

Simulations of inspiraling and merging double neutron stars using the Spectral Einstein Code

Roland Haas,^{1,2} Christian D. Ott,^{3,2} Bela Szilagyi,^{4,2} Jeffrey D. Kaplan,² Jonas Lippuner,²
Mark A. Scheel,² Kevin Barkett,² Curran D. Muhlberger,⁵ Tim Dietrich,¹ Matthew D. Duez,⁶
Francois Foucart,^{7,8} Harald P. Pfeiffer,⁷ Lawrence E. Kidder,⁵ and Saul A. Teukolsky⁵

¹Max Planck Institute for Gravitational Physics (Albert Einstein Institute), Am Mühlenberg 1, Potsdam-Golm, 14476, Germany

²TAPIR, Walter Burke Institute for Theoretical Physics, MC 350-17, California

Institute of Technology, 1200 E California Blvd, Pasadena CA 91125, USA

³Yukawa Institute for Theoretical Physics, Kyoto University, Kyoto, Japan

⁴Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr. Pasadena CA, 91109, USA

⁵Cornell Center for Astrophysics and Planetary Science, Cornell University, Ithaca, New York 14853, USA

⁶Department of Physics & Astronomy, Washington State University, Pullman, Washington 99164, USA

⁷Canadian Institute for Theoretical Astrophysics, University of Toronto, Toronto, Ontario M5S 3H8, Canada

⁸Lawrence Berkeley National Laboratory, 1 Cyclotron Rd, Berkeley, CA 94720, USA; Einstein Fellow

(Dated: April 5, 2016)

We present results on the inspiral, merger, and post-merger evolution of a neutron star – neutron star (NSNS) system. Our results are obtained using the hybrid pseudospectral–finite volume Spectral Einstein Code (SpEC). To test our numerical methods, we evolve an equal-mass system for ≈ 22 orbits before merger. This waveform is the longest waveform obtained from fully general-relativistic simulations for NSNSs to date. Such long (and accurate) numerical waveforms are required to further improve semi-analytical models used in gravitational wave data analysis, for example the effective one body models. We discuss in detail the improvements to SpEC’s ability to simulate NSNS mergers, in particular mesh refined grids to better resolve the merger and post-merger phases. We provide a set of consistency checks and compare our results to NSNS merger simulations with the independent BAM code. We find agreement between them, which increases confidence in results obtained with either code. This work paves the way for future studies using long waveforms and more complex microphysical descriptions of neutron star matter in SpEC.

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

I. INTRODUCTION

The Advanced Laser Interferometer Gravitational-Wave Observatory (LIGO) demonstrated its capability of measuring gravitational wave (GW) signals coming from compact binary systems and opened up a new window for astrophysical observations. Although the first detected GW signal [1, 2] was emitted by a binary made up of two black holes (BH), neutron star – neutron star (NSNS) systems are promising sources [3]. With further upgrades of LIGO and with the entire GW network consisting of LIGO, Virgo, and Kamioka Gravitational Wave Detector (KAGRA) [4, 5] operating, between 0.2 and 200 NSNS mergers per year [6] are expected to be observed.

The GWs emitted during the inspiral and merger contain unique information about the binary’s properties and about each binary constituent. In the case of NS systems, information about the equation of state (EOS) at supranuclear densities can be obtained that is not easily obtainable otherwise [4].

The GW signal of a NSNS merger can be roughly separated into three phases: the *inspiral* phase in which the neutron stars (NS) approach each other due to the emission of GWs, the *merger* phase in which the stars come

in contact and form a single object,¹ and the *post-merger* phase in which the remnant can (i) collapse promptly to a BH leading to a characteristic ringdown GW signal; (ii) form a hypermassive neutron star (HMNS) which is stabilized primarily by angular momentum over secular time scales before BH formation; (iii) form a long-term stable (supramassive or massive) NS if the total mass of the system is sufficiently low. The inspiral part of the signal sweeps through the most sensitive band of current GW detectors. The merger, post-merger, and ringdown parts of the signal are at high frequencies and difficult to observe unless the event is very close.

For most of the lifetime of the binary, the NSs are well separated and the signal is almost sinusoidal with slowly changing frequency. During this phase, the wave signal can be approximated by post-Newtonian (PN) theory to very high accuracy (see [7] and references therein). The binary orbit evolves adiabatically under the influence of the radiation reaction force. Close to merger and in the post-merger phase, PN theory is no longer valid and numerical relativity simulations are needed to correctly capture the fast dynamics and construct GW templates, see,

¹ We define the actual moment of merger as the time at which the GW strain has its maximum.

e.g., [8–10].

This late, high-frequency part of the GW signal is most interesting since it is directly affected by the star’s EOS. The EOS, via the tidal deformability of the NSs, leaves a clear imprint on the late inspiral and early merger phases of the GW signal [11–15]. In addition, also the post-merger frequency and the time evolution of the GW signal from the merger remnant can constrain the EOS via GW observations, see, e.g., [9, 10, 16–18].

While the original PN expansion breaks down before merger, the effective one body (EOB) model [19–22] provides techniques to extend the range of applicability of PN theory into the late inspiral phase, also including tidal effects [10, 23]. The tidal parameters of the model can be linked to the parameters of the EOS of the NSs, making it possible to determine EOS parameters from the GW signal. Similarly, there is interest in “universal” relations between observable quantities that are independent of the EOS [10, 14, 24–27]. Models accurately describing NS coalescences beyond the merger are still missing and only numerical simulations in full general relativity can give reliable information about this stage, but see [28] for a first attempt of a reduced-order model of the post-merger waveform.

Numerical simulations are needed to validate and calibrate EOB models, e.g. [29, 30]. This is possible only if (i) numerical waveforms have a sufficient length and span multiple orbits, so that PN approximations are valid at the beginning of the simulated period and (ii) are sufficiently accurate, i.e. having small eccentricities and small phase errors. Several authors [14, 30–33] have studied the detectability of tidal effects in detail, investigating the errors and uncertainties and the effect of different equations of state on the wave signal. A careful treatment of numerical errors and extrapolation procedures using the numerical data was found to be required to correctly estimate the quality of the simulation results. Some of us recently presented the results for spinning NSNS inspirals [34] and merger simulations of NS binaries including neutrino transport [35] using the Spectral Einstein Code (SpEC). SpEC simulations involving black hole – neutron star (BHNS) were performed [36–40] and binary black hole (BBH) [41–46] simulations using SpEC have a long history. This paper follows the line of work focusing on the accuracy and feasibility of constructing sufficiently long and accurate GW templates. For this purpose, we extended SpEC’s NSNSs simulating capabilities. SpEC employs a hybrid approach using pseudospectral methods for the spacetime evolution and finite volume or finite differencing methods for the hydrodynamical variables. This allows us to achieve very high phase accuracy at low computational costs for the spacetime part of the evolution and to exploit well-tested, stable high-resolution shock capturing methods for the fluid variables. SpEC uses a comoving coordinate system which reduces movement of the NSs on the grid and therefore reduces possible errors accumulated in moving-box mesh-refined schemes by the restriction and prolongation op-

eration as well as Eulerian advection errors. This paper provides numerical tests of these methods in SpEC, paving the way for a more systematic comparison with existing codes to simulate NSNS [47–51] systems and long NSNS inspiral simulations involving more realistic EOS. A first step toward such a comparison is made by comparing our data with a BAM waveform of [52]. We find very good agreement and a phase differences below 0.25 rad up to the end of the inspiral phase.

As outlined in the paragraph above, not only sufficient accuracy is needed in NSNS simulations, the waveforms also need to span multiple orbits before merger to be useful for semi-analytical waveform modeling. Here we consider a binary system with an initial separation of 81 km, which results in more than 22 orbits before merger, i.e., 44 GW cycles before merger. This is to date the longest NSNS merger simulation and the resulting waveform has already been used for the analysis of [53]. Such long simulations can be achieved due to the small computational expense for evolving the metric variables with our pseudospectral approach as well as the small fluid grids which cover the regions around the NSs, instead of the whole simulation volume.

This paper is organized as follows. Sec. II presents the methods used to evolve the spacetime and hydrodynamics sectors of the simulations. Sec. III describes how we construct initial data for NSNS systems. Sec. IV presents results on the convergence and diagnostics on the quality of the computed waveforms. We conclude in Sec. V. In the appendices Sec. A and B, we present convergence tests for a single Tolman-Oppenheimer-Volkoff (TOV) star and investigate the collapse dynamics of isolated NSs, respectively.

Unless stated otherwise, all results use $G = c = 1$ and masses are given in multiples of the solar mass M_{\odot} . ∇_{α} is used to denote the covariant derivative compatible with the 4-metric $g_{\alpha\beta}^{(4)}$.

II. METHODS

A. Two-Domain Approach to General-Relativistic Hydrodynamics

In SpEC we use a mixed approach to solve Einstein’s equations in the generalized harmonic formulation coupled to matter [36, 38]. We solve the evolution equations for the spacetime metric $g_{\alpha\beta}^{(4)}$ using spectral methods as described in [41, 42, 54–61] while the fluid equations are solved using high-resolution shock-capturing methods described in [36, 38, 62]. The NS material is modeled as a perfect fluid with rest mass density ρ_0 , pressure P , specific internal energy ϵ , and 4-velocity u^{α} , so that the stress-energy tensor is given by

$$T_{\alpha\beta} = \rho_0 h^e u_{\alpha} u_{\beta} + P g_{\alpha\beta}^{(4)}, \quad (1)$$

where $h^e = 1 + \epsilon + P/\rho_0$ is the relativistic specific enthalpy. The evolution equations for the conserved hydrodynamical variables $D = \sqrt{-g^{(4)}}u^0\rho_0$, $\tau = E - D$, $S_k = \sqrt{-g^{(4)}}T^0_k$ follow from conservation of stress-energy $\nabla_\alpha T^{\alpha\beta} = 0$ and conservation of baryon number $\nabla_\alpha (Du^\alpha) = 0$, where $g^{(4)}$ is the determinant of the 4-metric and $E = \sqrt{-g^{(4)}}T^{00}$. In this paper, we split the pressure and specific internal energy in cold and thermal pieces due to cold (nuclear force) and thermal contributions, respectively,

$$\epsilon = \epsilon_{\text{cold}}(\rho_0) + \epsilon_{\text{thermal}}, \quad (2)$$

$$P = P_{\text{cold}}(\rho_0) + (\Gamma - 1)\rho_0\epsilon_{\text{thermal}}, \quad (3)$$

$$P_{\text{cold}} = \kappa\rho_0^\Gamma, \quad (4)$$

$$\epsilon_{\text{cold}} = P_{\text{cold}}/[\rho_0(\Gamma - 1)], \quad (5)$$

where κ , Γ are the polytropic constant and the adiabatic index, respectively. In this paper we choose $\Gamma = 2$, $\kappa = 123.6$. Note, however, that SpEC can handle more general EOS than the one presented here [37, 39, 40, 63, 64]. We use the fifth-order weighted essentially non-oscillatory (WENO) reconstruction method of [63, 65–67], a Harten-Lax-van Leer (HLL) Riemann solver [68] to compute numerical fluxes at cell interfaces, and a 2D nonlinear root finding algorithm [63] to recover the primitive variables from the conserved variables at the beginning of each timestep.

Time integration is performed using the method of lines [69] and a 3rd order Runge-Kutta (RK3) method with adaptive step size control based on errors in both spacetime and hydrodynamical variables estimated by comparing second and third order accurate time stepper updates. The metric evolution couples to the hydrodynamical evolution via the stress-energy tensor $T_{\mu\nu}$ and the hydrodynamical evolution equations directly involve the metric and its first derivatives. We interpolate spacetime and hydrodynamical variables between spectral and finite volume grids at the end of each full timestep as described in detail in [36] using almost-spectral interpolation [70] to interpolate from the spectral grid to the finite volume grid and monotonicity preserving polynomial interpolation to interpolate between finite volume and spectral grids. Values for intermediate Runge-Kutta substeps are obtained via extrapolation in time.

Since [62] was published, we optimized this communication scheme to reduce the amount of data sent between compute nodes resulting in improved speed and scalability of the code. These changes include a reduction in the number of times SpEC’s internal dependency tracking recomputes quantities, limit copying of data in memory, and significantly speed up the transformation between collocation point and spectral coefficient representations of data. SpEC now interleaves communication and computation when interpolating data between the spectral and the finite volume grids after each timestep. The number of explicit MPI barriers has also been reduced. Beyond these infrastructure changes, the basic setup described in [36, 62] remains the same.

B. Comoving Coordinate System

SpEC uses a dual frame method [58, 60] to solve Einstein’s equations and the fluid equations. It uses explicit coordinate transformations to map between a set of inertial (physical) coordinates in which the NSs orbit and approach each other and a set of grid coordinates in which the NSs remain at a fixed coordinate location. Once a BH is formed the coordinate transformation is also used to map the excision surface inside of the apparent horizon (AH) to an excision sphere of constant radius in grid coordinates.

The finite volume grid is linked to the spectral grid via a final, piecewise constant in time coordinate transformation. During the inspiral phase of the simulation, the NSs are separated by a vacuum region that gradually shrinks as the stars approach each other and that does not need to be evolved with the matter evolution equations. See Sec. IV A for a detailed description of the grid setup during the individual phases of the simulation. A single finite volume grid which covers both NSs as used in the BH – NS simulations in [38] would thus be inefficient. Instead we cover each star with a separate cubical finite volume grid. In this paper, we choose the boxes to be initially 1.25 times the diameter of the NS. We track the motion of each star and follow the inspiraling stars with the grids by adjusting the grid locations. For the purpose of tracking the stars, we define each star’s position to be the centroids of the rest mass distribution,

$$X_{\text{CM}}^i = \int x^i D \, d^3x, \quad (6)$$

in each of the disjoint finite volume grids. Since the grid patches follow the stars and all fluxes are expressed in the frame comoving with the grid, the fluid velocities are small, which improves the accuracy of Eulerian finite volume codes [71–73].

We employ the remapping criteria outlined in [74] to control the volume covered by the grid. As the NSs spiral inwards, we map the center of each NS from its current position in the physical (inertial) frame to a fixed location in the grid frame in which the spectral basis functions are defined. Since the physical separation between the NSs shrinks with time, but their separation in the grid frame stays constant, the NSs appear to grow (cover more grid points) in the grid frame. This gradually brings the surface of the NSs closer to the grid boundaries and eventually the outer layers of the NSs reach the grid boundaries and matter flows off the grid.

We measure the flux of matter,

$$\dot{M}_0^{(i)} = \int_{S_{i,A}} D \frac{u^i}{u^t} \, d^2x, \quad (7)$$

through each of the outer boundaries $S_{i,\text{outer}}$ of the finite volume grid and through a set of interior surfaces $S_{i,\text{inner}}$, located approximately 44% the distance from the center to the outer boundary. If too much ($\dot{M}_0^{(i)} > 2 \times 10^{-8}$)

matter flows out of the grid through the i -direction outer surface $S_{i,\text{outer}}$, we expand the grid along the i -direction, if not enough matter ($|\dot{M}_0^{(i)}| < 3 \times 10^{-7}$, where $|\dot{M}_0^{(i)}|$ counts the total amount of matter passing through the surface ignoring direction) flows through the i -direction inner surface $S_{i,\text{inner}}$, we contract the grid. We then interpolate the evolved variables onto the new grid. During the inspiral this procedure typically keeps the outer layers of the NS with a rest mass density $\gtrsim 10^{-4}$ of the density in the center of the star $\rho_{0,\text{central}}$ inside of the grid, while lower density material may flow off the grid. In summary, this remapping procedure ensures that the total amount of matter leaving the domain is controlled and that the grid stays as small as possible. A side effect of this procedure is that the effective resolution changes during the simulation. When we expand or shrink the grid in response to the remapping criteria, the number of grid points is kept fixed but the physical volume covered by the grid changes, which leads to a discrete jump in the effective resolution. In our simulation the typical jump in resolution due to grid changes is approximately 10% which can be seen in the discrete jumps in Fig. 11. In addition to these discrete jumps the inspiral of the NSs toward each other causes a continuous increase of resolution since the physical area covered by the fixed number of grid points shrinks as the NSs approach each other in the inertial frame.

Finally, while the interpolation algorithm used is not strictly mass conservative, the remapping happens infrequently enough and with small enough incremental change in the grid size that the effect on the total rest mass is $< 10^{-5}$ of the total rest mass of the system over the course of the simulation. This is several orders of magnitude lower than the amount of material ($\approx 10^{-3}M_\odot$) lost through the outer grid boundaries until an apparent horizon forms, see Fig. 3. After horizon formation any matter outside of the horizon accretes rapidly onto the BH.

C. Mesh Refinement

SpEC employs adaptive mesh refinement in the spectral sector of the evolution equations, adjusting both the order of the spectral basis functions used as well as splitting subdomains into smaller subdomains as required to achieve a desired truncation error. Spectral adaptive mesh refinement (AMR) is described in [75] to which we refer the reader for details.

In the finite volume sector, in order to resolve both the region around each NS and cover a large enough volume to capture outflows and the remnant disk that forms after merger of the NSs and BH formation, we employ a variant of the mesh refinement techniques commonly used in compact binary merger simulations [40]. Often mesh refinement [48, 76–78] is used not only to increase the resolution in regions of interest but also to move the region of higher resolution along with the object. In this approach,

as the NSs move through the grid and get close to the current edge of the high resolution grid, new grid points are created and populated with data interpolated from the coarse grid in front of the NS and no longer needed points are destroyed once the NS has passed. Such an interpolation step necessarily leads to a loss of accuracy and great care needs to be taken to preserve physically conserved quantities such as rest mass as well as—in the absence of general relativity—energy and momentum [79–84].

In SpEC, on the other hand, due to its comoving coordinate system, the NSs are stationary on the grid during the inspiral phase and no mesh motion is required. During most of the inspiral we use only a single resolution in the grid patches that surround each NS and no mesh refinement. Eventually, however, the NSs approach each other close enough such that their individual grid patches overlap. Rather than continue evolving in the presence of overlapping grids as, e.g. in [48], we create a single refined grid hierarchy that contains both NSs. Since we continue to track the rotation of the NSs around each other but stop tracking their separation, the NSs appear to move directly toward each other on the grid and we again avoid having to create and destroy grid points to follow the NSs with a high resolution grid patch.

Our current implementation of mesh refinement in SpEC uses vertex-centered grid points such that for a factor of 2 difference in resolution on coarse and fine grids, every second fine grid point coincides with a coarse grid point. For the current set of simulations we employ only 2 refinement levels. The code, however, is not restricted to this and supports an arbitrary number of refinement levels in an arbitrary number of grid patches. We currently do not employ sub-cycling in time: all refinement levels step forward in time with the same timestep size. This is similar to, e.g. [85], but differs from the approach in [76]. In choosing to not implement sub-cycling in time we sacrifice efficiency of the simulation for code simplicity. In the current approach, we are able to leverage the existing multi-domain infrastructure in SpEC to implement the three required data movement operations [86]: “synchronization”, “restriction” and “prolongation” at grid boundaries and where fine and coarse grids overlap. In SpEC, synchronization is the exchange of grid point data between grid patches that make up a single refinement level. It provides data in neighboring grid cells required for WENO reconstruction of the variables to cell interfaces. It never moves data between different refinement levels, and since all grid patches are aligned, synchronization is a straightforward copy of values between grid patches. Restriction is the injection of data from a fine grid into a coarse grid in regions where fine and coarse grids overlap. In our vertex centered mesh refinement code, coarse points coincide with fine grid points and restriction is just a copy operation of data between grid patches on different refinement levels. Finally, prolongation refers to the interpolation of data from coarse grids into the outer boundaries of fine grids to provide boundary data for the WENO reconstruction

in the outermost grid points. Since fine grid points are more densely spaced than coarse grid points, this operation requires interpolation of values for which we use a simple linear interpolation method. Synchronization, restriction, and prolongation are applied in this order after each Runge-Kutta substep to ensure consistent data between the mesh-refined grids.

In the current implementation in SpEC, prolongation is not mass conservative and thus leads to mass non-conservation at the refinement level boundary. In practice, we find this effect to be very small since the matter density at the grid boundaries is small and mesh refinement is used only in the very late stages of the simulation. We have not found any noticeable increase in rest mass non-conservation (see Figure 3) once we turn on mesh refinement, since the core of the NSs stay inside the finest refinement level.

D. Gauge conditions

We evolve the spacetime metric $g_{\alpha\beta}^{(4)}$ using the generalized harmonic formulation of [36, 59, 87] in which the coordinate x^α satisfies the covariant scalar wave equation

$$\nabla^\beta \nabla_\beta x^\alpha = H^\alpha, \quad (8)$$

for a freely specifiable gauge source function H^α . The initial data is constructed in a gauge $H_{\text{initial}}^\alpha$ that assumes quasi-equilibrium and the existence of a helical Killing vector. At the beginning of the simulation, we use $H^\alpha = \hat{H}^\alpha$, where \hat{H}^α is defined to be a tensor that agrees with $H_{\text{initial}}^\alpha$ in a frame comoving with the grid, and is constant in time in this frame. Note that H^α is not a tensor. We smoothly transition from this initial gauge to a purely harmonic gauge $H_\alpha \equiv 0$ using a transition function

$$\mathcal{F}(t; t_0, \Delta T) = \begin{cases} 1 & t < t_0, \\ \exp(-(\frac{t-t_0}{\Delta T})^4) & t_0 \leq t. \end{cases} \quad (9)$$

For the transition to harmonic gauge, we choose $t_0 = 0$, $\Delta T = 2\sqrt{d^3/(2M_0)}$, where d is the initial coordinate separation of the stars and M_0 is the baryonic mass of each star. ΔT is approximately 2 orbital periods which is slow enough to avoid gauge artifacts in the numerical waveforms. This differs from what is typically done in BBH and BHNS simulations using SpEC, where the simulation directly transitions to damped harmonic gauge [59, 88] around each of the BH,

$$H_\alpha = \mu_L \ln \frac{\sqrt{g}}{N} t_\alpha - \mu_S g_{\alpha\beta}^{(4)} N^{-1} N^\beta, \quad (10)$$

where g is the determinant of the spatial 3-metric g_{ij} , $t_\alpha = -N\partial_\alpha t$ is the future directed unit normal to the constant- t surfaces, N is the lapse function and N^α is the shift vector, near the BH(s). We find that employing the damped harmonic gauge condition reduced the simulation speed to $\approx 20\%$ compared to the harmonic gauge,

due to damped harmonic gauge inducing a reduction of the allowed timestep size to $\approx 20\%$ of the value allowed in harmonic gauge. We therefore delay changing into a fully damped harmonic gauge as long as possible. On the other hand, a pure harmonic gauge condition can lead to coordinate singularities due to caustics near AH formation and we found that a “mild” version of the damped harmonic gauge condition lets us avoid caustics while still achieving good evolution speeds.

The simulations discussed in this paper stay in harmonic gauge until $t_0 = 22410 M_\odot$ (approx 2.5 ms before we find an AH) at which time we transition to the mild version of the damped harmonic gauge condition Eq. (10). For the mildly damped harmonic gauge, we set $\mu_L = \mu_S = 0.2 M_\odot/M_{\text{ADM}}$, with M_{ADM} denoting the Arnowitt-Deser-Misner (ADM) mass. This matches the value chosen in [89], and we smoothly transition to the new gauge using Eq. (9) with $t_0 = 22410 M_\odot$, $\Delta T = 200 M_\odot$.

Finally, just before we expect the AH to form, at $t = 22890 M_\odot$ ($t - t_{\text{horizon}} \approx 0.24$ ms) we add a fully damped harmonic instance of Eq. (10) with $\mu_L = \mu_S = [\ln(\sqrt{g}/N)]^2$ to the already active mild damped harmonic gauge source. This gauge change is very rapid with $\Delta T = 30 M_\odot$.

The complete gauge source term at the time of AH formation is thus

$$H_\alpha = (1 - \mathcal{F}(t; 22410, 200))H_\alpha^{\text{mild}} + (1 - \mathcal{F}(t; 22890, 30))H_\alpha^{\text{full}}, \quad (11)$$

$$H_\alpha^{\text{mild}} = \frac{0.2 M_\odot}{M_{\text{ADM}}} \left(\ln \frac{\sqrt{g}}{N} t_\alpha - g_{\alpha\beta}^{(4)} N^{-1} N^\beta \right), \quad (12)$$

$$H_\alpha^{\text{full}} = \left(\ln \frac{\sqrt{g}}{N} \right)^2 \left(\ln \frac{\sqrt{g}}{N} t_\alpha - g_{\alpha\beta}^{(4)} N^{-1} N^\beta \right). \quad (13)$$

E. Excision and post-collapse evolution

As described in greater detail in Sec. IV A, during the evolution, we use a complex spectral grid setup surrounding each NS with a set of concentric spherical shells, transitioning to a set of shared spherical shells in the wave zone. Such a grid setup is not well adapted to a single BH having formed after collapse of the merged NSs since it does not allow us to excise the interior of the BH from the grid.

Therefore soon after we turn on the full damped harmonic gauge, we begin searching for an AH using an iterative fast flow algorithm based on [90]. Once we have detected an AH and followed its evolution through several time steps, we construct a new spectral grid consisting only of nested spherical shells, with the innermost boundary of the innermost shell slightly inside the AH, so that the interior of the BH is excised. We then interpolate the spacetime variables from the old spectral grid to the new one, and continue the simulation, keeping the finite volume grid unchanged.

The algorithm for transitioning to a new spectral grid with a single excision boundary is almost the same as described in [41, 59, 60] for treating the merger and ring-down of a BBH system after a new common AH forms around the two individual AHs. In particular, the excision boundary of the spectral grid changes its shape and size dynamically to conform to the size and shape of the AH and to ensure that all characteristic fields of the evolution system are outgoing (into the horizon), so that excision is well-posed without a boundary condition.

The main difference between BH evolution for BBH ringdowns versus NSNS remnants is the form of the function that maps grid coordinates x^i to the coordinates $x^{\bar{i}}$ in which the excision boundary distorts to match the shape of the AH [60]. This map is

$$x^{\bar{i}} = x^i \left(1 - f_C(r) \sum_{\ell m} Y_{\ell m}(\theta, \phi) \lambda_{\ell m}(t) \right), \quad (14)$$

where $Y_{\ell m}$ are spherical harmonics, $\lambda_{\ell m}$ are coefficients, $f_C(r)$ is a prescribed function, and (r, θ, ϕ) are spherical polar coordinates computed in the usual way from x^i . For BBH ringdowns, $f_C(r)$ is chosen as a simple piecewise linear function that has discontinuous derivatives at spectral subdomain boundaries [60]. For NSNS remnants, $f_C(r)$ is a Gaussian that is smooth everywhere so that the Jacobian of the map is continuous over the finite volume domain, which overlaps subdomain boundaries of the spectral domain.

On the finite volume grid on which we evolve the fluid variables, we “mask” the excised region. Within the excised region, the metric variables are set to their value for Minkowski spacetime, while the density is set to its minimum allowed value ρ_{floor} . As we want to avoid any dependence of the evolution of the system on that arbitrarily (and unphysical) choice, we also use modified interpolation stencils when reconstructing the variables at cell faces close to the excised region, and when interpolating the fluid variables from the finite volume grid to the pseudospectral grid. For the reconstruction of the fluid variables on faces, we use the WENO5 algorithm in the bulk of the simulation, whenever the required 5-points stencil is available. We drop to the second-order MC2 algorithm [91] when only a 3-points stencil is available. Finally, we simply copy the value from the neighboring cell center when we do not have enough points to perform the MC2 reconstruction. On the face directly neighboring the excised region, the left and right fluxes are both set to their value at the nearest cell center. The metric variables at cell faces are similarly interpolated from a 3-point symmetric stencil, a 2-point symmetric stencil, or by copying from the only non-excised cell center, depending on the number of non-excised points available around a given face. Finally, interpolation from the finite volume grid to the pseudospectral grid is performed using, when possible, a polynomial fit to a 3-points stencil, with the additional constraint that the interpolation cannot create new extrema (i.e. the interpolated value is limited by

the minimum and maximum values of the function at the grid points used in the stencil). When we do not have 2 points available on each side of the desired interpolation location, we drop to linear interpolation, or copying from the nearest non-excised cell center when extrapolation is required.

F. Wave Extraction

We use the Cauchy-characteristic extraction (CCE) method described in [92–95] to evolve the gravitational waves emitted by the system from a finite radius to future null infinity \mathcal{I}^+ . Details on the characteristic method and its use in `SpEC` can be found in Sec. II.3.B of [95]. We compute the (2, 2)-mode $\Psi_4^{2,2}$ of the Newman-Penrose scalar [96, 97] at \mathcal{I}^+ decomposed into spin-weighted scalar spherical harmonic modes. We then use the fixed frequency integration method of [98] to compute the gravitational wave strain $h^{2,2}$ from $\Psi_4^{2,2}$ without taking into account possible drift effects described in [99]. Since the total drift of the BH at AH formation and at the end of the simulation is less than $0.1 M_\odot$ and $0.2 M_\odot$, respectively, the effect of drift is expected to be less than 1% on the dominant (2, 2) mode. Details on the extraction setup for our simulations are given in Sec. IV C.

III. INITIAL DATA

Initial data for this simulation was produced by a new NSNS initial data solver based on the work of Foucart et al. for BHNS systems [100]. As in that work, we start by considering systems in quasi-equilibrium, where time derivatives vanish in a corotating frame (this neglect of the small radial velocity will be addressed later). We take the metric to be conformally flat and solve for the lapse, shift, and conformal factor using the extended conformal thin sandwich (XCTS) equations [101]. The matter in the stars is modeled as a cold ($T = 0$) perfect fluid with an irrotational velocity profile, which is a special case of the more generic framework used in [34]. The irrotational limit allows a straightforward solution for the velocity and is a more realistic approximation than the corotating limit, as the effective viscosity of NS matter is insufficient to synchronize the stars’ spins with their orbital frequency [102, 103].

A particular NSNS system is specified in terms of the equation of state of NS matter, the baryon masses of both stars, and their coordinate separation. The solver then uses the above assumptions of quasi-equilibrium and cold irrotational flow to determine the metric and matter content of the corresponding spacetime. Since the initial data problem consists of several coupled nonlinear equations, the solver takes an iterative approach, with each iteration composed of a number of substeps (this procedure closely follows Sec. III.C of [100], which should be consulted for additional details).

First, given a trial matter distribution, we find an approximate solution to the elliptic XCTS equations by taking a single step of a nonlinear solver. By imposing force balance at the centers of the stars, we then adjust the orbital frequency of the binary. We also modify the enthalpy of the matter to drive the locations of its maxima to the specified stellar centers, thus controlling the stars' separation. Finally, we approximately solve the elliptic equations imposing irrotational flow (constrained to preserve the baryon masses of the stars) and feed the output to the next step of the iterative procedure. All of these updates are made using a relaxation scheme to aid convergence.

Throughout the solution process, the numerical data are represented on a spectral grid composed of hexahedra, cylindrical shells, and spherical shells, and approximate solutions to the elliptic equations are provided by the SPELLS framework [104]. We periodically evaluate the grid and adjust it to better conform to the stars' surfaces. By placing subdomain boundaries close to the surfaces, the discontinuities there do not strongly affect the spectral convergence of the method for the resolutions used in our simulations.

Additionally, we occasionally perturb the centers of the stars to reduce the ADM linear momentum of the system. During this procedure, the centers are *not* constrained to be colinear with the center of revolution, and the separation of the stars may deviate slightly from the initially specified value. Separations reported here are therefore measured from the final solution.

When constructing strictly quasi-equilibrium data, the solver chooses the orbital angular velocity Ω by requiring force balance at the centers of the stars. Later, when subsequently refining the initial data to reduce eccentricity, Ω is fixed. By adding an initial radial velocity, we relax the quasi-circular approximation in order to more accurately model inspiral conditions and reduce the initial eccentricity. The magnitude of this velocity is chosen by evolving each trial set of initial data for a short time in order to measure the eccentricity of the orbits, then adjusting the (fixed) orbital frequency and radial velocity according to a heuristic procedure based on the work of [105] and repeating until that eccentricity is below 10^{-3} . Similar approaches were also used by [106, 107], who achieved comparable results.

Results from our code closely match those of the Lorene solver by [108]. In particular, we can accurately reproduce the quasi-equilibrium sequences of [109, 110].

For this study, initial data are generated using a polytropic EOS of the form

$$P = \kappa \rho_0^\Gamma, \quad (15)$$

$$\epsilon = \frac{1}{\Gamma - 1} \frac{P}{\rho_0}, \quad (16)$$

with $\Gamma = 2$ and $\kappa = 123.6 M_\odot^2$. Both NSs have a baryon mass of $M_0 = 1.779 M_\odot$, corresponding to an isolated TOV star with an ADM mass of $M_\infty = 1.64 M_\odot$, circumferential radius of $R_{\text{areal}} = 15.1 \text{ km}$ and a compact-

ness of $M_\infty/R_{\text{areal}} = 0.16$. Due to the large initial coordinate separation of 81 km, the binding energy is small, $E_b = 6.7 \cdot 10^{-3} M_\odot$ and the total ADM mass of the system is $M_{\text{ADM}} \approx 2 M_\infty$. In the binary configuration, the stars extend to an isotropic coordinate radius of 12 km ($8.1 M_\odot$), and their centers are separated by a coordinate distance of 81 km ($54.5 M_\odot$). This system has an orbital frequency of $\Omega/2\pi = 133 \text{ Hz}$ and an eccentricity of less than 9×10^{-4} . Due to the total mass of the system exceeding the maximum mass of a hypermassive star for a $\Gamma = 2$ EOS we expect the merged NS to collapse to a BH very quickly [111–113].

IV. RESULTS

A. Dynamics and Grid setup

The evolution of the NSNS system proceeds through a series of stages, each of which is reflected by specific settings used by the simulations during this phase. Starting from large separation, the NSNS are initially in the inspiral phase characterized by quasi-circular motion around each other. During this phase, the separation changes very slowly compared to orbital timescales. Eventually, the NSs approach each other and come into contact, which leads to the development of a shear layer in the contact region and eventually a single merged object. Finally, the central core of the merged object collapses and forms a BH, which accretes the remaining material and eventually settles down to a stationary Kerr BH. Fig. 1 and Table I show the settings and grid structures used during the different phases.

For the spectral grid during the inspiral phase, we use an adapted version of the two-spheres domain used in [38]. We perform simulations using three different resolution levels: *Lev0*, *Lev1*, *Lev2*. During grid setup and during evolution, we use the spectral mesh refinement method of [75], decreasing the allowed truncation error in the spectral expansion of the solution as e^{-k} for resolution level k . Thus, the actual number of collocation points used differs from subdomain to subdomain and is based on features of the matter distribution and metric variables inside each subdomain. In contrast to [38], we replace the half of the grid that covers the BH in [38] by a second copy of the grid covering the NS. Thus, for resolution level *Levk*, our domain consists of two filled spheres covering the center of each neutron star using spherical harmonic basis functions in the angular directions, while the radial dependence is decomposed into one-sided Jacobi polynomials [114].

Each filled sphere is surrounded by 8 spherical shells of similar angular and radial resolution as the inner sphere. Initially the surface of the star is located in the third shell. The far field region around the binary is covered by 20 spherical shells starting at 1.5 times the initial separation of the stars to 40 times the initial separation, or approximately $2200 M_\odot$. These shells have slightly lower

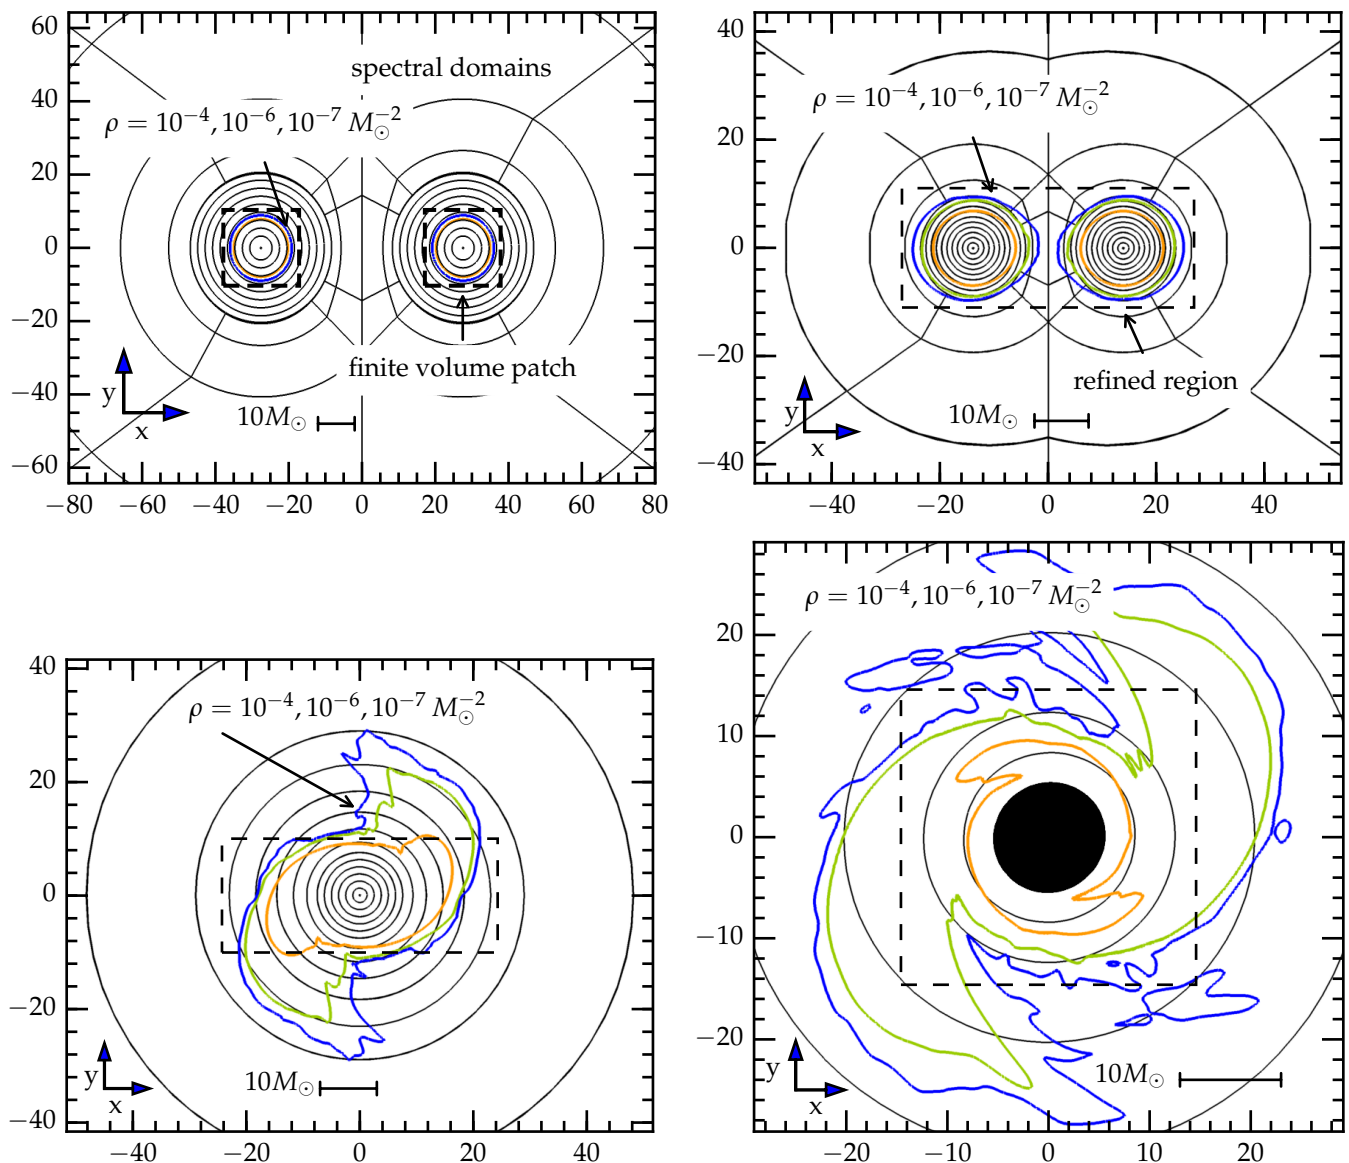


FIG. 1: Grid structure for the Lev0 run in the $z = 0$ plane during the simulation. We show the spectral grid (solid black lines) and the finite volume grid patches that surround the NSs. We also show contour lines of the rest mass density ρ_0 for $\rho_0 = 10^{-4}, 10^{-6}, 10^{-7} (M_\odot)^{-2}$ ($\rho_0 \approx 6 \times 10^{13}, 10^{11}, 10^{10} \text{ g cm}^{-3}$) in orange, green and blue and a scale bar indicating the size of $10 M_\odot$ or approximately 14.8 km. Top left: Long-dashed boxes outline the finite volume grids at the beginning of the simulation. Top right: Grid structure after creation of a single finite volume grid near merger. Bottom left: Simulation shortly after switching to a single set of nested spherical shells. Bottom right: Simulation after the AH has formed. The innermost spherical shell coincides with the AH. All but the top left plot show the entire region covered by the coarse (outer) finite volume mesh, with the fine (inner) mesh outlined using dashed lines.

angular and radial resolution than the spherical shells around the NSs. The region between the innermost shell and the stars is covered by a set of deformed cylindrical shells and filled cylinders interpolating between the spheres. There are a total of 48 subdomains in the initial setup. During the simulation, we measure the truncation error in each subdomain and adjust the subdomain structure by adding and removing points as well as splitting and joining subdomains such that the measured trunca-

tion error is close to the requested accuracy. Due to the presence of junk radiation at the beginning of the simulation we keep the grid structure in the outer spherical shells fixed for one light crossing time of the simulation domain to avoid the junk radiation triggering mesh refinement and leading to very high resolution when attempting to resolve the junk radiation.

The finite volume grid during the inspiral phase consists of two halved cubes with 48, 61, 77 grid points per

half-length of the cube for the three resolution levels *Lev0*, *Lev1*, *Lev2*. Initially the cube's sides are approximately 1.25 times the diameter of the stars. This corresponds to approximately 30, 38, 48 points across the radius of the NS and a linear resolution of 326 m, 252 m, and 192 m for resolution level *Lev0*, *Lev1*, and *Lev2*, respectively. We take advantage of the reflection symmetry across the $z = 0$ plane present in the system to only evolve the $z > 0$ half-space. Function values in the $z < 0$ half-space are computed using the symmetry condition when needed. The region outside of each NS but covered by the finite volume grid is filled with a low density atmosphere with rest mass density $\rho_{0,\text{atmosphere}} = 10^{-13} / M_{\odot}^2 \approx 10^{-10} \rho_{0,\text{central}}$, as is common in grid based simulations of NSs. For the majority of the simulation the spectral grid structure consists of sets of spherical shells around each of the NS, and the finite volume grid consists of one individual grid patch around each NS as shown in the upper left pane of Fig. 1.

During the simulation we monitor the dephasing in the orbital phase between resolution levels *Lev0*, *Lev1*, *Lev2* and interpolate onto a higher resolution grid once the phase difference increases too rapidly. For the set of simulations presented in this paper, this increase in resolution occurs at $t \approx 1.5 \times 10^4 M_{\odot}$ (≈ 38 ms before the horizon is formed). At this point, we increase the resolution of the *Lev0* run to that of the *Lev1* run (252 m), that of the *Lev1* run to *Lev2* (192 m) and finally that of the *Lev2* run to 144 m, which would be the resolution of a *Lev3* run. We also adjust the requested truncation error in a similar manner such that the *Lev0* simulation requests a truncation error corresponding to the truncation error originally requested by the *Lev1* simulation and similar for the higher resolution simulations. Once the NSs are close enough together so that the individual grids touch, we replace the two grids by a single rectangular grid that covers both NSs. This is shown in the top right panel of Fig. 1. At this time, we increase the resolution further such that three resolution levels *Lev0*, *Lev1*, *Lev2* use resolutions of 207 m, 148 m and 115 m. We surround this inner grid by a coarser grid of twice the size but half the resolution as described in Sec. II C. The coarse grid captures ejecta and the disk left behind after BH formation. At this point we no longer adjust the domain to contain all matter, instead we hold the physical size of the finite volume grid constant.

When the stars approach each other, they gradually deform, cf. Fig. 2 to see the deformation of the stars. At some point, the nested spherical shells that are used in the spectral grid are no longer a good approximation of the stellar shapes, causing the simulation speed to drop rapidly as more and more spectral resolution is required to resolve the deformed stellar shape. This in turn reduces the allowed timestep via the Courant-Friedrichs-Lewy (CFL) factor. We replace the nested spherical shells and cylinders around each star by a single set of concentric spherical shells centered at the center of the merging binary. During the merger phase, the matter dis-

tribution is very distorted while the metric terms gradually become centered around the origin. Hence, a spectral grid centered around the origin deals best with the lack of regularity in the data. The bottom left panel of Fig. 1 shows the grid layout at this point. The 3D density distribution shortly afterwards is shown in the bottom panel of Fig. 2.

When switching to a single set of spherical shells, we turn off tracking the orbital separation of the stars and smoothly transition to a constant scaling factor between the comoving grid coordinates and inertial coordinates. The numerical grid is thus no longer contracting with the stars and the stars move toward each other in the grid frame. Allowing the stars to move on the grid avoids strong grid deformation in the region between the stars where a constant volume in grid space is used to represent the shrinking separation between the stars while at the same time covering the region far from the stars with almost constant resolution. We continue tracking rotation of the object until an AH is detected at which point we transition to a coordinate frame that is at rest with respect to an observer at infinity.

Finally, once the merged object collapses to a BH, we excise the inner filled sphere and tie the inner boundary to the AH instead (see Sec. II E). At this point the setup is identical to what was used in [38]. The bottom right panel of Fig. 1 shows the grid structure at this point.

Table I lists the computational domains used in this work.

B. Diagnostics

During the course of the simulation we monitor several constraints and conserved quantities to assess the quality of the simulation.

1. Rest mass conservation

We evolve the relativistic rest mass density D using a conservative scheme [36]: SpEC is therefore expected to exactly conserve total rest mass,

$$M_0(t) = \int D d^3x, \quad (17)$$

during the evolution. However, there are several effects that introduce non-conservative changes to the rest mass density: (i) we employ (see Sec. IV A for details) a low density numerical atmosphere of density $\rho_{0,\text{atmosphere}}$ which can lead to matter creation in the region outside the NS when the density would drop below $\rho_{0,\text{atmosphere}}$ during the evolution. We employ an atmosphere density that is sufficiently small compared to the density at the center of the NSs, $\rho_{0,\text{central}}$, that this effect is expected to be small. (ii) Matter can reach grid boundaries and flow off the grid. We employ the remapping procedure described in Sec. II B to control the amount of matter

	Inspiral	Late Inspiral	Single Box
spectral grid	2 sets of spheres	2 sets of spheres	2 sets of spheres
finite volume grid	2 uniform boxes	2 uniform boxes	mesh-refined rectangular box
control outflows	yes	yes	no
track orbital separation	yes	yes	yes
finest finite-volume resolution at beginning of segment [M_\odot]	0.22, 0.17, 0.13	0.17, 0.13, 0.097	0.14, 0.10, N/A ^a
gauge condition	harmonic	harmonic	harmonic
no. of orbits	14	8	< 1
orbital angular frequency $M_{\text{ADM}}\omega$ at beginning of segment	0.014	0.020	0.042
	Plunge	Pre-Collapse	Post-Collapse
spectral grid	2 sets of spheres	spherical shells	excised spherical shells
finite volume grid	mesh refined rectangular box	mesh refined rectangular box	mesh refined square box
control outflows	no	no	no
track orbital separation	yes	no	no
finite volume resolution at beginning of segment [M_\odot]	0.14, 0.099, 0.078	0.13, 0.097, 0.076	0.11, 0.073, 0.058
gauge condition	mildly damped harmonic	fully damped harmonic	fully damped harmonic
no. of orbits	1	< 1	N/A
orbital angular frequency $M_{\text{ADM}}\omega$ at beginning of segment	0.045	0.092	N/A

^aWe transition to a mildly damped harmonic gauge at fixed evolution time while transition to a single box is triggered by the finite volume grids touching. Since the size of the finite volume grids differs between **Levs**, it so happens that the transition to a plunge occurs before the NSs approach each other close enough to force a single finite volume box.

TABLE I: Stages in the inspiral simulation. For each stage we list the type of spectral and finite volume grid used, whether we control the amount of matter flowing off the finite volume grid, whether we monitor the orbital separation of the NS, the minimum resolution for resolution levels **Lev0**, **Lev1**, **Lev2** on the finite volume grid, the gauge condition used, the approximate number of orbits the system spends in this phase in the medium resolution (**Lev1**) simulation, and the orbital angular frequency at the beginning of each segment in the medium resolution (**Lev1**) simulation.

leaving the system, limiting the matter loss rate through the boundary to $\dot{M}_0^{(i)} < 2 \times 10^{-8}$ through any of the grid boundaries. (iii) The interpolation algorithm used in the remapping procedure is not mass conservative and introduces a relative mass change of order $< 10^{-5}$ when interpolating to a new grid. (iv) Our current mesh refinement implementation (see Sec. II C for details) uses non-conservative interpolation operators to interpolate data between the different refined regions. The error introduced by this interpolation is very small, since the matter density near the refinement boundaries is very low and does not contribute much to the total mass.

We find (ii) to be the most important effect during the simulation. Figure 3 displays the conserved rest mass over the course of the simulation. We stop monitoring the total rest mass once an AH is found. At this point rest mass is lost from the simulated spacetime as matter falls into the AH. The accretion happens over a short time and matter is rapidly falling into the BH. Approximately 2 – 10 ms after BH formation, the mass left outside the BH falls below $10^{-3}M_\odot$, which corresponds roughly to the mass conservation error of our code for **Lev0** and

Lev1.

2. Constraints

In the following section, the L_2 volume norm $\|\cdot\|_2$ of a rank n tensor $T_{\alpha_1\dots\alpha_n}$ is defined as

$$\|T_{\alpha_1\dots\alpha_n}\|_2 = \sqrt{\frac{\int |\delta^{\alpha_1\beta_1}\dots\delta^{\alpha_n\beta_n} T_{\alpha_1\dots\alpha_n} T_{\beta_1\dots\beta_n}|^2 d^3x}{\int d^3x}}. \quad (18)$$

a. ADM constraints Figure 4 (top panel and center panel) show the violation of the ADM Hamiltonian and momentum constraints, respectively. The constraints are evaluated as the L_2 norm over the simulation volume. The numerical data contain a large amount of noise and we use a moving average of window width $\Delta t = 1.92$ ms to smooth out this high frequency noise. We observe a spike in the constraint violations around the time of merger and BH formation when the spacetime becomes




FIG. 2: Volume rendering of the rest-mass density in the late-inspiral, merger, and post-merger phases. Top left panel: Late-inspiral part of the simulation. Shortly before the two NSs touch, tidal deformations are clearly visible. Top right panel: Shortly after contact, a characteristic shear-layer is formed between the two NSs. Bottom left panel: Shortly before collapse, the mass is centered around the origin. Low-density spiral arms of the merger remnant have formed. Bottom right panel: 1.4 ms after merger the BH has settled down to an almost stationary BH.

highly dynamic. After excision of the BH, the constraint violations shrink since the inner part of the BH is no longer part of the numerical or physical domain of dependence and is no longer included in the computation of the constraint violations.

b. Generalized harmonic constraint energy The generalized harmonic formulation of Einstein's equations contains several constrained quantities. Monitoring these constraints during the simulation provides us with a useful measure of the faithfulness of our simulations. Figure 4 (bottom panel) shows the evolution of the L_2 norm of the generalized harmonic constraint energy as defined in Eq. (71) of [57].

At the beginning of the simulation we see clear convergence and the constraint decreases for increasing resolu-

tion. When AMR as described in Sec. II C is activated, i.e. after emission of the junk radiation, the clear convergence is lost, as constraint violations are no longer dominated by errors in the spectral domain but instead contain contributions due to matter which converge much more poorly as resolution increases from *Lev0* to *Lev1* and *Lev2*. Furthermore, as for the ADM constraints, we observe a spike in the constraint violation around the time of merger and BH formation when the spacetime becomes highly dynamic. Fortunately, the constraint violating numerical data are concentrated in the region that will be interior to the newly formed BH. This is seen as the sudden drop in the constraint energy once we excise the interior of the AH from the simulation domain.




FIG. 3: Conservation of rest mass for the three resolution levels *Lev0*, *Lev1*, *Lev2*. Mass conservation is very good until the horizon forms at which point numerical errors near the center of the forming BH lead to spurious mass creation. The vertical dashed line indicates when we switch on the refined meshes. We do not track mass conservation after the AH is formed, since material escapes from the simulation domain through the AH.

3. ADM integrals

We also monitor how well the code conserves the total ADM mass of the system during evolution. We approximate ADM mass conservation as conservation of the ADM mass surface integrals in the simulation domain corrected by the radiated energy. To this end, we evaluate the ADM surface integrals [115, 116]

$$M^{\text{NR}} = \frac{1}{16\pi} \oint_{r=r_{\text{ADM}}} \left(\frac{\partial g_{ik}}{\partial x^j} - \frac{\partial g_{ij}}{\partial x^k} \right) \delta^{ij} n^k dA, \quad (19)$$

$$J_i^{\text{NR}} = \frac{1}{8\pi} \epsilon_{ijk} \oint_{r=r_{\text{ADM}}} (K_l^k - \delta_l^k K) x^j n^l dA, \quad (20)$$

where g_{ij} is the 3-metric, K_{ij} is the extrinsic curvature, K is its trace, n^i is the outward pointing unit normal vector to the integration sphere of radius $r_{\text{ADM}} = 2090 M_\odot$, and ϵ_{ijk} is the Levi-Civita symbol. We keep track of the radiated energy and angular momentum in the GW modes up to $\ell = 8$ passing through the sphere [117]:

$$M^{\text{rad}} = \frac{1}{16\pi} \sum_{\ell, m} \int_0^t dt' \left| \frac{dh_{\ell m}(t')}{dt'} \right|^2, \quad (21)$$

$$J_z^{\text{rad}} = \frac{1}{16\pi} \sum_{\ell, m} \int_0^t dt' m \Im \left(h_{\ell m}(t') \left[\frac{dh_{\ell m}(t')}{dt'} \right]^* \right), \quad (22)$$

where $h_{\ell m}$ is the spin weighted spherical harmonic (ℓ, m) mode of the gravitational waveform, $\Im(z)$ is the imaginary part of z , and $*$ denotes complex conjugation. Since

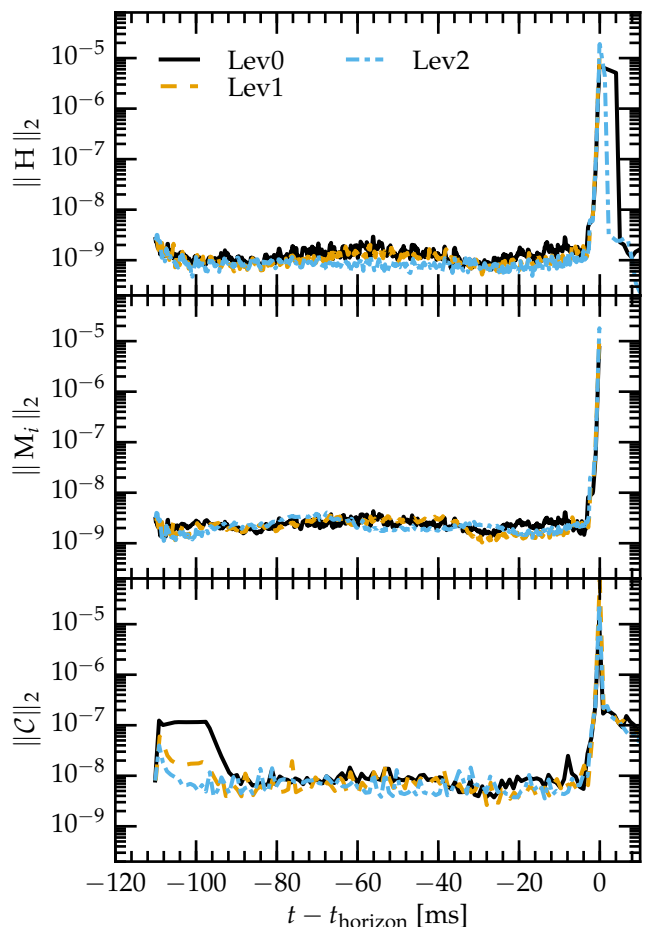


FIG. 4: L_2 volume norm of the Hamiltonian constraint violation (top panel), the square magnitude of the momentum constraint violation (middle panel), and the generalized harmonic constraint violation energy (bottom panel). The origin in time corresponds to AH formation. Clearly visible is the increase in constraint violation around this time.

we evaluate the ADM integral Eq. (19) on a surface at a finite radius, the integrated value depends on time. We correct the value by the amount of energy radiated through the integration surface and verify that the sum $M^{\text{NR}} + M^{\text{rad}}$ (and $J_z^{\text{NR}} + J_z^{\text{rad}}$) is constant over time.

The top panel of Figure 5 shows the deficit $M^{\text{ADM}} - (M^{\text{NR}} + M^{\text{rad}})$ i.e. the failure of the simulation to conserve the ADM mass. The lowest resolution simulation *Lev0* shows an unphysical decrease in the observed ADM mass integral between $-120 \text{ ms} \lesssim t \lesssim -80 \text{ ms}$, whose minimum occurs at approximately the same time as the spike in the generalized harmonic constraint energy. This decrease is caused by low resolution in the region $r > 165 M_\odot$ containing the surface used to evaluate the ADM mass integral and vanishes once spectral adaptive mesh refinement increases the resolution in this region.

The bottom panel of Figure 5 shows the fractional deficit of the total angular momentum in the z direction,




FIG. 5: Top panel: Mass deficit $M^{\text{ADM}} - (M^{\text{NR}} + M^{\text{rad}})$ computed using a sphere of coordinate radius $r_{\text{detector}} = 2090 M_{\odot}$. The initial value of M^{NR} for simulation Lev2 is $M^{\text{NR}} = 3.278 M_{\odot}$. The initial decrease in the Lev0 curve is due to constraint violations in the region $r > 165 M_{\odot}$, which are damped away once the spectral AMR [75] increases the resolution in this region. The higher resolution levels Lev1 and Lev2 are of sufficiently high resolution such that this issue does not arise. Bottom panel: Relative angular momentum deficit $1 - (J_z^{\text{NR}} + J_z^{\text{rad}}) / J_z^{\text{ADM}}$ computed on the same sphere. Shown are results for the three resolution levels Lev0, Lev1 and Lev2. The initial value of J_z^{NR} for simulation Lev2 is $J_z^{\text{NR}} = 12.41 M_{\odot}^2$ and the total radiated angular momentum is $J_z^{\text{rad}} = 3.680 M_{\odot}^2$.

$1 - (J_z^{\text{NR}} + J_z^{\text{rad}}) (J_z^{\text{ADM}})^{-1}$. Angular momentum non-conservation is well below 1% until just before AH formation. At this point, the increased constraint violations affect the measurement of the total angular momentum as well.

4. Black hole mass and spin

The total mass of the system is well above the maximum mass that can be supported by post-merger differential rotation and the chosen EOS. A BH forms less than 1 ms after the two NS come into contact and merge. Within 1 ms after the formation of the AH, almost all material has fallen into the BH. The final mass surrounding the BH is below $10^{-3} M_{\odot}$ for Lev2.

The final BH settles down to a Kerr BH with Christodoulou mass of $3.226 \pm 0.007 M_{\odot}$ and spin of $8.743 \pm 0.029 M_{\odot}^2$, where uncertainties are estimated as

the difference between Lev2 and Lev1. This corresponds to a dimensionless spin magnitude of $\chi \approx 0.84$ and is thus well below the extremal Kerr solution.

C. Gravitational wave signal

Figure 6 displays the (2,2)-component of the spherical harmonic decomposition of the GW strain h at \mathcal{I}^+ as obtained via the CCE method. The waveform lasts for more than 40 GW cycles. We emphasize that this is the longest waveform of a NSNS obtained from a fully general-relativistic simulation.

In addition to presenting the whole waveform, we also show a zoom-in around merger in Fig. 7. We compare our waveform with a shorter waveform obtained by [27, 52] for the same NSNS system using the finite differencing (spacetime) / finite-volume (hydrodynamics) BAM code [49, 107, 118]. We align the waveforms in a time interval $-25.8 \text{ ms} \leq t \leq -13.3 \text{ ms}$. This first comparison between the two codes suffers from the fact that different initial numerical datasets describing the same physical system were used. Nevertheless, we observe that after time and phase alignment, the phase difference stays below 0.25 rad up to merger. This is well within the uncertainty of the BAM waveform of ± 0.9 rad and shows that also around merger, where our error estimate becomes problematic (see the discussion in Sec. IV D), consistent results can be obtained.

The merged object collapses to a BH within less than 1 ms after merger and the BH then emits a characteristic ringdown GW signal. Our results allow us to estimate the quasi-normal mode frequencies. We obtain a frequency of $M_{\text{BH}}\omega = 0.613$ for the (2,2)-mode. This corresponds to within 0.5% to the result obtained via BH perturbation theory and to the value of 0.61454 stated in [119] for a BH with a dimensionless spin of $\chi = 0.840$. Figure 8 shows the (2,2)-mode of the ringdown signal in $\Psi_4^{2,2}$ (a similar plot can be obtained for the GW strain mode $h_{2,2}$). We observe very clean exponential decay of the dominant mode over more than 3 orders of magnitude of $\Psi_4^{2,2}$.

D. Convergence

The physics observable using gravitational wave detectors such as LIGO is primarily encoded in the phase ϕ of a GW signal $h(t) = A(t) \cos \phi(t)$ and thus the phase accuracy of a simulation is of primary importance when assessing the quality of a simulation. Since GW are quadrupolar in nature, for circular orbits, the (2,2)-mode captures the dominant GW signal and its complex phase ϕ can be used to compute (a proxy for) the GW phase ϕ .

SpEC uses a hybrid spectral – finite volume scheme that makes it difficult to assign a unique convergence order to simulations. Figure 9 shows the phase difference between




FIG. 6: Real part and magnitude of the (2,2)-mode of the gravitational wave strain $r h_{2,2}$ during the inspiral and merger. The time of maximum amplitude is labeled as $t - t_{\text{peak}} = 0$.




FIG. 7: Real part and magnitude of the (2,2)-mode of the gravitational wave strain $r h_{2,2}$ during the late inspiral and merger. The time of maximum amplitude is labeled as $t - t_{\text{peak}} = 0$. This figure displays a zoom-in of the last ≈ 10 ms of the signal shown in Figure 6, focusing on the final few cycles of the inspiral, merger, and ringdown GW signal. We compare the waveform obtained for an identical NSNS system by [52]. We align in time and phase in the interval $-25.8 \text{ ms} \leq t \leq -13.3 \text{ ms}$, minimizing Eq. (24). Both waveforms agree reasonably well. During the inspiral the phase difference due to the eccentricity of the BAM data is around 0.1 rad and even up to merger the phase difference stays below 0.25 rad.

different resolution levels `Lev0`, `Lev1`, `Lev2` in the (2,2)-mode of the GW strain at future null infinity \mathcal{I}^+ . We expect the polynomial error of the finite volume scheme to dominate the error budget, and thus model the phase error at each instant in time using a second order poly-




FIG. 8: Dominant 2,2-mode of the ringdown signal observed in Ψ_4 at \mathcal{I}^+ . We observe the ringdown signal for over 3 magnitudes in amplitude before the numerical noise overwhelms the signal. The dashed line shows the fitted decay behavior $\exp(-t/\tau_{\text{QNM}})$. We find a mode frequency of $M_{\text{BH}}\omega = 0.613$.

nomial of the form

$$\phi = \phi_0 + a_1 \Delta x + a_2 \Delta x^2. \quad (23)$$

Here, Δx is a measure of the finite volume resolution. Eq. (23) is able to capture second order convergence of the code in smooth regions of the flow and first order convergence across shocks and surfaces. In this model, ϕ_0 is the continuum value of the phase and the term $a_1 \Delta x + a_2 \Delta x^2$ is the phase error for a simulation using finite resolution Δx . Obviously the model neglects higher order error terms and the infinite-resolution extrapolated




FIG. 9: Phase difference $\Delta\phi_{n,m}(t) = \phi_n(t) - \phi_m(t)$ in the (2,2)-mode of the GW strain $r h^{2,2}$ during the inspiral phase of the simulation. The solid (black) line shows the phase difference between the low and medium resolution runs, while the dashed (orange) line shows the phase difference between the medium and high resolution runs. Dotted (black) and dash-dotted (orange) line segments indicate time intervals during which the phase difference $\Delta\phi_{n,m}(t)$ is negative. The upper x -axis is labeled by the GW frequency $\omega_{2,2}$ of the (2,2)-mode of the GW strain of the highest resolution (Lev2) run. We observe convergent behavior in the GW phase. However, no clear convergence order can be assigned. This is most likely due to interactions between numerical errors in the finite volume hydrodynamics part and in the adaptively refined spectral metric part of the code.

value ϕ_0 of the finite resolution GW phases obtained from it is only an approximation to the true phase. At very high resolution ($\Delta x \ll 1$) we expect to recover second order convergence away from the shock surfaces, which are of lower dimension than the bulk domain. Yet at the resolutions used here, the number of grid points affected by shocks and surfaces is not negligible compared to the total number of grid points, and a single monomial model for the error estimate cannot describe the simulation data. A further complication arises from the fact that the GW strain at \mathcal{I}^+ is given as a function of Bondi time whose relation to simulation time is complex and depends on both spatial location and time [94, 95]. This makes the assignment of a single resolution Δx for each timestep difficult. Instead of attempting to extract a value of Δx as a function of Bondi time, we instead use the fact that CCE introduces negligible error compared to the error in the evolution in the simulation domain [95]. Ignoring the small CCE error, we employ the phase error of the (2,2)-mode of the Newman-Penrose scalar Ψ_4 evaluated on a coordinate sphere of radius $2090 M_\odot$ as a proxy for the phase error in the GW strain at \mathcal{I}^+ so that there exists a unique resolution $\Delta x(t)$ as a function of simulation time. We find that the change in resolution to control dephasing during the inspiral at $t \approx 15 \times 10^3 M_\odot$ is abrupt, and different between the




FIG. 10: Combined error estimate for the phase of $\Psi_4^{2,2}$ after aligning at $t = 37$ ms before the AH forms. We define $\Delta\phi_{2,2}$ as $\Delta\phi_{2,2} = a_1\Delta x + a_2\Delta x^2$ according to Eq. (23).

different Levs. The differences introduced by this change are large enough such that the estimated error at merger is very large (> 0.1 rad at $t \approx 15 \times 10^3 M_\odot$ and multiple radians before an AH is detected) if the change of resolution is included in the dataset. Thus we align the Lev0 and Lev1 waveforms to the Lev2 waveform in the interval $t_{\min} = 15.2 \times 10^3 M_\odot \leq t \leq 20 \times 10^3 M_\odot = t_{\max}$ (-37 ms $< t - t_{\text{horizon}} < -14$ ms), corresponding to 5 inspiral wave cycles of Lev2, minimizing the root-mean-square of the phase difference,

$$\left(\int_{t_{\min}}^{t_{\max}} |\phi_{\text{LevN}}(t - \Delta t_N) - \phi_{\text{Lev2}} + \Delta\phi_N|^2 dt \right)^{1/2}, \quad (24)$$

by varying Δt_N and $\Delta\phi_N$ [120].

Figure 10 shows the error estimate Eq. (23) for the highest resolution run Lev2 from the point of alignment onward. Without alignment the estimated phase error $\Psi_4^{2,2}$ in Lev2 is significantly larger than 1 rad. The alignment procedure allows us to estimate the phase error in a hypothetical simulation that started approximately 40 ms or 11 orbits before AH formation. The estimated error is quite small until the last few orbits when approximately 5 ms before AH formation the error estimate becomes unreliable.

The jump in estimated error coincides with the time we enable fixed mesh refinement on the finite volume grid, which leads to a situation where temporarily the lowest resolution run Lev0 uses a higher resolution than Lev1. This is easily visible in Fig. 11 which shows the finite volume resolution during the final part of the simulations. The inset depicts a zoom-in view of the last 5 ms before AH formation. The slow increase in resolution over time is due to the inspiral of the NSs and the jumps are




FIG. 11: Resolution of the finite volume grid covering the NSs during the simulation. The inset depicts the resolution in the interval during which we enable fixed mesh refinement. Refined regions are added based on the NSs’ separation thus the refined region appears first in the Lev0 simulation leading to a short time interval during which the ordering of resolutions is inverted. This leads to problems when attempting to estimate phase errors in the GW strain if a single monomial dependence of the phase error on resolution is assumed.

due to remapping of the finite volume grid once material starts to leave the simulation box. During the period $-5 \text{ ms} \lesssim t \lesssim -2 \text{ ms}$ before AH formation, while Lev0 is of higher resolution than Lev1, the phase evolution between Lev0, Lev1, and Lev2 is also not proceeding as naively expected and a straightforward error estimate assuming that $|\Delta\phi_{2,2}(\text{Lev1}, \text{Lev2})| < |\Delta\phi_{2,2}(\text{Lev0}, \text{Lev1})|$ yields only an inaccurate estimate for the actual phase error.

V. CONCLUSIONS

In this paper, we presented simulation methods for NSNS mergers in SpEC and discussed the first long NSNS inspiral and merger simulation carried out with SpEC.

The advantages of SpEC compared to other codes are (i) the use of a hybrid pseudospectral – finite-volume approach, which reduces computational costs for the evolution of the spacetime, and, (ii) the use of comoving coordinates, which eliminates the movement of the NSs across the numerical domain during the inspiral. Currently, NSNS simulations using SpEC are not yet as robust as BBH simulations and require careful monitoring. This is particularly true for the phase error whose behavior is not yet fully understood. Further work is required to compute a robust error estimate for the GW phase.

As an example of SpEC’s capabilities, we presented the longest NSNS inspiral simulation performed to date. Two

NSs modeled with a $\Gamma = 2$ EOS and a compactness of 0.16 were evolved for ≈ 22 orbits (44 wave cycles). We demonstrated consistency of our results with another numerical code and found remarkable agreement. A more detailed study comparing results from multiple different numerical codes is planned for the future. Our results show that SpEC is capable of computing consistent long waveforms for NSNS systems up to and beyond merger. Such simulations are of great interest, because long and accurate numerical waveforms are urgently needed in the new field of GW astronomy. They are essential for calibrating and validating simpler waveform models employed to detect GWs and extract information about astrophysics and fundamental physics from observed GWs.

Acknowledgments

We acknowledge helpful discussions with Sebastiano Bernuzzi, Michael Boyle [121], Alessandra Buonanno, M. Brett Deaton, Sarah Gossan, Tanja Hinderer, Kenta Kiuchi, Luis Lehner, Geoffrey Lovelace, Maria Okounkova, David Radice, Jocelyn Read, Masaru Shibata, Nick Tacik, and members of our Simulating extreme Spacetimes (SXS) collaboration (<http://www.black-holes.org>). This research is partially supported by NSF grant nos. PHY-1068881, CAREER PHY-1151197, PHY-1306125, PHY-1404569, PHY-1402916, AST-1205732, AST-1333129, and AST-1333520, by the Alfred P. Sloan Foundation, the Max-Planck Society, by the Sherman Fairchild Foundation, and by the International Research Unit of Advanced Future Studies, Kyoto University. Support for F. Foucart was provided by NASA through Einstein Postdoctoral Fellowship grant numbered PF4-150122 awarded by the Chandra X-ray Center, which is operated by the Smithsonian Astrophysical Observatory for NASA under contract NAS8-03060. The simulations were performed on the Caltech compute cluster *Zwicky* (NSF MRI award No. PHY-0960291), the *Datura* cluster of the AEI, on machines of the Louisiana Optical Network Initiative under grant Ioni_numrel07, and on *Stampede* at TACC under NSF XSEDE allocations TG-PHY990007N and TG-PHY100033. All 2D graphs were generated with the Python-based *matplotlib* [122] and *ipython* [123] packages. *VisIt* [124, 125] was used for 3D and 2D sliced plots. This paper has been assigned Yukawa Institute for Theoretical Physics report number YITP-16-39.

Appendix A: Convergence analysis of TOV stars

In order to verify our numerical method and implementation, we present a convergence analysis of an isolated TOV star, re-examining the convergence study of [36]. Since both spacetime and matter are stationary, any non-trivial evolution is due to numerical error and in partic-




FIG. 12: L_2 volume norm of the difference between the density profile during the simulation and the initial time slice for a stable TOV star for six different resolutions. The inset shows the error computed over the entire hydrodynamical domain (black circles) and restricted to the inner region of the star with a radius $< 5M_\odot$ (orange diamonds) at $t = 1000 M_\odot$.

ular the presence of an atmosphere and sharp surface of the NS influence the observed evolution. This limits the ability of this test to verify the expected order of convergence of SpEC as the observed convergence order depends on the unresolved dynamics at the stellar surface. This fact is evidenced by finding different convergence orders when including or excluding the NS surface from the region in which we compute the convergence order. Nevertheless such a test provides a basic sanity check for the code and we include it here for this reason.

The initial star is a $\Gamma = 2$ polytrope with a total gravitational mass of $M_\infty = 1.40M_\odot$ and a radius of $8.1M_\odot$. The star is evolved with a Γ -law ($\Gamma = 2$) EOS. To quantify the numerical error, we compare the density profile during the evolution with the initial density profile obtained by solving the TOV equation, which can be used as the exact reference solution. Our comparison is only meaningful in the case of the frozen gauge condition described in Sec. B in which the TOV solution is stationary in the simulation frame. See App. B for a more detailed discussion of the influence of the gauge. Employing this gauge, the density profile should stay close to the initial configuration. In particular, we compute $\log_{10} \|\rho(t) - \rho(t=0)\|_2$, where the L_2 -norm is either computed within the entire domain or computed inside the central region of the star, which we define here as $R < 5M_\odot$. Figure 12 summarizes our results. The main plot shows the time evolution of the error $\log_{10} \|\rho(t) - \rho(t=0)\|_2$ computed over the entire star. The overall error clearly decreases with increasing resolution.

The inset of Fig. 12 shows errors at $t = 1000M_\odot$ for

the entire domain (black circles) and inner region (orange diamonds). Theoretically, our finite volume method is limited to second order because of the choice of using flux values at face centers for the averaged fluxes when evaluating the right-hand-side values for the time stepper as well as not distinguishing between averaged values and reconstructed values when computing the primitive variables from the conserved ones. Due to the hybrid grid approach and discontinuities at the surface of the star, a single expected convergence order is difficult to define. We observe that, triggered by the artificial atmosphere, the density profile already deviates at early times $t \approx 50M_\odot$ from the exact solution. Integrating the error over the entire star, we observe a convergence order of ≈ 1.75 (black dashed line). In fact we cannot expect to obtain high order or spectral convergence near the surface of the NS, since the hydrodynamical variables are not differentiable. Restricting the convergence computation to the inner region of the star, we observe a convergence order around ≈ 2.55 . We have verified that the observed convergence order does not change when the atmosphere density is varied by an order of magnitude. There is no obvious reason for observing a convergence order higher than 2. A possible reason is the fact that the *spatial* WENO reconstruction is of higher order and for a system that is stationary and very smooth, the error in smooth regions in the spatially integrated fluxes and inversion is small compared to the error incurred during reconstruction of cell averaged data to cell boundaries this displaying a convergence order between 2 and 5. Over-convergence is also typically observed if the resolution is not yet in the convergent regime. However, in Fig. 12, we observe a convergence order of 2.55 over a large range of resolutions, which makes this explanation less likely. Our analysis shows that when considering the whole domain, the dominant error comes from the surface of the NS, but the error stays localized and does not spoil the convergence in the inner region of the star.

Appendix B: Stellar collapse to a black hole in the generalized harmonic formulation

In this appendix, we address the general problem of simulating the collapse of a single NS to a BH in the generalized harmonic formulation used by SpEC as investigated in [126]. The methods presented here are a prerequisite for following the postmerger evolution to BH formation and ringdown. We demonstrate convergence and accuracy for nonrotating, and uniformly rotating test cases.

1. Initial conditions

We evolve three cases chosen from Baiotti et al. [127, 128]: a TOV case, and two uniformly rotating cases. The initial stars are modeled by a $\Gamma = 2$ polytrope and

Case	M_∞	M_0	$\rho_{0,\text{central}}$	R_{iso}	$r_{p/e}$
D0	1.636	1.770	3.325×10^{-3}	7.54	1.0
D2	1.728	1.913	3.189×10^{-3}	8.21	0.85
D4	1.861	2.059	3.116×10^{-3}	9.65	0.65

TABLE II: Cases evolved in this study. All units are given in terms of solar masses. M_∞ is the ADM (gravitational mass), M_0 is the baryonic mass, $\rho_{0,\text{central}}$ is the maximum baryon density of the configuration, R_{iso} is the coordinate radius in isotropic coordinates, and $r_{p/e}$ is the ratio of polar to equatorial coordinate radii.

evolved with a Γ -law ($\Gamma = 2$) EOS. Rotating equilibria are generated using the code of [129, 130]. The parameters specifying the cases are listed in Table II.

Truncation error alone will cause an unstable stellar equilibrium to evolve either to a stable equilibrium state or to collapse. In order to demonstrate convergence, we prefer to induce a resolved evolution toward gravitational collapse to a BH. We therefore deplete the fluid pressure by a constant factor f_d . A pressure depletion can be thought of as a change in the one-parameter EOS $P(\rho_0)$ used to construct the equilibrium.

In order to avoid violating the constraint equations at the initial time, this must be done carefully. Fortunately, the Hamiltonian and momentum constraints depend on the matter distribution only through the conserved variables E and S_i , defined in Sec. II A. If the primitive variables are adjusted but E and S_i are unchanged, the constraints will be unaffected. For a rotating star, there are two constraint source variables (E and S_ϕ or S^2) and two independent fluid variables (density and rotational velocity). This, suggests the following recipe:

1. Construct a constraint-satisfying equilibrium for the EOS $P = P_0(\rho_0)$, $h^e = h_0^e(\rho_0)$.
2. Take for one's actual EOS $P = f_d P_0(\rho_0)$, $h^e = 1 + f_d[h_0^e(\rho) - 1]$. Since this is the EOS actually used, one may prefer to think that step 1 uses a pressure-enhanced EOS.
3. At each point, re-solve for ρ_0 so that E and S_i are the same as before.

For a perfect fluid

$$E/\sqrt{g} = \rho_0 h W^2 - P, \quad (\text{B1})$$

$$S_i/\sqrt{g} = h W \rho_0 u_i, \quad (\text{B2})$$

where W is the Lorentz factor.

One eliminates W using

$$W^2 = \frac{E/\sqrt{g} + P}{\rho_0 h^e}. \quad (\text{B3})$$

The new density is obtained by solving for the root of the 1D equation

$$S^2 = (E + \sqrt{g}P)[E + \sqrt{g}(P - \rho h^e)]. \quad (\text{B4})$$

Hence, one solves for ρ_0 using Eq. (B4) and uses this to find the new rotation rate via Eq. (B3).

2. Gauge conditions and dynamics

We investigated a series of five different gauge conditions in order to study the coordinate dynamics during gravitational collapse and attempt to determine what condition leads to the most robust simulation of BH formation. All simulations start from the same gauge, set by the initial conditions, but we have the option of transitioning to another gauge, as described in Sec. II D, during the simulation. The gauge conditions are denoted

Frozen for a frozen gauge $H_\alpha(t) = H_\alpha^{\text{initial}}$;

Harm for transition to a pure harmonic gauge $H_\alpha = 0$;

Full for transition to the damped harmonic gauge given by Eq. (10);

Shift, which transitions to a gauge where only the spatial components of Eq. (10) are imposed;

Slice, which only imposes the damped harmonic condition on the t component of Eq. (10).

In all cases but the frozen gauge, we transition away from the initial gauge using the rolloff function Eq. (9) choosing a value for $\Delta T = 10.0 M_\odot$. This results in the $H_\alpha^{\text{initial}}$ contribution being driven to zero within roundoff precision by $t = 30.0 M_\odot$. For the shift only, slicing only, and fully damped harmonic gauge conditions, we transition to (“roll on”) the new gauge with $\Delta T = 25.0 M_\odot$, which is about half of the time to BH formation, so that our damped harmonic gauge condition has fully “kicked in” by the time of BH formation.

We will begin our discussion with results for D0 and since all cases display a similar behavior, we only point out differences between cases where appropriate.

The first observation is that the evolution of the central baryon density during the collapse simulations before AH formation proceeds differently for the different gauge choices. Although the proper time for the central density to reach any value is gauge-independent, the gauge choice affects the evolution of the lapse α as seen in Fig. 13, and therefore the central density growth curve as a function of coordinate time. This observation shows that we can expect different gauge choices to affect the dynamics of AH formation since more rapid evolution is harder for the code to resolve when the flow of time, given by the lapse, differs strongly between different regions of the simulation domain. This eventually leads to steep spatial gradients as different fluid regions evolve apart from each other.

Fig. 13 summarizes the dynamics of case D0 as a function of gauge choice. The top panel shows the maximum of the spatial volume element, \sqrt{g} . This quantity determines how much physical volume is represented per

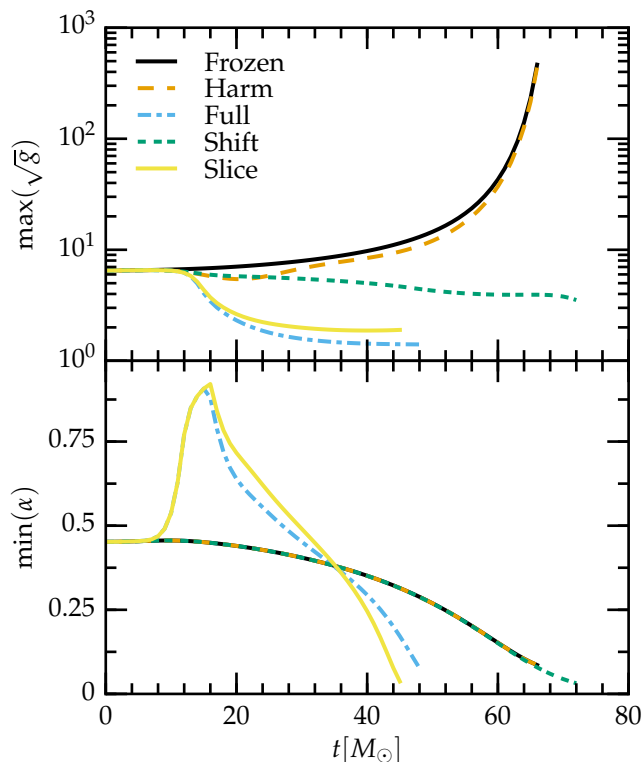


FIG. 13: Evolution of the maximum of the 3-volume element, \sqrt{g} , (top panel) and minimum of the lapse, α , (bottom panel) for the TOV case D0.

unit of coordinate volume. Thus the larger the value of \sqrt{g} is, the lower is the effective resolution, since a larger amount of physical volume is represented by a unit of coordinate volume. For a well resolved simulation, \sqrt{g} must not increase drastically. Otherwise, the coordinate evolution is de-resolving the simulation (effectively the grid is being fatally stretched out and distorted in physical space). One can see in Fig. 13 that this grid stretching is exactly what happens during collapse in pure harmonic and frozen gauges. The damped harmonic gauge is designed to dynamically damp $\log(\sqrt{g}/\alpha)$ to zero, and thus drive \sqrt{g}/α to order unity. This can be understood by looking at the evolution of the lapse function in the bottom panel of Fig. 13. It is interesting to note that the damped harmonic shift condition exhibits a lapse evolution similar to harmonic or frozen gauge, but a distinct evolution for \sqrt{g} . Imposing the damped harmonic condition on the t component (the Slice gauge choice) leads to an evolution of $\max(\sqrt{g})$ and $\min(\alpha)$ qualitatively similar to that produced by the damped harmonic condition. In general, we find that the damping of the coordinate dynamics imposed by the Full, Shift or Slice condition is enough to prevent the divergence of the volume element as the BH forms for case D0, and similar (not shown in the paper) for the rotation cases.

In practice we find that the damped harmonic gauge leads to the most robust BH formation simulations. Be-

cause of this, we use it in our NSNS simulations when the merged object is about to collapse to a BH, yet use harmonic gauge during the earlier phase since it yields faster simulations as described in Section II D. For the D0 case, an AH is first found at coordinate time $t = 48 M_\odot$ for the evolution in damped harmonic gauge. At coordinate time $t = 50 M_\odot$, after the AH has been found successfully a total of 8 times, the collapse evolution terminates. At this point, enough information is available to properly excise the BH and initialize the ringdown simulation. At the time of AH formation, the constraint violation has increased only by a factor of 10 in damped harmonic gauge. In contrast, by the time the constraints have increased by the same factor in the harmonic and frozen gauge runs, an AH has yet to be found and the code eventually crashes due to large constraint violations. Imposing the damped harmonic condition on the t component (Slice choice) or the spatial component (Shift choice) is also sufficient for following black hole formation. In our simulations, when we impose the damped harmonic condition only on the t component, the AH forms earlier in terms of coordinate time, but at the time the AH is found, the constraint violations have already increased by one order of magnitude compared to the damped harmonic gauge. When the damped harmonic shift condition is imposed, the AH is found at later times, and the constraints are at the same order as for the damped harmonic gauge.

Our finding that evolving in damped harmonic gauge is advantageous to resolve BH formation extends and confirms the results of Sorokin [131], who found that the damped harmonic gauge is particularly robust when forming BHs from a complex scalar field in axisymmetric simulations.

3. Convergence of simulations

To study the convergence properties of SpEC for the single star case studied here, we conduct three simulations of the same physical setup using three different resolutions. We use $\Delta x = 300$ m, 250 m, and 200 m in the finite volume grid. In the spectral grid, we use $N_r = 16, 18, 20$ grid points in the radial direction of each spherical shell and an angular resolution including up to $\ell = 10, 12$, and 14 spherical harmonic modes for all but the spherically symmetric case D0 for which we do not increase angular resolution with increasing Lev number. No spectral AMR is used for this test before an AH is found to simplify the convergence behavior. After an AH forms, we use AMR to adjust the number of grid points in the radial direction N_r , but not the spherical harmonic multipole number ℓ . This ensures that the region around the AH is resolved well enough to avoid code simulation failures due to large numerical errors. In Figure 14 we show plots which demonstrate the convergence of the simulations with resolution. We show the L_2 norm of the generalized harmonic constraints for the TOV case D0 and the rotating case D4. Case D2 shows a similar behavior. The




FIG. 14: L_2 norm of the normalized generalized harmonic constraints for the TOV case D0 (top) and the uniformly rotating case D4 (bottom panel). Note that the maximum in the constraints corresponds to the time of BH formation.

maximum in constraint violation corresponds to the time of BH formation, after which we excise the interior of the BH from the numerical domain. This reduces the amount of constraint violation on the grid. Both plots show clear evidence of convergence of the constraints with increasing resolution level before BH formation. After BH formation, case D0 and similarly cases D2 and D4 show an overall decrease of constraint violation with increasing resolution level, yet the detailed evolution with time varies slightly between resolution levels. Partially this is due to AMR which occasionally chooses identical resolution for individual subdomains for different resolution levels. This happens when the estimated truncation error in the affected subdomains is just above/below the threshold for derefinement/refinement for two resolution levels. We also observe a pulse of non-convergent constraint violation in the outer spherical shells which eventually leaves the simulation domain, yet contributes to the observed constraint violation. Case D4 shows a much stronger non-convergent behavior for Lev2 for which we unfortunately are not able to provide a simple explanation.

4. Gravitational Waveforms

Finally discuss briefly the GWs emitted from the collapse of our single NS cases. As pointed out in, e.g., [132–135] the GWs emitted during the collapse of a rotat-




FIG. 15: Gravitational wave emitted during the collapse of the uniformly rotating NS (case D4). We show different resolutions and compare our waveform with the published results of [136]. The BAM waveform of [136] is shifted in time such that amplitude maxima coincide.

ing NS have a particular simple structure consisting of a precursor-burst-ringdown pattern. We find this characteristic structure in our simulations. See Figure 15 for visualization of the $(2,0)$ -mode of Ψ_4 for the D4 case. In addition to our results, we present the waveform of [136] for the same case. As for the comparison of the NSNS waveform in Sec. IV C, Ref. [136] uses the BAM code. The extraction radii are slightly different between the waveforms, while our waveforms are extracted at a fixed coordinate radius of $r = 259M_\odot$, the BAM waveform is extracted at $r = 250M_\odot$. However, the main difference is caused by the artificial pressure perturbation. Here, we perturb the NS according to the discussion in Appendix B 1 and set $f_d = 0.9$ for a pressure depletion of 10%. In [136], the pressure is simply decreased by 0.5%. Although the ansatz of [136] does not ensure that the constraint equations are satisfied on the initial slice, the pressure perturbation is smaller than in our setup. This explains differences in the early part of the waveform at times $t \lesssim 350M_\odot$. In fact, this part of the waveform is unphysical and solely caused by the perturbation of the rotating NS. After $t = 350M_\odot$, the SpEC and BAM waveforms agree well and the maximum amplitude difference is $\lesssim 2 \times 10^{-5}$.

The results presented in these appendices show that SpEC is well suited to study the collapse of a NS into a BH. This is of great importance since in most realistic astrophysical scenarios, the merger remnant formed after the merger of two NSs will eventually collapse to a BH either on a dynamical or secular timescale, depending on its mass and on the nuclear EOS.

-
- [1] B. P. Abbott et al. (LIGO Scientific Collaboration, Virgo Collaboration), *Phys. Rev. Lett.* **116**, 061102 (2016).
- [2] B. P. Abbott et al. (LIGO Scientific Collaboration, Virgo Collaboration), arXiv:1602.03839 [gr-qc] (2016), submitted to PRD.
- [3] J. Abadie, B. P. Abbott, R. Abbott, M. Abernathy, T. Accadia, F. Acernese, C. Adams, R. Adhikari, P. Ajith, B. Allen, et al. (LIGO Scientific Collaboration, Virgo Collaboration), *Class. Quantum Grav.* **27**, 173001 (2010).
- [4] W. G. Anderson and J. D. E. Creighton, in *Astrophysics and Space Science Library*, edited by E. F. Milone, D. A. Leahy, and D. W. Hobill (Springer, Berlin, Germany, 2008), vol. 352 of *Astrophysics and Space Science Library*, p. 23.
- [5] M. Dominik, E. Berti, R. O’Shaughnessy, I. Mandel, K. Belczynski, C. Fryer, D. Holz, T. Bulik, and F. Panarale, *Astrophys. J.* **806**, 263 (2015).
- [6] B. P. Abbott et al. (LIGO Scientific Collaboration, Virgo Collaboration), *Living Rev. Rel.* **19** (2016).
- [7] L. Blanchet, *Living Rev. Rel.* **17**, 2 (2014).
- [8] K. Takami, L. Rezzolla, and L. Baiotti, *Phys. Rev. D* **91**, 064001 (2015).
- [9] A. Bauswein, N. Stergioulas, and H.-T. Janka, arXiv:1508.05493 [astro-ph.HE] (2015), submitted to EPJA.
- [10] S. Bernuzzi, T. Dietrich, and A. Nagar, *Phys. Rev. Lett.* **115**, 091101 (2015).
- [11] J. S. Read, C. Markakis, M. Shibata, K. Uryū, J. D. E. Creighton, and J. L. Friedman, *Phys. Rev. D* **79**, 124033 (2009).
- [12] T. Hinderer, B. D. Lackey, R. N. Lang, and J. S. Read, *Phys. Rev. D* **81**, 123016 (2010).
- [13] C. Markakis, J. S. Read, M. Shibata, K. Uryu, J. D. E. Creighton, and J. L. Friedman, in *On recent developments in theoretical and experimental general relativity, astrophysics and relativistic field theories. Proceedings, 12th Marcel Grossmann Meeting on General Relativity, Paris, France, July 12-18, 2009. Vol. 1-3* (2010), pp. 743–745, arXiv:1008.1822 [gr-qc].
- [14] J. S. Read, L. Baiotti, J. D. E. Creighton, J. L. Friedman, B. Giacomazzo, K. Kyutoku, C. Markakis, L. Rezzolla, M. Shibata, and K. Taniguchi, *Phys. Rev. D* **88**, 044042 (2013).
- [15] A. Maselli, L. Gualtieri, and V. Ferrari, *Phys. Rev. D* **88**, 104040 (2013).
- [16] K. Takami, L. Rezzolla, and L. Baiotti, *Phys. Rev. Lett.* **113**, 091104 (2014).
- [17] W. Kastaun and F. Galeazzi, *Phys. Rev. D* **91**, 064027 (2015).
- [18] S. Dall’Osso, B. Giacomazzo, R. Perna, and L. Stella, *Astrophys. J.* **798**, 25 (2015).
- [19] A. Buonanno and T. Damour, *Phys. Rev. D* **59**, 084006 (1999).
- [20] A. Buonanno and T. Damour, *Phys. Rev. D* **62**, 064015 (2000).
- [21] T. Damour, P. Jaranowski, and G. Schaefer, *Phys. Rev. D* **62**, 084011 (2000).
- [22] T. Damour, *Phys. Rev. D* **64**, 124013 (2001).
- [23] T. Hinderer, A. Taracchini, F. Foucart, A. Buonanno, J. Steinhoff, M. Duez, L. E. Kidder, H. P. Pfeiffer, M. A. Scheel, B. Szilagyi, et al., arXiv:1602.00599 [gr-qc] (2016).
- [24] K. Yagi and N. Yunes, *Science* **341**, 365 (2013).
- [25] K. Chatziioannou, K. Yagi, A. Klein, N. Cornish, and N. Yunes, *Phys. Rev. D* **92**, 104008 (2015).
- [26] D. D. Doneva, S. S. Yazadjiev, and K. D. Kokkotas, *Phys. Rev. D* **92**, 064015 (2015).
- [27] S. Bernuzzi, A. Nagar, S. Balmelli, T. Dietrich, and M. Ujevic, *Phys. Rev. Lett.* **112**, 201101 (2014).
- [28] J. A. Clark, A. Bauswein, N. Stergioulas, and D. Shoemaker, arXiv:1509.08522 [astro-ph.HE] (2015), submitted to CQG.
- [29] L. Baiotti, T. Damour, B. Giacomazzo, A. Nagar, and L. Rezzolla, *Phys. Rev. D* **84**, 024017 (2011).
- [30] S. Bernuzzi, A. Nagar, M. Thierfelder, and B. Brügmann, *Phys. Rev. D* **86**, 044030 (2012).
- [31] K. Hotokezaka, K. Kyutoku, and M. Shibata, *Phys. Rev. D* **87**, 044001 (2013).
- [32] M. Favata, *Phys. Rev. Lett.* **112**, 101101 (2014).
- [33] L. Wade, J. D. E. Creighton, E. Ochsner, B. D. Lackey, B. F. Farr, et al., *Phys. Rev. D* **89**, 103012 (2014).
- [34] N. Tacik, F. Foucart, H. P. Pfeiffer, R. Haas, S. Ossokine, J. Kaplan, C. Muhlberger, M. D. Duez, L. E. Kidder, M. A. Scheel, et al., *Phys. Rev. D* **92**, 124012 (2015).
- [35] F. Foucart, R. Haas, M. D. Duez, E. O’Connor, C. D. Ott, L. Roberts, L. E. Kidder, J. Lippuner, H. P. Pfeiffer, and M. A. Scheel, *Phys. Rev. D* **93**, 044019 (2016).
- [36] M. D. Duez, F. Foucart, L. E. Kidder, H. P. Pfeiffer, M. A. Scheel, and S. A. Teukolsky, *Phys. Rev. D* **78**, 104015 (2008).
- [37] M. D. Duez, F. Foucart, L. E. Kidder, C. D. Ott, and S. A. Teukolsky, *Class. Quantum Grav.* **27**, 114106 (2010).
- [38] F. Foucart, M. D. Duez, L. E. Kidder, M. A. Scheel, B. Szilágyi, and S. A. Teukolsky, *Phys. Rev. D* **85**, 044015 (2012).
- [39] M. B. Deaton, M. D. Duez, F. Foucart, E. O’Connor, C. D. Ott, L. E. Kidder, C. D. Muhlberger, M. A. Scheel, and B. Szilagyi, *Astrophys. J.* **776**, 47 (2013).
- [40] F. Foucart, M. B. Deaton, M. D. Duez, E. O’Connor, C. D. Ott, R. Haas, L. E. Kidder, H. P. Pfeiffer, M. A. Scheel, and B. Szilagyi, *Phys. Rev. D* **90**, 024026 (2014).
- [41] M. A. Scheel, M. Boyle, T. Chu, L. E. Kidder, K. D. Matthews and H. P. Pfeiffer, *Phys. Rev. D* **79**, 024003 (2009).
- [42] G. Lovelace, M. A. Scheel, and B. Szilágyi, *Phys. Rev. D* **83**, 024010 (2011).
- [43] I. MacDonald, A. H. Mroué, H. P. Pfeiffer, M. Boyle, L. E. Kidder, M. A. Scheel, B. Szilágyi, and N. W. Taylor, *Phys. Rev. D* **87**, 024009 (2013).
- [44] A. H. Mroue, M. A. Scheel, B. Szilagyi, H. P. Pfeiffer, M. Boyle, D. A. Hemberger, L. E. Kidder, G. Lovelace, S. Ossokine, N. W. Taylor, et al., *Phys. Rev. Lett.* **111**, 241104 (2013).
- [45] J. Blackman, S. E. Field, C. R. Galley, B. Szilágyi, M. A. Scheel, M. Tiglio, and D. A. Hemberger, *Phys. Rev. Lett.* **115**, 121102 (2015).
- [46] S. Ossokine, M. Boyle, L. E. Kidder, H. P. Pfeiffer, M. A. Scheel, and B. Szilágyi, *Phys. Rev. D* **92**, 104028 (2015).

- [47] B. Giacomazzo and L. Rezzolla, *Class. Quantum Grav.* **24**, 235 (2007).
- [48] T. Yamamoto, M. Shibata, and K. Taniguchi, *Phys. Rev. D* **78**, 064054 (2008).
- [49] M. Thierfelder, S. Bernuzzi, and B. Bruegmann, *Phys. Rev. D* **84**, 044012 (2011).
- [50] D. Radice, L. Rezzolla, and F. Galeazzi, *Mon. Not. Roy. Astr. Soc.* **437**, L46 (2014).
- [51] P. Mösta, B. C. Mundim, J. A. Faber, R. Haas, S. C. Noble, T. Bode, F. Löffler, C. D. Ott, C. Reisswig, and E. Schnetter, *Class. Quantum Grav.* **31**, 015005 (2014).
- [52] S. Bernuzzi, A. Nagar, T. Dietrich, and T. Damour, *Phys. Rev. Lett.* **114**, 161103 (2015).
- [53] K. Barkett, M. A. Scheel, R. Haas, C. D. Ott, S. Bernuzzi, D. A. Brown, B. Szilágyi, J. D. Kaplan, J. Lippuner, C. D. Muhlberger, et al., *Phys. Rev. D* **93**, 044064 (2016).
- [54] L. T. Buchman, H. P. Pfeiffer, M. A. Scheel, and B. Szilágyi, *Phys. Rev. D* **86**, 084033 (2012).
- [55] G. Lovelace, M. Boyle, M. A. Scheel, and B. Szilágyi, *Class. Quantum Grav.* **29**, 045003 (2012).
- [56] L. E. Kidder, M. A. Scheel, S. A. Teukolsky, E. D. Carlson, and G. B. Cook, *Phys. Rev. D* **62**, 084032 (2000).
- [57] L. Lindblom, M. A. Scheel, L. E. Kidder, R. Owen, and O. Rinne, *Class. Quantum Grav.* **23**, 447 (2006).
- [58] M. A. Scheel, H. P. Pfeiffer, L. Lindblom, L. E. Kidder, O. Rinne, and S. A. Teukolsky, *Phys. Rev. D* **74**, 104006 (2006).
- [59] B. Szilágyi, L. Lindblom, and M. A. Scheel, *Phys. Rev. D* **80**, 124010 (2009).
- [60] D. A. Hemberger, M. A. Scheel, L. E. Kidder, B. Szilágyi, G. Lovelace, N. W. Taylor, and S. A. Teukolsky, *Class. Quantum Grav.* **30**, 115001 (2013).
- [61] S. Ossokine, L. E. Kidder, and H. P. Pfeiffer, *Phys. Rev. D* **88**, 084031 (2013).
- [62] F. Foucart, M. B. Deaton, M. D. Duez, L. E. Kidder, I. MacDonald, C. D. Ott, H. P. Pfeiffer, M. A. Scheel, B. Szilágyi, and S. A. Teukolsky, *Phys. Rev. D* **87**, 084006 (2013).
- [63] C. D. Muhlberger, F. H. Nouri, M. D. Duez, F. Foucart, L. E. Kidder, et al., *Phys. Rev. D* **90**, 104014 (2014).
- [64] F. Foucart, E. O'Connor, L. Roberts, M. D. Duez, R. Haas, L. E. Kidder, C. D. Ott, H. P. Pfeiffer, M. A. Scheel, and B. Szilágyi, *Phys. Rev. D* **91**, 124021 (2015).
- [65] X.-D. Liu, S. Osher, and T. Chan, *J. Comput. Phys.* **115**, 200 (1994).
- [66] G.-S. Jiang and C.-W. Shu, *J. Comput. Phys.* **126**, 202 (1996).
- [67] As described in [63], instead of adding a fixed $\epsilon = 10^{-6}$ to each smoothness indicator β , we instead add $\epsilon(1 + \sum_i y_i)$ with $\epsilon = 10^{-17}$.
- [68] B. v. L. A. Harten, P. D. Lax, *SIAM Rev.* **25**, 35 (1983).
- [69] J. M. Hyman, Tech. Rep. COO-3077-139, ERDA Mathematics and Computing Laboratory, Courant Institute of Mathematical Sciences, New York University (1976).
- [70] J. P. Boyd, *J. Comput. Phys.* **103**, 243 (1992).
- [71] A. Tchekhovskoy, J. C. McKinney, and R. Narayan, *Mon. Not. Roy. Astr. Soc.* **379**, 469 (2007).
- [72] V. Springel, *Mon. Not. Roy. Astr. Soc.* **401**, 791 (2010).
- [73] J. W. Wadsley, G. Veeravalli, and H. M. P. Couchman, *Mon. Not. Roy. Astr. Soc.* **387**, 427 (2008).
- [74] F. Foucart, M. D. Duez, L. E. Kidder, and S. A. Teukolsky, *Phys. Rev. D* **83**, 024005 (2011).
- [75] B. Szilágyi, *Int. J. Mod. Phys. D* **23**, 1430014 (2014).
- [76] E. Schnetter, S. H. Hawley, and I. Hawke, *Class. Quantum Grav.* **21**, 1465 (2004).
- [77] B. Fryxell, K. Olson, P. Ricker, F. Timmes, M. Zingale, et al., *Astrophys.J.Suppl.* **131**, 273 (2000).
- [78] M. Anderson, E. Hirschmann, S. L. Liebling, and D. Neilsen, *Class. Quantum Grav.* **23**, 6503 (2006).
- [79] M. J. Berger and P. Colella, *J. Comput. Phys.* **82**, 64 (1989).
- [80] L. Baiotti, M. Shibata, and T. Yamamoto, *Phys. Rev. D* **82**, 064015 (2010).
- [81] Z. B. Etienne, Y. T. Liu, and S. L. Shapiro, *Phys. Rev. D* **82**, 084031 (2010).
- [82] W. E. East, F. Pretorius, and B. C. Stephens, *Phys. Rev. D* **85**, 124010 (2012).
- [83] C. Reisswig, R. Haas, C. D. Ott, E. Abdikamalov, P. Mösta, D. Pollney, and E. Schnetter, *Phys. Rev. D* **87**, 064023 (2013).
- [84] T. Dietrich, S. Bernuzzi, M. Ujevic, and B. Brügmann, *Phys. Rev. D* **91**, 124041 (2015).
- [85] P. MacNeice, K. M. Olson, C. Mobarri, R. de Fainchtein, and C. Packer, *Comput. Phys. Commun.* **126**, 330 (2000).
- [86] M. J. Berger and J. Olinger, *J. Comput. Phys.* **53**, 484 (1984).
- [87] <http://www.black-holes.org/SpEC.html>.
- [88] L. Lindblom and B. Szilágyi, *Phys. Rev. D* **80**, 084019 (2009).
- [89] W. E. East, F. Pretorius, and B. C. Stephens, *Phys. Rev. D* **85**, 124009 (2012).
- [90] C. Gundlach, *Phys. Rev. D* **57**, 863 (1998).
- [91] B. van Leer, *J. Comput. Phys.* **23**, 276 (1977).
- [92] N. T. Bishop, R. Gómez, R. A. Isaacson, L. Lehner, B. Szilágyi, and J. Winicour, in *Black Holes, Gravitational Radiation and the Universe*, edited by B. R. Iyer and B. Bhawal (Kluwer, Dordrecht, 1998), chap. 24.
- [93] C. Reisswig, N. T. Bishop, D. Pollney, and B. Szilágyi, *Class. Quantum Grav.* **27**, 075014 (2010).
- [94] M. C. Babiuc, B. Szilágyi, J. Winicour, and Y. Zlochower, *Phys. Rev. D* **84**, 044057 (2011).
- [95] N. W. Taylor, M. Boyle, C. Reisswig, M. A. Scheel, T. Chu, L. E. Kidder, and B. Szilágyi, *Phys. Rev. D* **88**, 124010 (2013).
- [96] C. Reisswig, N. T. Bishop, D. Pollney, and B. Szilágyi, *Phys. Rev. Lett.* **103**, 221101 (2009).
- [97] E. Newman and R. Penrose, *J. Math. Phys.* **3**, 566 (1962).
- [98] C. Reisswig and D. Pollney, *Class. Quantum Grav.* **28**, 195015 (2011).
- [99] M. Boyle, arXiv:1509.00862 [gr-qc] (2015), arXiv:1509.00862.
- [100] F. Foucart, L. E. Kidder, H. P. Pfeiffer, and S. A. Teukolsky, *Phys. Rev. D* **77**, 124051 (2008).
- [101] H. P. Pfeiffer and J. W. York, *Phys. Rev. D* **67**, 044022 (2003).
- [102] C. S. Kochanek, *Astrophys. J.* **398**, 234 (1992).
- [103] L. Bildsten and C. Cutler, *Astrophys. J.* **400**, 175 (1992).
- [104] H. P. Pfeiffer, L. E. Kidder, M. A. Scheel, and S. A. Teukolsky, *Comput. Phys. Commun.* **152**, 253 (2003).
- [105] H. P. Pfeiffer, D. A. Brown, L. E. Kidder, L. Lindblom, G. Lovelace, and M. A. Scheel, *Class. Quantum Grav.* **24**, S59 (2007).
- [106] K. Kyutoku, M. Shibata, and K. Taniguchi, *Phys. Rev. D* **90**, 064006 (2014).

- [107] T. Dietrich, N. Moldenhauer, N. K. Johnson-McDaniel, S. Bernuzzi, C. M. Markakis, B. Bruegmann, and W. Tichy, *Phys. Rev. D* **92**, 124007 (2015).
- [108] E. Gourgoulhon, P. Grandclément, K. Taniguchi, J.-A. Marck, and S. Bonazzola, *Phys. Rev. D* **63**, 64029 (2001).
- [109] K. Taniguchi and E. Gourgoulhon, *Phys. Rev. D* **66**, 104019 (2002).
- [110] K. Taniguchi and E. Gourgoulhon, *Phys. Rev. D* **68**, 124025 (2003).
- [111] L. Baiotti, B. Giacomazzo, and L. Rezzolla, *Class. Quantum Grav.* **26**, 114005 (2009).
- [112] J. A. Faber and F. A. Rasio, *Living Rev. Rel.* **15** (2012).
- [113] T. W. Baumgarte, S. L. Shapiro, and M. Shibata, *Astrophys. J. Lett.* **528**, L29 (2000).
- [114] T. Matsushima and P. S. Marcus, *J. Comput. Phys.* **120**, 365 (1995).
- [115] R. M. Wald, *General Relativity* (University of Chicago Press, Chicago and London, 1984).
- [116] T. W. Baumgarte and S. L. Shapiro, *Numerical Relativity: Solving Einstein's Equations on the Computer* (Cambridge University Press, New York, 2010).
- [117] T. Damour, A. Nagar, D. Pollney, and C. Reisswig, *Phys. Rev. Lett.* **108**, 131101 (2012).
- [118] B. Brügmann, J. A. González, M. Hannam, S. Husa, U. Sperhake, and W. Tichy, *Phys. Rev. D* **77**, 024027 (pages 25) (2008).
- [119] E. Berti, V. Cardoso, and A. O. Starinets, *Class. Quantum Grav.* **26**, 163001 (2009).
- [120] I. MacDonald, S. Nissanke, and H. P. Pfeiffer, *Class. Quantum Grav.* **28**, 134002 (2011).
- [121] *Triton: Code for waveform manipulations*, <https://github.com/moble/Triton>.
- [122] J. D. Hunter, *Comput. Sci. Eng.* **9**, 90 (2007).
- [123] F. Pérez and B. E. Granger, *Comput. Sci. Eng.* **9**, 21 (2007).
- [124] H. Childs, E. S. Brugger, K. S. Bonnell, J. S. Meredith, M. Miller, B. J. Whitlock, and N. Max, in *Proceedings of IEEE Visualization 2005* (2005), pp. 190–198.
- [125] *VisIt, a free interactive parallel visualization and graphical analysis tool*, <https://wci.llnl.gov/codes/visit/>.
- [126] J. D. Kaplan, Ph.D. thesis, California Institute of Technology (2014), URL <http://thesis.library.caltech.edu/7912/>.
- [127] L. Baiotti, I. Hawke, P. J. Montero, F. Löffler, L. Rezzolla, N. Stergioulas, J. A. Font, and E. Seidel, *Phys. Rev. D* **71**, 024035 (2005).
- [128] L. Baiotti, B. Giacomazzo, and L. Rezzolla, *Phys. Rev. D* **78**, 084033 (2008).
- [129] G. B. Cook, S. L. Shapiro, and S. A. Teukolsky, *Astrophys. J.* **398**, 203 (1992).
- [130] G. B. Cook, S. L. Shapiro, and S. A. Teukolsky, *Astrophys. J.* **422**, 227 (1994).
- [131] E. Sorkin, *Phys. Rev. D* **81**, 084062 (2010).
- [132] R. F. Stark and T. Piran, *Phys. Rev. Lett.* **55**, 891 (1985).
- [133] E. Seidel and T. Moore, *Phys. Rev. D* **35**, 2287 (1987).
- [134] E. Seidel, E. Myra, and T. Moore, *Phys. Rev. D* **38**, 2349 (1988).
- [135] E. Seidel, *Phys. Rev. D* **42**, 1884 (1990).
- [136] T. Dietrich and S. Bernuzzi, *Phys. Rev. D* **91**, 044039 (2015).