ell, m) = (2, 2)$ mode, we find a total dephasing on the order of $\Delta\phi \sim 1$ rad and an amplitude difference of about $\simeq 10\%$ between the waveform obtained at $R = 250 M_{\odot}$ and the one obtain at \mathcal{J}^+ . Waveforms extracted at smaller radii naturally yield larger differences to the result at \mathcal{J}^+ . While the amplitude error is rather large, the dephasing is comparable to the dephasing due to numerical error of the orbital evolution of the two neutron stars. Since this numerical error is convergent, but the systematic finite radius-error is not, the finite-radius error becomes a non-negligible effect as the numerical resolution is increased. As shown in [22, 23] for the case of binary black hole mergers, extrapolation to infinity using finite-radius data can reduce the errors to a tolerable level in cases where CCE is not available.

Finally, we investigate the influence of the outer boundary when it is *not* causally disconnected from the wave extraction region and interior evolution. We compare a setup with a causally connected outer boundary

located at $R_{\text{B}} = 2000 M_{\odot}$ and a causally disconnected boundary located at $R_{\text{B}} = 2800 M_{\odot}$. The former setup is in causal contact with the interior and wave-extraction region during the merger and ring-down phases. We find a difference in GW phase and amplitude, and final spin and mass of about $\simeq 7\%$. More details are given in Appendix F.

By comparing our results with those of [136, 137], we conclude that the accuracy of the orbital evolution of the two neutron stars is very similar. The errors in satisfying the Hamiltonian constraint and conserving baryonic mass are of comparable size. This is not surprising, since we find little difference between the new cell-centered AMR scheme compared with the vertex-centered AMR scheme that was also used in [136, 137]. Due to our multipatch grids, causally disconnected outer boundaries, and CCE, however, the waveforms that are extracted from our simulations are more accurate than what has been shown in previous studies.

IV. SUMMARY AND CONCLUSIONS

We have presented a new GR hydrodynamics scheme using multiple Cartesian/curvi-linear grid patches and flux-conservative cell-centered adaptive mesh refinement (AMR) to allow for a more efficient and accurate spatial discretization of the computational domain. This is the first study enabling GR hydrodynamic simulations with multipatches and AMR. Our multipatch scheme consists of a set of curvi-linear spherical “inflated-cube” grids with fixed angular resolution and variable radial spacing, and a central Cartesian grid with AMR. High-order Lagrange interpolation is used to fill ghost zones at patch boundaries for variables that are smooth, and second-order essentially non-oscillatory (ENO) interpolation for variables that contain discontinuities and shocks.

Apart from the successful implementation of multipatches and flux-conservative cell-centered AMR, we have introduced a number of additional improvements to the publicly available code `GRHydro`: (i) We have applied the enhanced piecewise-parabolic method (ePPM) to ensure high-order reconstruction at smooth maxima, a property that we have found to be crucial for cell-centered AMR. (ii) To speed up the computation, we have applied a multirate Runge-Kutta time integrator that exploits the less restrictive Courant-Friedrich-Lewy (CFL) condition for the hydrodynamic evolution by switching the time integration to second order and thus reducing the number of intermediate steps by a factor of two. Since the hydrodynamic evolution dominates the curvature evolution in terms of computational walltime when complex microphysics and neutrinos are included, the scheme can yield a speedup of $\gtrsim 30\%$ (e.g. [90]).

We have presented stable and convergent evolutions for binary neutron star mergers, stellar collapse to a neutron star, neutron star collapse to a black hole, and evolutions of isolated unperturbed and perturbed neutron

stars. For each test case, due to the more efficient domain discretization, we have been able to enlarge the domain sufficiently so that the outer boundary is causally disconnected from the interior evolution and wave-extraction zone. This has allowed us to remove the systematic error that arises from the lack of constraint preserving boundary conditions for the Einstein equations in the BSSN formulation. In the case of the binary neutron star merger problem, we have found that this error is on the order of a few percent, and thus limits the accuracy of the simulation and GW extraction.

In addition to enlarging the domain, multipatches have also allowed us to significantly increase the resolution in the GW extraction zone compared to previous studies. For the neutron star merger problem, we have been able to extract convergent spherical harmonic modes of the GW strain Dh up to $\ell = 6$. Previous studies have only considered the dominant $(\ell, m) = (2, 2)$ wave mode for this problem.

Furthermore, we have been able to remove the systematic error inherent in finite-radius wave extraction by application of Cauchy-characteristic extraction (CCE). This wave-extraction method computes gauge-invariant radiation at future null infinity \mathcal{J}^+ using boundary data from a worldtube at finite radius. This method has previously been applied in simulations of binary black holes and stellar collapse [2, 21–24, 75, 76]. Here, we have applied CCE also to simulations of binary neutron star mergers, neutron star collapse to a black hole, and isolated excited neutron stars. We have found that the error due to finite-radius extraction can be as large as 10%.

Finally, for each test case, we have compared the original vertex-centered AMR scheme using original PPM with the new flux-conservative cell-centered AMR scheme using enhanced PPM. The accuracy has been investigated and compared to results from previous studies. We have found that simulations of stellar collapse greatly benefit from flux-conservative cell-centered AMR with enhanced PPM compared to the original vertex-centered AMR scheme with original PPM. Conservation of mass and the satisfaction of the Hamiltonian constraint are significantly better with the new scheme. The isolated neutron star and binary neutron star test cases, on the other hand, are not much affected by the choice of cell-centered or vertex-centered AMR. This is mainly due to the choice of grid setup: no matter is crossing any refinement boundaries so that flux-conservation is not important. It can become important, however, in the post-merger phase of binary neutron star coalescence, especially in cases where a massive accretion torus forms.

The multipatch infrastructure, the associated curvature and hydrodynamics evolution codes, and all other computer codes used in this paper will be made (or are already) publicly available via the `EinsteinToolkit` [91].

ACKNOWLEDGMENTS

We acknowledge helpful discussions with Peter Diener, Frank Löffler, Uschi C. T. Gamma, and members of our Simulating eXtreme Spacetimes (SXS) collaboration (<http://www.black-holes.org>). This research is partially supported by NSF grant nos. AST-0855535, AST-1212170, PHY-1212460, PHY-1151197, and OCI-0905046, by the Alfred P. Sloan Foundation, and by the Sherman Fairchild Foundation. CR acknowledges support by NASA through Einstein Postdoctoral Fellowship grant number PF2-130099 awarded by the Chandra X-ray center, which is operated by the Smithsonian Astrophysical Observatory for NASA under contract NAS8-03060. RH acknowledges support by the Natural Sciences and Engineering Council of Canada. The simulations were performed on the Caltech compute cluster *Zwicky* (NSF MRI award No. PHY-0960291), on supercomputers of the NSF XSEDE network under computer time allocation TG-PHY100033, on machines of the Louisiana Optical Network Initiative under grant `loni_numrel07`, and at the National Energy Research Scientific Computing Center (NERSC), which is supported by the Office of Science of the US Department of Energy under contract DE-AC02-05CH11231. All figures were generated with the Python-based `matplotlib` package (<http://matplotlib.org/>).

Appendix A: Shock-tube Tests

We perform a number of basic Sod shock tube and spherical blast wave tests on fixed backgrounds to ensure correctness and convergence of our scheme at mesh-refinement and inter-patch boundaries.

In this appendix, we restrict our attention to a simple Sod test to show convergence of the primitive variables across inter-patch boundaries (see Sec. II C 3), and to demonstrate mass, energy, and momentum conservation at refinement boundaries when refluxing (see Sec. II D) is used.

The Sod shock-tube test consists of setting the initial fluid state according to [139]. The shock front is located at a position x_0 . The background metric is set to the flat space Minkowski metric. The tests below use a gamma-law equation of state $P = (\Gamma - 1)\rho\epsilon$ with $\Gamma = 1.4$.

If not stated otherwise, the tests below use cell-centered AMR with refluxing, ePPM reconstruction, second-order ENO inter-patch interpolation, RK4 time integration with $\Delta t/\Delta x = 0.4$, and the HLLC Riemann solver.

1. Inter-patch Interpolation

In this particular test, we check that shock fronts are correctly transported across inter-patch boundaries by maintaining convergence, and without introducing local




FIG. 23. L_1 -norm of the difference between exact and evolved fluid state for a Sod shock tube problem on low r_0 , medium r_1 , and high r_2 resolutions. As the resolution is increased, the error in primitive density ρ (upper panel), specific internal energy ϵ (middle panel), and x -component of the 3-velocity v^x (lower panel) correctly decrease by a factor of two in accordance with first-order convergence.




FIG. 24. Conservation of mass (top panel), energy (middle panel), and momentum (bottom panel) as a function of time for a shock front crossing a refinement boundary. The solid (red/blue/green) lines are from a simulation with refluxing, while the dashed (black) curves show the case without refluxing. With refluxing, mass, energy, and momentum are exactly conserved (to machine precision). Without refluxing, conservation of mass, energy, and momentum is violated.

oscillations at the shock, even in the presence of non-trivial Jacobians and coordinate transformations. We setup a multipatch grid consisting of a central Cartesian grid surrounded by the spherical inflated-cube grids. The outer boundary extends to $R_B = 2.5M_\odot$. The boundary between Cartesian and spherical grids is located at $R_S = 0.5M_\odot$. No AMR is employed. For the coarsest resolution (r_0), we set the Cartesian and radial resolution to $\Delta x = \Delta r = 0.05$, and use $(N_\rho, N_\sigma) = (20, 20)$ cells per spherical patch per direction. Medium (r_1) and high (r_2) resolutions double and quadruple, respectively, the resolution with respect to the coarsest resolution.

We set Sod initial data with $x_0 = 0$ and evolve the system for sufficiently long so that the shock propagates across inter-patch boundaries. At each timestep, we compare the evolved fluid state with a solution from an exact special relativistic Riemann solver [140].

In Fig. 23, we show the L_1 -norm of the difference between exact and evolved primitive density ρ , specific internal energy ϵ , and the x -component of the 3-velocity v^x . All quantities are plotted for the three resolutions r_0 , r_1 , and r_2 . As the resolution is increased, the error correctly decreases by a factor of two between successive resolutions, thus indicating first-order convergence. This is consistent with the ENO operator, which reduces to first-order at shocks.

2. Refluxing

In this simple test, we check the correctness of our refluxing scheme with a shock front crossing a refinement boundary. As the shock crosses the boundary, mass, momentum and energy must be conserved to machine precision.

The numerical grid consists of two levels of 2:1 AMR. The coarse level extends from $x = 0$ to $x = 1$. The fine level has a refinement half-width of $r = 0.1$ and is located at $x = 0.4$. We set the Sod shock front [139] at location $x_0 = 0.48$. Thus, the shock starts off on the fine grid and propagates onto the coarse grid.

A measure of conservation of energy and mass is given by the sum of the *conserved* internal energy τ and the *conserved* density D over the entire simulation domain, respectively. Both sums must be constant for all times t . A measure for conservation of momentum is given by the balance between the *conserved* momentum and the pressure gradient force F_{press} per unit time. The balance as a sum over the entire simulation domain must be constant as a function of time. In Fig. 24, we show the sums of conserved density, energy, and momentum when refluxing is used (solid lines). Without refluxing (dashed lines), the conserved mass, energy, and momentum grow significantly at time $t \approx 0.025$ when the shock front crosses the refinement boundary.




FIG. 25. TOV star (from Sec. III A): the effect of original PPM (oPPM) versus enhanced PPM (ePPM) on the Hamiltonian constraint as a function of x at time $t = 0.76$ ms. on cell-centered (cc) and vertex-centered (vc) AMR grids. The star’s radius is $R_e = 14.16$ km. The original PPM results in large constraint violations on the cell-centered grid. The enhanced PPM clearly outperforms oPPM. For ePPM, the error is dominated by the first-order error at the neutron star surface, where the scheme reduces to first order.

Appendix B: Enhanced PPM Scheme

The PPM scheme seeks to find “left” and “right” interpolated values, $a_{i,L}$ and $a_{i,R}$ at the left and right cell interfaces of a primitive quantity a_i defined on cell centers labeled by $i = 0, \dots, N - 1$. The left and right states are defined on cell interfaces labeled by $a_{i \pm \frac{1}{2}}$. Rather than assuming a constant value for a cell-averaged quantity within a given cell, the PPM scheme uses parabolas to represent cell averages within a given cell.

The *enhanced* PPM reconstruction proceeds in three steps: (i) Compute an approximation to a at cell interfaces using a high-order interpolation polynomial, (ii) limit the interpolated cell-interface values obtained in (i) to avoid oscillations near shocks and other discontinuities, (iii) constrain the parabolic profile so that no new artificial maximum is created within one single cell. The main difference to the original PPM scheme is in steps (i) and (ii). Both, the limiter and the constraining of the parabolic profiles is more restrictive in the original PPM scheme, thus reducing the order of accuracy in cases where it is not necessary.

a. First Step: Interpolation We compute an approximation to a at cell interfaces, which, assuming a uniform grid, is obtained via fourth order polynomial interpolation

$$a_{i+\frac{1}{2}} = \frac{7}{12}(a_{i+1} + a_i) - \frac{1}{12}(a_{i-1} + a_{i+2}), \quad (\text{B1})$$

using the cell center values of a from neighboring cells. Ref. [86] also suggests to use a sixth-order polynomial. This, however, requires more ghost points. In our tests, we find no significant difference between fourth and sixth-order interpolation. Hence, we stick to the fourth-order

interpolant.

b. Second Step: Limiting We require that the values $a_{i+\frac{1}{2}}$ satisfy

$$\min(a_i, a_{i+1}) \leq a_{i+\frac{1}{2}} \leq \max(a_i, a_{i+1}), \quad (\text{B2})$$

i.e., the interpolated value $a_{i+\frac{1}{2}}$ must lie between adjacent cell values [86]. This is enforced by the following conditions. If (B2) is not satisfied, then we define the second derivatives,

$$(D^2a)_{i+\frac{1}{2}} := 3(a_i - 2a_{i+\frac{1}{2}} + a_{i+1}), \quad (\text{B3})$$

$$(D^2a)_{i+\frac{1}{2},L} := (a_{i-1} - 2a_i + a_{i+1}), \quad (\text{B4})$$

$$(D^2a)_{i+\frac{1}{2},R} := (a_i - 2a_{i+1} + a_{i+2}). \quad (\text{B5})$$

If $(D^2a)_{i+\frac{1}{2}}$ and $(D^2a)_{i+\frac{1}{2},L,R}$ all have the same sign $s = \text{sign}((D^2a)_{i+\frac{1}{2}})$, we further define

$$(D^2a)_{i+\frac{1}{2},\text{lim}} := s \min(C|(D^2a)_{i+\frac{1}{2},L}|, C|(D^2a)_{i+\frac{1}{2},R}|, |(D^2a)_{i+\frac{1}{2}}|). \quad (\text{B6})$$

where C is a constant that we set according to [86] to $C = 1.25$. Otherwise, if one of the signs is different¹⁰, we set $(D^2a)_{i+\frac{1}{2},\text{lim}} = 0$. Then, we recompute (B1) by

$$a_{i+\frac{1}{2}} = \frac{1}{2}(a_i + a_{i+1}) - \frac{1}{3}(D^2a)_{i+\frac{1}{2},\text{lim}}. \quad (\text{B7})$$

c. Third Step: Constrain Parabolic Profiles Here, we apply the refined procedure from [87]. We begin by initializing left and right states according to the interpolated (and possibly limited) $a_{i+\frac{1}{2}}$ via

$$a_{i,R} = a_{i+1,L} = a_{i+\frac{1}{2}}, \quad (\text{B8})$$

so that the Riemann problem is trivial initially. The conditions below potentially alter $a_{i,R}$ and $a_{i+1,L}$, so that the Riemann problem becomes non-trivial.

First, we check whether we are at a smooth local maximum. A condition for local smooth maxima is given by

$$(a_{i,L} - a_i)(a_i - a_{i,R}) \leq 0, \quad \text{or} \\ (a_{i-2} - a_i)(a_i - a_{i+2}) \leq 0. \quad (\text{B9})$$

If (B9) holds, we compute, similar to (B3),

$$(D^2a)_i = -12a_i + 6(a_{i,L} + a_{i,R}), \\ (D^2a)_{i,C} = a_{i-1} - 2a_i + a_{i+1}, \\ (D^2a)_{i,L} = a_{i-2} - 2a_{i-1} + a_i, \\ (D^2a)_{i,R} = a_i - 2a_{i+1} + a_{i+2}. \quad (\text{B10})$$

¹⁰ For the specific internal energy ϵ , we also set $(D^2a)_{i+\frac{1}{2},\text{lim}} = 0$, in cases when $(D^2a)_{i+\frac{1}{2},\text{lim}} > \frac{1}{2}(a_i + a_{i+1})$. This is different from the procedure in [87], but is necessary at very strong contact discontinuities such as the surface of a neutron star to prevent ϵ from becoming negative for equations of state that do not allow $\epsilon < 0$. In practice, this additional limiter has no effect on the measured accuracy.

If $(D^2a)_{i,[C,L,R]}$ all have the same sign $s = \text{sign}((D^2a)_i)$, we compute

$$(D^2a)_{i+\frac{1}{2},\text{lim}} = s \min(C|(D^2a)_{i+\frac{1}{2},L}|, C|(D^2a)_{i+\frac{1}{2},R}|, C|(D^2a)_{i+\frac{1}{2},C}|, |(D^2a)_{i+\frac{1}{2}}|). \quad (\text{B11})$$

Otherwise, if one of the signs is different, we set $(D^2a)_{i+\frac{1}{2},\text{lim}} = 0$. If

$$|(D^2a)_i| \leq 10^{-12} \cdot \max(|a_{i-2}|, |a_{i-1}|, |a_i|, |a_{i+1}|, |a_{i+2}|) \quad (\text{B12})$$

then we define and set $\rho_i \equiv 0$. Otherwise, we define

$$\rho_i \equiv \frac{(D^2a)_{i+\frac{1}{2},\text{lim}}}{(D^2a)_i}. \quad (\text{B13})$$

To avoid limiting at small oscillations induced by round-off errors, we do not apply any limiter if $\rho_i \geq 1 - 10^{-12}$. Otherwise, we compute the third derivative according to

$$(D^3a)_{i+\frac{1}{2}} = (D^2a)_{i+1,C} - (D^2a)_{i,C}. \quad (\text{B14})$$

We set

$$(D^3a)_i^{\min} = \min((D^3a)_{i-\frac{3}{2}}, (D^3a)_{i-\frac{1}{2}}, (D^3a)_{i+\frac{1}{2}}, (D^3a)_{i+\frac{3}{2}}) \quad (\text{B15})$$

and

$$(D^3a)_i^{\max} = \max((D^3a)_{i-\frac{3}{2}}, (D^3a)_{i-\frac{1}{2}}, (D^3a)_{i+\frac{1}{2}}, (D^3a)_{i+\frac{3}{2}}). \quad (\text{B16})$$

Then, we test if

$$C_3 \cdot \max(|(D^3a)_i^{\max}|, |(D^3a)_i^{\min}|) \leq (D^3a)_i^{\max} - (D^3a)_i^{\min}, \quad (\text{B17})$$

holds. In the expression above, $C_3 = 0.1$, according to Ref. [87]. If (B17) does not hold, a limiter is not applied. Otherwise, we test the following conditions: (i) if $(a_{i,L} - a_i)(a_i - a_{i,R}) < 0$, we set

$$\begin{aligned} a_{i,L} &= a_i - \rho_i(a_i - a_{i,L}), \\ a_{i,R} &= a_i + \rho_i(a_{i,R} - a_i). \end{aligned} \quad (\text{B18})$$

Otherwise, (ii) if $|a_i - a_{i,L}| \geq 2|a_{i,R} - a_i|$, we set

$$a_{i,L} = a_i - 2(1 - \rho_i)(a_{i,R} - a_i) - \rho_i(a_i - a_{i,L}) \quad (\text{B19})$$

or (iii) if $|a_{i,R} - a_i| \geq 2|a_i - a_{i,L}|$, we set

$$a_{i,R} = a_i + 2(1 - \rho_i)(a_i - a_{i,L}) + \rho_i(a_{i,R} - a_i). \quad (\text{B20})$$

In the conditions (i)-(iii) above, we introduce a special treatment for the specific internal energy ϵ . If $|a_i - a_{i,L}| > |a_i|$ or $|a_{i,R} - a_i| > |a_i|$, we set $a_{i,L,R} = a_i$

instead of using the full expressions, respectively. This is different from the original procedure of Ref. [87]. It essentially reduces the reconstruction of ϵ to first order in cases when the correction becomes larger than the value of the reconstructed quantity itself. This is similar to the limiter step further above and is necessary at very strong contact discontinuities such as the surface of a neutron star. Without this additional limiter, ϵ may become ill-conditioned. This typically happens when ϵ is very small and the correction becomes larger than ϵ itself potentially leading to negative ϵ . For some equations of state, $\epsilon < 0$ is ill-defined, causing the HLLC Riemann solver to fail. In practice, this reduction does not affect the overall accuracy of the scheme. We also note that this special treatment does not *forbid* ϵ from becoming negative.

Finally, we recompute $a_{i,L}$ ($a_{i,R}$) according to

$$a_{i,L(R)} = a_i + (a_{i,L(R)} - a_i) \frac{(D^2a)_{i+\frac{1}{2},\text{lim}}}{(D^2a)_i}. \quad (\text{B21})$$

In case the denominator becomes zero in the expression above, we set the last term to zero.

Finally, if (B9) does not hold, we test whether $|a_{i,R(L)} - a_i| \geq 2|a_{i,L(R)} - a_i|$ holds. In that case, we set

$$a_{i,R(L)} = a_i - 2(a_{i,L(R)} - a_i) \quad (\text{B22})$$

for either $a_{i,L}$ or $a_{i,R}$, respectively. In the case of reconstructing the specific internal energy ϵ , if $|a_i - 2a_{i,L(R)}| > |a_i|$, we simply set $a_{i,R(L)} = a_i$. This is for the same reason that has been mentioned above already.

After having obtained $a_{i,L}$ and $a_{i,R}$, we apply the ‘‘standard’’ flattening procedure discussed in the Appendix of [98]. This completes the enhanced PPM scheme applied in our code. Note that Ref. [87] (in contrast to [86]) suggests to skip the second step. In our experiments with an excited neutron star and a collapsing stellar core, however, we find that when skipping this step, the scheme becomes too dissipative.

The enhanced PPM scheme requires four ghost points. For efficiency reasons, it may be desirable to use only three ghost points, since less memory and interprocessor communication is required. In order to reduce the number of required stencil points to three, we use fourth-order polynomial interpolation (B1) instead of sixth-order interpolation [86] in the first step, and we skip the check (B17) involving the third derivatives $(D^3a)_i$, assuming that it *does* hold. We also use a modified flattening scheme which allows us to use only three ghost points. This modified flattening scheme is the same as the one presented in the Appendix of [98], but we drop the maximum and neglect the term \tilde{f}_{i+s_i} in Eqn. (A.2) of [98], i.e., we directly use $f_i = \tilde{f}_i$. Since \tilde{f}_i involves derivatives of the pressure, by dropping \tilde{f}_{i+s_i} (where s_i can be +1 or -1), we are able to reduce the number of required stencil points by one. Effectively, this weakens the amount of flattening that is applied, and thus potentially results

in more oscillations at shocks. In our tests, however, we have found only small differences between the four- and three-point scheme. All simulations reported in this work use the three-point scheme.

In Fig. 25, we show the effect of ePPM compared with oPPM on the Hamiltonian constraint H along the x -axis for the example of an isolated TOV star (Sec. III A) on cell-centered and vertex-centered AMR grids. Clearly, ePPM results in a significantly lower error compared to oPPM on vertex-centered, and especially on cell-centered AMR grids.

Appendix C: Atmosphere Treatment

In vacuum, obviously, the equations describing the fluid dynamics break down. When simulating isolated neutron stars or binary neutron star mergers, a large fraction of the simulation domain is physically vacuum. At the surface of the fluid bodies where a sharp transition to vacuum occurs, the Riemann solver breaks down.

As a simple and common solution to this problem, we keep a very low and constant density fluid (the *atmosphere*) in the cells which would be vacuum otherwise. We also keep track of where the evolution of the fluid variables fails to produce a physical state and reset these cells to atmosphere. Typically, there are few such cells, which cluster around the surface of the star. The atmosphere density ρ_{atmo} is usually chosen to be 8 to 10 orders of magnitudes lower than the central density of the fluid body. This ensures that the atmosphere does not contribute noticeably to the total rest mass and energy in the simulation.

Whether a given fluid cell is set to atmosphere values is decided depending on the local fluid density. If it drops below atmosphere density ρ_{atmo} , the cell is set to atmosphere density with zero fluid velocity.

More specifically, we proceed in the following way.

1. During each intermediate time step, we set an “atmosphere” flag in an atmosphere mask M_A if $\tau + \Delta t R_\tau < 0$ or $D + \Delta t R_D < 0$, where R_τ and R_D are the right-hand sides of the τ and D equations (5), respectively and Δt is the temporal timestep size. In addition to setting the atmosphere flag, we also set all fluid right-hand sides for that cell to zero, in effect freezing the further evolution of this cell. In that case, we also skip conversion of conserved to primitive variables of that cell.
2. After a full time step, we set all variables of those cells to atmosphere values that are flagged as atmosphere.
3. Finally, we clear the atmosphere mask M_A .

Furthermore, we perform the following operations involving atmosphere checks:

1. After reconstruction, we check whether the reconstructed primitive density is below atmosphere density. If this is the case, we enforce first order reconstruction, i.e. we set left and right cell face $a_{i,L} = a_{i,R} = a_i$ to the cell average a_i for all primitive variables.
2. At the end of conservative to primitive conversion, we check whether the new set of primitive variables is below atmosphere level for a given cell. If this is the case, we reset that cell to atmosphere level.

In the two cases above, the atmosphere mask is not set.

To limit high-frequency noise in cells slightly above atmosphere level, we set cells to atmosphere value if they are within a given tolerance δ above atmosphere density, i.e. we test whether

$$\rho \leq \rho_{\text{atmo}} (1 + \delta). \quad (\text{C1})$$

In the cases considered here, we set $\delta = 0.001$.

The particular treatment of vacuum regions by enforcing a low density atmosphere is not ideal and has several drawbacks. If a cell is forced to be not lower than a particular minimum density, small amounts of baryonic mass can be created or removed. This breaks the strictly conservative nature of our hydrodynamics scheme and can thus lead to small errors. As noted in [141], introducing an artificial atmosphere may also change the local wave structure of the solution. An artificial low density atmosphere can be avoided by modifying the Riemann solver at those cells adjacent to vacuum cells [141]. In practice, however, if the atmosphere level is sufficiently low, the negative influence on the fluid evolution can be neglected.

Appendix D: Scheduling of Ghost-Zone Synchronization

We find that excessive inter-processor and inter-patch synchronization of ghost zone information can lead to significant performance drawbacks, especially on large numbers of processing units ($\gtrsim 1000$). We have thus optimized our ghost-zone update pattern and reduced the number of necessary synchronization calls.

We distinguish between three different synchronization update operations: (i) inter-processor and inter-patch synchronizations performed after each intermediate time step, and (ii) AMR buffer-zone prolongation performed after each full time step, and (iii) AMR prolongation after regridding (see [81] on the latter two cases for details).

We distinguish between two sets of variables. One set is comprised of the spacetime variables $\{\phi, \tilde{\gamma}_{ij}, K, \tilde{A}_{ij}, \tilde{\Gamma}^i, \alpha, \beta^i, B^i\}$ describing the curvature evolution and gauge (Sec. II B), and the other set is comprised of variables $\{D, \tau, S_i, \rho, \epsilon, v^i, \tilde{v}^i, P, W, Y_e, Y_e^{\text{con}}, T, s\}$ describing the evolution of the fluid elements (Sec. II A). The

TABLE VII. Required synchronizations for each quantity for the three synchronization operations. See text for more details.

Operation	inter-processor/inter-patch sync.	prolongation (buffer zone)	prolongation (regridding)
Quantities	$\{\phi, \tilde{\gamma}_{ij}, K, \tilde{A}_{ij}, \tilde{\Gamma}^i, \alpha, \beta^i, B^i\}$ $\{D, \tau, S_i, Y_e^{\text{con}}\}$ $\{M_A\}$	$\{\phi, \tilde{\gamma}_{ij}, K, \tilde{A}_{ij}, \tilde{\Gamma}^i, \alpha, \beta^i, B^i\}$ $\{D, \tau, S_i, Y_e^{\text{con}}, \rho, \epsilon, v^i, T\}$ $\{M_A\}$	$\{\phi, \tilde{\gamma}_{ij}, K, \tilde{A}_{ij}, \tilde{\Gamma}^i, \alpha, \beta^i, B^i\}$ $\{D, \tau, S_i, Y_e^{\text{con}}, \rho, \epsilon, v^i, Y_e, T, s\}$

primitive electron fraction Y_e , the conserved electron fraction Y_e^{con} , the temperature T , and the specific entropy s are only necessary when microphysical finite-temperature equations of state are used. In addition to these two sets of variables, we also need to consider the "pseudo-evolved" atmosphere mask M_A described in Appendix C. Thus, in total, we have $24 + 19 + 1 = 44$ evolved components that potentially need to be synchronized.

As described in Sec. IIB, the update terms for the spacetime variables are computed via finite differences and thus require ghost-zone synchronization after each intermediate step. In addition, they are also subject to AMR buffer-zone synchronization via prolongation to obtain valid ghost data from the coarse grid in the buffer zone.

As described in Sec. IIA, the update terms for the evolved conserved fluid variables are computed from reconstructed primitive variables at cell interfaces and thus also require ghost and buffer-zone synchronization in the same way as the spacetime variables. The conservative to primitive conversion requires the conserved variables and valid initial guesses for the primitive variables. Typically, these initial guesses are taken from the last valid time step on the given cell. Since cells located in the buffer zone become invalid during time integration substeps and need to be refilled via buffer-zone prolongation after a full time step, we also need to synchronize those primitive variables that are used as initial guesses in the conservative to primitive conversion. In our case, these are ρ , ϵ , v^i , and T . Note that we do not need to synchronize the *global* primitive velocity \tilde{v}^i since it is later obtained from a coordinate transformation.

Furthermore, we need to update the atmosphere mask M_A in each intermediate step via inter-processor and inter-patch synchronization, and also via buffer-zone prolongation after each full time step. This is necessary because the atmosphere mask is only set on cells of the evolved grid (i.e. all cells excluding ghost zones). Operations like conservative to primitive conversion, which depend on the atmosphere mask, are performed on the entire grid, including ghost zones. Thus they require a synchronized atmosphere mask. In addition, the synchronization order of the atmosphere mask is important during buffer-zone prolongation. Before prolongating all other required quantities, we first prolongate the atmosphere mask. Immediately afterwards, cells are set to atmosphere values according to the atmosphere mask. The atmosphere mask itself is cleared (also see Appendix C).

This completes the evolution step and all variables are in their final state for the given evolution step. Now, it is possible to prolongate also all remaining variables as discussed above.

Finally, we need to synchronize all variables (except for the atmosphere mask¹¹) via prolongation after regridding. A subsequent conservative to primitive conversion ensures that the two conservative and primitive sets of hydrodynamical variables are consistent with each other. Even though regridding requires all variables to be synchronized and is thus rather expensive, fortunately, this operation usually does occur only infrequently, say every 64 iterations, when moving the fine grids during binary neutron star evolution, and only very infrequently, say every couple of thousands of iterations, when adding additional refinement levels during stellar collapse or neutron star collapse.

In Table VII, we explicitly list all quantities that must be updated during one of the three possible synchronization operations. The most frequent operation, inter-processor and inter-patch synchronization require the least number of variables to be updated. Prolongation during regridding, which is the least frequent synchronization operation, requires the full set of variables (except for the atmosphere mask M_A which is invalid outside of a full time integration step). Also note that the global primitive velocity \tilde{v}^i never needs to be synchronized because it is obtained from the local primitive velocity v^i via a coordinate transformation after each synchronization step. Similarly, the Lorentz factor W and the pressure P are never synchronized since they are computed in the conservative to primitive routine, which is executed after each synchronization operation.

Appendix E: Volume Integration

Several quantities in our code require volume integration over the entire numerical grid. For instance, the total baryonic mass is given by

$$M_B = \int d^3x D(x, y, z) \quad (\text{E1})$$

¹¹ The atmosphere mask does not need to be synchronized because it is not valid during regridding. As explained in Appendix C, it is only valid during time integration substeps where regridding is not allowed. We clear it in any new grid region.

in terms of the conserved density D in the Cartesian tensor basis¹². In Cartesian coordinates, this can be approximated numerically by

$$M_B = \Delta x \Delta y \Delta z \sum_{ijk} D_{ijk}, \quad (\text{E2})$$

where Δx , Δy , and Δz is the grid spacing and the indices i, j, k , in this context, denote grid indices. In generic curvi-linear coordinates, the global grid spacing is not constant anymore. In order to compute the volume integral with respect to global coordinates, we make use of the local volume element

$$d^3u = \Delta u \Delta v \Delta w, \quad (\text{E3})$$

where Δu , Δv , and Δw denote the local uniform grid spacing, and we make use of the relation between local volume form d^3u and global volume form

$$d^3x = d^3u \left| \det \frac{\partial x^i}{\partial u^j} \right|. \quad (\text{E4})$$

The volume form d^3x is introduced as an additional grid function which can be computed once the coordinates and grids are set up.

Next, we need to take into account the non-trivial overlap between neighboring grid patches. For instance, the spherical boundary of the spherical outer grid (Fig. 1) cuts through cells of the central Cartesian patch, i.e., parts of the Cartesian cells reach into the nominal domain of the spherical grid. Consequently, the volume associated with each of those cells is only a fraction of the volume of the entire cell. In practice, we set up a weight mask \mathcal{W}_{ijk} defining the contribution of each cell to the total volume. A cell fully contained on the nominal grid has a weight of $\mathcal{W}_{ijk} = 1$. Correspondingly, a cell completely outside of the nominal grid has a weight of $\mathcal{W}_{ijk} = 0$. Cells, whose vertices are not all on the nominal grid, carry a weight $0 < \mathcal{W}_{ijk} < 1$. In that case, we determine the weight by using 3D Monte-Carlo integration [e.g., 112] of the volume fraction of the overlapping regions. The weights need to be calculated only once after the grids have been setup and therefore the cost of Monte Carlo volume integration is negligible compared to the total cost of the simulation.

For simplicity, we absorb the weight mask into the volume form (E4), i.e., we effectively store

$$(d^3x)_{ijk} = \Delta u \Delta v \Delta w \left| \det \frac{\partial x^l}{\partial u^m} \right|_{ijk} \mathcal{W}_{ijk}, \quad (\text{E5})$$

¹² We remark that our code uses the conserved density D in the *local* coordinate basis. Since D is a *densitized* scalar, (E1) requires an additional Jacobian factor to transform D to the global basis. For simplicity of discussion, we omit this here and temporarily assume that D is given in the global basis.




FIG. 26. Binary neutron stars: influence of the outer boundary on the accuracy of the wave extraction and evolution. The upper panel shows the “+” polarization of the Weyl scalar $D\Psi_4$ extracted via CCE for the two setups with different outer boundary locations. At time $t \simeq 7.5$ ms, when the outer boundary in the setup with $R_B = 2000M_\odot$ comes in causal contact with the interior evolution, differences start to become visible for the $R_B = 2000M_\odot$ setup: the amplitude of $D\Psi_4$ deviates by $\simeq 7\%$, the phase ϕ deviates by $\simeq 0.2$ rad, and the L_1 -norm of the Hamiltonian constraint $\|H\|_1$ is larger by $\simeq 15\%$.

where the indices i, j, k label grid points and are not subject to the Einstein sum convention. Similar to the Jacobians introduced for computing global Cartesian derivatives from local finite differences, any volume integration needs to take into account (E5). For instance (E2) takes the form

$$M_B = \sum_{ijk} D_{ijk} (d^3x)_{ijk}. \quad (\text{E6})$$

Appendix F: Influence of the Outer Boundary

All GR binary neutron star merger simulations to date employ grids which are too small to allow for causally disconnected outer boundaries. Since no constraint preserving boundary conditions are known for the BSSN evolution system, the simulations may be affected by incoming constraint violations. Thus, it is interesting to investigate the influence of the outer boundary condition on the interior evolution and extracted GWs of the binary neutron star merger problem considered in Sec. III D when the boundary is *not* causally disconnected.

We compare a simulation with outer boundary at

$R_B = 2000M_\odot$ to the simulations in Sec. III D, which use an outer boundary at $R_B = 2800M_\odot$. The setup with $R_B = 2000M_\odot$ has an outer boundary which is in causal contact with the interior evolution and the wave-extraction region during the merger and ring-down phases. All simulations impose an approximate and non-constraint preserving radiative boundary condition (e.g. [21]). We focus on baseline resolution $r1$. We expect the simulations to be very similar at least up to the point when the constraint violations from the outer boundary reach the wave-extraction region which happens at $t \simeq 7.5$ ms.

In Fig. 26, we show the “+” polarization of the leading order harmonic $(\ell, m) = (2, 2)$ mode of the complex Weyl scalar $D\Psi_4$ computed via CCE. The difference in amplitude are on the order of $\simeq 7\%$. The effects on the phase are more subtle and not clearly visible from a simple inspection of the waveform itself. Therefore, in the two panels below, we plot the phase ϕ of the $(\ell, m) = (2, 2)$ mode. The maximum dephasing in the two simulations is $\simeq 0.2$ rad and thus, the systematic dephasing due to the influence from the outer boundary is only slightly below the one due to the convergent numerical error. This indicates that when the resolution is further increased, the error due to constraint violations from the outer boundary cannot be neglected anymore.

In the same figure, we also show the L_1 -norm¹³ of the Hamiltonian constraint $\|H\|_1$ for the two simulations. We find that the difference of $\simeq 15\%$ is smaller than the difference of $\simeq 25\%$ between the numerical resolutions $r1$ and $r2$, but not so small that it can be ignored.

Finally, we also compare mass and spin of the merger remnant, and find that the differences are on the order of the numerical error between resolutions $r1$ and $r2$.

Overall, we find that causally disconnected outer boundaries have a non-negligible impact on the accuracy of the binary neutron star simulation presented in Sec. III D. It is thus likely that longer inspiral simulations are even more strongly affected.

¹³ We show here the L_1 -norm since it does not require a log-scaling. Thus, subtle differences are better visible. We note, however, that the L_2 -norm $\|H\|_2$ shows similar differences.

-
- [1] C. D. Ott, E. Abdikamalov, E. O'Connor, C. Reisswig, R. Haas, P. Kalmus, S. Drasco, A. Burrows, and E. Schnetter, *Phys. Rev. D.* **86**, 024026 (2012), arXiv:1204.0512 [astro-ph.HE].
- [2] C. D. Ott, C. Reisswig, E. Schnetter, E. O'Connor, U. Sperhake, F. Löffler, P. Diener, E. Abdikamalov, I. Hawke, and A. Burrows, *Phys. Rev. Lett.* **106**, 161103 (2011).
- [3] T. Kuroda, K. Kotake, and T. Takiwaki, *Astrophys. J.* **755**, 11 (2012), arXiv:1202.2487 [astro-ph.HE].
- [4] T. Kuroda and H. Umeda, *Astrophys. J. Supp. Ser.* **191**, 439 (2010).
- [5] L. Baiotti, I. Hawke, and L. Rezzolla, *Class. Quantum Grav.* **24**, 187 (2007).
- [6] S. Bernuzzi, A. Nagar, M. Thierfelder, and B. Bruegmann, arXiv:1205.3403 (2012).
- [7] L. Baiotti, T. Damour, B. Giacomazzo, A. Nagar, and L. Rezzolla, *Phys. Rev. D.* **84**, 024017 (2011).
- [8] B. Giacomazzo, L. Rezzolla, and L. Baiotti, *Phys. Rev. D.* **83**, 044014 (2011).
- [9] L. Rezzolla, B. Giacomazzo, L. Baiotti, J. Granot, C. Kouveliotou, and M. A. Aloy, *Astrophys. J. Lett.* **732**, L6 (2011).
- [10] Z. B. Etienne, Y. T. Liu, V. Paschalidis, and S. L. Shapiro, *Phys. Rev. D.* **85**, 064029 (2012).
- [11] K. Kiuchi, Y. Sekiguchi, K. Kyutoku, and M. Shibata, *Class. Quantum Grav.* **29**, 124003 (2012).
- [12] W. E. East and F. Pretorius, *Astrophys. J. Lett.* **760**, L4 (2012).
- [13] M. Anderson, E. W. Hirschmann, L. Lehner, S. L. Liebling, P. M. Motl, D. Neilsen, C. Palenzuela, and J. E. Tohline, *Phys. Rev. Lett.* **100**, 191101 (2008).
- [14] J. A. Faber and F. A. Rasio, *Liv. Rev. Rel.* **15**, 8 (2012), arXiv:1204.3858 [gr-qc].
- [15] C. D. Ott, H. Dimmelmeier, A. Marek, H.-T. Janka, I. Hawke, B. Zink, and E. Schnetter, *Phys. Rev. Lett.* **98**, 261101 (2007).
- [16] C. D. Ott, H. Dimmelmeier, A. Marek, H.-T. Janka, B. Zink, I. Hawke, and E. Schnetter, *Class. Quantum Grav.* **24**, 139 (2007).
- [17] S. Scheidegger, T. Fischer, S. C. Whitehouse, and M. Liebendörfer, *Astron. Astrophys.* **490**, 231 (2008).
- [18] S. Scheidegger, R. Käppeli, S. C. Whitehouse, T. Fischer, and M. Liebendörfer, *Astron. Astrophys.* **514**, A51 (2010).
- [19] J. Winicour, *Liv. Rev. Rel.* **12** (2009), <http://www.livingreviews.org/lrr-2009-3>.
- [20] N. T. Bishop, R. Gómez, L. Lehner, M. Maharaj, and J. Winicour, *Phys. Rev. D.* **56**, 6298 (1997).
- [21] C. Reisswig, C. D. Ott, U. Sperhake, and E. Schnetter, *Phys. Rev. D.* **83**, 064008 (2011).
- [22] C. Reisswig, N. T. Bishop, D. Pollney, and B. Szilágyi, *Phys. Rev. Lett.* **103**, 221101 (2009).
- [23] C. Reisswig, N. T. Bishop, D. Pollney, and B. Szilágyi, *Class. Quantum Grav.* **27**, 075014 (2010).
- [24] M. C. Babiuc, B. Szilágyi, J. Winicour, and Y. Zlochower, *Phys. Rev. D.* **84**, 044057 (2011).
- [25] J. Winicour, *Class. Quantum Grav.* **29**, 113001 (2012).
- [26] D. Brown, O. Sarbach, E. Schnetter, M. Tiglio, P. Diener, I. Hawke, and D. Pollney, *Phys. Rev. D* **76**, 081503(R) (2007), arXiv:0707.3101 [gr-qc].
- [27] D. Brown, P. Diener, O. Sarbach, E. Schnetter, and M. Tiglio, *Phys. Rev. D* **79**, 044023 (2009).
- [28] T. Zwerger and E. Müller, *Astron. Astrophys.* **320**, 209 (1997).
- [29] H. Dimmelmeier, J. A. Font, and E. Müller, *Astron. Astrophys.* **393**, 523 (2002).
- [30] H. Dimmelmeier, J. A. Font, and E. Müller, *Astron. Astrophys.* **388**, 917 (2002).
- [31] H. Dimmelmeier, C. D. Ott, H.-T. Janka, A. Marek, and E. Müller, *Phys. Rev. Lett.* **98**, 251101 (2007).
- [32] E. B. Abdikamalov, C. D. Ott, L. Rezzolla, L. Dessart, H. Dimmelmeier, A. Marek, and H. Janka, *Phys. Rev. D.* **81**, 044012 (2010).
- [33] B. Müller, H.-T. Janka, and H. Dimmelmeier, *Astrophys. J. Supp. Ser.* **189**, 104 (2010), arXiv:1001.4841 [astro-ph.SR].
- [34] F. Hanke, A. Marek, B. Müller, and H.-T. Janka, *Astrophys. J.* **755**, 138 (2012), arXiv:1108.4355 [astro-ph.SR].
- [35] T. Takiwaki, K. Kotake, and Y. Suwa, *Astrophys. J.* **749**, 98 (2012), arXiv:1108.3989 [astro-ph.HE].
- [36] H. Dimmelmeier, N. Stergioulas, and J. A. Font, *Mon. Not. Roy. Astron. Soc.* **368**, 1609 (2006).
- [37] E. B. Abdikamalov, H. Dimmelmeier, L. Rezzolla, and J. C. Miller, *Mon. Not. Roy. Astron. Soc.* **392**, 52 (2009).
- [38] S. S. Komissarov, *Mon. Not. Roy. Astron. Soc.* **367**, 19 (2006), arXiv:astro-ph/0510310.
- [39] J. A. Font and J. M. A. Ibanez, *Astrophys. J.* **494**, 297 (1998).
- [40] J. A. Font and F. Daigne, *Mon. Not. Roy. Astron. Soc.* **334**, 383 (2002), arXiv:astro-ph/0203403.
- [41] O. Zanotti, L. Rezzolla, and J. A. Font, *Mon. Not. Roy. Astron. Soc.* **341**, 832 (2003), arXiv:gr-qc/0210018.
- [42] R. Narayan, A. Sądowski, R. F. Penna, and A. K. Kulkarni, *Mon. Not. Roy. Astron. Soc.* **426**, 3241 (2012), arXiv:1206.1213 [astro-ph.HE].
- [43] M. Alcubierre and M. D. Mendez, *Gen.Rel.Grav.* **43**, 2769 (2011).
- [44] P. J. Montero and I. Cordero-Carrión, *Phys. Rev. D.* **85**, 124037 (2012).
- [45] I. Cordero-Carrión and P. Cerdá-Durán, arXiv:1211.5930 (2012).
- [46] T. W. Baumgarte, P. J. Montero, I. Cordero-Carrión, and E. Müller, arXiv:1211.6632 (2012).
- [47] O. Korobkin, E. B. Abdikamalov, E. Schnetter, N. Stergioulas, and B. Zink, *Phys. Rev. D.* **83**, 043007 (2011).
- [48] O. Korobkin, E. Abdikamalov, N. Stergioulas, E. Schnetter, B. Zink, S. Rosswog, and C. D. Ott, ArXiv e-prints (2012), arXiv:1210.1214 [astro-ph.HE].
- [49] J. Thornburg, *Class. Quant. Grav.* **21**, 743 (2004).
- [50] E. Pazos *et al.*, *Class. Quantum Grav.* **24**, S341 (2007).
- [51] J. Thornburg, *Class. Quantum Grav.* **21**, 3365 (2004), arxiv:gr-qc/0404059.
- [52] E. N. Dorband, E. Berti, P. Diener, E. Schnetter, and M. Tiglio, *Phys. Rev. D* **74**, 084028 (2006), arXiv:gr-qc/0608091.
- [53] E. Pazos, M. Tiglio, M. D. Duez, L. E. Kidder, and S. A. Teukolsky, *Phys. Rev. D.* **80**, 024027 (2009).
- [54] B. Zink, E. Schnetter, and M. Tiglio, *Phys. Rev. D.* **77**, 103015 (2008).
- [55] M. Ansorg, *Class. Quantum Grav.* **24**, S1 (2007).

- [56] H. P. Pfeiffer, L. E. Kidder, M. A. Scheel, and S. A. Teukolsky, *Comp. Phys. Comm.* **152**, 253 (2003).
- [57] E.ourgoulhon, P. Grandclément, K. Taniguchi, J.-A. Marck, and S. Bonazzola, *Phys. Rev. D* **63**, 064029 (2001), arXiv:gr-qc/0007028.
- [58] E.ourgoulhon, P. Grandclément, and S. Bonazzola, *Phys. Rev. D* **65**, 044020 (2002).
- [59] P. Grandclément, S. Bonazzola, E.ourgoulhon, and J.-A. Marck, *J. Chem. Phys.* **170**, 231 (2001).
- [60] C. Reisswig, N. T. Bishop, C. W. Lai, J. Thornburg, and B. Szilagyi, *Class. Quantum Grav.* **24**, 327 (2007).
- [61] R. Gómez, W. Barreto, and S. Frittelli, *Phys. Rev. D* **76**, 124029 (2007).
- [62] M. A. Scheel, M. Boyle, T. Chu, L. E. Kidder, K. D. Matthews, and H. P. Pfeiffer, *Phys. Rev. D* **79**, 024003 (2009).
- [63] B. Szilágyi, L. Lindblom, and M. A. Scheel, *Phys. Rev. D* **80**, 124010 (2009).
- [64] T. Chu, H. P. Pfeiffer, and M. A. Scheel, *Phys. Rev. D* **80**, 124051 (2009), arXiv:0909.1313.
- [65] G. Lovelace, M. Boyle, M. A. Scheel, and B. Szilágyi, *Class. Quantum Grav.* **29**, 045003 (2012).
- [66] L. T. Buchman, H. P. Pfeiffer, M. A. Scheel, and B. Szilagyi, arxiv:1206.3015 (2012).
- [67] I. MacDonald, A. H. Mroue, H. P. Pfeiffer, M. Boyle, L. E. Kidder, M. A. Scheel, B. Szilagyi, and N. W. Taylor, arXiv:1210.3007 (2012).
- [68] M. A. Scheel, H. P. Pfeiffer, L. Lindblom, L. E. Kidder, O. Rinne, and S. A. Teukolsky, *Phys. Rev. D* **74**, 104006 (2006), arXiv:gr-qc/0607056.
- [69] F. Foucart, M. D. Duez, L. E. Kidder, M. A. Scheel, B. Szilagyi, and S. A. Teukolsky, *Phys. Rev. D* **85**, 044015 (2012).
- [70] F. Foucart, M. D. Duez, L. E. Kidder, and S. A. Teukolsky, *Phys. Rev. D* **83**, 024005 (2011), arXiv:1007.4203 [astro-ph.HE].
- [71] M. D. Duez, F. Foucart, L. E. Kidder, H. P. Pfeiffer, M. A. Scheel, and S. A. Teukolsky, *Phys. Rev. D* **78**, 104015 (2008).
- [72] F. Foucart, L. E. Kidder, H. P. Pfeiffer, and S. A. Teukolsky, *Phys. Rev. D* **77**, 124051 (2008).
- [73] D. Pollney, C. Reisswig, E. Schnetter, N. Dorband, and P. Diener, *Phys. Rev. D* **83**, 044045 (2011).
- [74] D. Pollney, C. Reisswig, N. Dorband, E. Schnetter, and P. Diener, *Phys. Rev. D* **80**, 121502 (2009).
- [75] D. Pollney and C. Reisswig, *Astrophys. J. Lett.* **732**, L13 (2011).
- [76] T. Damour, A. Nagar, D. Pollney, and C. Reisswig, *Phys. Rev. Lett.* **108**, 131101 (2012).
- [77] L. Santamaria, F. Ohme, P. Ajith, B. Bruegmann, N. Dorband, *et al.*, *Phys. Rev. D* **82**, 064016 (2010).
- [78] N. Bishop, D. Pollney, and C. Reisswig, *Class. Quantum Grav.* **28**, 155019 (2011).
- [79] P. Ajith, M. Hannam, S. Husa, Y. Chen, B. Brügmann, N. Dorband, D. Müller, F. Ohme, D. Pollney, C. Reisswig, L. Santamaria, and J. Seiler, *Phys. Rev. Lett.* **106**, 241101 (2011).
- [80] T. Goodale, G. Allen, G. Lanfermann, J. Massó, T. Radke, E. Seidel, and J. Shalf, in *Vector and Parallel Processing – VECPAR’2002, 5th International Conference, Lecture Notes in Computer Science* (Springer, Berlin, 2003).
- [81] E. Schnetter, S. Hawley, and I. Hawke, *Class. Quantum Grav.* **21**, 1465 (2004), gr-qc/0310042.
- [82] E. Schnetter, P. Diener, E. N. Dorband, and M. Tiglio, *Class. Quantum Grav.* **23**, S553 (2006).
- [83] F. Löffler, J. Faber, E. Bentivegna, T. Bode, P. Diener, R. Haas, I. Hinder, B. C. Mundim, C. D. Ott, E. Schnetter, G. Allen, M. Campanelli, and P. Laguna, *Class. Quantum Grav.* **29**, 115001 (2012), arXiv:1111.3344 [gr-qc].
- [84] M. J. Berger and J. Olinger, *J. Chem. Phys.* **53**, 484 (1984).
- [85] W. E. East, F. Pretorius, and B. C. Stephens, *Phys. Rev. D* **85**, 124010 (2012).
- [86] P. Colella and M. D. Sekora, *J. Chem. Phys.* **227**, 7069 (2008).
- [87] P. McCorquodale and P. Colella, *Comm. Appl. Math. Comp. Sci.* **6**, 1 (2011).
- [88] M. Schlegel, O. Knoth, M. Arnold, and R. Wolke, *J. Comp. Appl. Math.* **226**, 345 (2009).
- [89] E. Constantinescu and A. Sandu, *SIAM J. Sci. Comput.* **33**, 239 (2007).
- [90] C. D. Ott, E. Abdikamalov, P. Moesta, R. Haas, S. Drasco, E. O’Connor, C. Reisswig, C. Meakin, and E. Schnetter, Submitted to the *Astrophys. J.*, arXiv:1210.6674 (2012).
- [91] EinsteinToolkit: A Community Toolkit for Numerical Relativity.
- [92] B. Baiotti, I. Hawke, P. J. Montero, F. Löffler, L. Rezzolla, N. Stergioulas, J. A. Font, and E. Seidel, *Phys. Rev. D* **71**, 024035 (2005).
- [93] J. W. York, Jr., in *Gravitational radiation*, edited by N. Deruelle and T. Piran (North-Holland Publishing Company, 1983) pp. 175–201.
- [94] F. Banyuls, J. A. Font, J. M. Ibáñez, J. M. Martí, and J. A. Miralles, *Astrophys. J.* **476**, 221 (1997).
- [95] J. A. Font, *Liv. Rev. Rel.* **11**, 7 (2008).
- [96] L. Baiotti, B. Giacomazzo, and L. Rezzolla, *Phys. Rev. D* **78**, 084033 (2008).
- [97] B. Einfeldt, in *Shock tubes and waves; Proceedings of the Sixteenth International Symposium, Aachen, Germany, July 26–31, 1987. VCH Verlag, Weinheim, Germany* (1988) p. 671.
- [98] P. Colella and P. R. Woodward, *J. Comp. Phys.* **54**, 174 (1984).
- [99] J. M. Hyman, *The Method of Lines Solution of Partial Differential Equations*, Tech. Rep. (ERDA Mathematics and Computing Laboratory, Courant Institute of Mathematical Sciences, New York University, 1976).
- [100] T. Nakamura, K. Oohara, and Y. Kojima, *Prog. Theor. Phys. Suppl.* **90**, 1 (1987).
- [101] M. Shibata and T. Nakamura, *Phys. Rev. D* **52**, 5428 (1995).
- [102] T. W. Baumgarte and S. L. Shapiro, *Phys. Rev. D* **59**, 024007 (1999).
- [103] M. Alcubierre, B. Brügmann, T. Dramlitsch, J. A. Font, P. Papadopoulos, E. Seidel, N. Stergioulas, and R. Takahashi, *Phys. Rev. D* **62**, 044034 (2000).
- [104] M. Alcubierre, *Introduction to 3+1 Numerical Relativity, by Miguel Alcubierre. ISBN 978-0-19-920567-7 (HB). Published by Oxford University Press, Oxford, UK, 2008.*, edited by Alcubierre, M. (Oxford University Press, 2008).
- [105] C. Bona, J. Massó, E. Seidel, and J. Stela, *Phys. Rev. Lett.* **75**, 600 (1995).
- [106] M. Alcubierre, B. Brügmann, P. Diener, M. Koppitz, D. Pollney, E. Seidel, and R. Takahashi, *Phys. Rev. D*

- 67**, 084023 (2003).
- [107] E. Schnetter, *Class. Quantum Grav.* **27**, 167001 (2010).
- [108] D. Müller and B. Brügmann, *Class. Quantum Grav.* **27**, 114008 (2010).
- [109] P. Diener, E. N. Dorband, E. Schnetter, and M. Tiglio, *J. Sci. Comput.* **32**, 109 (2007).
- [110] C.-W. Shu, *Lecture Notes in Mathematics* **1697**, 325 (1998).
- [111] M. J. Berger and P. Colella, *J. Comp. Phys.* **82**, 64 (1989).
- [112] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes, 3rd. edition* (Cambridge University Press, Cambridge, U. K., 2007).
- [113] K. S. Thorne, *Rev. Mod. Phys.* **52**, 299 (1980).
- [114] M. Shibata and Y. Sekiguchi, *Phys. Rev. D* **69**, 084024 (2004).
- [115] H. Dimmelmeier, C. D. Ott, A. Marek, and H.-T. Janka, *Phys. Rev. D* **78**, 064056 (2008).
- [116] A. Nagar and L. Rezzolla, *Class. Quantum Grav.* **22**, 167 (2005).
- [117] R. Penrose, *Phys. Rev. Lett.* **10**, 66 (1963).
- [118] E. T. Newman and R. Penrose, *J. Math. Phys.* **3**, 566 (1962).
- [119] C. Reisswig, N. T. Bishop, and D. Pollney, arXiv:1208.3891 (2012).
- [120] C. Reisswig and D. Pollney, *Class. Quantum Grav.* **28**, 195015 (2011).
- [121] C. O. Lousto and Y. Zlochower, *Phys. Rev. D* **76**, 041502 (2007).
- [122] M. Ruiz, M. Alcubierre, D. Núñez, and R. Takahashi, *Gen. Rel. Grav.* **40**, 2467 (2008).
- [123] J. Thornburg, *Phys. Rev. D* **54**, 4899 (1996).
- [124] O. Dreyer, B. Krishnan, E. Schnetter, and E. Shoemaker, *Phys. Rev. D* **67**, 024018 (2003).
- [125] R. C. Tolman, *Phys. Rev.* **55**, 364 (1939).
- [126] J. R. Oppenheimer and G. M. Volkoff, *Phys. Rev.* **55**, 374 (1939).
- [127] L. Baiotti, S. Bernuzzi, G. Corvino, R. de Pietri, and A. Nagar, *Phys. Rev. D* **79**, 024002 (2009).
- [128] H. Komatsu, Y. Eriguchi, and I. Hachisu, *Mon. Not. Roy. Astron. Soc.* **237**, 355 (1989).
- [129] H. Komatsu, Y. Eriguchi, and I. Hachisu, *Mon. Not. Roy. Astron. Soc.* **239**, 153 (1989).
- [130] H.-T. Janka, T. Zwerger, and R. Mönchmeyer, *Astron. Astrophys.* **268**, 360 (1993).
- [131] E. Berti, V. Cardoso, and A. O. Starinets, *Class. Quantum Grav.* **26**, 163001 (2009).
- [132] M. Thierfelder, S. Bernuzzi, and B. Brügmann, *Phys. Rev. D* **84**, 044012 (2011).
- [133] S. Bernuzzi, M. Thierfelder, and B. Brügmann, *Phys. Rev. D* **85**, 104030 (2012).
- [134] R. Gold, S. Bernuzzi, M. Thierfelder, B. Bruegmann, and F. Pretorius, arXiv:1109.5128 (2011).
- [135] LORENE, “LORENE: Langage Objet pour la RELativité Numérique,”.
- [136] L. Baiotti, M. Shibata, and T. Yamamoto, *Phys. Rev. D* **82**, 064015 (2010), arXiv:1007.1754 [gr-qc].
- [137] L. Baiotti, B. Giacomazzo, and L. Rezzolla, *Classical and Quantum Gravity* **26**, 114005 (2009).
- [138] M. Boyle, A. Buonanno, L. E. Kidder, A. H. Mroue, Y. Pan, H. P. Pfeiffer, and M. A. Scheel, *Phys. Rev. D* **78**, 104020 (2008).
- [139] G. A. Sod, *Journal of Computational Physics* **27**, 1 (1978).
- [140] J. M. Martí and E. Müller, *Liv. Rev. Rel.* **6**, 7 (2003).
- [141] E. F. Toro, *Riemann Solvers and Numerical Methods for Fluid Dynamics* (Springer, Berlin, 1999).