

Improving the convergence order of binary neutron star merger simulations in the Baumgarte- Shapiro-Shibata-Nakamura formulation

Carolyn A. Raithel^{1,2,3} and Vasileios Paschalidis^{4,5}

¹*School of Natural Sciences, Institute for Advanced Study,
1 Einstein Drive, Princeton, New Jersey 08540, USA*

²*Princeton Center for Theoretical Science, Jadwin Hall, Princeton University,
Princeton, New Jersey 08540, USA*

³*Princeton Gravity Initiative, Jadwin Hall, Princeton University, Princeton, New Jersey 08540, USA*

⁴*Department of Astronomy and Steward Observatory, University of Arizona,
933 N. Cherry Avenue, Tucson, Arizona 85721, USA*

⁵*Department of Physics, University of Arizona, 1118 E. Fourth Street, Arizona 85721, USA*



(Received 1 April 2022; accepted 3 June 2022; published 13 July 2022)

High-accuracy numerical relativity simulations of binary neutron star mergers are a necessary ingredient for constructing gravitational-waveform templates to analyze and interpret observations of compact object mergers. Numerical convergence in the postmerger phase of such simulations is challenging to achieve with many modern codes. In this paper, we study two ways of improving the convergence properties of binary neutron star merger simulations within the Baumgarte-Shapiro-Shibata-Nakamura formulation of Einstein's equations. We show that discontinuities in a particular constraint damping scheme in this formulation can destroy the postmerger convergence of the simulation. A continuous prescription, in contrast, ensures convergence until late times. We additionally study the impact of the equation of state parametrization on the premerger and postmerger convergence properties of the simulations. In particular, we compare results for a piecewise polytropic parametrization, which is commonly used in merger simulations but suffers unphysical discontinuities in the sound speed, with results using a “generalized” piecewise polytropic parametrization, which was designed to ensure both continuity and differentiability of the equation of state. We report on the differences in the gravitational waves and any spurious premerger heating, depending on which equation of state parametrization is used.

DOI: [10.1103/PhysRevD.106.023015](https://doi.org/10.1103/PhysRevD.106.023015)

I. INTRODUCTION

With the advent of gravitational-wave astronomy, binary neutron star mergers offer an exciting new laboratory for studying a wide range of topics, from the dense matter equation of state (EoS) to the production of heavy elements via the r-process, and can even be used as a standard siren for cosmological measurements. Although each of these pursuits is sensitive to different information contained within an observed merger, they all rely (at least in part) on precise modeling of the gravitational waveform to determine the source properties of the binary.

Constructing robust gravitational waveforms for inferring source properties is a challenging problem. While analytic and semianalytic waveform templates suffice during the early inspiral of a neutron star coalescence [1–4], nonlinear hydrodynamical effects become significant in the final orbits. As a result, from the late inspiral onwards, analytic waveforms must be calibrated with fully nonlinear numerical relativity calculations, which solve the Einstein equations coupled to the equations of (magneto) hydrodynamics and neutrinos (e.g., [5–8]). For recent

reviews on the status of numerical relativity simulations of binary neutron star mergers, see e.g., [9–13].

During the inspiral, numerical relativity simulations generally show rigorous self-convergence, typically with errors converging to zero at second or third order with resolution, depending on the details of the numerical scheme (e.g., [14–19]). However, numerical convergence is challenging to establish in the postmerger regime (but see [19]). The challenge is in part due to the turbulent nature of the postmerger evolution, but may also be related to aspects of the spacetime evolution, as has been found for binary black holes [20,21].

In this paper, we investigate two ways of improving numerical convergence in the postmerger phase of a binary neutron star merger, within the Baumgarte-Shapiro-Shibata-Nakamura (BSSN) formulation of Einstein's equations. First, we study how the treatment of parabolic damping of the Hamiltonian constraint can affect the global numerical convergence of the Hamiltonian constraint and of the gravitational-wave signals. One might expect that because the Hamiltonian constraint must be zero in the continuum limit, it might not matter how the damping of

this constraint is implemented. However, the challenge of treating parabolic damping arises specifically for codes with adaptive-mesh-refinement grids, in which adjacent refinement levels may have different degrees of numerical dissipation due to differences in the grid spacing. As we will describe in further detail below, the jumps in effective damping at the refinement level boundaries introduce discontinuities in the associated evolution equations, which can reduce the convergence order of global quantities.

In a separate context, it was recently reported that replacing the standard, discontinuous Kreiss-Oliger dissipation prescription with a continuous prescription was necessary to achieve long-term stability in numerical simulations of charged binary black holes [22]. In this paper, we will demonstrate that discontinuities in a parabolic damping of the Hamiltonian constraint across different refinement levels can also spoil the convergence of binary neutron star merger simulations in the postmerger phase, using the dynamical spacetime and (magneto)hydrodynamics code of Refs. [23–26], as it was most recently extended in [27]. We show that by adopting a continuous prescription for the parabolic term that is used to damp Hamiltonian constraint violations in this framework, we recover convergent behavior in several quantities until late times postmerger.

We stress that the results found in this paper are specific to the BSSN formulation, but we expect them to hold also for more general types of parabolic constraint damping at the analytic level of formulations of the Einstein equations [28,29]. On the other hand, improved constraint damping may also be achieved with different formulations which implement lower-order constraint damping terms, such as the conformal Z4 [30–32] or the generalized harmonic [33–35] families of formulations.

As a second main goal of this paper, we investigate how the smoothness of the dense-matter equation of state affects the convergence order of these simulations. In particular, we perform simulations using the same EoS model, parameterized with either piecewise polytropes (PWP) [36,37] or with the recently proposed generalized piecewise polytrope parametrization [38]. Although piecewise polytropes are commonly used in the merger literature, they are discontinuous in the first derivative of the pressure and hence in the sound speed. These discontinuities are purely artifacts of the construction, and are of potential concern for two reasons. First, as discussed in [24,39,40], when the fluxes of hydrodynamic variables are nonsmooth, the equations of hydrodynamics are nonconvex and this can lead to unphysical solutions such as split waves and composite structures. Additionally, in the context of spectral methods, Ref. [41] recently showed that using a smooth EoS can significantly improve the accuracy of the inspiral waveform, compared to simulations that use a piecewise polytropic EoS. In that work, the authors used a spectral expansion of the EoS to ensure smoothness in the pressure. The generalized PWP construction, which we explore in this work, in a finite-difference/finite-volume numerical

approach, likewise ensures that the sound speed remains continuous at all densities [38].

In this paper, we perform the first numerical simulations of binary neutron star mergers with the generalized PWP construction and we show robust self-convergence in the inspiral gravitational waves using this parametrization, as well as convergence of the Hamiltonian constraint until late times postmerger. We find similar gravitational-wave phase errors during the inspiral with both the standard and smoothed PWP parametrizations. In contrast, we find a difference of ~ 130 Hz in the postmerger gravitational-wave spectra between these parametrizations, which cannot easily be explained by differences in the radii or tidal deformabilities of the neutron stars. This could point to a subtle sensitivity of the postmerger gravitational waves on the smoothness of the underlying EoS that we plan to explore further in forthcoming work.

The outline of the paper is as follows. In Sec. II, we describe our numerical methods, starting with a brief overview of the evolution equations in the BSSN formulation. In Sec. II A, we describe three different prescriptions for damping the Hamiltonian constraint, while in Sec. II B we describe the two EoS parametrizations that we explore in this work. We present our findings in Sec. III, starting in Sec. III A with a comparison of global quantities for each of the Hamiltonian constraint damping treatments and EoS parametrizations. In Sec. III B, we discuss the convergence of the Hamiltonian constraint for each study, and in Sec. III C, we analyze the gravitational-wave signals from all of our simulations. Appendix provides additional details about the smoothed EoS parametrization.

Unless otherwise stated, we use geometrized units in which $G = c = 1$.

II. NUMERICAL METHODS

All simulations in this paper were performed with the dynamical spacetime, general-relativistic (magneto)hydrodynamics code with adaptive mesh refinement of [23–26] as it was most recently extended in [27]. The code is built within the Cactus/Carpet framework [42–44]. Here, we review only a few key aspects of the code, in order to highlight the changes introduced in this paper.

In the code, the spacetime is evolved using the BSSN formulation of the Einstein equations [45,46], in which the evolution equations are given by

$$(\partial_t - \mathcal{L}_\beta)\tilde{\gamma}_{ij} = -2\alpha\tilde{A}_{ij}, \quad (1a)$$

$$(\partial_t - \mathcal{L}_\beta)\phi = -\frac{1}{6}\alpha K, \quad (1b)$$

$$(\partial_t - \mathcal{L}_\beta)K = -\gamma^{ij}D_j D_i \alpha + \frac{1}{3}\alpha K^2 + \alpha\tilde{A}_{ij}\tilde{A}^{ij} + 4\pi\alpha(\rho + S), \quad (1c)$$

$$(\partial_t - \mathcal{L}_\beta)\tilde{A}_{ij} = e^{-4\phi}[-D_i D_j \alpha + \alpha(R_{ij} - 8\pi S_{ij})]^{TF} + \alpha(K\tilde{A}_{ij} - 2\tilde{A}_{il}\tilde{A}_j^l), \quad (1d)$$

and

$$(\partial_t - \mathcal{L}_\beta)\tilde{\Gamma}^i = -\partial_j(2\alpha\tilde{A}^{ij} + \mathcal{L}_\beta\tilde{\gamma}^{ij}), \quad (1e)$$

where $\tilde{\gamma}_{ij}$ is the conformally-related 3-metric, α is the lapse, ϕ is the conformal exponent, K is the trace of the extrinsic curvature K_{ij} , $\tilde{A}_{ij}$ is the conformal traceless part of the extrinsic curvature, R_{ij} is the Ricci tensor associated with the 3-metric of spacelike hypersurfaces γ_{ij} , and $\tilde{\Gamma}^i$ are the conformal connection functions. The matter source terms ρ , S_i , and S_{ij} are the usual projections of the stress-energy tensor. Finally, D_i is the covariant derivative operator associated with γ_{ij} , and $\mathcal{L}_\beta$ is the Lie derivative with respect to the shift, β^i . For calculation of the Lie derivative terms and for further details, see [46].

In terms of these variables, the Hamiltonian and momentum constraints are given by

$$0 = \mathcal{H} = \tilde{\gamma}^{ij}\tilde{D}_i\tilde{D}_j e^\phi - \frac{e^\phi}{8}\tilde{R} + \frac{e^{5\phi}}{8}\tilde{A}_{ij}\tilde{A}^{ij} - \frac{e^{5\phi}}{12}K^2 + 2\pi e^{5\phi}\rho, \quad (2)$$

and

$$0 = \mathcal{M}^i = \tilde{D}_j(e^{6\phi}\tilde{A}^{ji}) - \frac{2}{3}e^{6\phi}\tilde{D}^i K - 8\pi e^{6\phi}S^i. \quad (3)$$

In the continuum limit, both $\mathcal{H}$ and $\mathcal{M}^i$ must be zero for the solution to be physical. Thus, during the numerical evolution of the BSSN equations, these quantities must show convergence to zero, in order to remain consistent with the Einstein equations. We monitor the L_2 norm of the Hamiltonian constraint, $\|\mathcal{H}\|$, as defined e.g., in [47], to test for convergence in our simulations. A fourth-order accurate finite difference method for the BSSN equations is employed as described in [47].

A. Treatment of constraint damping

In simulations involving strongly gravitating matter, the growth of the Hamiltonian constraint can be minimized by the addition of damping terms to the BSSN equations. For example, it has been found that adding in such a damping term to the evolution equation for ϕ can improve the numerical stability and accuracy of the code [48,49]. As in [48], we add a multiple of the Hamiltonian constraint to Eq. (1b) to damp Hamiltonian constraint-violating degrees of freedom, according to

$$(\partial_t - \mathcal{L}_\beta)\phi = -\frac{1}{6}\alpha K + c_H \mathcal{H}, \quad (4)$$

where c_H governs the amplitude of the damping. We note that Eq. (4) is very similar to one of the modified evolution equations proposed, at the analytic level, in [50] to control constraint violations. This version of the evolution equation for ϕ is the standard implementation within the code of Refs. [23–26] for damping the Hamiltonian constraint, and has been adopted in all binary neutron star simulations with this code. The principal part of the new term is $c_H \mathcal{H} \propto c_H \tilde{\gamma}^{ij} e^\phi \partial_i \partial_j \phi$. When $c_H > 0$, it is a second-order parabolic operator which parabolically diffuses the Hamiltonian constraint. As such, it is subject to a Courant stability condition of the form $c_H \Delta t / (\Delta x)^2 \leq A$ (with A of order 1/6 in 3 + 1 dimensional explicit time integrations in flat spacetime), where Δt is the time step and Δx is the grid spacing. Depending on the implementation of c_H , the damping on different levels of the grid can vary due to different refinement levels having different grid spacing, leading to potential complications as we discuss in more detail below. Understanding how the prescription for this parabolic constraint damping term affects the convergence of the code is one of the main goals of this paper.

In [48], and more generally in the standard implementation of Eq. (4) in the code of Refs. [23–26], the damping coefficient is set as

$$c_H = c_{H1} \Delta t, \quad (5)$$

where c_{H1} is a constant. We refer to this prescription as “ Δt -scaled damping”. The rationale behind this form is that the stability condition from the parabolic operator now becomes $c_{H1}(\Delta t / \Delta x)^2 = c_{H1} \text{CFL}^2 \leq A$, where $\text{CFL} = (\Delta t / \Delta x)$ is the Courant factor. Therefore, by choosing c_{H1} one can satisfy the parabolic Courant stability condition independently of the grid spacing, since in these integrations the Courant-Friedricks-Lewy (CFL) factor is often also fixed. However, when both c_{H1} and the CFL factor are fixed, the degree of damping (i.e., c_H) will decrease when moving to evolutions with higher spatial resolution, since the time step on a given refinement level gets smaller as the overall resolution of the simulation is increased. As we will show, using a baseline set of simulations with a constant $c_{H1} = 0.08$, this leads to undesirable behavior in the growth of the Hamiltonian constraint, where the reduced damping leads to larger constraint violations for evolutions at higher resolution. The Δt -scaled damping prescription is what has been used in all prior simulations with the code of Refs. [23–26].

In an attempt to reduce this behavior, we additionally implement a modified version of this parabolic damping treatment, in which c_{H1} is scaled by the global resolution. In particular, we use the standard prescription of Eq. (5) for our highest resolution evolution, while each lower-resolution evolution, i , uses a scaled damping according to

$$c_H = c_{H1}(\Delta x_{\text{HR}} / \Delta x_i) \Delta t. \quad (6)$$

We note that the index i above does not label refinement levels, but rather simulations with the same grid hierarchy and different spatial resolution. We scale the damping coefficient with respect to the highest resolution performed, so that the Courant stability condition is always satisfied. For the three resolutions used in this work (which vary in grid spacing by factors of 1.25; see Sec. II C), the scaled coefficients are thus 0.0512, 0.064, and 0.08. In this prescription, which we refer to as “modified Δt -scaled damping”, the degree of damping still varies across the refinement levels (due to the dependence of c_H on Δt), but the amount of damping on a *given* level is the same for evolutions with different resolutions.

Both of these prescriptions suffer from discontinuities in the damping between refinement levels with different timesteps Δt . Such discontinuities may become problematic when integrating global quantities or as waves propagate across the grid. For example, in a recent study, Ref. [22] found that requiring the Kreiss-Oliger dissipation to be continuous everywhere on the grid (i.e., eliminating the discontinuities of the standard prescription) was essential to achieving long-term, stable evolutions of charged binary black holes.

Motivated by this finding, we thus also implement a third prescription, which ensures a continuous parabolic damping operator for the Hamiltonian constraint. We define this case simply as

$$c_H = c_{H1}, \quad (7)$$

where we use the finest-time level of our highest-resolution grid ($\Delta t = 0.05625$ km) to set $c_H = 0.0045$ km in the simulations presented below. This choice ensures that the Courant stability condition is satisfied on all refinement levels for all resolutions explored in this work. Additionally, because c_H no longer scales with the local time step, the damping is the same everywhere on the grid, thereby ensuring that the same set of partial differential equations will be solved across all refinement levels and independently of the time step. In other words, this is the only prescription for which the continuum formulation is fixed independently of the grid. We refer to this final prescription as “constant c_H damping”.

B. Equation of state parametrization

While discontinuities in constraint damping can cause numerical issues when integrating quantities across the grid or as waves propagate across refinement levels, discontinuities in the sound speed resulting from the equation of state can introduce unphysical solutions in the hydrodynamics of the evolution. This, in turn, may also affect the convergence properties of the simulation. For example, in the context of spectral/finite-volume methods, it has been shown that discontinuities in the sound speed can lead to increased phase errors in the inspiral of a binary neutron

star merger simulation, compared to results found with a smooth EoS [41]. This is of particular interest, given the prevalence in the merger literature of the piecewise polytropic method for modeling the EoS [36,37], in which the EoS is nonsmooth. This leads us to the second main goal of the present paper: to investigate how the smoothness of the EoS affects the convergence of a binary neutron star merger simulation in the context of finite-difference/volume methods. The finite-volume method adopted in this work uses the Harten-Lax-Van Leer scheme in conjunction with the piecewise parabolic reconstruction method as described in [51,52].

To that end, we perform simulations with two different EoS parametrizations: standard piecewise polytropes [36,37] and the generalized piecewise polytrope (GPP) framework introduced recently by [38]. The GPP parametrization was designed to share the accuracy and flexibility of the PWP approach, while also ensuring a smooth (i.e., continuous and differentiable) pressure function. In this section, we briefly summarize the two approaches.

For a PWP EoS, the pressure between two dividing densities, ρ_{i-1} and ρ_i , is defined as

$$P(\rho) = K_i \rho^{\Gamma_i}, \quad \rho_{i-1} < \rho \leq \rho_i \quad (8)$$

where the polytropic constant, K_i , is determined by imposing continuity in the pressure between adjacent polytropic segments, according to

$$K_i = \frac{P_{i-1}}{\rho_{i-1}^{\Gamma_i}} = \frac{P_i}{\rho_i^{\Gamma_i}}. \quad (9)$$

The polytropic index, Γ_i , is then given by

$$\Gamma_i \equiv \frac{\partial \ln P}{\partial \ln \rho} = \frac{\log(P_i/P_{i-1})}{\log(\rho_i/\rho_{i-1})}. \quad (10)$$

In this formulation, it is clear that the pressure is continuous but not differentiable.

In the generalized PWP formulation of [38], the pressure is instead defined as

$$P(\rho) = \bar{K}_i \rho^{\bar{\Gamma}_i} + \Lambda_i, \quad \rho_{i-1} < \rho \leq \rho_i \quad (11)$$

where $\bar{\Gamma}_i \equiv \partial \ln P / \partial \ln \rho$ is the new adiabatic index, which now differs from the right-hand side of Eq. (10). Imposing *differentiability* requires the polytropic constant to be

$$\bar{K}_i = \bar{K}_{i-1} \left(\frac{\bar{\Gamma}_{i-1}}{\bar{\Gamma}_i} \right) \rho_{i-1}^{\bar{\Gamma}_{i-1} - \bar{\Gamma}_i} \quad (12)$$

Finally, the new parameter, Λ_i , is introduced to impose continuity in the pressure, such that

$$\Lambda_i = \Lambda_{i-1} + \left(1 - \frac{\bar{\Gamma}_{i-1}}{\bar{\Gamma}_i} \right) \bar{K}_{i-1} \rho_{i-1}^{\bar{\Gamma}_{i-1}}. \quad (13)$$

TABLE I. Summary of tests run. The resolutions correspond to finest-level grid spacings of $\Delta x_{\text{LR}} = 195$ m, $\Delta x_{\text{MR}} = 156$ m, $\Delta x_{\text{HR}} = 125$ m.

Name	Damping prescription	Parametrization	Resolutions
Δt -scaled	$c_H = 0.08\Delta t$	GPP	LR, MR, HR
Modified Δt -scaled	$c_H = 0.08[\Delta x_{\text{HR}}/\Delta x_i]\Delta t$	GPP	LR, MR, HR
Constant c_H	$c_H = 0.0045$ km	GPP	LR, MR, HR
		PWP	LR, MR, HR

With this construction, the pressure is continuous and differentiable, ensuring that the sound speed will also be continuous [38].

In this paper, we adopt the nuclear EoS ENG, which predicts a radius of 12.06 km for a $1.4 M_\odot$ cold, non-rotating neutron star, and a corresponding maximum mass of $2.24 M_\odot$ [53,54]. These properties are both consistent with latest astrophysical constraints [55–62]. For the PWP case, we use the coefficients provided by [37] for the three-polytrope approximation of ENG at high densities. We additionally use the PWP coefficients from that paper for the SLy EoS [63] to describe the low-density crust. For the GPP formulation, the parameters cannot be calculated directly, but rather must be fit numerically. To do so, we use the fit coefficients for SLy provided by [38] to describe the low-density crust. We then perform a Markov Chain Monte Carlo simulation to find the best-fit parameters for a three-polytrope approximation to ENG at higher densities. We report the details of this fit and compare the results to the PWP approximation in the Appendix.

C. Initial conditions and numerical setup

For all simulations, we construct binary initial data with LORENE [64]. The initial configurations describe two unmagnetized, irrotational, equal-mass neutron stars in a quasicircular orbit, with an Arnowitt-Deser-Misner (ADM) mass of $\sim 2.76 M_\odot$ and an initial separation of 32.5 km. The neutron stars start at zero temperature and are described by either the PWP representation of the ENG EoS [37,53,54], or the generalized PWP approximation of ENG described further in the Appendix.

We extend the cold parametric EoSs to finite temperatures using the M^* framework of [65], which is based on a two-parameter approximation of the particle effective mass and includes the leading-order effects of degeneracy in the thermal prescription. The implementation and validation of this framework into our code was recently described in [27]. In this work, we use an intermediate set of M^* parameters to describe the matter in the degenerate regime, corresponding to $n_0 = 0.12 \text{ fm}^{-3}$ and $\alpha = 0.8$. For additional details, see [27,65].

For each binary evolution, we use nine spatial refinement levels, which are separated by a 2:1 refinement ratio. We additionally use six temporal refinement levels, such that the coarsest four levels are evolved with identical time

steps, while the remaining levels are separated by a 2:1 ratio. The CFL factor on the finest refinement level is 0.45. The outer boundary of the computational domain is located at 3200 km, and we impose equatorial symmetry to reduce computational costs. We study three different resolutions, with grid spacings on the finest level of $\Delta x_{\text{LR}} = 195$ m, $\Delta x_{\text{MR}} = 156.25$ m, and $\Delta x_{\text{HR}} = 125$ m, where the subscripts indicate low, medium, and high resolution, respectively. These grid spacings correspond to ~ 100 , 125, and 156 points across the diameter of each initial neutron star in the x-direction (i.e., along the line connecting the two stars).

D. Summary of simulations performed

All together, we perform a total of four new convergence studies. This includes three convergence studies to study the three parabolic constraint damping approaches described in Sec. II A. For each of these studies, we use the new GPP EoS parametrization, in order to validate its implementation into our code. We additionally perform a fourth study with the PWP EoS parametrization together with the continuous parabolic constraint damping which, as we will show, achieves the best convergence properties of the various damping prescriptions studied. This fourth study enables us to compare the impact of the EoS parametrization on the convergence properties, starting from a robust baseline. We summarize these tests in Table I.

III. RESULTS

We now turn to the results of the simulations for each of the three parabolic damping treatments, as well as for the two EoS parametrizations studied in this work. In all cases except one, we evolve the mergers from approximately the final two orbits, through the merger itself, and for ~ 10 ms – 20 ms postmerger, with the longer range of evolutions possible for the lower resolutions studied. The one exception is the PWP evolution at the highest resolution studied, which we evolve only in the inspiral phase, as this was sufficient for the comparisons we will make below.

A. Comparison of global quantities

In order to confirm the stability of the evolutions, we start by comparing their global spacetime properties, including the L_2 norms of the Hamiltonian constraint, $\|\mathcal{H}\|$, and momentum constraint, $\|\mathcal{M}\|$ (for definitions, see Eqs. 40–43 of [47]), as well as the ADM mass. We show




FIG. 1. Comparison of global properties for the four cases studied in this paper. From left to right: the L_2 norm of the Hamiltonian constraint, the L_2 norm of the momentum constraint, and the fractional change to the ADM mass. For the latter quantity, we compute the fractional change with respect to M_{ADM} at $t = 0$, accounting for the losses from due to gravitational waves leaving the computational domain. All quantities are plotted from the medium resolution evolutions. The vertical dashed line indicates the time of merger, which corresponds to the time of maximum gravitational-wave strain. By reducing the damping parameter c_H to its smallest value to make it continuous across the grid (i.e., constant c_H damping, shown in blue and orange), the Hamiltonian constraint violations grow to slightly larger values over time. The momentum constraint violations and ADM mass conservation are comparable for all three damping prescriptions and for both the PWP and GPP EoSs.

the evolution of these three quantities for all four of our tests in Fig. 1. All results correspond to the evolutions at our intermediate resolution. The Hamiltonian constraint jumps at merger, as is commonly found in compact binary mergers, and continues to slowly grow at late times. The magnitude of this late-time growth is directly governed by the degree of damping in Eq. (4). For example, for the case of the Δt -scaled damping, the damping is enhanced on refinement levels with larger time steps; thus, there is overall more damping in this prescription, and the growth of the Hamiltonian constraint is correspondingly reduced. On the other hand, the constant c_H prescription has the smallest degree of damping ($c_H = 0.0045$ km on every refinement level), and thus the Hamiltonian constraint exhibits slightly larger growth over time.

We additionally find some difference in $\|\mathcal{H}\|$ during the inspiral between the PWP evolution (which uses the constant c_H prescription) and all of the GPP evolutions, with the PWP result being $\sim 50\%$ larger. This difference enters already at $t = 0$, suggesting that the difference in the smoothness of the sound speed between these parametrizations affects the construction of the initial data.

The momentum constraint violations, $\|\mathcal{M}\|$, shown in the middle panel of Fig. 1, are comparable for all tests considered here. Finally, the right panel of Fig. 1 shows the ADM mass (M_{ADM}), i.e., the contribution to the ADM mass integral in our computational domain, compared to the $t = 0$ value for each test, with losses due to gravitational waves added back in. Overall, we find small errors ($\lesssim 0.7\%$) in the ADM mass conservation over the course of the simulation. We find no significant differences in the conservation of M_{ADM} between the different parabolic constraint damping treatments. Finally, in comparing between the two EoS parametrizations, we again find worse conservation of M_{ADM} for the PWP parametrization during the inspiral. However, at late times, the conservation is marginally better

for the PWP parametrization and, overall, we conclude that the difference between the conservation of M_{ADM} for the PWP and GPP parametrizations is small.

All together, these results confirm that the three parabolic damping treatments and two EoS parametrizations studied in this paper can be used to stably evolve binary neutron stars, with robust preservation of the global spacetime properties reported here.

B. Convergence of the Hamiltonian constraint

Although we find similar conservation of the Hamiltonian constraint for the various parabolic constraint damping treatments and EoS parametrizations, we find significant differences in the numerical convergence of $\|\mathcal{H}\|$ between these treatments.

We show the evolution of $\|\mathcal{H}\|$ at all three resolutions for each of our studies in the top row of Fig. 2. The bottom row shows the corresponding convergence order, assuming $\|\mathcal{H}\|$ approaches zero with increasing resolution. This convergence order is defined as

$$n = \log \left(\frac{\|\mathcal{H}(\Delta x_1)\|}{\|\mathcal{H}(\Delta x_2)\|} \right) / \log \left(\frac{\Delta x_1}{\Delta x_2} \right), \quad (14)$$

for two grid spacings, $\Delta x_{1,2}$.

In general, we find similar qualitative behavior in the growth of $\|\mathcal{H}\|$ at all resolutions to what was described for the intermediate resolution in Sec. III A. The one notable exception is the low-resolution evolution with Δt -scaled damping ($c_H = 0.08\Delta t$; red dashed line in Fig. 2), which shows a reduction in $\|\mathcal{H}\|$ at late times. The reason for this turnover is related to the fact that the low-resolution evolution has the largest time steps, Δt (located on the lowest-resolution refinement levels). Because c_H scales with Δt for this prescription, this leads to the highest degree of damping and, hence, the late-time decrease in $\|\mathcal{H}\|$.




FIG. 2. Top panel: L_2 norm of the Hamiltonian constraint for each of the three parabolic constraint damping treatments and two parametrizations studied in this work. From left to right, we show the results for the Δt -scaled damping, the modified Δt -scaled damping, and the constant c_H damping prescription, respectively. Bottom panel: Convergence order, n , for the Hamiltonian constraints, assuming $\|\mathcal{H}\|$ approaches zero for infinite resolution. The left three columns all use the GPP parametrization, while the rightmost column corresponds to the PWP parametrization. Only with the constant c_H treatment do we find positive convergence at late times postmerger, across all resolutions studied.

The Δt -scaled damping prescription also exhibits a troublesome trend, in which $\|\mathcal{H}\|$ in the highest resolution actually *exceeds* the value at lower resolutions, starting shortly after merger. This is due to the fact that at higher resolutions, Δt becomes smaller (for fixed CFL number) and hence the damping also becomes smaller in this prescription. This results in a negative convergence order, as shown in the bottom, left panel of Fig. 2. These findings are consistent with our previous results from [27], in which we used the same $c_H = 0.08\Delta t$ damping prescription and similar initial conditions, but with the standard PWP parametrization of ENG.

Our second parabolic constraint-damping prescription is scaled by the resolution such that $c_H = 0.08(\Delta x_{\text{HR}}/\Delta x_i)\Delta t$. Although c_H still scales with the local timestep, and thus can vary depending on the refinement level, this prescription ensures that the amount of damping on a *given* refinement level is the same for each of the three resolutions studied in this work. The results from evolutions with this treatment, which start from identical initial data as the previous case, are shown in the second column of Fig. 2. With this modified Δt -scaled damping, we find improved convergence behavior, with positive (but decaying) convergence order achieved for nearly 10 ms postmerger. However, at later times ($t - t_{\text{merger}} > 12$ ms), the convergence order again becomes negative between the two highest resolutions studied, indicating loss of convergence.

Finally, the right two columns of Fig. 2 show the results from the evolutions with the constant $c_H = 0.0045$ km damping, for both the GPP and PWP parametrizations. This prescription ensures equal damping on different refinement levels, so that the parabolic damping operator added to the BSSN equations is continuous across the entire grid.

With this continuous damping prescription, we find a positive and stable convergence order until the end of our simulations (~ 15 ms postmerger), for either EoS parametrization. In particular, the medium and high resolution GPP evolutions indicate convergence-to-zero at approximately first order. This is consistent with the order of our hydrodynamical numerical scheme which becomes first-order accurate when shocks arise, e.g., postmerger. We note that the convergence order is lower when calculated with the low-resolution data ($n \approx 0.5$ – 0.6), which suggests that the lowest resolution in our study is not be in the convergent regime for all times, despite covering the diameter of each initial neutron star with ~ 100 grid points. Nevertheless, the finding of $n_{\text{MR,HR}} \approx 1$ convergence only with the constant c_H prescription provides strong evidence that using a constant parabolic damping parameter is necessary to achieve the expected convergence order.

In other words, we find that discontinuities in the parabolic damping term in Eq. (4) between adjacent refinement levels lead to a decay in the convergence of $\|\mathcal{H}\|$. The reasons for this decay are twofold. First, if there are different degrees of damping on different refinement levels, the set of partial differential equations that is being solved will also change depending on the level. In addition, the discontinuities at the refinement level boundaries may propagate in ways that further spoil the numerical convergence. These effects accumulate when integrating $\mathcal{H}$ across the grid to compute the L_2 -norm and the convergence order is correspondingly reduced.

We also compare between the PWP and GPP parametrizations of the EoS in the right two columns of Fig. 2. For both parametrizations, we use the constant c_H damping, to ensure optimal convergence in the comparison. We find

similar convergence results between the two cases, with small differences in the initial convergence order. In particular, we achieve a slightly higher initial convergence order, $n \simeq 2.25$ (2.50), with the GPP parametrization, compared to $n \simeq 1.90$ (1.96) for the standard PWP parametrization, for the convergence order calculated between the high and medium (low) resolution evolutions. As discussed in Sec. III A, where we found larger $\|\mathcal{H}\|$ at $t = 0$ for the PWP parametrization, the slight reduction in the initial convergence order seems to originate at the initial data level. For both parametrizations, the postmerger convergence order asymptotes to $n \simeq 0.5$ – 0.6 , calculated between the low and intermediate resolutions. Given the similarity between the results at low and medium resolutions for the GPP and PWP parametrizations, we did not evolve the highest resolution past the merger for the PWP case. Thus, in summary, the PWP and GPP parametrizations lead to overall similar convergence of $\|\mathcal{H}\|$.

Finally, we note that it is possible that the Kreiss-Oliger dissipation, which is added (at fifth order) to all evolved BSSN and gauge variables at the refinement level boundaries and which has a diffusive coefficient that is different on different refinement levels, may also suffer the same issues as the parabolic term in Eq. (4). Implementing a continuous prescription for the Kreiss-Oliger dissipation may further improve the numerical convergence of our results.

In summary, we find that a continuous parabolic constraint-damping prescription is necessary to recover convergence of $\|\mathcal{H}\|$ in the postmerger phase. We note that this conclusion was exhibited in the context of parabolic constraint damping in the BSSN formulation, but it should be the case for more general types of parabolic constraint damping at the analytic level of formulations of the Einstein equations [28,29]. However, this is not a problem for low-order constraint damping terms such as those of the conformal Z4 [30–32] or the generalized harmonic [33–35] families of formulations. It is not obvious how, or to what extent, discontinuities across refinement levels in Kreiss-Oliger dissipation operators might affect such formulations. We leave the investigation of this question to future work.

C. Comparison of gravitational-wave signals

Finally, we turn to the gravitational-wave signals extracted from each of our simulations. We show the $\ell = m = 2$ mode of the plus-polarized gravitational-wave strain, $h_{2,2}^+$, for each of our simulations in Fig. 3. These signals correspond to face-on mergers, located at a distance of 40 Mpc. We find that the gravitational-wave signals leading up to merger are nearly identical for all parabolic constraint damping treatments and EoS parametrizations considered in this work. However, small differences emerge following the merger. We will spend the remainder of this section discussing these differences in detail. For the specifics of the analysis methods used to extract the




FIG. 3. The $\ell = m = 2$ mode of the plus-polarized gravitational-wave strain for a merger at a distance of 40 Mpc, for each of the parabolic constraint damping treatments and EoS parametrizations considered in this work. The strain is plotted as a function of the retarded time. All results are shown for the medium resolution simulations. The inspirals in all cases are nearly identical, but small differences emerge in the postmerger phase.

gravitational-wave strains and to calculate the corresponding spectra for our simulations, see Appendix C of [66].

Before discussing these signals in detail, we first briefly comment on the convergence properties of the gravitational wave-strain. We find approximately second-order self-convergence in the phase of the gravitational wave strain, $\phi_{2,2}$, during the inspiral, for all four of our tests. Second-order self-convergences requires that

$$\frac{\phi_{\text{LR}} - \phi_{\text{HR}}}{\phi_{\text{MR}} - \phi_{\text{HR}}} = \frac{(\Delta x_{\text{LR}}/\Delta x_{\text{HR}})^2 - 1}{(\Delta x_{\text{MR}}/\Delta x_{\text{HR}})^2 - 1}, \quad (15)$$

where we have suppressed the (2,2) subscript on ϕ for clarity. We can rearrange this expression to alternatively obtain

$$\begin{aligned} & (\phi_{\text{LR}} - \phi_{\text{HR}})[(\Delta x_{\text{MR}}/\Delta x_{\text{HR}})^2 - 1] \\ &= (\phi_{\text{MR}} - \phi_{\text{HR}})[(\Delta x_{\text{LR}}/\Delta x_{\text{HR}})^2 - 1]. \end{aligned} \quad (16)$$

We plot these scaled, differential quantities in Fig. 4, where $\kappa \Delta \phi_{2,2}$ corresponds to the left- or right-hand side of Eq. (16), depending on the color of the line. The alignment of the scaled $\Delta \phi_{2,2}$ curves during the inspiral indicates that second-order self-convergence is indeed achieved. However, after the merger, the self-convergence is lost in all four cases, regardless of the damping treatment or EoS parametrization. Although the curves in Fig. 4 look very different during the postmerger evolution, in all four cases the order of self-convergence becomes negative at late times. As a result, the late-time differences in Fig. 4




FIG. 4. Self-convergence of the phase of the $\ell = m = 2$ mode of the gravitational-wave signal, for (from left to right) the GPP parametrization with Δt -scaled damping, the GPP parametrization with the modified Δt -scaled damping, the GPP and PWP parametrizations with constant c_H damping, respectively. All results are scaled for second-order convergence, according to Eq. (16).

between damping treatments do not hold much significance. We find similar results for the amplitude of the $h_{2,2}$ and the instantaneous gravitational-wave frequency, as well as for other quantities, such as the maximum rest-mass density and the maximum conformal exponent ϕ . That is, for all of these quantities, we find approximate, second-order self-convergence during the inspiral, but that the self-convergence disappears at late times postmerger.

The lack of self-convergence at late times may be a consequence of the lowest resolution evolution not being in the convergent regime (see Sec. III B), or it may have another origin. In either case, without this convergence, we cannot draw strong conclusions about the differences in the waveforms seen in Fig. 3. For example, Fig. 3 indicates a possible difference in the damping rate of the postmerger gravitational-wave amplitudes, depending on the parabolic damping treatment that is adopted. However, whether this behavior is physical or numerical in nature requires further investigation.

Nevertheless, we can take advantage of the approximate second-order self-convergence during the inspiral, to calculate the Richardson extrapolation of the inspiral phase using the low and high resolutions. We then estimate the numerical errors during the inspiral as the difference between this Richardson extrapolation and the phase extracted from our highest resolution evolution. We show the resulting error estimates, $\Delta\phi_{2,2}$, in Fig. 5.

We find that the inspiral phase errors are slightly larger for the evolutions with constant c_H damping, but that, over the course of the (relatively short) inspiral, the differences between the various parabolic damping prescriptions are generally small.

Interestingly, we do not find significant differences in the phase errors between the GPP and PWP parametrizations. This is in contrast to a previous study by [41], which showed that a continuous sound speed significantly reduced the inspiral phase errors, compared to a piecewise polytropic approximation of the same EoS. The authors in that study evolved the spacetime using spectral methods, and used a spectral expansion of the EoS to ensure smoothness (which has a different construction than the

GPP approximation used here), and they also used a different nuclear EoS (SLy). Additionally, in that study, the authors start their simulations at larger initial binary separation than we do, thereby allowing more gravitational-wave cycles for differences in phase to accumulate. Any of these reasons could explain the smaller effect that we find here. We leave further study of this point to future work.

As a further check on our results, we also compute the approximate errors during the inspiral for the maximum rest-mass density, $\rho_{b,\max}$, and the maximum conformal exponent, $\phi_{\max}$. The calculation is identical to that described above, and we find errors in these quantities that are similar to what is shown in Fig. 5. In particular, we find that the approximate (Richardson-extrapolated) errors in $\rho_{b,\max}$ and $\phi_{\max}$ are comparable for each of the different parabolic-damping treatments as well as for the two EoS




FIG. 5. Difference in the gravitational-wave phase, $\phi_{2,2}$, between the phase extracted from the highest resolution evolution and the Richardson extrapolation of the phase, assuming second-order convergence during the inspiral. The standard and modified Δt -scaled parabolic damping treatments lead to similar phase errors during the inspiral, while we find slightly larger phase errors for the constant c_H prescription. The phase errors are nearly identical for the PWP and GPP parametrizations.

parametrizations. In other words, both the self-convergence properties and the approximate inspiral errors for $\rho_{b,\text{max}}$ and ϕ_{max} are similar to what is reported here for the gravitational waves.

In order to compare the gravitational waves in the postmerger phase, we compute the characteristic strain for each evolution, according to

$$h_c = 2f\tilde{h}(f), \quad (17)$$

where f is the frequency and $\tilde{h}(f)$ represents the Fourier transform of the strain. To compute $\tilde{h}(f)$, we first window each signal between $t = t_{\text{merger}}$ and 12 ms postmerger. This window length corresponds to our shortest-duration evolution, and ensures that all spectra have the same effective resolution. Additionally, we include all $\ell = 2, 3$ modes in these spectra and assume a face-on merger located at 40 Mpc. For additional details on our calculation of h_c , see Appendix C of [66]. We show the resulting spectra in Fig. 6.

In general, we find that the postmerger spectra roughly agree with one another, for all four tests considered. However, we find a few notable differences. First, we find a small difference in the approximate errors in the peak frequency, f_{peak} , of the spectra, for the different parabolic damping treatments. We estimate the errors (σ_f) in f_{peak} simply as the difference between the low and medium resolutions; this provides a rough estimate of the numerical error. We find that the error is largest for the noncontinuous parabolic damping treatments: $\sigma_f \approx 80$ Hz and 45 Hz for the standard and modified Δt -scaled damping treatments, respectively. These differences correspond to fractional errors of $\sim 2\%$ and 1% , calculated with respect to the medium resolution. In contrast, the errors are significantly reduced with the constant c_H damping: we find $\sigma_f \lesssim 5$ Hz (0.2%) and 10 Hz (0.4%) with this damping prescription evolved with either the GPP or PWP parametrization, respectively. This is consistent with the overall worse convergence behavior found for $\|\mathcal{H}\|$ with the discontinuous parabolic damping prescriptions (see Sec. III B). In other words, we find that using a continuous damping helps to reduce the approximate errors in the postmerger spectra, at least as defined here.

We find that f_{peak} agrees to within this approximate error for all three damping treatments with the GPP EoS parametrization. In particular, $f_{\text{peak}} \approx 3200$ Hz in all three cases, suggesting that the details of the parabolic-damping method do not significantly influence the peak frequency of the postmerger gravitational waves.

In contrast, we find a systematic difference of ~ 130 Hz (fractional difference of 4%) in f_{peak} between the GPP and PWP parametrizations with the continuous c_H prescription (bottom panel of Fig. 6), which is much larger than our estimated error in f_{peak} for either parametrization. Many




FIG. 6. Characteristic strain for each of the parabolic damping treatments and EoS parametrizations considered in this work, for a face-on merger at 40 Mpc. These spectra include all $\ell = 2, 3$ modes of the gravitational-wave strain. The top panel shows the results for the standard Δt -scaled damping; the middle panel shows the results for the modified Δt -scaled damping, in which the damping coefficient is scaled by the global resolution; and the bottom panel shows the results for the evolutions with constant c_H damping, for both the GPP (blue) and PWP (orange) parametrizations of the EoS. In all cases, the low resolution results are shown in the light dash-dot lines, while the medium resolution results are shown as the darker, solid lines.

studies have shown that f_{peak} is a sensitive probe of the neutron star structure, e.g., of the radius or tidal deformability of the initial stars [10,12,67–69]. However, as we show in the Appendix, the mass-radius and tidal

deformability curves for the GPP and PWP parametrizations of the ENG EoS used in this paper are nearly identical. Indeed, the differences in the radius and tidal deformability for a $1.4 M_{\odot}$ star between the two parametrizations are only $\Delta R_{1.4} = 0.01$ km (0.1%) and $\Delta \Lambda_{1.4} = 4.6$ (1%), respectively. Thus, the difference in f_{peak} cannot be easily explained away in terms of differences in these macroscopic properties.

Instead, this difference may point to a dependence of h_c on the smoothness of the underlying EoS; for example, through the different sound speed gradients that must be resolved for each parametrization. However, it remains possible that this difference in f_{peak} is simply a result of nonconvergence in the postmerger gravitational waves (as found in Fig. 4). Additionally, as we discussed in Sec. III B, we find hints that our lowest-resolution evolution is not yet in the convergent regime for all times. Thus, it is possible that the observed difference in f_{peak} between the GPP and PWP models may go away with increased resolution; however, such a study is beyond the feasibility of the present work.

D. Spurious heating with the PWP and GPP parametrizations

Finally, we briefly discuss the impact of the EoS parametrization on the spurious heating of the neutron star surface during the inspiral. Spurious heating at the neutron star surface is a common feature of merger simulations, and is caused by spurious shocks that form across the steep density gradient of the stellar surface. This heating can raise the premerger stellar surface to significant temperatures, but the heating is typically limited to low-density regions and thus has little effect on the inspiral dynamics. In addition, the shock heating that develops at first contact during merger is much stronger than the spurious inspiral heating. As a result, the inspiral heating is often simply ignored, even though it can hide interesting physical processes, such as the melting of the neutron star crust [70].

One might naturally wonder whether using a smooth EoS can reduce this spurious heating during the inspiral, compared to simulations that use a PWP EoS with large jumps in the sound speed. We find that this is indeed the case.

Figure 7 shows the 2D thermal profiles of the neutron stars at two times during the inspiral and at one time near the end of our simulations, for the GPP and PWP EoS parametrizations. In both cases, the constant c_H damping prescription is used. Figure 7 shows the thermal pressure, P_{th} , relative to the cold pressure, P_{cold} , for all matter with densities above $0.01 n_{\text{sat}}$, where $n_{\text{sat}} = 0.16 \text{ fm}^{-3}$ is the nuclear saturation density. Matter at lower densities (where P_{cold} drops precipitously) is masked in black, so as not to saturate the color scale. Finally, we note that the colorbar varies for the three snapshots, in order to account for the increased spurious heating throughout the evolution.




FIG. 7. Equatorial snapshots showing the thermal pressure relative to the cold pressure, at two times before merger and at the end of the simulations. All times are labeled relative to the time of merger, t_{mer} . The left column shows the results from the evolution with the generalized PWP EoS and the right column shows the results with the standard piecewise polytropic version of the EoS. In both cases, the constant c_H damping prescription is used. Only matter with densities above $0.01 \times n_{\text{sat}}$ (where $n_{\text{sat}} = 0.16 \text{ fm}^{-3}$ is the nuclear saturation density) are plotted; the black background is included for visual clarity. The PWP parametrization leads to more significant spurious heating near the neutron star surface during the inspiral.

In the two premerger snapshots, we indeed find differences in the spurious heating of the neutron star surface, with the piecewise polytropic parametrization leading to higher values of $P_{\text{th}}/P_{\text{cold}}$ in the surface layers of each neutron star. The spurious heating is reduced, though still present, with the use of the generalized PWP parametrization. In the postmerger snapshot, the thermal profiles of the remnants are similar.

In order to explore these spatial profiles in more detail, Fig. 8 shows the median thermal properties of the matter as a function of density, where the median values are calculated within density bins that are log-uniformly sampled. The top row of Fig. 8 shows the median values of $P_{\text{th}}/P_{\text{cold}}$, while the bottom row shows the median temperatures. We focus in particular on the surface region in these plots, which we define as the range of densities from $n/n_{\text{sat}} \approx 0.01$ – 0.5 . This range corresponds to a volume containing 0.99 to 0.9999 of the total mass for a TOV star with the same EoS. During the inspiral, $P_{\text{th}}/P_{\text{cold}}$




FIG. 8. Median thermal properties, calculated within density bins that are log-uniformly spaced, at two times premerger and one time postmerger. These times correspond to the same snapshots as shown in Fig. 7. The top panel shows the median value of $P_{\text{th}}/P_{\text{cold}}$ at each density, while the bottom panel shows the corresponding, median temperature. The PWP parametrization (orange lines) leads to greater heating of the neutron star surface during inspiral compared to the GPP parametrization (shown in blue), where the surface corresponds to densities of $n/n_{\text{sat}} \approx 0.01\text{--}0.5$ (defined as the volume containing 0.99–0.9999 of the total mass for a TOV star with this EoS). Both parametrizations give similar thermal profiles for the late time remnant (far right column).

is approximately twice as large over the neutron star surface when the piecewise polytropes are used, compared to the GPP parametrization. During the inspiral, the surface temperatures likewise vary by a factor of ~ 1.5 , with the PWP parametrization again leading to higher temperatures. In contrast, at late times, the thermal profiles of the remnant are very similar between the two parametrizations. In the postmerger remnant, we only find differences at the very highest densities, where the temperature varies by a maximum factor of ~ 2 . In this inner core region, the thermal pressure is subdominant to the cold pressure (i.e., $P_{\text{th}} < 0.01 P_{\text{cold}}$), so the dynamical impact of this difference is likely to be small. However, such differences in the high-density temperature profiles could be important in determining when bulk-viscous effects become important [e.g., [71,72]], or in determining the long-term cooling of the remnant and neutrino irradiation of the disk.

IV. DISCUSSION AND CONCLUSIONS

In this work, we have investigated two ways of improving the convergence order in several quantities of a binary

neutron star merger simulation within the BSSN formulation. First, we have studied how discontinuities in a constraint damping approach between different refinement levels can spoil the convergence in the postmerger phase. In particular, we studied three different treatments for the parabolic term added to the evolution equation for the BSSN variable ϕ , which is used to damp the growth of the Hamiltonian constraint. One might expect that because this term is supposed to be zero in the continuum limit, it may not matter for the simulations how the constraint damping parameter in this approach is chosen. We showed that the standard “ Δt -scaled” damping ($c_H = 0.08 \Delta t$), which always respects the strict Courant parabolic stability condition as resolution increases, leads to nonconvergence in $\|\mathcal{H}\|$, starting shortly after merger. Scaling the damping coefficient by the resolution ($c_H = 0.08 [\Delta x_{\text{HR}}/\Delta x_i] \Delta t$) to ensure the same (refinement level-dependent) damping for all three resolutions improves the convergence of $\|\mathcal{H}\|$ somewhat, but the convergence order still decays with time with this prescription and is completely lost by the end of our simulations. We find that using a constant prescription for the damping ($c_H = 0.0045$ km at all resolutions and at

all points on the grid) enables us to achieve convergence until late times postmerger. With this continuous damping prescription, we recover the expected first-order convergence of $\|\mathcal{H}\|$ with our numerical scheme for the postmerger phase until the end of our simulations.

In other words, we find that continuity in the parabolic term used to damp the Hamiltonian constraint in the BSSN formulation is necessary to ensure convergence in the postmerger phase of our simulations.

In addition to studying how discontinuities in the parabolic damping of the Hamiltonian constraint affect the convergence of $\|\mathcal{H}\|$, we have also investigated the role of discontinuities in the sound speed, using two different parametrizations of the dense-matter EoS. We compared simulation results for the standard piecewise polytropic parametrization of the nuclear EoS ENG, as well as a “generalized” piecewise polytropic approximation of the same EoS [38], which ensures smoothness in the EoS.

In general, we find only small differences between the evolutions with these two parametrizations. Interestingly, we find negligible differences between the inspiral gravitational waves for these two parametrizations, in contrast to previous work [41]. As discussed further in Sec. III C, this may be due to the different initial binary separation in our simulations, the different nuclear EoS that is being approximated, the different smooth EoS parametrization that we use, or the fact that [41] adopts a spectral method for the spacetime evolution instead of the finite-difference method used here. If the difference in parametrization turns out to be the source of the different results, it would be interesting to understand what feature of the EoS is causing the reduction of the phase errors, as this would imply that continuity in the sound speed alone is insufficient. Further study will be needed to address these questions.

We find a small difference in the peak frequency of the postmerger gravitational waves between the PWP and GPP parametrizations of the EoS, which is larger than our approximate estimate of the numerical errors for these evolutions. This finding may point to a (small) dependence of the postmerger gravitational-wave signal on the smoothness of the microscopic sound speed, which may be important to take into account when estimating the errors for numerical simulations that use parametric EoSs. However, we cannot rule out the possibility that this is simply an artifact of nonconvergence in the postmerger gravitational-wave signals. We leave further exploration of this issue to future work.

Finally, we explored the impact of the EoS parametrization on the spurious heating of the neutron star surfaces during the inspiral. Both the GPP and PWP parametrizations have multiple transition densities at low densities (for details, see the Appendix and Refs. [37,38]). We find that by using an EoS that is smooth across these transitions—i.e., an EoS with no artificial jumps in the sound speed—the spurious heating during the inspiral can be reduced. In the postmerger phase, we do not find significant differences in the remnant thermal profiles, except at the very highest

densities, where the thermal pressure is subdominant to the cold pressure.

As a final remark, we note that it is unclear how discontinuities at the refinement level boundaries in Kreiss-Oliger dissipation affect our results, especially with regards to the gravitational-wave signals, since these are extracted at large distances from the center of mass and hence travel across multiple refinement levels to reach the extraction point. In [22], it was demonstrated that ensuring continuity of the Kreiss-Oliger dissipation operator was crucial for the stability of simulations of charged black hole binaries. Such discontinuities could affect other formulations of the Einstein equations as well, such as the families of Z4 or generalized harmonic formulations. By making the Kreiss-Oliger dissipation continuous, it may be possible to further improve the postmerger convergence in our simulations. This will be the subject of a future investigation of ours.

ACKNOWLEDGMENTS

We are grateful to Michael O’Boyle for providing us with a high-accuracy table of fitting coefficients for the low-density GPP representation of SLy and for his correspondence on fitting for GPP coefficients. We thank William East and Frans Pretorius for providing us with their generic primitives recovery routine. C. R. would like to thank Elias Most for useful conversations related to this work. C. R. gratefully acknowledges support from a joint postdoctoral fellowship at the Princeton Center for Theoretical Science, the Princeton Gravity Initiative, and as a John N. Bahcall Fellow at the Institute for Advanced Study. This work was in part supported by NSF Grants No. PHY-1912619 and No. PHY-2145421 to the University of Arizona. The simulations presented in this work were carried out in part with the Princeton Research Computing resources at Princeton University, which is a consortium of groups led by the Princeton Institute for Computational Science and Engineering (PIC-SciE) and Office of Information Technology’s Research Computing. The simulations were additionally made possible thanks to the Stampede2 cluster at the Texas Advanced Computing Center, under XSEDE Allocation No. PHY190020.

APPENDIX: COMPARISON OF GPP AND PWP EoS APPROXIMATIONS

In this Appendix, we describe our numerical procedure for creating a generalized piecewise polytrope parametrization of the ENG EoS.

For the low-density portion of the EoS, we use the GPP representation of SLy, the coefficients for which are provided in Table II of [38].¹ SLy is used to describe

¹For this work, we use coefficients from a higher-accuracy version of Table II which was graciously provided by M. O’Boyle (priv. comm.).

TABLE II. Generalized piecewise polytrope-fit parameters for the ENG EoS. $R_{1.4}$ and $\Lambda_{1.4}$ indicate the radius and tidal deformability of a $1.4 M_\odot$ neutron star, predicted by the GPP EoS. The columns $\%_{\text{PWP}}$ report the percent differences between the predictions for each of these values, computed by the GPP and the PWP approximations. The column $\%_{\text{full}}$ reports the percent differences between the predictions of the GPP approximation and the full EoS.

EoS	ρ_0 ($\times 10^{14}$ g/cm 3)	$\log_{10} K_1$	Γ_1	Γ_2	Γ_3	$R_{1.4}$	$\%_{\text{PWP}}$	$\%_{\text{full}}$	$\Lambda_{1.4}$	$\%_{\text{PWP}}$	$\%_{\text{full}}$
ENG	1.064	-34.7162	3.277	2.863	3.272	11.95	0.09	1.5	362.5	1.2	11.3

the crust of the EoS, up to the density at which the crust and high-density EoSs intersect.

For the high-density EoS, we follow [38] in using three piecewise polytropic segments, which are divided at fiducial densities $\rho_1 = 10^{14.87}$ g/cm 3 and $\rho_2 = 10^{14.99}$ g/cm 3 . This leaves us with four free variables; $\{\bar{K}_1, \bar{\Gamma}_1, \bar{\Gamma}_2, \bar{\Gamma}_3\}$. From these quantities and the crust coefficients, all other $\bar{K}_i, \Lambda_i$ in Eqs. (11)–(13) are uniquely determined.

We perform a Markov Chain Monte Carlo simulation to find the set of $\{\bar{K}_1, \bar{\Gamma}_1, \bar{\Gamma}_2, \bar{\Gamma}_3\}$ that minimizes the differences between the GPP pressure and the pressure predicted by the tabulated EoS. We bound the fit between densities of 10^{14} g/cm 3 and the density corresponding to the maximum mass. We calculate an initial guess for the parameters by performing a piecewise power-law fit to

$\partial P / \partial \rho$, and iterate from these values for a total of 50,000 iterations. The resulting best-fit parameters are reported in Table II, while the agreement between the tabulated ENG pressure, the standard PWP approximation, and our new GPP fit is shown in the top left panel of Fig. 9. We find only very small differences in the pressures predicted by either approximation of the EoS and the full model.

The top right panel of Fig. 9 shows the adiabatic index for each of these representations of ENG. The adiabatic index of the complete model is smooth, as expected for this nuclear EoS which does not undergo any physical phase transition at high densities. The PWP approximation, in contrast, shows large discontinuities at each of the fiducial densities dividing the polytropic segments. The adiabatic index with the GPP approximation is again smooth, as intended.




FIG. 9. Comparison of properties for the full ENG EoS model, the generalized PWP approximation (GPP), and the PWP approximation. Clockwise from top left: pressure P as a function of density ρ ; the polytropic index, defined as $\Gamma = \partial \ln P / \partial \ln \rho$; the tidal deformability as a function of mass; and the mass-radius relation. The GPP and PWP give nearly indistinguishable results for the global properties of the neutron star, with small deviations from the full ENG model due to the different crust descriptions.

Finally, we also calculate the mass-radius and mass-tidal deformability curves for each representation of the EoS, which we show in the bottom row of Fig. 9. The PWP and GPP representations are nearly identical to one another, and show only a small offset from the full model. This offset stems from the differences in the crust EoS: both approximations switch to the SLy EoS at low densities, whereas the “full model” in Fig. 9 corresponds to ENG at all densities.

We also report the radius and tidal deformability of a $1.4 M_{\odot}$ neutron star predicted by the GPP approximation with our best-fit parameters in Table II. For the purposes of this

paper, the most relevant comparison is that between the PWP and GPP approximations, but we also report the differences between the GPP approximation and full EoS for completeness. We find a fractional difference in $R_{1.4}$ of only 0.09% between the two parametrizations, and a fractional difference of 1.2% in $\Lambda_{1.4}$. Thus, the global properties predicted by the GPP and PWP approximations are nearly indistinguishable, which allows us to directly study the impact of the sound speed treatment on our numerical simulations, without the confounding variable of changes to the macroscopic stellar structure.

-
- [1] T. Hinderer, *Astrophys. J.* **677**, 1216 (2008).
 - [2] T. Binnington and E. Poisson, *Phys. Rev. D* **80**, 084018 (2009).
 - [3] T. Damour and A. Nagar, *Phys. Rev. D* **80**, 084035 (2009).
 - [4] T. Damour, The general relativistic two body problem and the effective one body formalism, in *General Relativity, Cosmology and Astrophysics*, edited by J. Bičák and T. Ledvinka, Fundamental Theories of Physics Vol. 177 (Springer, Cham, 2014), p. 111.
 - [5] T. Dietrich, S. Bernuzzi, and W. Tichy, *Phys. Rev. D* **96**, 121501 (2017).
 - [6] K. Kawaguchi, K. Kiuchi, K. Kyutoku, Y. Sekiguchi, M. Shibata, and K. Taniguchi, *Phys. Rev. D* **97**, 044044 (2018).
 - [7] A. Nagar *et al.*, *Phys. Rev. D* **98**, 104052 (2018).
 - [8] T. Dietrich, A. Samajdar, S. Khan, N. K. Johnson-McDaniel, R. Dudi, and W. Tichy, *Phys. Rev. D* **100**, 044003 (2019).
 - [9] V. Paschalidis, *Classical Quantum Gravity* **34**, 084002 (2017).
 - [10] L. Baiotti and L. Rezzolla, *Rep. Prog. Phys.* **80**, 096901 (2017).
 - [11] M. D. Duez and Y. Zlochower, *Rep. Prog. Phys.* **82**, 016902 (2019).
 - [12] D. Radice, S. Bernuzzi, and A. Perego, *Annu. Rev. Nucl. Part. Sci.* **70**, 95 (2020).
 - [13] R. Ciolfi, *Gen. Relativ. Gravit.* **52**, 59 (2020).
 - [14] D. Radice, L. Rezzolla, and F. Galeazzi, *Mon. Not. R. Astron. Soc.* **437**, L46 (2014).
 - [15] D. Radice, L. Rezzolla, and F. Galeazzi, *Classical Quantum Gravity* **31**, 075012 (2014).
 - [16] S. Bernuzzi and T. Dietrich, *Phys. Rev. D* **94**, 064062 (2016).
 - [17] W. E. East, V. Paschalidis, F. Pretorius, and S. L. Shapiro, *Phys. Rev. D* **93**, 024011 (2016).
 - [18] K. Kiuchi, K. Kawaguchi, K. Kyutoku, Y. Sekiguchi, M. Shibata, and K. Taniguchi, *Phys. Rev. D* **96**, 084060 (2017).
 - [19] E. R. Most, L. J. Papenfort, and L. Rezzolla, *Mon. Not. R. Astron. Soc.* **490**, 3588 (2019).
 - [20] Y. Zlochower, M. Ponce, and C. O. Lousto, *Phys. Rev. D* **86**, 104056 (2012).
 - [21] Z. B. Etienne, J. G. Baker, V. Paschalidis, B. J. Kelly, and S. L. Shapiro, *Phys. Rev. D* **90**, 064032 (2014).
 - [22] G. Bozzola and V. Paschalidis, *Phys. Rev. D* **104**, 044004 (2021).
 - [23] M. D. Duez, Y. T. Liu, S. L. Shapiro, and B. C. Stephens, *Phys. Rev. D* **72**, 024029 (2005).
 - [24] V. Paschalidis, Z. Etienne, Y. T. Liu, and S. L. Shapiro, *Phys. Rev. D* **83**, 064002 (2011).
 - [25] Z. B. Etienne, V. Paschalidis, Y. T. Liu, and S. L. Shapiro, *Phys. Rev. D* **85**, 024013 (2012).
 - [26] Z. B. Etienne, V. Paschalidis, R. Haas, P. Mösta, and S. L. Shapiro, *Classical Quantum Gravity* **32**, 175009 (2015).
 - [27] C. A. Raithel, V. Paschalidis, and F. Ä-zel, *Phys. Rev. D* **104**, 063016 (2021).
 - [28] V. Paschalidis, *Phys. Rev. D* **78**, 024002 (2008).
 - [29] V. Paschalidis, J. Hansen, and A. Khokhlov, *Phys. Rev. D* **78**, 064048 (2008).
 - [30] S. Bernuzzi and D. Hilditch, *Phys. Rev. D* **81**, 084003 (2010).
 - [31] D. Hilditch, S. Bernuzzi, M. Thierfelder, Z. Cao, W. Tichy, and B. Brügmann, *Phys. Rev. D* **88**, 084057 (2013).
 - [32] D. Alic, W. Kastaun, and L. Rezzolla, *Phys. Rev. D* **88**, 064049 (2013).
 - [33] C. Gundlach, G. Calabrese, I. Hinder, and J. M. Martín-García, *Classical Quantum Gravity* **22**, 3767 (2005).
 - [34] F. Pretorius, *Classical Quantum Gravity* **22**, 425 (2005).
 - [35] L. Lindblom, M. A. Scheel, L. E. Kidder, R. Owen, and O. Rinne, *Classical Quantum Gravity* **23**, S447 (2006).
 - [36] F. Özel and D. Psaltis, *Phys. Rev. D* **80**, 103003 (2009).
 - [37] J. S. Read, B. D. Lackey, B. J. Owen, and J. L. Friedman, *Phys. Rev. D* **79**, 124032 (2009).
 - [38] M. F. O’Boyle, C. Markakis, N. Stergioulas, and J. S. Read, *Phys. Rev. D* **102**, 083027 (2020).
 - [39] R. J. LeVeque, *Finite Volume Methods for Hyperbolic Problems* (Cambridge University Press, Cambridge, England, 2002).
 - [40] A. Voss, Exact Riemann solution for the Euler equations with nonconvex and nonsmooth equation of state, Ph.D. thesis, RWTH Aachen University, 2005.
 - [41] F. Foucart, M. Duez, A. Gudinas, F. Hébert, L. Kidder, H. Pfeiffer, and M. Scheel, *Phys. Rev. D* **100**, 104048 (2019).

- [42] G. Allen, D. Angulo, I. Foster, G. Lanfermann, C. Liu, T. Radke, E. Seidel, and J. Shalf, *Int. J. High Perform. Comput. Appl.* **15**, 345 (2001).
- [43] E. Schnetter, S. H. Hawley, and I. Hawke, *Classical Quantum Gravity* **21**, 1465 (2004).
- [44] E. Schnetter, P. Diener, E. N. Dorband, and M. Tiglio, *Classical Quantum Gravity* **23**, S553 (2006).
- [45] M. Shibata and T. Nakamura, *Phys. Rev. D* **52**, 5428 (1995).
- [46] T. W. Baumgarte and S. L. Shapiro, *Phys. Rev. D* **59**, 024007 (1998).
- [47] Z. B. Etienne, J. A. Faber, Y. T. Liu, S. L. Shapiro, K. Taniguchi, and T. W. Baumgarte, *Phys. Rev. D* **77**, 084002 (2008).
- [48] M. D. Duez, P. Marronetti, S. L. Shapiro, and T. W. Baumgarte, *Phys. Rev. D* **67**, 024004 (2003).
- [49] M. D. Duez, S. L. Shapiro, and H.-J. Yo, *Phys. Rev. D* **69**, 104016 (2004).
- [50] G. Yoneda and H.-a. Shinkai, *Phys. Rev. D* **66**, 124003 (2002).
- [51] Z. B. Etienne, Y. T. Liu, V. Paschalidis, and S. L. Shapiro, *Phys. Rev. D* **85**, 064029 (2012).
- [52] Z. B. Etienne, V. Paschalidis, Y. T. Liu, and S. L. Shapiro, *Phys. Rev. D* **85**, 024013 (2012).
- [53] L. Engvik, M. Hjorth-Jensen, E. Osnes, G. Bao, and E. Østgaard, *Phys. Rev. Lett.* **73**, 2650 (1994).
- [54] L. Engvik, E. Osnes, M. Hjorth-Jensen, G. Bao, and E. Ostgaard, *Astrophys. J.* **469**, 794 (1996).
- [55] F. Özel and P. Freire, *Annu. Rev. Astron. Astrophys.* **54**, 401 (2016).
- [56] J. Antoniadis *et al.*, *Science* **340**, 448 (2013).
- [57] E. Fonseca *et al.*, *Astrophys. J.* **832**, 167 (2016).
- [58] H. T. Cromartie *et al.*, *Nat. Astron.* **4**, 72 (2020).
- [59] L. Baiotti, *Prog. Part. Nucl. Phys.* **109**, 103714 (2019).
- [60] C. A. Raithel, *Eur. Phys. J. A* **55**, 80 (2019).
- [61] K. Chatziioannou, *Gen. Relativ. Gravit.* **52**, 109 (2020).
- [62] E. Annala, T. Gorda, E. Katerini, A. Kurkela, J. Nättilä, V. Paschalidis, and A. Vuorinen, *Phys. Rev. X* **12**, 011058 (2022).
- [63] F. Douchin and P. Haensel, *Astron. Astrophys.* **380**, 151 (2001).
- [64] E. Gourgoulhon, P. Grandclement, J.-A. Marck, J. Novak, and K. Taniguchi, Lorene, <https://lorene.obspm.fr/>.
- [65] C. A. Raithel, F. Özel, and D. Psaltis, *Astrophys. J.* **875**, 12 (2019).
- [66] E. R. Most and C. A. Raithel, *Phys. Rev. D* **104**, 124012 (2021).
- [67] V. Paschalidis and N. Stergioulas, *Living Rev. Relativity* **20**, 7 (2017).
- [68] A. Bauswein and N. Stergioulas, *J. Phys. G* **46**, 113002 (2019).
- [69] S. Bernuzzi, *Gen. Relativ. Gravit.* **52**, 108 (2020).
- [70] P. Hammond, I. Hawke, and N. Andersson, *Phys. Rev. D* **104**, 103006 (2021).
- [71] M. G. Alford, L. Bovard, M. Hanauske, L. Rezzolla, and K. Schwenzer, *Phys. Rev. Lett.* **120**, 041101 (2018).
- [72] E. R. Most, S. P. Harris, C. Plumberg, M. G. Alford, J. Noronha, J. Noronha-Hostler, F. Pretorius, H. Witek, and N. Yunes, *Mon. Not. R. Astron. Soc.* **509**, 1096 (2021).