


Thermal aspects of neutron star mergers

P. Hammond¹, I. Hawke¹, and N. Andersson¹*Mathematical Sciences and STAG Research Centre, University of Southampton,
Southampton SO17 1BJ, United Kingdom* (Received 6 September 2021; accepted 7 October 2021; published 5 November 2021)

In order to extract maximal information from neutron-star merger signals, both gravitational and electromagnetic, we need to ensure that our theoretical models/numerical simulations faithfully represent the extreme physics involved. This involves a range of issues, with the finite temperature effects regulating many of the relevant phenomena. As a step toward understanding these issues, we explore the conditions for β -equilibrium in neutron star matter for the densities and temperatures reached in a binary neutron star merger. Using the results from our out-of-equilibrium merger simulation, we consider how different notions of equilibrium may affect the merger dynamics, raising issues that arise when attempting to account for these conditions in future simulations. These issues are both computational and conceptual. We show that the effects lead to, in our case, a softening of the equation of state in some density regions, and to composition changes that affect processes that rely on deviation from equilibrium, such as bulk viscosity, both in terms of the magnitude and the equilibration timescales inherent to the relevant set of reactions. We also demonstrate that it is difficult to determine exactly which equilibrium conditions are relevant in which regions of the matter due to the dependence on neutrino absorption, further complicating the calculation of the reactions that work to restore the matter to equilibrium.

DOI: [10.1103/PhysRevD.104.103006](https://doi.org/10.1103/PhysRevD.104.103006)

I. INTRODUCTION

In terms of cosmic fireworks, neutron star mergers rank (alongside supernovae) as the most spectacular. Reaching temperatures similar to those recorded in intermediate-energy collider experiments, with matter compressed to a density that cannot be reached in the laboratory, these events involve physics at the limit of our understanding—and beyond. The excitement of observing such events in the full glory of multimessenger astronomy, including gravitational waves, is evident. The breakthrough observations of GW170817 [1–3]—including the late inspiral gravitational-wave signal, the short gamma-ray burst that followed the merger and the kilonova emission associated with nuclear reactions in the matter outflow [4]—raised the stakes by providing a wealth of information. However, the modeling of neutron star mergers still presents a severe challenge [5,6], which can only be met by large-scale numerical simulations with an increasing level of realism—conceptually difficult and computationally expensive.

In this study, we focus our attention on the thermal aspects of the problem. This is important for a number of reasons. First of all, recent simulations have established that neutron star mergers may reach higher temperatures than one might naively expect (in some cases rising to almost 100 MeV [7–9]) and it makes sense to ask to what extent this impacts on our understanding of the problem. Second, recent theoretical work points to thermal effects impacting on key physics aspects, like the meaning of chemical

equilibrium [10,11]. Again, we need to explore what impact this new understanding may have on the simulations and the interpretation of the results. This is a natural step toward higher level issues which aim to sharpen the questions we address with numerical relativity simulations. Which aspects of the rich physics are within reach and which are not?

II. THE CURRENT STATE OF THE ART

A. The equation of state

While it may be tempting to focus the discussion on different aspects of our numerical simulations, the key points we aim to make relate to the thermodynamics and the microphysics encoded in the matter equation of state. This means that we need to keep a keen eye on the input physics—and to what extent this physics is encoded in the simulation. Given this, and to provide the proper context for the discussion, it makes sense to first take a closer look at the assumptions associated with the simulated matter model.

Neutron stars are complex physical systems, involving a range of aspects from nuclear physics and condensed matter physics through to electromagnetism and relativistic gravity. As the different issues are (inevitably) linked, the assembly of even moderately realistic models is a challenge. Unless we want to get stuck at the very first step, we have to make simplifications. A natural starting point is to

describe the star as a multifluid system. There are two good reasons for this. First, we have a good intuitive understanding of (at least the basics of) the relevant fluid dynamics. Second, the formalism required to model relativistic multifluid systems is fairly well developed [12]. It is then relatively easy to argue [13] that we need to keep track of (at least) four fluid degrees of freedom for a moderately cold neutron star. We have to consider the usual bulk flow along with the electromagnetic charge current, a heat flux and a relative drift of superfluid neutrons in the star's core. If we focus on neutron star mergers, the thermal aspects dominate and we inevitably have to consider the role of the neutrinos that are generated as the matter composition evolves. As we will see, this complicates the problem in a number of ways.

In view of these different issues, the assumptions required to reach the starting point for all current neutron star simulations may seem rather drastic. We ignore all the multifluid features and reduce the problem to a single (perfect) fluid model (or, in the case of magnetohydrodynamics, a single fluid alongside some simple closure assumption for the charge current). This retains aspects related to the matter composition and the heat, but ignores intricate details regarding the expected relative flows. A pragmatist might argue that this step is necessary to make progress—which may well be true—but a realist might want to urge some level of caution. We need to make sure that we are not trying to draw (too many) conclusions about physics that is not faithfully represented in the model.

The immediate advantage of the single-fluid model is that we can make a direct connection with the thermodynamics and the small scale physics. The natural starting point for this discussion is the first law of thermodynamics, which then states that the change in energy dE of a many-particle-species system (with different species labeled by x , all moving together) is given by

$$dE = TdS - pdV + \sum_x \mu_x dN_x, \quad (1)$$

where T is the temperature, dS is the change in entropy, p is the pressure, dV is the change in volume and μ_x and dN_x are the chemical potential (the energy associated with adding or removing a particle from the system) and change in particle number, respectively, of species x . The first term, TdS , may be identified as the heat put into the system; the second, $-pdV$, corresponds to the work done on the system; and the final term, $\mu_x dN_x$, is associated with a change in number of each particle species, e.g., due to nuclear reactions. Effectively, we have an equation of state for the system, $E(S, V, \{N_x\})$, but in practice it is more convenient to express the thermodynamics in terms of densities. This leads to a Gibbs relation (essentially the integrated and densitized version of (1))

$$e = sT - p + \sum_x n_x \mu_x, \quad (2)$$

where e is the total energy density, s is the entropy density and n_x is each species' number density. Noting that

$$T = \left(\frac{\partial e}{\partial s} \right)_{\{n_x\}}, \quad \mu_x = \left(\frac{\partial e}{\partial n_x} \right)_{s, \{n_y\}}, \quad (3)$$

with $y \neq x$, we see that the pressure is completely determined from an equation of state of form $e = e(n_x, s)$. So far, the arguments are quite general.

The specific simulation we will consider implements the APR equation of state from [14] (as implemented in the CompOSE library [15]). This means that we have matter composed of neutrons, protons and electrons¹ (n , p and e). Neutrinos, although key to many of the aspects we will consider later, are not included in the equation of state prescription (and hence will not be considered in the arguments we sketch at this point). If the system retains local charge neutrality (as we will assume) we have $n_p = n_e$ and if we introduce the baryon number density $n_b = n_n + n_p$ and the electron fraction $Y_e = n_e/n_b$ we have the three-parameter model $e = e(n_b, Y_e, s)$ or, equivalently (after a Legendre transform), $f = f(n_b, Y_e, T)$ where $f = e - Ts$ is the free energy density. The required information is typically provided as an equation of state table.

As much of our discussion will focus on issues relating to chemical equilibrium, let us explore what this involves. In general, the system is in chemical equilibrium with respect to the relevant reactions when the chemical potentials balance (we will discuss the details later in Sec. IV). This typically leads to an equation which can be solved for the equilibrium matter composition, which then reduces the number of parameters by one. We may, for example, use the equilibrium relation to determine the electron fraction. In essence, we are then assuming that nuclear reactions act to reinstate equilibrium on a timescale much faster than the dynamics we intend to study [16]. The upshot is that the equilibrium equation of state for our system has only two parameters, $e = e(n_b, s)$ and if we also assume that the matter is cold—not a useful assumption for neutron star mergers—then we may ignore the entropy and work with a simple barotropic model $e = e(n_b)$.

The arguments we have provided outlines why realistic neutron-star simulations have to be based on a three-parameter equation of state prescription. Reality may be even more complicated, with additional particle species and so on, but the three-parameter model is the minimum we require to describe the hot merger dynamics.

¹For clarity, we ignore the muons even though they are expected to be present throughout the bulk of a neutron star's core. Including them would not be very difficult as they enter the problem in the same way as the electrons.

B. Simulations in context

Having considered the different assumptions and approximations associated with the matter description we are well placed to put existing simulations in context. In the early days of numerical relativity little attention was paid to the matter model. The first hurdle to overcome was to simulate matter moving in a dynamical spacetime—a challenging problem in itself. As a result, early simulations were based on barotropic single-parameter models (polytropes being particularly common). As these simulations were (eventually) demonstrated to be robust, the first efforts to account for thermal aspects were made. Much of this work took the barotropic model as starting point and introduced an effective model for the internal energy and the heat. As this approach is still used in simulations (see for example [17,18]), it is useful to explain the logic involved.

Noting that the barotropic model leads to

$$p = n_b \mu_n - e, \quad (4)$$

we may introduce

$$\mu_n = m + \bar{\mu}, \quad (5)$$

where m is the baryon mass. In terms of the mass density $\rho = mn_b$, we may use Eqs. (4) and (5) to define the (specific) internal energy ϵ through

$$p = \rho + (n_b \bar{\mu} - e) = \rho(1 - \epsilon), \quad (6)$$

where finite temperature effects can now be encoded in ϵ . Phenomenologically (see [19] for a more detailed argument), comparing to the ideal gas law

$$p = nk_B T, \quad (7)$$

where k_B is Boltzmann's constant (later we use units such that $k_B = 1$), we see that $\epsilon = C_v T$, with C_v the specific heat capacity (at fixed volume) while Mayer's relation

$$\frac{k_B}{C_v} = m(\Gamma - 1), \quad (8)$$

with Γ is the adiabatic index, leads to

$$p = \rho\epsilon(\Gamma - 1). \quad (9)$$

This argument motivates the so-called Gamma-law equation of state [20]. It has the advantage of being easy to implement, but there is no reason to expect it to be particularly realistic—at least not for neutron star mergers.

The next step up in complexity involves replacing algebraic models like the polytrope with a tabulated equation of state for cold nuclear matter in equilibrium.

Tabulated data must be used as the microphysics calculations required to evaluate the equation of state at any given point are prohibitively expensive to use live, hence interpolation within a precalculated table becomes a necessity. However, one has to pay careful attention because the interpolation has to be thermodynamically consistent in order to avoid modeling errors. The end result is a fairly significant increase in computational cost (as well as memory footprint, which impacts on the potential use of GPUs for large scale simulations). The most simplistic tabulated equation of state tables are barotropic, however many modern simulations (see [21] for references to the recent literature) use a cold tabulated equation of state in combination with a Gamma-law prescription to approximate thermal effects. The table then includes pressure $p_{\text{cold}}(\rho)$ and specific internal energy $\epsilon_{\text{cold}}(\rho)$ from a zero-temperature nuclear physics derived equation of state, and is augmented with a thermal contribution given by

$$p_{\text{th}}(\rho, \epsilon) = \rho\epsilon_{\text{th}}(\rho, \epsilon)(\Gamma_{\text{th}} - 1), \quad (10)$$

where ϵ_{th} is the thermal component of the specific internal energy $\epsilon_{\text{th}}(\rho, \epsilon) = \epsilon - \epsilon_{\text{cold}}(\rho)$. The total pressure is given by the sum of the cold and thermal components. While this approach is not expected to be realistic (see the discussion in [22]), it is definitely a step in the right direction.

The key step toward realism involves working with a true finite temperature equation of state. There are a number of such models on the market [15], typically developed for supernova simulations. The implementation of these models involves a more complex inversion from evolved to “primitive” variables (as we discuss in Appendix Sec. 2) but it does not add conceptual issues, unless we consider problems involving phase transitions [9,23] or possible mixed matter phases.

When it comes to the matter, the simplest option is to still insist on local equilibrium. The question is if this is a reasonable representation of reality. Are the nuclear reactions fast enough for the equilibrium assumption to be appropriate? Estimates suggest that the answer may be no [24]—at least in parts of the simulated domain—so we need to consider deviations from equilibrium. This raises the complexity level as we need to keep track of the local nuclear reaction rates and implement the associated changes in composition as the simulation proceeds. As an alternative, we may consider the slow-reaction limit, in which the composition is effectively frozen. In the case of npe-matter, this involves advecting the lepton fraction along with the fluid flow. This is the assumption we make in our simulation. This may not be an appropriate assumption either, but it is a useful starting point for discussions of out-of-equilibrium issues. We need to do better, but (as we will argue in the following) there are important lessons to be learned already at this level, and we

need to pay attention to these lessons when developing the next generation of simulations.

C. Going further

The discussion to this point has assumed a single fluid mixture. That is, the different particle species are assumed to move together with a single four velocity, and in addition the fluctuations in each species must relax to the average on negligible scales [25]. However, there is a range of regions in spacetime and phase space relevant for neutron stars where these assumptions are expected to fail. When dealing with neutrinos the fluid approximation fails. When dealing with superfluidity the single-fluid approximation fails [12]. When considering transport properties, or with phase transitions, quantitative calculations must be done outside the single fluid limit, and the application of these calculations on top of single fluid simulations needs care to ensure the approximations remain reasonable. As an example we may consider the estimate for the thermal conductivity from [24], which suggests a relaxation timescale of order

$$t_{\text{relax}} \sim \left(\frac{z}{\text{km}}\right)^2 \left(\frac{T}{10 \text{ MeV}}\right)^2 \text{ s}, \quad (11)$$

where z represents the scale of thermal gradients. This rough estimate suggests that gradients on a scale of a few tens of meters would relax on about a millisecond, a timescale similar to that of the anticipated postmerger oscillations [26–28]. This estimate is interesting because; on the one hand, it suggests that we could probably get away with ignoring the heat flux (the inclusion of which would require us to go beyond the perfect fluid assumption) on the timescale of a typical merger simulations (especially noting that thermal gradients of order 10 meters may not be resolved in the simulation anyway). On the other hand, the timescale resolvable by a typical simulation is of order a fraction of micro-second (combine a grid resolution of $\Delta x \sim 100$ m with a typical time step determined by $\Delta t/\Delta x \simeq 0.1$ to see that the resolvable time variation is of order $\Delta t \sim 10^{-7}$ s). That is, the thermal relaxation may take place on a timescale that should (at least in principle) be resolvable. The issue of small scale thermal features is also closely linked to questions of turbulence and the need to consider a large-eddy approach to the simulations in the first place [25,29].

To highlight the differences, in a multifluid context the expression for the first law extends from Eq. (1) to

$$de = \Theta_a ds^a + \sum_x \mu_a^x dn_x^a, \quad (12)$$

where the chemical potentials have been extended to conjugate momenta μ_a^x (with Θ_a for the entropy, x the species label and a a spacetime index) and the particle

number extends to the fluxes n_x^a [12]. In situations where the different fluxes are not aligned, substantial changes in the equilibrium state can be driven by, for example, entrainment. It is important to note that these differences are expected, and may in fact be essential, on some scales. For example, on microscopic scales the difference in charged species flows drives the charge current and hence is required for the neutron star magnetic field. Similarly, on larger scales, the difference between fluxes is central to the hydrodynamical description of superfluids.

A fully quantitative discussion of transport properties and the interaction of different phases of matter within a neutron star will need to start from a micro-scale theory, such as a multispecies Boltzmann equation based model [30]. However, the dimensionality of these kinetic-theory based calculations make them computationally prohibitive in many dynamical situations. To estimate to what extent the (single) fluid approximation is valid, the interaction terms between and within each species need to be estimated. Unfortunately, once the fluid approximation is made the information required to confirm its validity is lost, and hence the estimates must be used with care.

Having outlined how different equation of state models are used in numerical simulations and sketched some of the main points of concern, let us return to the elephant in the room—the neutrino. It is well known that the neutrinos are central to any discussion of hot neutron star merger dynamics: They remove energy from the system, which may drive matter outflows, and facilitate the reactions that lead to a changing composition which may, in turn, lead to dissipation acting on the fluid dynamics. Yet, current simulations tend to either not include the neutrinos at all or account for them approximately through some (suitably simple) leakage scheme. In fact, many recent discussions [7,8] are based on postprocessing simulation results. The key point is that the neutrinos do not naturally lend themselves to a fluid description. Instead, the strategy for dealing with the neutrino problem is formally the same as for photons. The relevant radiation transfer problem must be solved, taking into account frequency dependence and directionality. In current models this is done at different levels of sophistication, starting from some kind of leakage scheme alongside a suitable closure condition for the moment expansion of the radiation stresses (recent examples include [31–33] with, in particular [34] paying attention to the multifrequency aspects), increasing realism by considering the angular dependency (as in [35,36]) or taking a full-blown Monte-Carlo approach (as in [37–39]). In essence, the problem is complex and extremely expensive computationally.

However, and this is the main point we want to make here, there are issues we need address before we can introduce the neutrino aspects in a faithful fashion. Somewhat simplistically, in order to establish how the neutrinos impact on the dynamics we need to know to what

extent they are trapped in the simulated “fluid elements”. This is important as it, in turn, dictates the rate at which weak interactions work to establish local beta equilibrium [10,40]. The issue of deviation from equilibrium comes to the fore and we need to carefully consider what we mean by “equilibrium” in the first place.

III. PERFORMING AN OUT-OF-EQUILIBRIUM SIMULATION

As a step toward studying how the equilibrium processes affect the merger dynamics, we have performed a simulation in which the composition of the fluid was advected along with the fluid flow and reactions were ignored. However, before we cover the details of said simulation, let us discuss a few general issues that arise when simulating a continuous fluid on a discrete grid.

A. Analysis quantities

Numerical simulations typically provide a limited set of variables which may not be directly useful for extracting physically interesting signatures. Each step in the analysis process that leads to observable or interesting quantities will add numerical error, and may also introduce additional modeling errors in ways that are hard to disentangle.

The key quantities that are controlled in a simulation (based on the standard Valencia formulation, see Appendix Sec. 2 for a summary of the relevant evolution equations) are the conserved variables—the relativistic equivalents of rest mass density, spatial momenta, and energy. These are represented as cell integral averages in the simulation, and no information about the structure on subgrid scales is available. While some simulations use large-eddy-style approaches [25,29], these only give the (assumed) contribution from physics on smaller scales, and not the subgrid structure itself.

Many quantities of interest can be computed locally using purely algebraic relations, for example thermodynamical quantities from the equation of state. The accuracy of these quantities is limited by the numerical error as well as modeling errors. The modeling errors come from, in particular, neglecting subgrid fluctuations, and hence assuming that the average of the thermodynamical quantity is equivalent to the thermodynamical quantity of the average. This may be reasonable over the bulk of the spacetime, but the modeling error is hard to quantify in the (necessarily underresolved) turbulent region that forms in the merger process. Notably, this is the hottest region of the simulation.

Additional quantities can be calculated using nonlocal operations. Examples include vorticities (using derivatives of given or computed quantities) or optical depths (using integrals of given or computed quantities). In addition to the problems noted above, this will involve additional error due to the numerical approximation of the nonlocal

operator. Moreover, as nonlocal operators are typically less accurate at high frequencies (see, e.g., [41]), they will again be problematic in the turbulent region of interest. Different quantities of interest may be robust while others are not. For example, the gravitational-wave signal depends on bulk integrated quantities (in the sense of the quadrupole formula) which are likely to be robust as they mainly depend on low-frequency features. In contrast, emission spectra may be less reliable due to the dependence on high-frequency effects.

Finally, there are quantities that can be considered “local” but which depend on nonlocal constructions. The key example of this is Lagrangian tracers, which have been used to investigate the dynamics of phase transitions and other neutron star properties [9,23]. This relies on evaluating the thermodynamic potentials and velocity field at a “single particle point” which is then advected through the spacetime. Using a set of point quantities, rather than cell-integral-average quantities, one might hope to bypass the subgrid structure issues discussed above. However, all the problems of nonlocal operators and the modeling of subgrid structure discussed above will enter through the required interpolation scheme.

This is, in fact, an interesting problem. Lagrangian tracers as an analysis probe of fluid properties have been investigated in depth in the Newtonian literature (example reviews include [42]). In particular, the discussion in [42] emphasizes that three requirements must be met to get accurate results from Lagrangian tracers. First, the simulation must resolve sufficiently small scales in the instantaneous velocity field. This will severely limit the utility of tracers passing through turbulent regions of neutron star simulations. Second, the interpolation scheme used must be of sufficiently high accuracy and differentiability. Finally, a sufficient number of tracers must be used for statistical accuracy (e.g., so that the average Lagrangian velocity is sufficiently close to the average Eulerian velocity). This last point highlights that Lagrangian tracer information should be used statistically. The analysis in [42] also demonstrates that the number of tracers required scales with the Reynolds number, which may be impractically high for neutron star applications. These issues highlight why we focus on local grid-cell based analysis in the following.

B. Our setup

To take the first step toward quantifying the out of equilibrium behavior in binary neutron star mergers we focus on a single simulation performed using the Einstein Toolkit [43]. We used the BSSN formalism of the ADM equations as provided by McLachlan [44] for the spacetime evolution and GRHydro [45] for the fluid evolution. The microphysics is represented by the APR equation of state [14] from the CompOSE database [15]. This model can be seen as a “vanilla” representation of the relevant physics as it assumes matter composed only of neutrons, protons and

electrons and the obtained neutron star radii accord well with current observational constraints (see for example [46]). In order to perform the simulation, it was necessary to make modifications to both GRHydro and EOS_Omni—the general purpose equation of state interface available in the Toolkit—in order to support the three-parameter tabulated form of the APR equation of state from CompOSE. We describe these modifications in the Appendix.

Initial data was obtained using LORENE [47], with a fixed temperature of $T = 0.02$ MeV across the entire domain (we comment on this later in Sec. III C), and under the assumption of cold β -equilibrium (see Sec. IV A). It is important to note that, due to the low-temperature cut-off in the equation of state table we have to assume an initial temperature that is much higher than the expected core temperature of a mature neutron star (below 10^6 K $\sim 10^{-4}$ MeV or so). The reason this is an issue will become clear in the discussion of Fig. 3 in Sec. III C. Finally, the initial separation of the two stars was 40 km, and each star had an initial baryon mass of $M_b = 1.4 M_\odot$.

While the composition of the fluid is fixed at the initial data stage to be in β -equilibrium, this constraint was not imposed on the fluid during the evolution, instead the electron fraction Y_e was advected along with the fluid flow.

The simulation does not feature reactions between the constituent species of the fluid, neutrinos are not included, and no magnetic fields are present. The simulation domain uses adaptive mesh refinement for the grid setup, with the finest grids being centered on and completely encompassing each star, and the finest grid spacing being ~ 400 m, and we use $\Delta t/\Delta x = 0.25$.

C. Comments on the temperature

At a glance, the results of our simulation agree well with other models considered in the literature (indicating that our comparatively low numerical grid resolution has not impacted significantly on the results, see for example [7–9]) and we believe they provide a fair representation of current simulation technology. The results we discuss should not, in any sense, be “particular” to our chosen model. This is evident if we compare the snapshots in Fig. 1 to similar results in the literature.

The results in Fig. 1 show the temperature T , electron fraction Y_e , rest-mass density with respect to nuclear saturation density ρ/ρ_{nuc} , and the neutron chemical potential μ_n in the equatorial plane, both at the time of merger (left panels) and 5 ms later (right panels). The results are relevant for a range of questions one may want to answer




FIG. 1. Temperature T , electron fraction Y_e , rest-mass density with respect to nuclear saturation density ρ/ρ_{nuc} , and neutron chemical potential μ_n for a simulation of the merger of two $1.4 M_\odot$ neutron stars using the APR equation of state from [14]. The left column coincides with the merger time, while the right column snapshots are 5.0 ms postmerger.

with simulations. For example, while the shock heating associated with the merger heats the matter up to temperatures well above 10 MeV, the high-density core of the merger remnant remains relatively cold—the postmerger configuration is effectively a rotating, fairly cold, high-density “peanut” that lasts for some time (we have not tracked the evolution through to the stage where the remnant may eventually collapse to form a black hole as our main interest is in the dynamics immediately following the merger).

The temperature results impact on the expected postmerger oscillations and issues relating to, for example, bulk viscosity (which we will return to in Sec. IV B). It is notable that the thermal hotspots, which start out localized and corotating with the high-density core, spread out as the evolution proceeds. Finally, the results for the chemical potential demonstrate that the matter outflow involves relatively large values of μ_n , while the values for the electron fraction are moderate. These features influence the r-process (with lower values of Y_e required to form heavier elements) and the signature of the anticipated kilonova emission. It is also worth noting the spiral nature of the matter outflow, involving colder matter associated with the rarefaction wave following the merger shock. This feature is notable because it is yet another aspect that is affected by the low-temperature cut-off in the equation of state table. It is important to get this part right as it impacts directly on both the r-process and possible jet formation. There are (obviously) many interesting issues here, but we focus on questions relating to the temperature and the matter composition in the following.

Given this particular focus, it is worth remarking on the robustness of the numerical implementation. As a measure of the reliability of the thermal description we may consider the total entropy per baryon, as shown in Fig. 2. This is a

useful measure (as it should not be negative) which provides a good idea of to what extent the high-density matter is hot or cold (in a nuclear physics sense). The entropy results bring out the expectations—the high density core remains cold—but they also highlight an artefact of the simulations (also apparent in, for example, the results of [7–9]). The region close to the neutron star surface is artificially hot already from the early stages of the simulation. This is also apparent from the temperature result in Fig. 1. We see that, away from the shock region (where the two stars first touch) the low density matter near each star’s surface has reached a temperature well above 10 MeV. This feature is clearly “unfortunate” as it means that we cannot discuss the fine print of the thermal physics (e.g., the point at which the neutron star crust melts).

The artificially high surface temperature may, perhaps intuitively, be explained in terms of a numerical shock associated with the transition from the high-density neutron star interior to the low-density (artificial) atmosphere. Consider any perturbation in the interior that leads to a wave propagating outwards. As it travels “down” the density gradient, the speed of sound (and hence the characteristic speed) decreases, meaning that the perturbation steepens and will eventually lead to a shock. There is an assumption that this shock takes place before hitting the atmosphere (a relevant toy problem demonstrating the effect is discussed in [48]), which seems to match the numerical results, but the subsequent propagation through the atmosphere is certainly artificial. Such a feature may (to some extent) be unavoidable. One may argue that this is not a major problem, as the artificial heating is overwhelmed by the entropy generated in the merger and therefore does not feature prominently in the postmerger dynamics. Pragmatically, this may be true on average and hence sufficient for gravitational wave emission (which is




FIG. 2. Total entropy per baryon S for the simulation of the merger of two $1.4 M_{\odot}$ neutron stars using the APR equation of state from [14]. The left panel coincides with merger time, while the right panel is 5.0 ms postmerger. Note that $S \leq 1$ indicates degenerate matter while $S \gg 1$ indicates strongly nondegenerate matter.

determined by the bulk matter dynamics), but may not be true for determining correct local matter properties (which is needed for e.g., electromagnetic and neutrino emission). One would also like to be able to distinguish the (fairly large) modeling error due to this artificial heating from the robust merger physics.

The surface artefact warrants closer scrutiny since the feature may cause problems for more the detailed modeling of the neutrino aspects and the associated weak interactions. In particular, the temperature is required to decide if neutrinos are trapped (a point we return to in Sec. IV D). The artificial surface feature may “confuse” any automated neutrino treatment and hence it would seem important to understand if this makes a—qualitative or quantitative—difference.

It is also worth noting that there may be a link to the initial data prescription. Our initial model assumes a uniform temperature, which would not be a realistic representation of a mature neutron star (more likely represented by a uniform redshifted temperature). The choice is due to the initial data construction within LORENE, which requires a barotropic model. The immediate alternative to our choice would be to assume a uniform entropy distribution, but this does not represent the expected physics either. Recent work [49] has considered more realistic entropy profiles, but these have not yet been used in merger simulations. We are investigating this issue in more detail but are not yet in a position to comment on it further.

Finally, it is worth commenting on the temperature of the relatively cold high-density region. As is evident from the snapshots in Fig. 2, the temperature in the core of the merger remnant is at the level of 5 MeV. This may be a reasonable reflection of the physics but the result also has to be considered with caveats. As the evolved internal energy

τ only depends weakly on the temperatures (see Appendix Sec. 2 for a discussion of the evolved variables), the numerical inversion to extract the temperature may be associated with significant uncertainties. We can approximately quantify this by a *condition number* $\mathcal{K}_{A \rightarrow B}$ which gives, to first order, the relative error in B given a relative error in A . For the temperature, we assume an error $\Delta\tau$ in the energy τ , which induces an error in T as

$$\begin{aligned} \frac{\Delta T}{T} &= \frac{1}{T} \left(\frac{\partial \tau}{\partial T} \right)^{-1} \Delta \tau \\ &= \left[\frac{\tau}{T} \left(\frac{\partial \tau}{\partial T} \right)^{-1} \right] \frac{\Delta \tau}{\tau} \\ &= \mathcal{K}_{\tau \rightarrow T} \frac{\Delta \tau}{\tau}. \end{aligned} \quad (13)$$

We then compute $\mathcal{K}_{\tau \rightarrow T}$ for a range of densities, velocities, and temperatures, leading to the results shown in Fig. 3.

We see that the condition number, indicating the growth in the inevitable numerical error, is largely insensitive to the density and velocity. It is, however, extremely sensitive to the temperature and grows rapidly as the temperature decreases. While condition numbers only give a qualitative, rather than quantitative indication, this suggests that working directly with the temperature will be prone to large numerical errors. This is not a feature of all numerical operations converting between conserved and primitive variables, as the analogous condition number for the densities $\mathcal{K}_{D \rightarrow \rho} = 1$.

If we want to do (significantly) better, we may have to work with a different set of variables (e.g., the entropy). This kind of development may be required if we want to be able to explore astrophysically motivated questions associated with the neutron star crust and high-density




FIG. 3. The condition number $\mathcal{K}_{\tau \rightarrow T} = \frac{\tau}{T} \left(\frac{\partial \tau}{\partial T} \right)^{-1}$ which illustrates how an intrinsic numerical error in the energy τ is amplified to an error in the temperature. The behavior is robust to changes in density and velocity (as encoded in the Lorentz factor W explicitly defined in Appendix Sec. 2). For the APR equation of state from [14] we see the result generically diverges as the temperature decreases, indicating serious problems with numerical accuracy at low temperatures.




FIG. 4. Degeneracy parameter T/E_{F_x} for protons p , neutrons n , and electrons e . The baryons are assumed to be nonrelativistic while the electrons are assumed to be relativistic. The data relate to the simulation of the merger of two $1.4 M_{\odot}$ neutron stars using the APR equation of state from [14]. The left column coincides with merger time, and right column is 5.0 ms postmerger. Note that $T/E_{F_x} \ll 1$ indicates strongly degenerate matter, and $T/E_{F_x} \gg 1$ indicates strongly nondegenerate matter.

condensates, which come into play at temperatures below about 1 MeV. However, as it would either involve a different equation of state parametrization, or dealing with nonconserved fluxes, it is unlikely to be a preferred immediate direction of travel.

The simulation results also allow us to comment on issues related to the finite temperature equation of state. In particular, we may consider the degeneracy of the different matter components. This is crucial, because the evaluation of the thermodynamic integrals required to build the equation of state in the first place is relatively “straightforward” for both degenerate matter (where we are essentially dealing with a low-temperature expansion) and highly relativistic matter. In the intermediate regime, the calculation is more involved (see for example [50]). Considering the results from Figs. 2 and 4 we see that much of the high density core remains strongly degenerate. However, as the temperature varies significantly across the merger remnant we always have to consider the (more challenging) cross-over regions for the baryons, especially if we are interested

in the low-density matter dynamics. It is useful to keep this in mind.



IV. THE ISSUE OF EQUILIBRIUM

As we declared from the outset, our aim is not to consider the wider aspects of the neutron star merger problem (like the detectable signatures in different channels, electromagnetic and gravitational). Rather, we want to focus on the implementation of the physics. In particular, we want to better understand which aspects of the rich physics (that we would ideally like to probe) are realistically within reach of a simulation and which ones are not? This may seem a somewhat conservative strategy given that recent simulation work considers a much wider range of issues (matter outflows and nucleosynthesis, electromagnetism and jet launch leading to observed gamma-ray bursts, neutrino transport and so on). However, as will become evident, important questions remain to be resolved already at this more “basic” level. In particular, we will



focus on the issue of local equilibrium, an issue of immediate relevance for problems involving nuclear reactions (from bulk viscosity to neutrino emission).

A. Cold chemical equilibrium





It is well known that β -equilibrium in neutron star matter is dictated by a balance of the different electron capture (ec in the following) and neutron decay (nd) reactions. Schematically, these take the form



where one of the ellipses for each reaction contains a neutrino, and they may both contain a spectator nucleon to ensure momentum balance. When the fluid is cold, it is transparent to neutrinos, and thus they may only appear as products for each reaction. In fact, because of their long mean-free path, the neutrinos depart the systems as soon as they are formed. This leaves us with a total of 6 reactions that may take place at $T \rightarrow 0$: The direct Urca processes



and the modified Urca processes (which require a bystander nucleon)



Equivalent reactions involving muons will also be involved at sufficiently high temperatures and densities, but we ignore them here for simplicity. Moreover, as the electrons are relativistic across the simulation domain (we have $\mu_e \gg T$), positron occupation will be suppressed by a factor of $\exp(-\mu_e/T)$, so we can safely ignore reactions involving them as well [10] (although positrons do play an important role in other processes e.g., setting outflow composition [51]).

To obtain a condition for equilibrium between the relevant reactions, we need a relation of the form



$$N_{i_1}P_{i_1} + N_{i_2}P_{i_2} + \dots \leftrightarrow N_{j_1}P_{j_1} + N_{j_2}P_{j_2} + \dots, \quad (22)$$

where P_k are the particle species, and N_k are stoichiometric coefficients representing the balance between products and reactants in a given reaction. The principle of detailed

balance means that the reactions will be in equilibrium when the chemical potentials obey the relation

$$N_{i_1}\mu_{i_1} + N_{i_2}\mu_{i_2} + \dots = N_{j_1}\mu_{j_1} + N_{j_2}\mu_{j_2} + \dots, \quad (23)$$

where μ_k is the chemical potential of species P_k . However, none of the 6 allowed reactions above are the inverse of any other—as they all must feature the neutrino as a product—so we cannot write the reactions in the form suggested in Eq. (22). Instead, we assume that the neutrinos are kinematically negligible (their energy will be on the order of a few times the temperature [52]) allowing us to pair the electron capture and neutron decay reactions together (effectively assuming that they proceed at the same rate, $\Gamma_{\text{nd}} = \Gamma_{\text{ec}}$) to give



These all have the equilibrium condition

$$\mu_n = \mu_p + \mu_e, \quad (27)$$

which dictates the matter composition at $T = 0$ (and holds up to temperatures of around a few hundred keV [10]).

Given that our simulation does not enforce the equilibrium (the lepton fraction is advected with the fluid), we can use the results to quantify the local deviation from the cold beta-equilibrium. To measure this deviation we introduce

$$\mu_\Delta = \mu_n - \mu_p - \mu_e. \quad (28)$$

The results, which are provided (in the equatorial plane) in Fig. 5, show that there are regions—notably associated with the merger shock and the post-merger hotspots—where $\mu_\Delta \sim 100$ MeV. That is, the matter deviates significantly from the cold equilibrium. Having said that, it is important to note that the deviation is not uniform throughout the simulation domain.

B. A step toward bulk viscosity

The results for μ_Δ from Fig. 5 represent a first step toward a discussion of the reactions that serve to reinstate equilibrium and issues relating to, for example, bulk viscosity [10,53], as they allow us to quantify the rate at which the system equilibrates and assess the relevance of bulk viscosity on the fluid dynamics. This argument is notably different from the recent discussion in [54]. Our focus here is on the parameters required to determine the relevant equilibration rates, while [54] mainly considers the nature of the fluid flow during merger. The latter argument is complicated as bulk viscosity is typically thought of as a




FIG. 5. Deviation from cold β -equilibrium $\mu_\Delta = \mu_n - (\mu_p + \mu_e)$ for a simulation of the merger of two $1.4 M_\odot$ neutron stars using the APR equation of state from [14]. The left panel coincides with merger time, and right panel is 5.0 ms post merger.

resonant phenomenon [30], which is particularly important when the dynamical motion of the fluid and the rate of relaxation toward equilibrium occur on comparable time-scales. The usual logic is perturbative and the calculation is typically done in terms of a periodic solution for the dimensionless parameter

$$\mathcal{A} = \frac{\mu_\Delta}{T}. \quad (29)$$

The calculation of the viscosity from \mathcal{A} (see Alford *et al.* [55] for a detailed discussion) can be done in three regimes: a general solution and two limiting cases. First of all, the sub-thermal limit, where $\mathcal{A} \ll 1$, allows an analytic solution, essentially an expansion in small deviations from

equilibrium (where it may be sufficient to include only the linear term). Meanwhile, the general solution and the suprathreshold limit, where $\mathcal{A} \gg 1$, are more complicated as they both require numerical integration of the periodic solution. Given this understanding, it is interesting to quantify the regions in the merger simulation where each assumption applies. As an indication of this, we plot in Fig. 6 the value of $|\mu_\Delta|/T$ [using Eq. (28)] reached by the matter in our simulation. As far as we are aware, this is the first time this kind of information has been extracted from numerical merger data. The results paint a complex picture. While there are regions in the simulation that remain subthermal (where a low-temperature expansion and a perturbative analysis of bulk viscosity would suffice) there are also large regions where the matter is in the




FIG. 6. Bulk viscosity regime parameter $|\mu_\Delta|/T$ for the simulation of the merger of two $1.4 M_\odot$ neutron stars using the APR equation of state from [14]. The left panel coincides with the merger time, and the right panel is 5.0 ms postmerger. We see that there are regions in the simulation that remain subthermal (where a low-temperature expansion and a perturbative analysis of bulk viscosity would suffice), but there are also large regions where the matter is in the suprathreshold regime.




FIG. 7. Distribution of baryon mass M_b in the $\mu_\Delta - T$ plane for a simulation of the merger of two $1.4 M_\odot$ neutron stars using the APR equation of state from [14], where $\mu_\Delta = \mu_n - \mu_p - \mu_e$. The left panel coincides with the merger time, and the right panel is 5.0 ms post merger. The dashed line denotes $T = |\mu_\Delta|$. Matter above this line therefore has $\mathcal{A} < 1$ and is in the sub-thermal regime, whereas matter below the line has $\mathcal{A} > 1$ and is in the suprathermal regime. This distinction is replicated in the histograms to the side of and below each plot, with blue indicating subthermal, and red indicating suprathermal.

suprathermal regime. The upshot is that the analysis of the impact of bulk viscosity in a live simulation would require a careful on-the-fly identification of these regions.

We can also look at this in the $\mu_\Delta - T$ plane, as we show in Fig. 7. We see that the majority of the matter reaches the point of merger around $\mu_\Delta = 0$, but in the remnant we see that the modal μ_Δ shifts to around $\mu_\Delta = 40$ MeV. There is a strong correlation between T and μ_Δ , which is perhaps unsurprising as much of the temperature increase in the simulation is driven by compression heating, and changes in density are also able to drive changes in μ_Δ . In the left panel we see that before merger most of the matter is in the subthermal regime, whereas in the right panel most of the matter has crossed the $T = |\mu_\Delta|$ line to enter the suprathermal region. While most of this matter is at high temperature, meaning that if reactions were included it would have been quickly driven to a lower μ_Δ , we see that there is some matter (corresponding to the inner core of the remnant) with $T \lesssim 5$ MeV and $\mu_\Delta \sim 40$ MeV, meaning that it may well be necessary to consider suprathermal bulk viscosity.

In principle, we could proceed and use the results to estimate the bulk viscosity. There are, however, a couple of issues that prevent us from taking this step. First, as stated in previous section, the calculation of μ_Δ assumes that the fluid is at $T \rightarrow 0$ (we are measuring the deviation from the cold equilibrium). As $T \neq 0$ is required for the calculation of \mathcal{A} , we need to consider how the problem changes for finite temperatures. The second issue we need to acknowledge relates to the fact that the standard bulk viscosity prescription involves an estimate of the local dynamical timescale. The estimates from [53] suggest that the bulk viscosity resonance may be relevant for

neutrino-transparent matter, but the outcome depends on the local conditions and the required information is not easily extracted from a live simulation, at least not in a way that does not significantly ramp up the computational cost. A more practical approach would likely involve implementing the relevant reaction network directly, but this is neither cheap nor easy.

Before we move on, it is worth commenting on the timescales of the reactive problem (as estimated in for example [53]). To outline the argument, following [42], consider a general fluid quantity X varying in time and space. Working on the fluid element scale, we can construct Eulerian $\tau_{E,X}$ and Lagrangian $\tau_{L,X}$ timescales for this quantity from

$$\tau_{E,X} = \left(\frac{\text{Var}(X)}{\text{Var}\left(\frac{\partial X}{\partial t}\right)} \right)^{1/2}, \quad (30)$$

$$\tau_{L,X} = \left(\frac{\text{Var}(X)}{\text{Var}\left(\frac{DX}{Dt}\right)} \right)^{1/2}, \quad (31)$$

where D represents the appropriate convective derivative, and $\text{Var}(X)$ denotes the statistical variation of X .

It is typical in simulations to note that reactions typically take place on the fast timescale of the weak interaction (of order $10^{-8} - 10^{-10}$ s), and hence may be neglected on the timescales that are resolved in the simulation ($\sim 10^{-7}$ s). In effect, we assume that the quantity X depends *separately* on the slow timescale t and the fast timescale $\tau_{\text{fast}} = t/(t_{\text{reaction}}/\Delta t)$ and integral-average over the fast timescale. The quantities used in the simulation are then considered to be

$$\hat{X}(t) = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T d\tau_{\text{fast}} X(t, \tau_{\text{fast}}). \quad (32)$$

When the fluctuations induced by a given reaction decay exponentially on timescales $\sim \tau_{\text{fast}}$ (which follows when their statistics approximate a linear Markov process) then it is reasonable to assume that the reaction takes place “instantaneously” and so equilibrium should be imposed on the simulated quantities. When the fluctuations instead are oscillatory there can be $\mathcal{O}(1)$ differences between X and \hat{X} . This would be the case, for example, for bulk viscosity and it would no longer be appropriate to ignore the fast timescale dynamics.

Moreover, we need to keep in mind that the reactions apply to fluid elements reacting on Lagrangian timescales, while a typical simulation is working in an Eulerian frame on Eulerian timescales. Interestingly, it is noted in [42] that, in fully resolved Newtonian simulations, these two timescales need not match. In fact, they may be rather different. In turbulent regions with high Reynolds number, the hypothesis is that $\tau_{E,X}/\tau_{L,X} \sim \text{Re}^{1/2}$ (noting that while the turbulence and reaction timescales are not necessarily similar, the point to take away is that the Eulerian and Lagrangian timescales could be quite different). This also highlights that the assumptions leading to an imposed, instantaneous equilibrium may not hold in the most interesting regions of a neutron star merger. We clearly need to approach these issues with some care.

C. Equilibrium at finite temperature

Let us return to the issue of equilibrium. When working with cold matter, the momenta of the particles taking part in the reactions are assumed to be at their respective Fermi surfaces (this is known as the Fermi surface approximation). At low densities the Fermi momenta p_F of the three involved particles (neglecting the neutrinos) satisfy the relation

$$p_{F_n} > p_{F_p} + p_{F_e}. \quad (33)$$

This blocks the direct Urca processes as there is no way to balance the momenta, requiring the introduction of the spectator nucleon in the modified Urca processes. However, as the density increases, $p_{F_p} + p_{F_e}$ increases faster than p_{F_n} , and some equations of state allow the direct processes to proceed above a threshold density at which $p_{F_n} = p_{F_p} + p_{F_e}$ [10]. For cold matter, this condition typically requires a proton fraction of order 10% (in the case of the APR model we consider here, the direct Urca threshold is reached at a around $5\rho_{\text{nuc}}$, beyond the densities reached in our simulation).

By relaxing the Fermi surface approximation at finite temperature for the direct processes, one finds that instead of being forbidden, they are Boltzmann suppressed by a

factor depending on the single particle free energy γ_i , defined by

$$\gamma_i(p) = E_i(p) - \mu_i = E_i(p) - E_{F,i}, \quad (34)$$

where p is the particle momentum. $\gamma_i(p)$ therefore represents the energy difference between a particle on the Fermi surface and a particle with momentum p . This will result in Boltzmann suppression of the rates of electron capture and neutron decay by a factor of $\exp(-|\gamma_i|/T)$ where each reaction will be dominated by a particular γ_i . For neutron decay, the dominant factor is finding a hole in the electron Fermi sea below p_{F_e} by an amount in the region of $\gamma_e = 20\text{--}25$ MeV, whereas for electron capture the energy mismatch is lessened by the antialigning of the neutrino produced with the resultant neutron, reducing the increase of momentum needed on the initial proton to around $\gamma_p = 10\text{--}15$ MeV [10]. The upshot of this is that the reactions are no longer in balance when Eq. (27) holds. In fact, as the direct electron capture rate is suppressed to a lesser extent than the neutron decay rate, there will be a net production of neutrons when the cold equilibrium condition is satisfied. To account for this, one may introduce a chemical potential offset μ_δ , such that the fluid reaches its actual equilibrium when [10]

$$\mu_n = \mu_p + \mu_e + \mu_\delta. \quad (35)$$

The offset is determined by balancing the rates of the Urca processes, again setting $\Gamma_{\text{ec}} = \Gamma_{\text{nd}}$ but now for a finite temperature. This then allows us to work out the matter composition assuming that the temperature is held fixed.

Alford *et al.* [10] (see also the recent discussion in [11]) demonstrate that, for temperatures of order 10 MeV the required offset can be as large as $\mu_\delta = 20\text{--}25$ MeV. We get an immediate idea of the impact of this, by comparing to the results for μ_Δ from Fig. 5. Clearly, there are large regions in the simulation where the suggested effect would be significant, bringing the matter closer to a de facto equilibrium than one might have expected. The effect becomes less important after the merger as the thermal hotspots drive outflows and significant regions are driven much further ($\gtrsim 100$ MeV) from the cold beta-equilibrium.

Using the results for μ_δ from Alford *et al.* [10], we may estimate the degree to which the shift of equilibrium causes the equation of state to soften in equilibrated matter. This provides a useful insight into the level at which the effect may impact on, for example, the gravitational-wave signal. In Fig. 8 we plot the adiabatic index Γ for the APR equation of state and matter in “cold” equilibrium ($\mu_\Delta = 0$) as well as matter where the suggested temperature effects are taken into account (setting $\mu_\Delta = \mu_\delta$ and using the results from Fig. 4 in [10]). We see that there is a general trend toward lower values of Γ as the temperature increases, and the softening effect due to composition changes is most




FIG. 8. Illustrating the impact of the warm equilibrium on the stiffness of the equation of state. Dashed lines assume the “cold” β -equilibrium, while solid lines use the “warm” equilibrium prescription with μ_δ take from Fig. 4 in [10]. The results show that the effect tends to soften the equation of state by up to 5% at densities up to $2\rho_{\text{nuc}}$, but also that the impact is much less pronounced at higher temperatures.

pronounced at intermediate temperatures, peaking at around the 5% level for matter in the $1 - 2\rho_{\text{nuc}}$ density range. While we cannot quantify the extent to which this degree of softening impacts on observables, like the frequency of the postmerger oscillations, it is clear that this is a question worth returning to in future work (which would require an equation of state accounting for the temperature dependent μ_δ). Finally, it is worth noting that the softening associated with the warm equilibrium is much less pronounced at a temperature of 10 MeV (and above). This is likely due to the general softening associated with the thermal pressure, an effect that is evident in the top panel of Fig. 8, dominating at higher temperatures.

D. Equilibrium in matter with trapped neutrinos

The equilibrium relations given in Eqs. (27) and (35) both assume that the fluid is transparent to neutrinos, which can then only appear as products in the allowed reactions. However, at sufficiently high temperatures the neutrino mean free path is expected to become short enough that they begin to react with the fluid matter at a meaningful rate, participating in the inverses of the six aforementioned Urca reactions. Now that the processes are allowed to proceed in both directions, it follows from Eq. (23) that all six pairs of reactions have the equilibrium condition

$$\mu_n + \mu_{\nu_e} = \mu_p + \mu_e \quad (36)$$

where $\mu_{\nu_e} = -\mu_{\bar{\nu}_e}$ is the electron-neutrino chemical potential.

Under these conditions, the relaxation timescale is expected to be (see Fig. 7 in [40]) of order $10^{-8} - 10^{-10}$ s (faster than any current numerical simulation is likely to be able to resolve). This, in turn, leads to a

weakening of the bulk viscosity by several orders of magnitude. As this has an obvious impact on the matter dynamics—the new equilibrium that the matter will evolve toward is rather different from the cold and warm cases—we clearly need to consider the neutrinos. Unfortunately, this is problematic. As the neutrinos were not included in our simulation (or indeed the equation of state in the first place!), we cannot directly measure the neutrino chemical potential. We have to resort to approximations. For example, we get some idea of the likely magnitude of the neutrino chemical potential by assuming $\mu_\nu \sim \mu_\Delta$ and then considering the results shown in Fig. 5. Still, if we want to do better then we need to account for the neutrinos in the matter description. Similarly, we cannot reliably quantify the rate at which the neutrinos are involved in reactions to determine the regions of the simulation for which Eq. (36) is the appropriate statement of β -equilibrium. However, we can estimate these regions by postprocessing (as in [8]) the opacity κ (which is related to the mean free path ℓ through $\ell = 1/\kappa$).

Given a path λ of length L , the probability of a neutrino being transmitted along that path is

$$P_T(\lambda) = \exp(-\kappa_{\text{ab}}L), \quad (37)$$

where κ_{ab} is the absorption opacity. Scattering increases the effective length of the path taken through the fluid by the neutrinos, which we can account for by using the effective opacity κ_* given by

$$\kappa_* = \sqrt{\kappa_{\text{ab}}(\kappa_{\text{ab}} + \kappa_{\text{sc}})} \quad (38)$$

where κ_{sc} is the scattering opacity [56].

Hence, we can estimate the regions where neutrino absorption will be significant by examining the magnitude




FIG. 9. Energy averaged electron neutrino and antineutrino effective opacities $\tilde{\kappa}_*$ multiplied by grid cell spacing $\Delta x = 400$ m for a simulation of the merger of two $1.4 M_\odot$ neutron stars using the APR equation of state from [14]. The left panel coincides with the merger time, and the right panel is 5.0 ms post merger. The neutrino opacities were calculated using NuLib [57].

of $\kappa_* L$. In regions where $\kappa_* L \gg 1$, the probability of transmission is very low, and thus neutrinos are available to partake in reactions, whereas in the opposite limit they are likely to escape freely. The question then is, what should we take to be the relevant length scales? Pragmatically, in merger simulations there are two pertinent length scales: the grid spacing Δx , which for the simulation under discussion is $\Delta x \approx 400$ m, and the size of the region containing hot matter r_{hot} , which here is $\sim 30\text{--}50$ km $\approx 100\Delta x$, see Fig. 1.

In regions where $\kappa_* \Delta x \gg 1$ the neutrinos are unlikely to escape the computational cell in which they were emitted, so for our purposes they are definitely trapped and available for reactions. Similarly, in regions where $\kappa_* r_{\text{hot}} \ll 1 \Rightarrow \kappa_* \Delta x \ll 1/100$ the neutrinos are likely to escape the simulation without reacting, and so these regions should evolve toward the cold/warm equilibrium conditions on some relatively fast timescale. In the intermediate regime, the neutrinos emitted in one place are not likely to be absorbed locally (i.e., within the same computational cell), but they are also unlikely to escape the simulation completely, hence it is difficult to make a definite statement one way or the other—without directly simulating the neutrinos. These arguments show why the neutrino treatment is problematic. The logic that the matter is either neutrino transparent or not is too simplistic to describe the conditions throughout much of the simulation domain. Nevertheless, as one has to start somewhere, let us consider where the estimates we have suggested take us.

The opacity of the fluid to neutrinos depends on the equation of state parameters (ρ, T, Y_e), the neutrino energy E_ν , and the neutrino species x . We use NuLib [57] to calculate $\kappa_*(\rho, T, Y_e, E_\nu, x)$ for the two most relevant species, the electron neutrino/antineutrino. In order to obtain an estimate for the appropriate neutrino energies, we assume that there are sufficient neutrinos available to be

in chemical equilibrium, and that those neutrinos will be in thermal equilibrium with the fluid. Endrizzi *et al.* [8] have shown that this assumption is good enough to determine equilibrium surfaces, which depend strongly on the absorption, and thus this approximation should work well for the qualitative analysis here. The neutrinos then follow an isotropic Fermi-Dirac distribution

$$f_\nu(E_\nu) = \left[\exp\left(\frac{E_\nu - \mu_\nu}{T}\right) + 1 \right]^{-1}, \quad (39)$$

where $\mu_{\nu_e} = -\mu_{\bar{\nu}_e} = \mu_p + \mu_e - \mu_n$. We then calculate an energy averaged opacity $\tilde{\kappa}_*$ for each species through (as we focus on particle number transport [8])

$$\tilde{\kappa}_*(\rho, T, Y_e, x) = \frac{\int_0^\infty f_\nu(E_\nu) \kappa_*(\rho, T, Y_e, E_\nu, x) E_\nu^2 dE_\nu}{\int_0^\infty f_\nu(E_\nu) E_\nu^2 dE_\nu}. \quad (40)$$

The results in Fig. 9 show that $\tilde{\kappa}_* \Delta x$ is small in the dense core regions up to merger, and although some of the core matter in the remnant has been heated enough that absorption on the scale of grid cells has become significant, there remains a transparent central region. The hotspots mentioned in Sec. III C show large potential for absorption of both species of neutrino, as expected. The outflowing matter visible in the +5 ms panel is mostly in the intermediate regime, however it does highlight the difference between the values of $\tilde{\kappa}_*$ for the two different neutrino species, with the electron neutrinos being subjected to a stronger degree of absorption than their antiparticle. Hence we are likely to see regions where the inverses of the Urca processes are suppressed to different degrees, further complicating the calculation of the equilibrium condition.




FIG. 10. Distribution of baryon mass M_b in the $\mu_n - T$ plane for a simulation of the merger of two $1.4 M_\odot$ neutron stars using the APR equation of state from [14]. The left panel coincides with the merger time, and the right panel is 5.0 ms post merger. The results suggest that thermal pions, which may play a role at temperature above $T \gtrsim 25$ MeV [58], may impact on a significant fraction of the simulated matter.

E. Equilibration via the strong interaction

So far, we considered the electron to be the only negatively charged particle (muons were ignored for simplicity, but partake in similar reactions to the electrons), meaning that equilibration relied on the weak interaction. However, the situation may change as the temperature ramps up. For example, one may argue that we should account for the presence of thermal pions [58]. A small population of pions could drastically change the story. At high temperatures neutrons may decay to protons and pions on the strong interaction timescale (which at 10^{-23} s is instantaneous for all practical purposes). The system would then reach an equilibrium where

$$\mu_n - \mu_p = \mu_\pi \quad (41)$$

with the pions subsequently equilibrating with the electrons on the weak interaction timescale (through pion decay). In the first estimates of this effect, Fore and Reddy [58] show that pion-equilibrated matter is more proton rich than the weak equilibria we have discussed. As a result, the equation of state may soften by as much as 10%-15%, which could have a significant impact on the merger dynamics. In addition, the presence of the thermal pions increases the heat capacity, leading to the matter cooling which may self-regulate the process. These are interesting ideas, but in order to test them with simulations we would need an equation of state that accounts for the thermal pions. In the absence of such a model, the best we can do is estimate if (and where) the thermal pions are likely to come into play.

Drawing on the estimates from Fore and Reddy [58], which suggest that the effect becomes significant when $T \gtrsim 25$ MeV, we may consider the results in Fig. 10, which shows the distribution of matter in the $\mu_n - T$ phase space,

as indicative. We see that after merger the outflows from the core lead to a substantial migration of matter to higher chemical potentials, and the distribution of the matter has shifted to higher temperature, with most of the matter now above $T \gtrsim 15$ MeV—the spatial regions where this become important will be discussed in Sec. V. Evidently, the temperature reaches the level where the pions may play a role throughout much of the matter. This suggests that we should seriously consider the role of the thermal pions, which inevitably involves moving toward a more complex matter model.

V. SUMMARY AND OUTLOOK

In order to better understand how the extreme physics of neutron star mergers is represented by large scale simulations, we have explored issues associated with finite temperature effects. While it is well established that the matter equation of state softens when the temperature ramps up and that the thermal effects also regulate the extent to which the matter is transparent to neutrinos (with immediate impact on the ability of the neutrinos to remove energy from the system), the problem has more subtle aspects which have not (in our view) been considered at the level of detail they warrant. These aspects need attention if we aim toward a faithful representation of the physics, especially since they regulate the rates of nuclear reactions that enter discussion of bulk viscosity and neutrino leakage. Our particular focus was on the notion of β -equilibrium in neutron star matter for the densities and temperatures reached in a neutron star merger. By post-processing the results from an out-of-equilibrium merger simulation, for equal mass neutron stars described by the APR equation of state [14], we explored how different notions of equilibrium may affect the merger dynamics, and how this raises issues

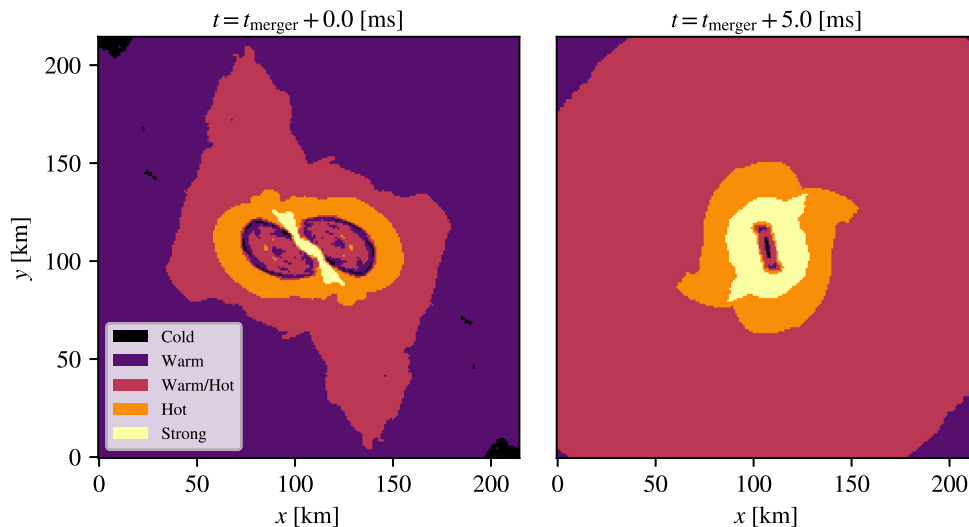


FIG. 11. Schematic of where the different equilibrium conditions discussed in Sec. IV are expected to occur in a simulation of the merger of two $1.4 M_{\odot}$ neutron stars using the APR equation of state from [14]. Left panel coincides with merger time, and right panel is 5.0 ms post-merger. The conditions on the matter are labelled as follows: “Cold” denotes matter with $T < 1$ MeV where the standard cold β equilibrium condition eq. (27) is relevant, “Warm” denotes matter with $T > 1$ MeV and $\tilde{\kappa}_* r_{\text{hot}} < 1$ for both neutrino species where the warm β equilibrium condition in Eq. (35) is relevant, “Hot” denotes matter where $\tilde{\kappa}_* \Delta x > 1$ for both neutrino species, thus the hot β equilibrium condition in Eq. (36) is relevant, “Warm/Hot” denotes matter with $T > 1$ MeV but with $\Delta x / r_{\text{hot}} < \tilde{\kappa}_* \Delta x < 1$ for at least one of the neutrino species, where we cannot distinguish cleanly between hot and warm equilibrium, and “Strong” denotes matter with $T > 25$ MeV where pions may be relevant.

when attempting to account for the relevant conditions in a live simulation. We argued that this leads to a set of problems, both computational and conceptual. In particular, our results suggest that the finite temperature effects lead to additional softening of the equation of state in some density regions, and to compositional changes that affect the processes that rely on deviation from equilibrium, such as bulk viscosity, both in terms of the magnitude and the equilibration timescales. We have demonstrated that it is far from straightforward to determine exactly which equilibrium conditions are relevant in which regions of the matter—the what and where of the problem—due to the dependence on neutrino absorption, further complicating the calculation of the reactions that work to restore the matter to equilibrium.

As an attempt to summarize the discussion we provide the schematic illustration in Fig. 11. The figure identifies regions where the different equilibrium conditions discussed in Sec. IV apply in our simulation (at merger and 5 ms later). We distinguish regions with “cold” matter ($T < 1$ MeV) where the condition (27) applies, “warm” ($T > 1$ MeV and $\tilde{\kappa}_* r_{\text{hot}} < 1$, the transparent regime, for both neutrino species) where we need to account for the imbalance of the nuclear reaction by adding μ_{δ} as in (35), a “warm/hot” region which involves matter with $T > 1$ MeV but with $\Delta x / r_{\text{hot}} < \tilde{\kappa}_* \Delta x < 1$ (neither locally trapped, nor transparent) for at least one of the neutrino species, a “hot” region where $\tilde{\kappa}_* \Delta x > 1$ (locally trapped) for both neutrino species, and finally a “strong” region with matter at $T > 25$ MeV where

one might anticipate thermal pions to play a role. As these regions evolve with time, a live identification of the appropriate local conditions is clearly far from trivial.

Admittedly, our discussion poses more problems than it solves. The main message is that we have to be careful when we consider how thermal effects enter the merger problem, especially if we aim toward a realistic implementation of the neutrinos with an on-the-fly calculation of the emission/absorption features. If the error bars associated with the inferred matter temperature are, indeed, as large as our results suggest then this issue is problematic. If we want to do better, the next steps are fairly clear. On the one hand we need to quantify the uncertainties associated with the temperature. While some of the aspects we are trying to model are likely sensitive to these errors, other are expected to be more robust. For example, the gravitational-wave signal may not be hugely affected by the issues we have discussed (although this remains to be quantified). The natural next step would be to work toward simulations that account for the different notions of equilibrium, either with a consistent model for μ_{δ} for the chosen equation of state or with the thermal pions explicitly included. Such simulations are clearly within reach of our computational technology, but require some effort on the nuclear physics side. Similarly, we can take further steps toward a quantitative understanding of the relevance of bulk viscosity on the problem. A recent discussion of this effect [54] focused on the bulk viscous pressure (associated with the divergence of the fluid flow). Such results are indicative, but it is

appropriate to raise a flag of caution. First of all, we need to keep in mind that the representation of the bulk viscosity as an effective contribution to the pressure is very much a perturbative notion. It is not at all clear that it remains appropriate for the nonlinear dynamics we are trying to model. Secondly, and perhaps more critically, the bulk viscosity arises as reactions strive to equilibrate the matter. Effectively, this leads to the different species fraction (like Y_e in our case) not being conserved. Now one has to be careful because the same effect is already accounted for in schemes aimed at modeling the neutrinos. Hence, adding a bulk viscous pressure may, in fact, lead to double counting. Clearly, these issues require further thinking.

Source files, parameter files, and plotting scripts for this work can be found at Ref. [63].

ACKNOWLEDGMENTS

Many colleagues have contributed useful discussion during this project. We would like to thank, in particular, Mark Alford, Constantinos Constantinou, Steven Harris, Madappa Prakash, Sanjay Reddy, Andreas Schmitt and Andre Schneider. We are also grateful for support from STFC via Grants No. ST/R00045X/1 and No. ST/V000551/1.

APPENDIX: Code Modifications

1. Modifications to EOS_Omni

The EOS_Omni module provides native support for a number of types of equation of state (including some tabulated), but there is no ability to use three-parameter tables in the format provided by the CompOSE database [15]. Therefore, we have written our own interface to tables in this form within the EOS_Omni framework, and added the temperature inversion step required by some conservative-to-primitive methods.

We use a set of tricubic interpolators (one for each tabulated variable, p , ϵ , and c_s^2) defined piece-wise for each cell in the table grid. In order to account for the large range of values of the table parameters ρ and T , and all of the tabulated variables we use the logarithm instead of the raw value. An issue arises with ϵ as, unlike p and c_s^2 , it is not everywhere positive. Hence, instead of using $\log(\epsilon)$ we use $\operatorname{arcsinh}(\epsilon)$, which behaves like \log for $|\epsilon| \gg 1$, but is also defined for $\epsilon \leq 0$.

The interpolators each take the form (for a general variable f to be interpolated)

$$f(\log \rho, \log T, Y_e) = \sum_{i=0}^3 \sum_{j=0}^3 \sum_{k=0}^3 a_{ijk} x^i y^j z^k, \quad (\text{A1})$$

where (x, y, z) are the shifted and normalized values of $(\log \rho, \log T, Y_e)$ with respect to the table cell that contains $(\log \rho, \log T, Y_e)$ scaled such that each cell runs from

$0 \rightarrow 1$ in all three dimensions. The polynomial coefficients a_{ijk} are calculated using the values of the variable f and all derivatives that contain at most first order terms in each of the three parameters, evaluated at each of the eight corners of the grid cell. These 64 values are arranged in a vector \mathbf{m} which we relate to a vector containing the elements a_{ijk} , \mathbf{a} , through a matrix \mathbf{B} , giving

$$\mathbf{m} = \mathbf{B}\mathbf{a}, \quad (\text{A2})$$

where we now need to find the elements of \mathbf{B} . This is done by applying the same derivative terms in \mathbf{m} to eq. (A1) and reading off the numerical coefficients that appear in front of each of the elements of \mathbf{a} for each of the elements of \mathbf{m} . The final step is to invert eq. (A2) to give

$$\mathbf{a} = \mathbf{B}^{-1}\mathbf{m}, \quad (\text{A3})$$

where the interpolation coefficients are now defined in terms of the table variable f and its derivatives with respect to the other table parameters. Constructing the table in this form also allows easy calculation of derivatives of the table variables with respect to the table values, however care must be taken to undo the transformations (taking the log or arcsinh) that were made when creating the table in the first place.

The temperature inversion step is required when the known values are (ρ, ϵ, Y_e) , and the value of T must be calculated. This is used by one of the conservative-to-primitive methods we mention below, in particular the more stable backup method. As this step is only required in the backup method, we do not need to ensure the inversion step is fast, it is significantly more important that it be robust. Hence we use an exhaustive search over all possible values of T (in this context meaning within the limits of the table).

As the values of ρ and Y_e are known, we can evaluate those two respective sums in eq. (A1) to leave a cubic polynomial in $\log T$ for each cell in the table. We then use a standard cubic polynomial solver to find which of these have solutions to $\epsilon = \epsilon(\rho, T, Y_e)$ in the correct limits by disregarding solutions for a given cell that lie outside said cell, which leaves us with the value of T that corresponds to the input ϵ given ρ and Y_e . There exists the possibility for multiple solutions given an arbitrary equation of state table, so if this does occur we choose T to minimise the change in T compared to the value on the simulation grid at the previous timestep.

2. Modifications to GRHydro

Modern GRMHD codes utilise the *Valencia formulation* [45,59] of the relativistic hydrodynamics equations for the evolution of the fluid, which take the form (neglecting the magnetic field component)

$$\partial_t(\sqrt{\gamma}\mathbf{q}) + \partial_i(\alpha\sqrt{\gamma}f^i(\mathbf{p})) = \mathbf{s}(\mathbf{p}), \quad (\text{A4})$$

where α and γ are taken from the 4-metric, \mathbf{q} denotes a vector of *conserved variables*

$$\mathbf{q} = (D, S_i, \tau, DY_e), \quad (\text{A5})$$

and \mathbf{p} denotes a vector of *primitive variables*

$$\mathbf{p} = (\rho, v^i, T, Y_e). \quad (\text{A6})$$

The conserved variables are given analytically in terms of the elements of \mathbf{p} through

$$\mathbf{q}(\mathbf{p}) = \begin{pmatrix} D \\ S_i \\ \tau \\ DY_e \end{pmatrix} = \begin{pmatrix} \rho W \\ \rho h W^2 v_i \\ \rho h W^2 - p - \rho W \\ \rho W Y_e \end{pmatrix}, \quad (\text{A7})$$

where the isotropic pressure p and specific enthalpy density $h = 1 + \epsilon + p/\rho$ are calculated from the equation of state, and the Lorentz factor $W = (1 - v^i v_i)^{-1/2}$. As the evolution of \mathbf{q} depends on \mathbf{p} we must calculate $\mathbf{q}(\mathbf{p})$ at each step. As this is not possible analytically, a number of numerical schemes have been developed (see Siegel *et al.* [60] for a recent summary).

The conservative-to-primitive inversion method used by GRHydro relies on the availability of the derivatives $\partial p/\partial \rho$ and $\partial p/\partial \epsilon$ from the EoS. However, most three-parameter EoSs are tabulated by ρ , T , and Y_e , with ϵ a function of these parameters. This means that $\partial p/\partial \epsilon$ must be calculated through

$$\begin{aligned} \frac{\partial p}{\partial \epsilon} &= \frac{\partial p}{\partial T} \frac{\partial T}{\partial \epsilon} + \frac{\partial p}{\partial \rho} \frac{\partial \rho}{\partial \epsilon} + \frac{\partial p}{\partial Y_e} \frac{\partial Y_e}{\partial \epsilon} \\ &= \frac{\partial p}{\partial T} \left(\frac{\partial \epsilon}{\partial T} \right)^{-1} + \frac{\partial p}{\partial \rho} \left(\frac{\partial \epsilon}{\partial \rho} \right)^{-1} + \frac{\partial p}{\partial Y_e} \left(\frac{\partial \epsilon}{\partial Y_e} \right)^{-1} \end{aligned} \quad (\text{A8})$$

which, when any of the $\partial \epsilon/\partial X$ terms approaches 0, can explode to infinity, leading to instability.

To improve upon this, we have implemented two methods from Siegel *et al.* [60] which are formulated in terms more suited for (ρ, T, Y_e) tables. The first is based on the 2-dimensional Newton-Raphson scheme presented by Anton *et al.* [61], and solves for W and T . The second method was suggested by Palenzuela *et al.* [62] and solves for $x = hW$ on a bracketed interval, where we use Dekker's method for the bracketed root finding. These two methods were chosen for a balance of speed and stability. The faster 2-D Newton-Raphson scheme is attempted first, and if it fails then the more stable 1-D Dekker method is used.

It is possible that the evolution of the conserved variables will cause them to no longer correspond to a physical solution within the bounds of the equation of state table. This typically occurs when the effective temperature of the fluid falls below the lowest temperature available in the

equation of state table T_{\min} . In this case, we assume $T = T_{\min}$ and solve for W . Due to this mismatch in temperatures, we know that the primitives we obtain will not correspond exactly to the conservatives with which we started, and so we must decide where we want this error to lie.

We experimented with trying to minimise the total relative error across the conservatives, but found this introduced spurious kinetic energy into the simulation. Instead, we do our best to conserve the 3-momentum density $p^i = \rho W v^i = D v^i$. As the value of W will not impact on the direction of p^i , we can instead choose to attempt to minimise the error in $p^2 = p_i p^i = D^2 v^2$, and because D does not depend on the equation of state we can safely ignore it. Therefore, we must solve to minimise error in v^2 .

The combination of conserved variables closest to v^2 is

$$\frac{S^2}{(\tau + D)^2} = \frac{z^2 v^2}{(z - p)^2} = v^2 \frac{1}{(1 - \frac{p}{z})^2}, \quad (\text{A9})$$

where $z = \rho h W^2$, which reduces to v^2 for $p \ll z$. We can also calculate this quantity from the primitive variables using

$$v^2 \frac{1}{(1 - \frac{p}{z})^2} = \frac{W^2 - 1}{W^2 (1 - \frac{p(W)}{z(W)})^2}, \quad (\text{A10})$$

where $p(W)$ and $z(W)$ are calculated from the EoS using $(\rho = D/W, T = T_{\min}, Y_e = DY_e/D)$. Combining eqs. (A9) and (A10) we define an error function f , which measures the difference between the input $S^2/(\tau + D)^2$ and the value calculated from our W guess, which takes the form

$$f(W) = \log \left(\frac{W^2 - 1}{W^2 (1 - \frac{p(W)}{z(W)})^2} \right) - \log \left(\frac{S^2}{(\tau + D)^2} \right), \quad (\text{A11})$$

where we have taken the log of each equation in order to improve behavior far from $f(W) = 0$. We also obtain better behavior where $v^2 \ll 1 \Rightarrow W \approx 1$ by using $w = \log(W - 1)$, which yields

$$\begin{aligned} f(w) &= \log \left(\frac{e^{2w} + 2e^w}{(e^{2w} + 2e^w + 1)(1 - \frac{p(w)}{z(w)})^2} \right) \\ &\quad - \log \left(\frac{S^2}{(\tau + D)^2} \right). \end{aligned} \quad (\text{A12})$$

W is physically constrained by the lower bound $1 \leq W$, which corresponds to $-\infty \leq w$. However the limits of floating point number representation limit meaningful values of $W > 1$ to

$$W = 1 + \epsilon, \quad (\text{A13})$$

where ε is the machine epsilon (in the case of double precision floating point numbers, $\varepsilon = 2^{-52} \approx 2.22 \times 10^{-16}$), thus giving a lower bound for w of

$$w_{\min} = \log \varepsilon. \quad (\text{A14})$$

We can obtain an upper bound for w by considering the fact that if ρ drops below an atmosphere density ρ_{atmo} , the simulation will consider that point to be part of the atmosphere and reset all of the fluid variables to pre-defined defaults. This means that for the point in question w must be such that

$$\frac{D}{(1 + e^w)} \geq \rho_{\text{atmo}}, \quad (\text{A15})$$

to avoid resetting, which we can rearrange to

$$w_{\max} = \log \left(\frac{D}{\rho_{\text{atmo}}} - 1 \right). \quad (\text{A16})$$

Using the two bounds for w and the function $f(w)$, we again use Dekker's method to solve for w such that $f(w) = 0$. If we find there is no solution for $f(w) = 0$ on the interval $w_{\min} \leq w \leq w_{\max}$, we use a modified objective function $g(w)$, defined by

$$g(w) = \left(\frac{e^{2w} + 2e^w}{(e^{2w} + 2e^w + 1)(1 - \frac{p(w)}{z(w)})^2} \right) - \left(\frac{S^2}{(\tau + D)^2} \right), \quad (\text{A17})$$

to determine whether there is a solution on the interval $-\infty \leq w \leq w_{\min}$ by observing that

$$\begin{aligned} \lim_{w \rightarrow -\infty} g(w) &= \left(\frac{0}{0 + 1} \frac{1}{(1 - \frac{p(w)}{z(w)})^2} \right) - \frac{S^2}{(\tau + D)^2} \\ &= -\frac{S^2}{(\tau + D)^2} \end{aligned} \quad (\text{A18})$$

If we find that $g(w_{\min}) \times g(-\infty) \leq 0$, then we accept $W = 1$ as the solution. We can also test for a solution on the interval $w_{\max} \leq w \leq \infty$ by observing that

$$\begin{aligned} \lim_{w \rightarrow \infty} g(w) &= \lim_{w \rightarrow \infty} \left(\frac{e^{2w} + 2e^w}{(e^{2w} + 2e^w + 1)} \times \left(1 - \frac{1}{W^2 \frac{\rho(1+\varepsilon)}{p} + 1} \right)^{-2} \right) \\ &\quad - \left(\frac{S^2}{(\tau + D)^2} \right), \end{aligned} \quad (\text{A19})$$

where we can immediately see that

$$\lim_{w \rightarrow \infty} \left(\frac{e^{2w} + 2e^w}{(e^{2w} + 2e^w + 1)} \right) = 1, \quad (\text{A20})$$

and we can use the dominant energy condition

$$p^2 \leq \rho^2(1 + \varepsilon)^2, \quad (\text{A21})$$

to show that

$$\frac{1}{\frac{\rho(1+\varepsilon)}{p} + 1} \leq \frac{1}{2}. \quad (\text{A22})$$

This leaves us with

$$\begin{aligned} \lim_{w \rightarrow \infty} g(w) &= \left(1 \times \left(1 - \lim_{w \rightarrow \infty} \frac{1}{2(1 + e^w)^2} \right)^{-2} \right) \\ &\quad - \left(\frac{S^2}{(\tau + D)^2} \right) \end{aligned} \quad (\text{A23})$$

$$= 1 - \left(\frac{S^2}{(\tau + D)^2} \right). \quad (\text{A24})$$

and so if $g(w_{\max}) \times g(\infty) \leq 0$ then the solution is sufficiently large that the point will be set to atmosphere, and this is accepted as a solution. If both of these tests fail, then Con2Prim fails, and the simulation is aborted.

Once we have a solution for W , we can use this and T_{\min} to calculate the values of the primitives through

$$\rho = \frac{D}{W}, \quad (\text{A25})$$

$$Y_e = \frac{DY_e}{D}, \quad (\text{A26})$$

$$p = p(\rho, T_{\min}, Y_e), \quad (\text{A27})$$

$$v^i = \frac{\gamma^{ij} S_j}{\tau + D + p}. \quad (\text{A28})$$

As we know the primitive variables obtained will be inconsistent with the conservative variables with which we started, we calculate the values of the conserved variables using Eq. (A7), and use these to replace the previous values.

- [1] B. P. Abbott *et al.*, *Astrophys. J. Lett.* **848**, L12 (2017).
- [2] B. P. Abbott *et al.*, *Phys. Rev. Lett.* **119**, 161101 (2017).
- [3] B. P. Abbott *et al.*, *Phys. Rev. X* **9**, 011001 (2019).
- [4] P. S. Cowperthwaite *et al.*, *Astrophys. J. Lett.* **848**, L17 (2017).
- [5] L. Baiotti and L. Rezzolla, *Rep. Prog. Phys.* **80**, 096901 (2017).
- [6] D. Radice, S. Bernuzzi, and A. Perego, *Annu. Rev. Nucl. Part. Sci.* **70**, 95 (2020).
- [7] A. Perego, S. Bernuzzi, and D. Radice, *Eur. Phys. J. A* **55**, 124 (2019).
- [8] A. Endrizzi, A. Perego, F. M. Fabbri, L. Branca, D. Radice, S. Bernuzzi, B. Giacomazzo, F. Pederiva, and A. Lovato, *Eur. Phys. J. A* **56**, 15 (2020).
- [9] A. Prakash, D. Radice, D. Logoteta, A. Perego, V. Nedora, I. Bombaci, R. Kashyap, S. Bernuzzi, and A. Endrizzi, *Phys. Rev. D* **104**, 083029 (2021).
- [10] M. G. Alford and S. P. Harris, *Phys. Rev. C* **98**, 065806 (2018).
- [11] M. G. Alford, A. Haber, S. P. Harris, and Z. Zhang, arXiv: 2108.03324.
- [12] N. Andersson and G. L. Comer, *Living Rev. Relativity* **24**, 3 (2021).
- [13] N. Andersson, I. Hawke, K. Dionysopoulou, and G. L. Comer, *Classical Quantum Gravity* **34** (2017).
- [14] A. S. Schneider, C. Constantinou, B. Muccioli, and M. Prakash, *Phys. Rev. C* **100**, 025803 (2019).
- [15] CompOSE Core Team, CompOSE (accessed 2021), <https://compose.obspm.fr/home>.
- [16] N. Andersson and P. Pnigouras, *Mon. Not. R. Astron. Soc.* **489**, 4043 (2019).
- [17] X. Xie, I. Hawke, A. Passamonti, and N. Andersson, *Phys. Rev. D* **102**, 044040 (2020).
- [18] R. De Pietri, A. Feo, J. A. Font, F. Löffler, M. Pasquali, and N. Stergioulas, *Phys. Rev. D* **101**, 064052 (2020).
- [19] C. Constantinou, B. Muccioli, M. Prakash, and J. M. Lattimer, *Phys. Rev. C* **92**, 025801 (2015).
- [20] L. Rezzolla and O. Zanotti, *Relativistic Hydrodynamics* (Oxford University Press, Oxford, 2013).
- [21] A. Figura, F. Li, J.-J. Lu, G. Burgio, Z.-H. Li, and H.-J. Schulze, *Phys. Rev. D* **103**, 083012 (2021).
- [22] A. Figura, J.-J. Lu, G. F. Burgio, Z.-H. Li, and H.-J. Schulze, *Phys. Rev. D* **102**, 043006 (2020).
- [23] E. R. Most, L. J. Papenfort, V. Dexheimer, M. Hanauske, H. Stoecker, and L. Rezzolla, *Eur. Phys. J. A* **56**, 59 (2020).
- [24] M. G. Alford, L. Bovard, M. Hanauske, L. Rezzolla, and K. Schwenzer, *Phys. Rev. Lett.* **120**, 041101 (2018).
- [25] T. Celora, N. Andersson, I. Hawke, and G. L. Comer,
- [26] A. Bauswein and H. T. Janka, *Phys. Rev. Lett.* **108**, 011101 (2012).
- [27] S. Bernuzzi, T. Dietrich, and A. Nagar, *Phys. Rev. Lett.* **115**, 091101 (2015).
- [28] L. Rezzolla and K. Takami, *Phys. Rev. D* **104**, 084090 (2021).
- [29] D. Radice, *Astrophys. J. Lett.* **838**, L2 (2017).
- [30] A. Schmitt and P. Shternin, Reaction rates and transport in neutron stars, in *The Physics and Astrophysics of Neutron Stars* (Springer International Publishing, Cham, 2018), Vol. 457, pp. 455–574, ISBN 978-3-319-97616-7.
- [31] J. C. McKinney, A. Tchekhovskoy, A. Sadowski, and R. Narayan, *Mon. Not. R. Astron. Soc.* **441**, 3177 (2014).
- [32] A. Sadowski, R. Narayan, J. C. McKinney, and A. Tchekhovskoy, *Mon. Not. R. Astron. Soc.* **439**, 503 (2014).
- [33] S. Fujibayashi, K. Kiuchi, N. Nishimura, Y. Sekiguchi, and M. Shibata, *Astrophys. J.* **860**, 64 (2018).
- [34] P. Anninos and P. C. Fragile, *Astrophys. J.* **900**, 71 (2020).
- [35] Y.-F. Jiang, J. M. Stone, and S. W. Davis, *Astrophys. J.* **796**, 106 (2014).
- [36] D. Radice, E. Abdikamalov, L. Rezzolla, and C. D. Ott, *J. Comput. Phys.* **242**, 648 (2013).
- [37] D. Kasen, R. C. Thomas, F. Röpke, and S. E. Woosley, *J. Phys. Conf. Ser.* **125**, 012007 (2008).
- [38] J. M. Miller, B. R. Ryan, and J. C. Dolence, *Astrophys. J. Suppl. Ser.* **241**, 30 (2019).
- [39] F. Foucart, M. D. Duez, F. Hebert, L. E. Kidder, H. P. Pfeiffer, and M. A. Scheel, arXiv:2103.16588.
- [40] M. Alford, A. Harutyunyan, and A. Sedrakian, *Phys. Rev. D* **100**, 103021 (2019).
- [41] S. K. Lele, *J. Comput. Phys.* **103**, 16 (1992).
- [42] P. K. Yeung, *Annu. Rev. Fluid Mech.* **34**, 115 (2002).
- [43] Z. Etienne *et al.*, The Einstein Toolkit (accessed 2021), <http://einsteintoolkit.org/index.html>.
- [44] P. Diener, E. Schnetter, and J. Tao, McLachlan, a public BSSN code (accessed 2021), <http://www.cct.lsu.edu/eschnett/McLachlan/>.
- [45] 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).
- [46] G. Raaijmakers, S. K. Greif, K. Hebeler, T. Hinderer, S. Nisanke, A. Schwenk, T. E. Riley, A. L. Watts, J. M. Lattimer, and W. C. G. Ho, *Astrophys. J. Lett.* **918**, L29 (2021).
- [47] E. Gourgoulhon, P. Grandclément, and J. Novak, LORENE (accessed 2021), <https://lorene.obspm.fr/>.
- [48] C. Gundlach and R. J. Leveque, *J. Fluid Mech.* **676**, 237 (2011).
- [49] G. Camelio, T. Dietrich, M. Marques, and S. Rosswog, *Phys. Rev. D* **100**, 123001 (2019).
- [50] S. Lalit, M. A. A. Mamun, C. Constantinou, and M. Prakash, *Eur. Phys. J. A* **55**, 10 (2019).
- [51] D. Radice, F. Galeazzi, J. Lippuner, L. F. Roberts, C. D. Ott, and L. Rezzolla, *Mon. Not. R. Astron. Soc.* **460**, 3255 (2016).
- [52] P. Haensel, *Astron. Astrophys.* **262**, 131 (1992), <https://ui.adsabs.harvard.edu/abs/1992A&A...262..131H>.
- [53] M. Alford, A. Harutyunyan, and A. Sedrakian, *Particles* **3**, 500 (2020).
- [54] E. R. Most, S. P. Harris, C. Plumberg, M. G. Alford, J. Noronha, J. Noronha-Hostler, F. Pretorius, H. Witek, and N. Yunes, arXiv:2107.05094.
- [55] M. G. Alford, S. Mahmoodifar, and K. Schwenzer, *J. Phys. G* **37**, 125202 (2010).
- [56] G. B. Rybicki and A. P. Lightman, *Radiative Processes in Astrophysics* (Wiley, New York, 1979), ISBN 978-0-471-82759-7.
- [57] E. O’Connor and C. Sullivan, NuLib open-source neutrino interaction library (accessed 2021), <https://github.com/evanocconnor/NuLib>.
- [58] B. Fore and S. Reddy, *Phys. Rev. C* **101**, 035809 (2020).

- [59] F. Banyuls, J. A. Font, J. M. Ibanez, J. M. Marti, and J. A. Miralles, *Astrophys. J.* **476**, 221 (1997).
- [60] D. M. Siegel, P. Mösta, D. Desai, and S. Wu, *Astrophys. J.* **859**, 71 (2018).
- [61] L. Anton, O. Zanotti, J. A. Miralles, J. M. Marti, J. M. Ibanez, J. A. Font, and J. A. Pons, *Astrophys. J.* **637**, 296 (2006).
- [62] C. Palenzuela, S. L. Liebling, D. Neilsen, L. Lehner, O. L. Caballero, E. O'Connor, and M. Anderson, *Phys. Rev. D* **92**, 044045 (2015).
- [63] P. Hammond, Source, parfile, and plotting scripts relating to “Thermal aspects of neutron star mergers”, [10.5281/zenodo.5469497](https://doi.org/10.5281/zenodo.5469497).