Monthly Notices
of the
ROYAL ASTRONOMICAL SOCIETY

MNRAS **512,** 1499–1521 (2022)
Advance Access publication 2022 March 4

https://doi.org/10.1093/mnras/stac589

A new moment-based general-relativistic neutrino-radiation transport code: Methods and first applications to neutron star mergers

David Radice [1,2,3★] Sebastiano Bernuzzi [4] Albino Perego [5,6] and Roland Haas [7]

[1] *Institute for Gravitation and the Cosmos, The Pennsylvania State University, University Park, PA 16802, USA*
[2] *Department of Physics, The Pennsylvania State University, University Park, PA 16802, USA*
[3] *Department of Astronomy and Astrophysics, The Pennsylvania State University, University Park, PA 16802, USA*
[4] *Theoretisch-Physikalisches Institut, Friedrich-SchillerUniversität Jena, D-07743 Jena, Germany*
[5] *Dipartimento di Fisica, Università di Trento, Via Sommarive 14, I-38123 Trento, Italy*
[6] *INFN-TIFPA, Trento Institute for Fundamental Physics and Applications, via Sommarive 14, I-38123 Trento, Italy*
[7] *National Center for Supercomputing Applications and Department of Physics, University of Illinois, Urbana, IL 61801, USA*

Accepted 2022 March 1. Received 2022 February 8; in original form 2021 December 1

ABSTRACT

We present a new moment-based energy-integrated neutrino transport code for neutron star merger simulations in general relativity. In the merger context, ours is the first code to include Doppler effects at all orders in v/c, retaining all non-linear neutrino–matter coupling terms. The code is validated with a stringent series of tests. We show that the inclusion of full neutrino–matter coupling terms is necessary to correctly capture the trapping of neutrinos in relativistically moving media, such as in differentially rotating merger remnants. We perform preliminary simulations proving the robustness of the scheme in simulating ab-initio mergers to black hole collapse and long-term neutron star remnants up to $\sim$70 ms. The latter is the longest dynamical space-time, 3D, general relativistic simulations with full neutrino transport to date. We compare results obtained at different resolutions and using two different closures for the moment scheme. We do not find evidences of significant out-of-thermodynamic equilibrium effects, such as bulk viscosity, on the post-merger dynamics or gravitational wave emission. Neutrino luminosities and average energies are in good agreement with theory expectations and previous simulations by other groups using similar schemes. We compare dynamical and early wind ejecta properties obtained with M1 and with our older neutrino treatment. We find that the M1 results have systematically larger proton fractions. However, the differences in the nucleosynthesis yields are modest. This work sets the basis for future detailed studies spanning a wider set of neutrino reactions, binaries, and equations of state.

Key words: neutrinos – methods: numerical – neutron star mergers.

1 INTRODUCTION

Neutrinos mediate the transport of energy and lepton number in dense and hot environments. As such, neutrinos play a crucial role in powering the explosion of massive stars as core-collapse supernovae (Lentz et al. 2015; Melson et al. 2015; O'Connor & Couch 2018b; Burrows et al. 2020; Mezzacappa et al. 2020; Bollig et al. 2021; Burrows & Vartanyan 2021), in the cooling of the protoneutron star (Roberts & Reddy 2017) and in the synthesis of heavy elements neutrino-driven winds (Arcones & Thielemann 2013). Neutrinos also determine the composition and the final r-process nucleosynthesis yields of the dynamical ejecta from neutron star (NS) mergers (Sekiguchi et al. 2015; Foucart et al. 2016a; Radice et al. 2016; Sekiguchi et al. 2016; Perego, Radice & Bernuzzi 2017b). Neutrinos directly drive winds from NS merger remnants (Dessart et al. 2009; Perego et al. 2014; Fujibayashi et al. 2017) and impact the composition of outflows driven by hydrodynamic or magnetic torques and nuclear processes (Metzger & Fernández 2014; Fujibayashi et al. 2018; Fernández et al. 2019; Miller et al. 2019a; Nedora et al. 2019; Fujibayashi et al. 2020a, b; Just et al. 2022; Li & Siegel 2021). Finally, neutrinos might participate in the launching of gamma-ray burst jets from these systems (Eichler et al. 1989; Rosswog & Ramirez-Ruiz 2002; Zalamea & Beloborodov 2011; Just et al. 2016; Perego, Yasin & Arcones 2017a).

In the context of NS mergers, the most popular approach to include neutrinos in simulations is the so-called neutrino leakage scheme. This method was first proposed by van Riper & Lattimer (1981) in the context of core-collapse supernovae, and then used to perform Newtonian simulations of NS mergers by Ruffert, Janka & Schaefer (1996) and Rosswog & Ramirez-Ruiz (2002). A general-relativistic (GR) extension of the leakage scheme was first proposed in Sekiguchi (2010) and was subsequently applied to NS mergers in Sekiguchi et al. (2011). Publicly available implementations of the relativistic leakage scheme are available in GR1D and ZELMANI (O'Connor & Ott 2010) and in the THC code (Radice et al. 2016). The latter uses a methodology first proposed by Neilsen et al. (2014) to compute the optical depth, which is able to capture the complex geometries of neutron star merger remnants (Endrizzi et al. 2020).

★ E-mail: david.radice@psu.edu

This approach has also been used by Siegel & Metzger (2017) and Murguia-Berthier et al. (2021). More sophisticated implementations include the Advanced Spectral Leakage scheme of Perego, Cabezón & Käppeli (2016), Gizzi et al. (2021), and the Improved Leakage-Equilibration-Absorption scheme of Ardevol-Pulpillo et al. (2019), Kullmann et al. (2022). Leakage schemes do not attempt to simulate the transport of neutrinos, but instead parametrize the rate of cooling of the remnant using of phenomenological formulas based on the optical depth. Specifically, they replace the emission rate of neutrinos with a scaling factor $O(e^{-\tau})$, where τ is the optical depth. As such, leakage schemes avoid stiff source terms in the hydrodynamics equations and are computationally inexpensive. Standard leakage schemes ignore the reabsorption of neutrinos, so they cannot model the deposition of heat and lepton number in the ejecta by neutrinos. Moreover, leakage schemes are not accurate over time-scales comparable with the cooling time-scale of optically thick source, that is several hundreds of milliseconds for NS merger remnants (Sekiguchi et al. 2011).

To include the effect of neutrino reabsorption, several groups have coupled leakage scheme, used to treat the optically thick regions, with schemes designed to treat the free streaming neutrinos (Perego et al. 2014; Sekiguchi et al. 2015; Radice et al. 2016; Fujibayashi et al. 2017; Radice et al. 2018b; Ardevol-Pulpillo et al. 2019). This approach is likely inspired by the isotropic diffusion source approximation developed in the context of core-collapse supernovae (Liebendoerfer, Whitehouse & Fischer 2009). The combination of leakage and transport schemes addresses some of the limitations of the formers, namely the inability to model reabsorption, while preserving the overall computational efficiency of the method, since no stiff source terms are present. However, the use of these methods is questionable when modelling optically thick sources on time-scales comparable to their cooling time-scale. This is an important limitation, since it is now well established that secular ejecta, launched on time-scales of several seconds, likely dominate the kilonova signal and the nucleosynthesis yield from mergers (Shibata & Hotokezaka 2019; Siegel 2019; Radice, Bernuzzi & Perego 2020; Nedora et al. 2021b; Shibata, Fujibayashi & Sekiguchi 2021). Moreover, most of these methods cannot model out-of-weak-equilibrium effects, which might impact the post-merger evolution and the gravitational wave (GW) signal of binary NS systems (Alford et al. 2018; Alford, Harutyunyan & Sedrakian 2020; Hammond, Hawke & Andersson 2021; Most et al. 2022).

On the opposite end of the spectrum, the most sophisticated GR radiation-(magneto)hydrodynamics simulations of NS mergers and their post-merger evolution use Monte Carlo schemes (Miller et al. 2019a; Miller, Ryan & Dolence 2019b; Foucart et al. 2020, 2021b). These schemes directly attempt to solve the 7D Boltzmann equation by sampling the distribution function of neutrinos at random points in phase space. While these methods can be very accurate, they become prohibitively expensive when optically thick media are present. This is because, in order to correctly capture the thermodynamic equilibrium of matter and radiation, Monte Carlo schemes need to resolve the mean free path of the neutrinos. To avoid this issue, the method of Foucart et al. (2021b) artificially alters emission, absorption, and scattering rates at high optical depth in a way that does not impact the energy distribution of neutrinos close to the neutrino sphere. This approach can accurately predict the neutrino distribution outside of the remnant, but it is only valid for short times compared to the diffusion time-scale. Moreover, this method does not correctly capture out-of-thermodynamic equilibrium effects for matter and neutrinos.

Other methods solving the full-Boltzmann equation of radiation transport equations in seven dimensions include the short characteristic method (Davis, Stone & Jiang 2012), the S_N schemes of Nagakura, Sumiyoshi & Yamada (2014) and Chan & Müller (2020), the FP_N approach (McClarren & Hauck 2010; Radice et al. 2013), the lattice Boltzmann method (Weih et al. 2020b), and the recently proposed method of characteristics moment closure (MOCMC) method (Ryan & Dolence 2020). All of these approaches can, in principle, model the full range of conditions and effects encountered in NS mergers. In practice, these methods are extremely computationally intensive, because high angular resolutions is required to obtain solutions that are competitive with those of moment based schemes (Richers et al. 2017). So while the continued development of such methods is important and full-Boltzmann simulations are necessary to validate NS merger models, simplified neutrino transport methods are necessary to perform systematic surveys of the binary and equation of state (EOS) parameter space.

The moment formalism casts the Boltzmann equation for classical neutrino transport in a form resembling the hydrodynamics equations (Thorne 1981; Shibata et al. 2011). The main advantage of moment-based approaches is that they reduce the 7D Boltzmann equation to a system of 3+1 equations. Unlike the hydrodynamics equations, however, the moment equations for radiative transfer cannot be closed with an EOS, because, in general, there is no frame in which radiation can be assumed to be isotropic. Consequently, although moment-based approaches can model all effects arising from the interaction between matter and radiation, their accuracy is limited by the accuracy of the adopted closures (Richers 2020). Moment-based approaches are currently becoming very popular in the context of core-collapse supernovae (Obergaulinger et al. 2014; O'Connor 2015; Kuroda, Takiwaki & Kotake 2016; Roberts et al. 2016; O'Connor & Couch 2018a; Glas et al. 2019; Rahman, Just & Janka 2019; Skinner et al. 2019; Laiu et al. 2021). Moment-based methods have been first introduced by Foucart et al. (2015, 2016a, b) in the context of NS mergers.

Here, we introduce `THC_M1`: a new moment-based radiation transport code designed to perform long-term merger and post-merger simulations of binary NS. We adopt a formalism similar to that of Foucart et al. (2016b), but with two important differences. First, we introduce a new numerical scheme able to capture the diffusion limit of radiative transfer without resorting to the use of the relativistic heat-transfer equation, which is known to be ill posed (Hiscock & Lindblom 1985; Andersson & Lopez-Monsalvo 2011). Secondly, we retain all terms appearing in the coupling of matter and radiation. To the best of our knowledge, the only other codes to include these terms are that of Anninos & Fragile (2020) and Kuroda et al. (2016), which have not been applied to NS mergers. We demonstrate that these terms are necessary to correctly capture the trapping of neutrinos in relativistically moving media. After having validated our code with a series of tests, we use it to perform inspiral, merger, and post-merger simulations of two binary NS systems, and we study the impact of neutrinos on their dynamics, GW signal, and nucleosynthesis yields.

The rest of this paper is organized as follows. We introduce the mathematical formalism for the moment-based treatment of radiation in Section 2. We give the details of our numerical implementation in Section 3. We validate our approach with a series of tests in Section 4. We present a first application to the study of the merger and post-merger evolution of binary NS systems in Sections 5 and 6. Finally, Section 7 is dedicated to discussion and conclusions. Unless otherwise specified, we use a system of units in which $G = c = 1$.

2 MATHEMATICAL FORMALISM

The M1 scheme describes the neutrino fields in term of their associated (energy integrated) stress energy tensors $T^{\alpha\beta}_{(\nu)}$, where $\nu \in \{\nu_e, \bar{\nu}_e, \nu_x\}$ and $\alpha, \beta \in \{0, 1, 2, 3\}$. Since the formalism we are going to discuss applies in the same way to all neutrino species, we will omit the $\cdot_{(\nu)}$ subscript in the following discussion.

We decompose the (neutrino) radiation stress energy tensor along and orthogonally to n^α, the future-oriented unit normal to the $t =$ const hypersurfaces, as

$$T^{\alpha\beta} = En^\alpha n^\beta + F^\alpha n^\beta + n^\alpha F^\beta + P^{\alpha\beta}, \tag{1}$$

with $F^\alpha n_\alpha = 0$ and $P^{\alpha\beta} n_\alpha = 0$. The quantities E, F^α, and $P^{\alpha\beta}$ appearing in this decomposition are the radiation energy density, the radiation flux, and the radiation pressure tensor in the Eulerian frame, respectively.

In an analogous way, we can decompose the radiation stress energy tensor using the fluid four-velocity u^α:

$$T^{\alpha\beta} = Ju^\alpha u^\beta + H^\alpha u^\beta + u^\alpha H^\beta + K^{\alpha\beta}, \tag{2}$$

with $H^\alpha u_\alpha = 0$ and $K^{\alpha\beta} u_\alpha = 0$. The new quantities J, H^α, and $K^{\alpha\beta}$ are, respectively, the radiation energy density, the radiation flux, and the radiation pressure tensor in the fluid rest frame.

Conservation of energy and angular momentum reads

$$\nabla_\beta T^{\alpha\beta} = -\nabla_\beta T^{\alpha\beta}_{\rm HD}, \tag{3}$$

where ∇ is the covariant derivative operator compatible with the space-time metric and $T^{\alpha\beta}_{\rm HD}$ is the matter stress–energy tensor. In 3+1 form equation (3) reads (Shibata et al. 2011)

$$\begin{aligned}
&\partial_t \left(\sqrt{\gamma} E\right) + \partial_i \left[\sqrt{\gamma}(\alpha F^i - \beta^i E)\right] \\
&\quad = \alpha\sqrt{\gamma}[P^{ik} K_{ik} - F^i \partial_i \log\alpha - \mathcal{S}^\mu n_\mu], \\
&\partial_t \left(\sqrt{\gamma} F_i\right) + \partial_k \left[\sqrt{\gamma}\left(\alpha P^k_{\ i} - \beta^k F_i\right)\right] \\
&\quad = \sqrt{\gamma}\left[-E\partial_i\alpha + F_k \partial_i \beta^k + \frac{\alpha}{2} P^{jk} \partial_i \gamma_{jk} + \alpha \mathcal{S}^\mu \gamma_{i\mu}\right],
\end{aligned} \tag{4}$$

where γ_{ik} is the three metric and γ is its determinant, α is the lapse function, β^i is the shift vector, and K_{ik} is the extrinsic curvature, not to be confused with the fluid frame radiation pressure tensor. $\mathcal{S}^\mu$ is the term representing the interaction between the neutrino radiation and the fluid. It can be written as

$$\mathcal{S}^\mu = (\eta - \kappa_a J) u^\mu - (\kappa_a + \kappa_s) H^\mu, \tag{5}$$

where η, κ_a, and κ_s are the neutrino emissivity, and absorption and scattering coefficients. Scattering is assumed to be isotropic and elastic. Inelastic scattering effects could, in principle, be treated within this formalism as absorption events immediately followed by emission.

It is important to remark that equation (4) are exact, but they are not closed, since P^{ik} cannot be expressed in terms of E and F^i. The key idea of the M1 scheme is to introduce an (approximate) analytic closure for these equations, that is a relation $P^{ik} = f(E, F^i)$. Clearly, if P^{ik} were known, then the M1 scheme would provide an exact solution of the transport equation. However, because P^{ik} depends on the global geometry of the radiation field, no closure in the form $P^{ik} = f(E, F^i)$ can be exact in general.

THC_M1 adopts the so-called *Minerbo closure*, which is exact in two limits: (1) the optically thick limit in which matter and radiation and in thermodynamic equilibrium and (2) the propagation of radiation from a point source (at large distances) in a transparent medium. We consider these two cases separately below.

2.1 Optically thick limit

In the optically thick limit, in which matter and radiation are in equilibrium, the radiation pressure tensor is isotropic in the fluid frame

$$K_{\alpha\beta} = \frac{1}{3} J (g_{\alpha\beta} + u_\alpha u_\beta), \tag{6}$$

where $g_{\alpha\beta}$ is the space-time metric. The stress energy tensor reads

$$T^\alpha_{\ \beta} = \frac{4}{3} J u^\alpha u_\beta + H^\alpha u_\beta + H_\beta u^\alpha + \frac{1}{3} J \delta^\alpha_{\ \beta}, \tag{7}$$

where $\delta^\alpha_{\ \beta}$ is the Kronecker delta. The radiation pressure tensor in the laboratory frame is written as

$$\begin{aligned}
P_{\alpha\beta} &= \gamma_{\alpha\gamma}\gamma_{\beta\delta} T^{\gamma\delta} = \frac{4}{3} J W^2 \upsilon_\alpha \upsilon_\beta \\
&\quad + \gamma_{\alpha\gamma} H^\gamma \upsilon_\beta W + \gamma_{\gamma\beta} H^\gamma \upsilon_\alpha W + \frac{1}{3} J \gamma_{\alpha\beta},
\end{aligned} \tag{8}$$

where $W = -u^\alpha n_\alpha$ is the fluid Lorentz factor and $\upsilon^\alpha = \frac{1}{W}\gamma^\alpha_{\ \beta} u^\beta$ is the fluid three velocity. Since M1 evolves (E, F^i), it is necessary to reformulate equation (8) in terms of these variables. To this aim, we exploit the decomposition of equation (1) to write

$$E = T_{\alpha\beta} n^\alpha n^\beta = \frac{4}{3} J W^2 - 2 H_\alpha n^\alpha W - \frac{1}{3} J, \tag{9}$$

$$F_\alpha = -\gamma_{\alpha\beta} n_\mu T^{\beta\mu} = \frac{4}{3} J W^2 \upsilon_\alpha + W H_\alpha + W H^\beta n_\beta (n_\alpha - \upsilon_\alpha). \tag{10}$$

Since H^α is orthogonal to u_α, it is possible to project equation (10) to find

$$\begin{aligned}
F_\alpha u^\alpha &= \frac{4}{3} J W (W^2 - 1) - W H^\beta n_\beta (W + W^{-1}(W^2 - 1)) \\
&= \frac{4}{3} J W (W^2 - 1) - H^\beta n_\beta (2W^2 - 1) \\
&= \left[\frac{4}{3} J W^2 - 2 H^\beta n_\beta W - \frac{1}{3} J\right] W - JW + H^\alpha n_\alpha.
\end{aligned}$$

The term in parenthesis in the last expression is the RHS of equation (9), so we conclude that

$$H^\alpha n_\alpha = F_\alpha u^\alpha - EW + JW. \tag{11}$$

Substituting this into equation (9) we find

$$\left(\frac{2}{3} W^2 + \frac{1}{3}\right) J = E(2W^2 - 1) - 2W^2 F_\alpha \upsilon^\alpha. \tag{12}$$

This equation can be used to evaluate J given the evolved fluid and radiation quantities. Determining H^α is more complex, but fortunately only its projection on the $t =$ const hypersurface is required. To find it, we use equation (10) to write

$$W H^\alpha = F^\alpha - \frac{4}{3} J W^2 \upsilon^\alpha - W H^\beta n_\beta (n^\alpha - \upsilon^\alpha) \tag{13}$$

and

$$\begin{aligned}
\gamma^\alpha_{\ \beta} H^\beta &= \frac{F^\alpha}{W} - \frac{4}{3} J W \upsilon^\alpha + \upsilon^\alpha H^\beta n_\beta \\
&= \frac{F^\alpha}{W} - \frac{4}{3} J W \upsilon^\alpha + \upsilon^\alpha \left[W F^\beta \upsilon_\beta - EW + JW\right].
\end{aligned} \tag{14}$$

We can thus evaluate the radiation pressure tensor by combining equations (8), (12), and (14).

2.2 Optically thin limit

In the optically thin limit, we assume that radiation is streaming at the speed of light in the direction of the radiation flux. This ansatz is well

verified for radiation propagating at large distances from a central source. In this case, the radiation pressure tensor can be written as

$$P_{\alpha\beta} = \frac{E}{F_\mu F^\mu} F_\alpha F_\beta. \tag{15}$$

We remark that, differently from the optically thick limit, the optically thin limit is not unique. It is instead determined by the global geometry of the radiation field. This choice of the optically thin limit is also responsible for the appearance of 'radiation shocks' in M1 calculations. These artefacts emerge when radiation beams from different directions intersect. In these cases, the M1 method will force radiation to stream in the direction of the total (weighted and averaged) radiation flux causing neutrinos to interact in an unphysical manner. To quantify the impact of such artefact, we perform calculations in which the optically thick closure is used throughout the simulation domain. This is the so-called *Eddington closure*. It is not affected by radiation shocks, since it preserves the linearity of the transport operator. However, it predicts a maximum propagation speed of neutrinos of $\frac{c}{\sqrt{3}}$ and leads to substantial artificial diffusion (radiation can diffuse past obstacles that would otherwise cause shadows to appear).

2.3 Minerbo closure

The Minerbo closure combines the optically thin and optically thick limits as

$$P_{\alpha\beta} = \frac{3\chi - 1}{2} P^{\rm thin}_{\alpha\beta} + \frac{3(1-\chi)}{2} P^{\rm thick}_{\alpha\beta}, \tag{16}$$

where $\chi \in \left[\frac{1}{3}, 1\right]$ is the so-called Eddington factor, which is taken to be

$$\chi(\xi) = \frac{1}{3} + \xi^2 \left(\frac{6 - 2\xi + 6\xi^2}{15}\right), \tag{17}$$

where

$$\xi^2 = \frac{H_\alpha H^\alpha}{J^2}. \tag{18}$$

In the optically thick regions of the flow $H_\alpha \simeq 0$ and $\chi \simeq \frac{1}{3}$, so $P_{\alpha\beta} \simeq P^{\rm thick}_{\alpha\beta}$. Conversely, in the optically thin regions $\xi \simeq 1$ and $\chi \simeq 1$, so $P_{\alpha\beta} \simeq P^{\rm thin}_{\alpha\beta}$. It is important to remark that ξ is computed using H_α and J, instead of F_α and E. This is because F_α is not guaranteed to be small in the optically thick limit if the background flow is moving. On the other hand, the knowledge of the M1 evolved quantities, E and F^α, is not immediately sufficient to calculate H_α: it is necessary to also know $P_{\alpha\beta}$. Equations (16), (17), and (18) need to be solved numerically for χ using a root finding scheme. To this purpose, we adopt the Brent–Dekker method as implemented in the GNU Scientific Library (Galassi 2009).

2.4 Neutrino number density

Weak reactions conserve the total lepton number of the system, but they can alter the electron fraction of the matter. For this reason, it is desirable to also evolve the number density of neutrinos. To this aim, we follow the phenomenological approach proposed by Foucart et al. (2016b) and, for each neutrino species, we introduce a neutrino number current $N^\alpha_{(\nu)}$, with $\nu \in \{\nu_e, \bar\nu_e, \nu_x\}$. The neutrino number density in the fluid frame is

$$n = -N^\alpha u_\alpha, \tag{19}$$

where we have suppressed once again the index $\cdot_{(\nu)}$.The continuity equation for neutrinos reads

$$\nabla_\alpha N^\alpha = \sqrt{-g}\left(\eta^0 - \kappa^0_a n\right), \tag{20}$$

where g is the determinant of the space-time metric and κ^0_a and η^0 are the neutrino number absorption and emission coefficients. Equation (20) is exact, but like the neutrino energy and momentum equations (4), it is also not closed. The closure we adopt for equation (20) is

$$N^\alpha = n f^\alpha = n\left(u^\alpha + \frac{H^\alpha}{J}\right). \tag{21}$$

Since $H^\alpha u_\alpha = 0$, this closure is consistent with equation (19). The closure assumes that the neutrino number and energy flux are aligned. While this closure would be exact if neutrinos had a single energy, it is not for the energy-integrated fluxes in general. The closure on the neutrino number flux (21) and neutrino pressure tensor (16), the simplified treatment of the energy dependence of neutrino absorption and scattering opacities (Section 3.2.3), and the fact that we neglect neutrino oscillations are the only modelling assumptions in `THC_M1`.

In 3+1 form equation (20) becomes

$$\partial_t\left(\sqrt{\gamma} n \Gamma\right) + \partial_i\left(\alpha\sqrt{\gamma} n f^i\right) = \alpha\sqrt{\gamma}\left(\eta^0 - \kappa^0_a n\right), \tag{22}$$

where

$$\Gamma = \alpha f^0 = W - \frac{1}{J} H^\alpha n_\alpha, \quad f^i = W\left(\upsilon^i - \frac{\beta^i}{\alpha}\right) + \frac{H^i}{J}. \tag{23}$$

When computing Γ, we follow Foucart et al. (2016b) and rewrite it as

$$\Gamma = -f^\alpha n_\alpha = W\left(\frac{E - F_\alpha \upsilon^\alpha}{J}\right), \tag{24}$$

where we have used the fact that

$$-H^\mu n_\mu = W(E - J - F^\alpha \upsilon_\alpha). \tag{25}$$

3 NUMERICAL IMPLEMENTATION

The M1 equations can be summarized as

$$\partial_t \boldsymbol{U} + \partial_i \boldsymbol{F}^i(\boldsymbol{U}) = \boldsymbol{G}(\boldsymbol{U}) + \boldsymbol{S}(\boldsymbol{U}), \tag{26}$$

where

$$\boldsymbol{U} = \begin{pmatrix} \sqrt{\gamma} n \Gamma \\ \sqrt{\gamma} E \\ \sqrt{\gamma} F_k \end{pmatrix}, \tag{27}$$

$$\boldsymbol{F}^i = \begin{pmatrix} \alpha\sqrt{\gamma} n f^i \\ \sqrt{\gamma}[\alpha F^i - \beta^i E] \\ \sqrt{\gamma}\left[\alpha P^i_{\ k} - \beta^i F_k\right] \end{pmatrix}, \tag{28}$$

$$\boldsymbol{S} = \begin{pmatrix} \alpha\sqrt{\gamma}[\eta^0 - \kappa^0_a n] \\ -\alpha\sqrt{\gamma}\mathcal{S}^\mu n_\mu \\ \alpha\sqrt{\gamma}\mathcal{S}^\mu \gamma_{k\mu} \end{pmatrix}, \tag{29}$$

and

$$\boldsymbol{G} = \begin{pmatrix} 0 \\ \alpha\sqrt{\gamma}[P^{ik} K_{ik} - F^i \partial_i \log\alpha] \\ \sqrt{\gamma}\left[F_i \partial_k \beta^i - E\partial_k \alpha + \frac{\alpha}{2} P^{ij} \partial_k \gamma_{ji}\right] \end{pmatrix}. \tag{30}$$

Among these terms, the coupling with matter $\boldsymbol{S}$ is stiff and cannot be treated using an explicit time integration strategy. Since $\mathcal{S}^\mu$ is a function of (E, F^i) through the (non-linear) closure of the M1 scheme, the matter coupling is not only stiff, but also non-linear. Our code is the first M1 code in GR to treat this term in full generality in the merger context. On the other hand, if the opacity coefficients are

kept fixed during the update of the radiation quantities, the number density equation formally decouples from the others, so it can be treated separately.

THC_M1 integrates equation (26) using a semi-implicit scheme. Given the solution $\boldsymbol{U}^{(k)}$ at time $t = k\Delta t$, we compute the solution at the next time-step $\boldsymbol{U}^{(k+1)}$ in two main steps:

$$\text{(i)}\ \frac{\boldsymbol{U}^{*} - \boldsymbol{U}^{(k)}}{\Delta t} = -\partial_i \boldsymbol{F}^i[\boldsymbol{U}^{(k)}] + \boldsymbol{G}[\boldsymbol{U}^{(k)}] + \boldsymbol{S}[\boldsymbol{U}^{*}],$$
$$\text{(ii)}\ \frac{\boldsymbol{U}^{(k+1)} - \boldsymbol{U}^{(k)}}{\Delta t} = -\partial_i \boldsymbol{F}^i[\boldsymbol{U}^{*}] + \boldsymbol{G}[\boldsymbol{U}^{*}] + \boldsymbol{S}[\boldsymbol{U}^{(k+1)}].$$

In particular, the advection terms and the metric sources are treated explicitly, as discussed below, while the coupling with matter is treated implicitly. Fluid quantities are kept fixed during the radiation update until the end of the second step, when matter energy and momentum densities, as well as the electron fraction, are updated according to energy, momentum, and lepton number conservation. Conservation is also enforced by limiting the changes in the radiation quantities that would correspond to negative matter energy density, or to electron fractions outside the boundaries of the EOS table (typically $0 \leq Y_e \leq 0.6$). The treatment of the advective and source terms are discussed in detail below. The derivative of the metric terms appearing in $\boldsymbol{G}$ are discretized using standard second order finite differencing.

3.1 Radiation advection

THC_M1 uses a second order flux-limited conservative finite-differencing scheme to evolve the radiation fields. In particular, numerical fluxes are computed separately for each variable and direction-by-direction. These are then combined in a directionally unsplit fashion. For simplicity, we discuss the treatment of the radiation fluxes for one of the evolved variables, say u, in the x-direction.

Let u_i be the evolved quantity at the coordinate position x_i. Then, THC_M1 approximates the derivative of the flux $f(u)$ at the location x_i as

$$\partial_x f(u) \simeq \frac{F_{i-1/2} - F_{i+1/2}}{\Delta x}, \tag{31}$$

where $F_{i-1/2}$ and $F_{i+1/2}$ are numerical fluxes defined at $x_i \mp \frac{\Delta x}{2}$, respectively. The fluxes are constructed as linear combination of a non-diffusive second order flux $F^{\rm HO}$ and a diffusive first order correction $F^{\rm LO}$:

$$F_{i+1/2} = F^{\rm HO}_{i+1/2} - A_{i+1/2}\varphi_{i+1/2}\left(F^{\rm HO}_{i+1/2} - F^{\rm LO}_{i+1/2}\right). \tag{32}$$

The term $\varphi_{i+1/2}$ is the so-called *flux limiter* (LeVeque 1992), while $A_{i+1/2}$ is a coefficient introduced to switch off the diffusive correction at high optical depth (more below). The role of the flux limiter is to introduce numerical dissipation in the presence of unresolved features in the solution u and ensure the non-linear stability of the scheme. In particular, if $A_{i+1/2}\varphi_{i+1/2} = 0$ the second-order flux is used, while if $A_{i+1/2}\varphi_{i+1/2} = 1$, then the low order flux is used. A standard second order non-diffusive flux is used for $F^{\rm HO}$, while the Lax–Friedrichs flux is used for $F^{\rm LO}$:

$$F^{\rm HO}_{i+1/2} = \frac{f(u_i) + f(u_{i+1})}{2}, \tag{33}$$

$$F^{\rm LO}_{i+1/2} = \frac{1}{2}[f(u_i) + f(u_{i+1})] - \frac{c_{i+1/2}}{2}[u_{i+1} - u_i]. \tag{34}$$

The characteristic speed in the Lax–Friedrichs flux c_i is taken to be the maximum value of the speed of light between the right and left cells

$$c_{i+1/2} = \max_{a\in\{i,i+1\}}\left\{\left|\alpha_a\sqrt{\gamma^{xx}_a} \pm \beta^x_a\right|\right\}. \tag{35}$$

We remark that it is known that the M1 system can, in some circumstances, lead to acausal (faster than light) propagation of neutrinos in GR (Shibata et al. 2011). For this reason, one might argue that a better choice of the characteristic velocity for the Lax–Friedrichs formula would have been given by the eigenvalue of the Jacobian of $\boldsymbol{F}$. These values are known analytically (Shibata et al. 2011), however in our preliminary tests we found that the use of the full eigenvalues resulted did not improve on the stability or accuracy of the M1 solver.

The flux limiter is computed using a standard minmod approach:

$$\varphi_{i+1/2} = \min\left[1, \min\left(\frac{u_i - u_{i-1}}{u_{i+1} - u_i}, \frac{u_{i+2} - u_{i+1}}{u_{i+1} - u_i}\right)\right]. \tag{36}$$

The resulting scheme is formally second-order accurate away from shocks or extrema in the solution.

The coefficient $A_{i+1/2}$ is computed as

$$A_{i+1/2} = \min\left(1, \frac{1}{\Delta x \kappa_{\rm ave}}\right), \tag{37}$$

where

$$\kappa_{\rm ave} = \frac{1}{2}[(\kappa_a)_i + (\kappa_a)_{i+1} + (\kappa_s)_i + (\kappa_s)_{i+1}]. \tag{38}$$

In particular, $A_{i+1/2} = 1$ in optically thin regions, while $A_{i+1/2} \ll 1$ at high optical depths ($\Delta x\kappa_{\rm ave}$ is the optical distance between x_i and x_{i+1}). In the optically thick limit $F_{i+1/2} \simeq F^{\rm HO}_{i+1/2}$ and the scheme reduces to a centred second order scheme, which is asymptotic preserving (Rider & Lowrie 2002). This means that THC_M1 can capture the optically thick limit without having to artificially replace the advective terms with the flux obtained from the diffusion equation, which is known to be ill posed[1] in special and general relativity (Hiscock & Lindblom 1985; Andersson & Lopez-Monsalvo 2011).

This can be shown easily for an optically thick stationary medium in flat space-time. To keep our notation simple, we also restrict ourselves to the discussion of the 1D case, however the generalization to 3D is straightforward. In this case, the radiative transfer equations reduce to

$$\partial_t E + \partial_x F^x = \kappa_a(B - E),$$
$$\partial_t F^x + \frac{1}{3}\partial_x E = -(\kappa_a + \kappa_s)F^x, \tag{39}$$

where B is the blackbody function. In the limit of $L(\kappa_a + \kappa_s) \gg 1$, where L is a characteristic length scale of the system, the radiation flux becomes

$$F_x = \frac{-1}{3(\kappa_a + \kappa_s)}\partial_x E, \tag{40}$$

and the energy equation reduces to the heat diffusion equation:

$$\partial_t E - \partial_x\left(\frac{1}{3(\kappa_a + \kappa_s)}\partial_x E\right) = \kappa_a(B - E). \tag{41}$$

A similar derivation applied to the THC_M1 discretization of equation (39), shows that the numerical discretization of the radiation energy

[1] Here we say that a mathematical problem is 'ill posed' if it is not well posed according to Hadamard. That is if it does not (1) admit a single solution that (2) depends continuously on the initial/boundary data.

flux reduces to a finite differencing scheme for the heat equation:

$$[\partial_x F^x]_i \simeq \frac{F^x_{i+1} - F^x_{i-1}}{2\Delta x} = \frac{-1}{3(\kappa_a + \kappa_s)} \left(\frac{E_{i+2} - 2E_i + E_{i-2}}{(2\Delta x)^2} \right). \tag{42}$$

In the last step, we have also assumed the absorption coefficients to be constant in space for simplicity. However, a valid scheme for the diffusion equation is also obtained for non-constant coefficients.

Although the scheme described by equation (42) is a valid discretization of the heat equation, it can suffer from an odd–even decoupling instability, as evident from the fact that the solution at x_i does not depend on the solution at x_{i-1} and x_{i+1}. To suppress this instability, we check if

$$(u_i - u_{i-1})(u_{i+1} - u_i) < 0 \text{ and } (u_i - u_{i-1})(u_{i+1} - u_i) < 0.$$

If this condition is satisfied, we set $A_{i+1/2} = 1$. We find this to be sufficient to obtain stable evolution in the scattering dominated regime.

3.2 Radiation-matter coupling

The implicit update of the neutrino number densities does not pose particular challenges and reads (in the first substep of the method):

$$N^* = \frac{N^{(k)} - \Delta t \partial_i \left[\alpha \sqrt{\gamma} n^{(k)} f^i\right] + \Delta t \left[\alpha \sqrt{\gamma} \eta^0\right]}{1 + \Delta t \alpha \kappa_a^0 \Gamma^{-1}}, \tag{43}$$

where $N = \sqrt{\gamma} n \Gamma$, and Γ is given by equation (23). n^* is obtained from N^* using the Γ recomputed with the updated neutrino fields (E, F_i). The flux terms are computed as discussed in the previous section.

The implicit part of the time update for the radiation energy quantities is significantly more complex, since it involves the solution of a 4×4 system of non-linear equations. These are in the form

$$\boldsymbol{U}^* = \boldsymbol{W} + \Delta t \boldsymbol{S}[\boldsymbol{U}^*], \tag{44}$$

where $\boldsymbol{W}$ contains the explicit terms of the scheme. For example, in the first substep of the update

$$\boldsymbol{W} = \boldsymbol{U}^{(k)} + \Delta t(-\partial_i \boldsymbol{F}^i[\boldsymbol{U}^{(k)}] + \boldsymbol{G}[\boldsymbol{U}^{(k)}]). \tag{45}$$

We employ the Powell's Hybrid method as implemented in the GNU scientific library (Galassi 2009) to solve (44). This algorithm requires the evaluation of the Jacobian of the system as well as a suitable initial guess. Both are discussed in detail below. Before diving into the details, we remark that equation (44) requires the solution of a nested non-linear equation for the closure. `THC_M1` is the first GR code to treat these terms without approximations and it is thus able to correctly captured the trapping of neutrinos in optically thick rapidly moving media. In some rare situations, the non-linear solver can fail to converge to the desired accuracy. This typically happens in the optically thick limit, since the source term become stiff only in this limit. In such cases, we linearize the equations by fixing $\chi = 1/3$. Finally, we remark that, to save computational resources, we treat the source term explicitly in the optically thin (non-stiff) limit.

3.2.1 Source Jacobian

The undensitized collisional source terms $\boldsymbol{S}[\boldsymbol{U}]$ are composed of the projections

$$-\alpha n_\alpha \mathcal{S}^\alpha = \alpha W \left[\eta + \kappa_s J - \kappa_{\rm as}(E - F_i \upsilon^i)\right], \tag{46}$$

$$+\alpha \gamma_{i\alpha} \mathcal{S}^\alpha = \alpha W(\eta - \kappa_a J)\upsilon_i - \alpha \kappa_{\rm as} H_i, \tag{47}$$

where $\kappa_{\rm as} = \kappa_a + \kappa_s$. For the computation of the Jacobian matrix $\mathcal{J}_{ab} = \partial \boldsymbol{S}_a / \partial \boldsymbol{U}_b$ $(a, b = 0,..., 3)$ the density and momentum in the laboratory frame must be expressed in terms of the Eulerian quantities:

$$J(E, F_i) = B_0 + d_{\rm thin} B_{\rm thin} + d_{\rm thick} B_{\rm thick}, \tag{48}$$

$$H_i(E, F_i) = -(a_{\upsilon\,0} + d_{\rm thin} a_{\upsilon\ \rm thin} + d_{\rm thick} a_{\upsilon\ \rm thick})\upsilon_i, \tag{49}$$

$$- d_{\rm thin} a_{f\ \rm thin} \hat{f}_i - (a_{F\,0} + d_{\rm thick} a_{F\ \rm thick}) F_i, \tag{50}$$

with $\hat{f}_i = F_i / \sqrt{F_k F^k} = F_i / F$, the definitions

$$d_{\rm thick} = \frac{3}{2}(1 - \chi), \quad d_{\rm thin} = 1 - d_{\rm thick}, \tag{51}$$

and the coefficients

$$B_0 = W^2[E - 2(\upsilon \cdot F)], \tag{52}$$

$$B_{\rm thin} = W^2 E(\upsilon \cdot \hat{f})^2, \tag{53}$$

$$B_{\rm thick} = \frac{W^2 - 1}{2W^2 + 1}[4W^2(\upsilon \cdot F) + (3 - 2W^2)E], \tag{54}$$

$$a_{\upsilon 0} = W^3[E - 2(\upsilon \cdot F)] = W B_0, \tag{55}$$

$$a_{\upsilon\ \rm thin} = W^3 E(\upsilon \cdot \hat{f})^2 = W B_{\rm thin}, \tag{56}$$

$$\begin{aligned} a_{\upsilon\ \rm thick} &= W \frac{W^2 - 1}{2W^2 + 1}[4W^2(\upsilon \cdot F) + (3 - 2W^2)E] \\ &\quad + \frac{W}{2W^2 + 1}[(2W^2 - 1)(\upsilon \cdot F) + (3 - 2W^2)E] \\ &= W B_{\rm thick} + \frac{W}{2W^2 + 1}[(2W^2 - 1)(\upsilon \cdot F) + (3 - 2W^2)E], \end{aligned} \tag{57}$$

$$a_{f\ \rm thin} = W E(\upsilon \cdot \hat{f}), \tag{58}$$

$$a_{F0} = -W, \tag{59}$$

$$a_{F\ \rm thick} = W \upsilon^2. \tag{60}$$

The contractions between the fluid's velocity and the radiation momentum are shortly indicated as, e.g. $F_i \upsilon^i = \upsilon \cdot F$. The Jacobian is then given by

$$\mathcal{J}_{00} = -\alpha W \left(\kappa_{\rm as} - \kappa_s \frac{\partial J}{\partial E}\right), \tag{61}$$

$$\mathcal{J}_{0j} = \alpha W \kappa_s \frac{\partial J}{\partial F_j} + \alpha W \kappa_{\rm as} \upsilon^j, \tag{62}$$

$$\mathcal{J}_{i0} = -\alpha \left(\kappa_{\rm as} \frac{\partial H_i}{\partial E} + W \kappa_a \frac{\partial J}{\partial E} \upsilon_i\right), \tag{63}$$

$$\mathcal{J}_{ij} = -\alpha \left(\kappa_{\rm as} \frac{\partial H_i}{\partial F_j} + W \kappa_a \upsilon_i \frac{\partial J}{\partial F_j}\right). \tag{64}$$

The necessary derivatives are

$$\frac{\partial J}{\partial E} = W^2 + d_{\rm thin}(\upsilon \cdot \hat{f})^2 W^2 + d_{\rm thick} \frac{(3 - 2W^2)(W^2 - 1)}{(1 + 2W^2)}, \tag{65}$$

$$\frac{\partial J}{\partial F_j} = J^\upsilon_F \upsilon^j + J^f_F \hat{f}^j, \tag{66}$$

$$\frac{\partial H_i}{\partial E} = H^\upsilon_E \upsilon_i + H^f_E \hat{f}_i, \tag{67}$$

$$\frac{\partial H_i}{\partial F_j} = H^\delta_F \delta^j_i + H^{\upsilon\upsilon}_F \upsilon_i \upsilon^j + H^{ff}_F \hat{f}_i \hat{f}^j + H^{\upsilon f}_F \upsilon_i \hat{f}^j + H^{f\upsilon}_F \hat{f}_i \upsilon^j, \tag{68}$$

where the factors X^z_Y in the derivatives $\partial X/\partial Y$ are the common terms multiplying the terms with indexes z^j_i. Specifically, they are

$$J^{\upsilon}_F = 2W^2\left(-1 + d_{\rm thin}\frac{E(\upsilon\cdot\hat{f})}{F} + 2d_{\rm thick}\frac{W^2-1}{1+2W^2}\right), \tag{69}$$

$$J^{f}_F = -2d_{\rm thin}\frac{W^2E(\upsilon\cdot\hat{f})^2}{F}, \tag{70}$$

$$H^{\upsilon}_E = W^3\left(-1 - d_{\rm thin}(\upsilon\cdot\hat{f})^2 + d_{\rm thick}\frac{2W^2-3}{1+2W^2}\right), \tag{71}$$

$$H^{f}_E = -d_{\rm thin}W(\upsilon\cdot\hat{f}), \tag{72}$$

$$H^{\delta}_F = W\left(1 - d_{\rm thick}\upsilon^2 - d_{\rm thin}\frac{E(\upsilon\cdot\hat{f})}{F}\right) \tag{73}$$

$$H^{\upsilon\upsilon}_F = 2W^3\left[1 - d_{\rm thin}\frac{E(\upsilon\cdot\hat{f})}{F} - d_{\rm thick}\left(\upsilon^2 + \frac{1}{2W^2(1+2W^2)}\right)\right], \tag{74}$$

$$H^{ff}_F = 2d_{\rm thin}\frac{WE(\upsilon\cdot\hat{f})}{F}, \tag{75}$$

$$H^{\upsilon f}_F = 2d_{\rm thin}\frac{W^3E(\upsilon\cdot\hat{f})^2}{F}, \tag{76}$$

$$H^{f\upsilon}_F = -d_{\rm thin}\frac{WE}{F}. \tag{77}$$

The calculation of the above terms proceed as follows. The Eulerian multipole (E, F_i) enter the term $-\alpha n_\alpha\mathcal{S}^\alpha$ both directly and via the fluid frame multipoles (J, H_i). In particular, F_i enters only via combinations $\upsilon\cdot F$ (directly and in B_0, $B_{\rm thin}$) and $(\upsilon\cdot\hat{f})^2$ (in $B_{\rm thin}$). The relevant derivatives are

$$\frac{\partial F_i}{\partial F_j} = \delta^j{}_i \tag{78}$$

$$\frac{\partial(\upsilon\cdot F)}{\partial F_j} = \upsilon^j \tag{79}$$

$$\frac{\partial\hat{f}_i}{\partial F_j} = \frac{1}{F}\delta^j_i - \frac{1}{F}\hat{f}_i\hat{f}^j \tag{80}$$

$$\frac{\partial(\upsilon\cdot\hat{f})}{\partial F_j} = \frac{1}{F}\upsilon^j - \frac{(\upsilon\cdot\hat{f})}{F}\hat{f}^j \tag{81}$$

$$\frac{\partial(\upsilon\cdot\hat{f})^2}{\partial F_j} = 2\frac{(\upsilon\cdot\hat{f})}{F}\upsilon^j - 2\frac{(\upsilon\cdot\hat{f})^2}{F}\hat{f}^j \tag{82}$$

Consequently, the derivatives $\partial J/\partial F_j$ have terms proportional to υ^j and to $\hat{f}^j$. The Eulerian multipoles (E, F_i) do not enter directly the terms $\alpha\gamma_{i\alpha}\mathcal{S}^\alpha$. The dependence on F_i of H_i is either in terms proportional to $(\upsilon\cdot\hat{f})^2$ (in $a_{\upsilon\ \rm thin}$), $\upsilon\cdot\hat{f}$ (in $a_{f\ \rm thin}$), $\upsilon\cdot F$ (in $a_{\upsilon\ \rm thick}$), or in the direct terms explicitly indicated in equation (15). Consequently, the derivatives $\partial H_i/\partial F_j$ have terms proportional to $\delta^j{}_i$, to $\upsilon_i\upsilon^j$ and to F_iF^j.

A particular cases of the above calculation is the linearization around the zero state $\boldsymbol{U}_0 = 0$ and the zero fluid's velocity limit $\upsilon_i = 0$. For the former case, the undensitized collisional term is

$$\boldsymbol{S}(0) = [\alpha\eta W,\ \alpha\eta W\upsilon_i], \tag{83}$$

and the Jacobian matrix simplifies: since $\hat{f}_i = 0$, the first column and first row are proportional to υ_i and υ^j, respectively, while the spatial block has a term proportional to δ_{ij} and a term proportional to $\upsilon_i\upsilon^j$. A simple analytical inversion can be calculated with any computer algebra software. For a static fluid $\upsilon^i = 0$ ($E = J$ and $F_i = H_i$), one obtains

$$\boldsymbol{S}(\boldsymbol{U}_0) = [\alpha\eta - \kappa_a E,\ -\alpha\kappa_{\rm as}F_i], \tag{84}$$

and the Jacobian matrix is diagonal

$$\mathcal{J}_{00} = -\alpha\kappa_a, \tag{85}$$

$$\mathcal{J}_{ij} = -\alpha\kappa_{\rm as}\delta_{ij}. \tag{86}$$

The THC_M1 implementation is different from most of the other M1 schemes in general relativity. In particular, THC_M1 and Anninos & Fragile (2020) are the only codes fully treating the non-linear terms in the radiation matter coupling. In Roberts et al. (2016) the linearization is performed about the zero state and only retains some of the $(\upsilon\cdot\hat{f})$ terms in the Jacobian matrix. In Foucart et al. (2015) and Foucart et al. (2016b), the linearization is also performed about the zero state and the angle between the velocity and the neutrino flux is kept fixed, i.e. $(\upsilon\cdot F) = \text{const}$ and $(\upsilon\cdot\hat{f}) = \text{const}$. In Weih, Olivares & Rezzolla (2020a), the linearization is also performed about the zero state and P^{ij} is assumed to be independent from $\boldsymbol{U}$. Hence the projections of P^{ij} appear explicitly in the $\boldsymbol{U}$ terms, but the P^{ij} (closure) is not included in the Jacobian matrix.

3.2.2 Blackbody function

Emissivity, absorption, and scattering coefficients are kept fixed throughout the implicit time integration. This can cause the numerical scheme to oscillate if matter is thrown out of equilibrium over a small time-scale compared with Δt. To avoid this problem, first we compute the blackbody function for neutrinos in two ways.

(i) When the radiation–matter equilibration time $\tau = (c\sqrt{\kappa_a(\kappa_a+\kappa_s)})^{-1}$ is larger than Δt, then we set

$$B_\nu = \frac{4\pi}{(ch)^3}(k_BT)^4F_3(\eta_\nu), \tag{87}$$

where F_3 is the Fermi function of the order of 3

$$F_k(\eta) = \int_0^\infty \frac{x^k}{e^{x-\eta}+1}{\rm d}x \tag{88}$$

and $\eta_\nu = \mu_\nu/(k_BT)$ is the degeneracy parameter of the neutrinos. The equilibrium number density of neutrinos is computed as

$$\mathcal{B}_\nu = \frac{4\pi}{(ch)^3}(k_BT)^3F_2(\eta_\nu). \tag{89}$$

The temperature T is taken to be the fluid temperature, while the neutrino chemical potential are evaluated at equilibrium using the EOS at the fluid density, temperature, and electron fraction Y_e, separately for each neutrino flavour. In particular, $\mu_{\nu_e} = \mu_e + \mu_p - \mu_n$, $\mu_{\bar{\nu}_e} = -\mu_{\nu_e}$, and $\mu_{\nu_x} = 0$.

(ii) If τ is smaller than $\Delta t/2$, then the blackbody function is computed again using (87), but now T and Y_e are taken to be the equilibrium temperature and electron fraction that matter would achieve under the assumption of weak equilibrium with neutrinos, and lepton number and energy conservation (Perego, Bernuzzi & Radice 2019). In particular, we solve the following equations

$$Y_l = Y_{e,\rm eq} + Y_{\nu_e}(Y_{e,\rm eq},T_{\rm eq}) - Y_{\bar{\nu}_e}(Y_{e,\rm eq},T_{\rm eq}), \tag{90}$$

$$u = e(Y_{e,\rm eq},T_{\rm eq}) + \frac{\rho}{m_{\rm b}}\Big[Z_{\nu_e}(Y_{e,\rm eq},T_{\rm eq}) + Z_{\bar{\nu}_e}(Y_{e,\rm eq},T_{\rm eq}) + 4Z_{\nu_x}(T_{\rm eq})\Big], \tag{91}$$

where Y_l is the total lepton fraction, inferred from both fluid and radiation quantities, u is the total (matter and neutrino-radiation) energy density, and $Z_\times$ denotes the energy fraction of the species

×. These equations are solved for $T_{\rm eq}$ and $Y_{e,\,\rm eq}$ under the assumption of weak equilibrium, that is $T_{\rm matter} = T_{\nu_x} = T_{\rm eq}$, $Y_e = Y_{e,\,\rm eq}$, $\mu_{\nu_e} = -\mu_{\bar\nu_e}$, $\mu_{\nu_x} = 0$, and $\mu_{\nu_e} = \mu_e + \mu_p - \mu_n$. The rationale for this choice is that it captures the correct equilibrium distribution for neutrinos, while the blackbody function of point (i) is valid for a mixture of matter and radiation in a thermal and lepton bath, or for short times compared with the equilibration time.

(iii) For intermediate values of τ we linearly interpolate between the prescriptions from points (i) and (ii).

Given the blackbody functions, we compute the ν_e and $\bar\nu_e$ emission coefficients and the ν_x absorption coefficients using Kirchhoff's law. That is, we set

$$\eta_{\nu_e} = \kappa_{a,\nu_e} B_{\nu_e}, \qquad \eta_{\bar\nu_e} = \kappa_{a,\bar\nu_e} B_{\bar\nu_e}, \qquad \kappa_{a,\nu_x} = \frac{B_{\nu_x}}{\eta_{\nu_x}}. \tag{92}$$

We apply the same treatment to the neutrino number emissivities and opacities, but using $\mathcal{B}$ instead of B.

3.2.3 *Opacity correction*

Following Foucart et al. (2016b), we correct absorption and scattering opacities by a factor

$$\left(\frac{\varepsilon_\nu}{\varepsilon_{\nu,\rm eq}}\right)^2,$$

where ε_ν is the average incoming neutrino energy and $\varepsilon_{\nu,\rm eq}$ is the average neutrino energy at the thermodynamic equilibrium (computed as in the previous section). This correction is applied prior to the imposition of Kirchhoff's law, to ensure the preservation of the correct equilibrium.

3.2.4 *Initial guess*

In order to initialize the implicit solver for equation (44) we proceed as follows.

(i) We update the radiation fields according to the non-stiff part of the equations. For the first substep this update reads:

$$\tilde{\boldsymbol{U}} = \boldsymbol{U}^{(k)} + \Delta t(\boldsymbol{G}[\boldsymbol{U}^{(k)}] - \partial_i \boldsymbol{F}^i[\boldsymbol{U}^{(k)}]). \tag{93}$$

A similar formula is used for the second substep, but using $\boldsymbol{U}^*$ to evaluate the terms in the parenthesis.

(ii) The M1 closure is updated and quantities are transformed to the fluid frame to obtain $\tilde J$ and $\widetilde{H_i}$.

(iii) We compute new values $\hat J$ and $\widehat{H_i}$ in the fluid rest-frame according to

$$\hat J = \tilde J + \frac{\Delta t}{W}(\eta - \kappa_a \hat J), \tag{94}$$

$$\widehat{H_i} = \widetilde{H_i} - \frac{\Delta t}{W}(\kappa_a + \kappa_s)\widehat{H_i}. \tag{95}$$

$\widehat{H_0}$ is obtained from the requirement that $\hat H_\alpha u^\alpha = 0$.

(iv) Finally, the initial guess for equation (44) is obtained by transforming the radiation quantities to the laboratory frame. For this transformation we take $\chi = 1/3$, since the initial guess becomes important only in the optically thick limit.

It is important to remark that $\hat J$ and $\hat H_\alpha$ are exact solution only at leading order in v/c, when $u^\alpha \partial_\alpha \simeq W\partial_t$. It would be incorrect to take the obtained $\hat E$ and $\hat F_i$ as the updated solution, even if we were to update the closure before boosting back the solution to the laboratory frame. However, `THC_M1` only uses $\hat E$ and $\hat F_i$ as initial guesses for the full non-linear solver. An exception, is the test in Section 4.3, where we show that using the $\hat E$ and $\hat F_i$ as the new states for the radiation fields, instead of performing a non-linear solve, result in large errors in the case of moving media.

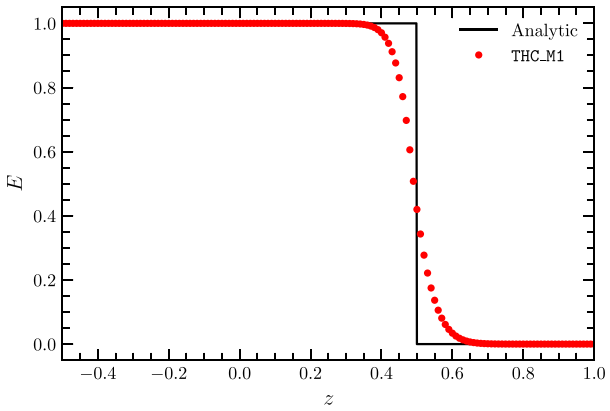

Figure 1. Optically thin advection of radiation through a large velocity discontinuity. The frame in which we compute the closure has a velocity 0.87 c for $z < 0$ and a velocity of −0.87 c for $z > 0$ (the relative Lorentz factor between left and right state is 7). No artefact appears as `THC_M1` advects a pulse of radiation through the interface at $z = 0$.

4 TEST PROBLEMS

We validate `THC_M1` by performing a series of demanding tests meant to independently verify different components of the code. This section describes the most representative tests we have performed. Most of the tests discussed here are fairly standard and have been considered by a number of authors, although with some differences in the setup (e.g. Audit et al. 2002; Vaytet et al. 2011; Radice et al. 2013; McKinney et al. 2014; Foucart et al. 2015; Skinner et al. 2019; Anninos & Fragile 2020; Weih et al. 2020a)

4.1 Optically thin advection through a velocity jump

As a first test we consider the propagation of beam of radiation in an optically thin medium. We assume slab geometry and consider initial data with $E(t = 0, z) = H(z + \frac{1}{2})$ (arbitrary units), where H is the Heaviside function, and $F^z = E$. The background fluid velocity is chosen to be:

$$v^z(z) = \begin{cases} 0.87, & z < 0, \\ -0.87, & z > 0. \end{cases}$$

That is, the medium is moving with Lorentz factor $W = 2$ in the grid frame and the two parts of the domain with $z < 0$ and $z > 0$ have a relative Lorentz factor of 7. The fluid is taken to be transparent. We set $\Delta z = 0.01$. The time-step is chosen so as to have a CFL of 0.5. It is important to emphasize that, although matter and radiation do not interact in this test, because our closure is defined in the fluid frame (equation 18), the equations become stiff in the limit in which $W \gg 1$, so this is actually a rather demanding test.

Fig. 1 shows the radiation energy density profile at time $t = 1$, after the beam has propagated through the velocity jump at $z = 0$. As it can be seen from the figure, `THC_M1` transports the radiation front through the shock without creating artificial oscillations. The discontinuity is spread over many grid cells, since `THC_M1` uses a rather dissipative central scheme to handle the transport operator in the M1 formalism (Section 3.1). However, since neutrino sources do

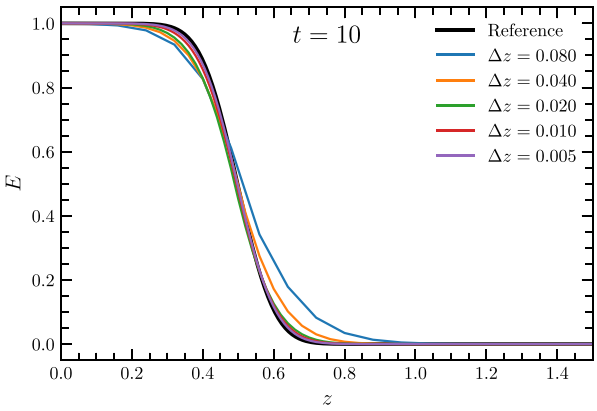

Figure 2. Diffusion of radiation in a purely scattering medium. Initially the radiation has a square profile. The reference profile is a semi-analytic solution of the diffusion equation. THC_M1 correctly capture the diffusion limit. We measure first order convergence for this test problem.

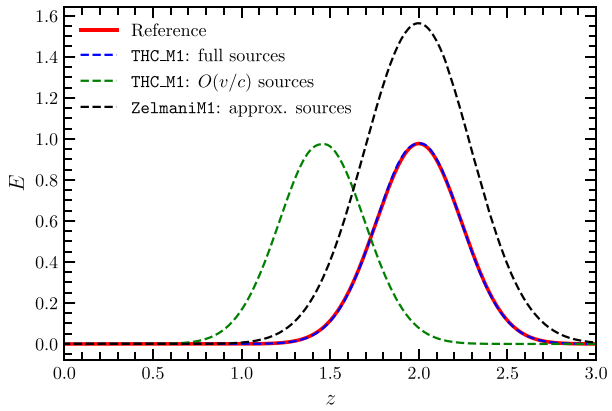

Figure 3. Diffusion and advection of Gaussian pulse of radiation in a purely scattering moving medium. The medium is moving with velocity $\upsilon = 0.5$. The reference profile is a translated semi-analytic solution of the diffusion equation. Our results show that it is essential to properly treat all of the source terms in the M1 equations to correctly capture the advection of trapped radiation.

not switch on abruptly, a sharp preservation of the radiation front is not a critical modelling requirement for our applications. Being able to handle transport through fast moving media is, instead, critical for NS merger applications, since the outflows produced in these events can be mildly relativistic ($W \lesssim 2$) (Hotokezaka et al. 2018; Nedora et al. 2021a). This test demonstrates that THC_M1 meets this requirement.

4.2 Diffusion limit

Another requirement for the modelling of NS mergers is to correctly handle the diffusion of neutrinos from the central remnant on secular time-scales. As discussed in Section 3.1, THC_M1 uses a numerical scheme designed to correctly capture the scattering dominated limit.

To validate it, we consider a purely scattering medium of constant density $\rho = 1$ (arbitrary units) and with scattering opacity $\kappa_s = 10^3$. As in the previous test, we assume slab geometry, so this is effectively a 1D problem. Initially, radiation is concentrated in the region $[-0.5, 0.5]$ and is spatially homogeneous and isotropic in this region. That is, $E(t = 0, z) = H(z + 0.5) - H(z - 0.5)$ and $F^i = 0$. In these conditions, when considering time-scales longer than the equilibration time, the radiative transfer equation can be well approximated by the diffusion equation:

$$\partial_t E = \frac{1}{3\kappa_s}\partial_x^2 E. \tag{96}$$

THC_M1 solves the equations in hyperbolic form (4). Typical hyperbolic solvers have numerical diffusion with an effective diffusion coefficient $\nu_{\rm num} \sim (\Delta z)^{-1}$. In essence, this means that standard numerical schemes fail to predict the correct diffusion of radiation in a scattering dominated region, unless the mean free path of the neutrinos (or photons) is well resolved on the grid (Rider & Lowrie 2002; McClarren & Lowrie 2008). Given that the mean free path of neutrinos at the centre of an NS merger remnant is of the order of a few meters or less, the resolution requirements for merger simulations would be extremely demanding. To work around this issue, THC_M1 uses a special numerical scheme for which $\nu_{\rm num} \to 0$ when $\kappa_s(\Delta z) \gtrsim 1$ (see Section 3.1). In this respect, our approach is different from that of Foucart et al. (2015), which instead solve the heat diffusion equation in the scattering regime.

Fig. 2 shows the radiation energy density profile at time $t = 10$ at different resolution. The CFL is set to 0.625 in all calculations. The reference solution is a semi-analytic solution of equation (96). We find that THC_M1 captures the correct diffusion rate for radiation even when $\kappa_s(\Delta z) \gg 1$. The numerical solutions are non-oscillatory, even though the initial radiation profile is discontinuous and slope limiting is essentially disabled in the scattering dominated regime. We measure first order convergence in this test, which is the expected order of convergence given that the initial data are discontinuous.

4.3 Diffusion limit in a moving medium

Matter in NS mergers is not only optically thick, but also moving at mildly relativistic velocities. Correctly capturing the advection of trapped radiation in moving media is one of the key challenges in radiation hydrodynamics and is of crucial importance for both mergers and core-collapse supernovae (Nagakura et al. 2014; Chan & Müller 2020). This requires a careful treatment of the radiation matter coupling in the stiff limit.

To demonstrate that our code can handle this, we consider a constant density, purely scattering medium with $\rho = 1$ and $\kappa_s = 10^3$, which we take to be moving towards the right with velocity $\upsilon^z = 0.5$. Once again, we assume slab geometry. We setup a Gaussian pulse of radiation centred around the origin:

$$E(t = 0, z) = e^{-9z^2}. \tag{97}$$

To initialize the radiation flux, we use equations (7), (9), and (10) under the assumption of fully trapped radiation ($H^\alpha = 0$) to write

$$J = \frac{3E}{4W^2 - 1}, \qquad F_i = \frac{4}{3} J\, W^2\, \upsilon_\alpha. \tag{98}$$

The exact solution corresponds to a slowly diffusing and translating pulse of radiation. The baseline grid spacing adopted for this problem is $\Delta z = 0.01$ and the CFL is fixed to 0.625.

Fig. 3 shows the results obtained using different schemes. The reference profile is a semi-analytic solution of the diffusion equation advected along the background fluid velocity. We find that THC_M1 reproduce the correct solution when all the non-linear terms in the sources are consistently treated. This ensures that neutrinos will not be 'left behind' as the NS merger remnant, typically deformed into a bar (Shibata 2005), rotates.

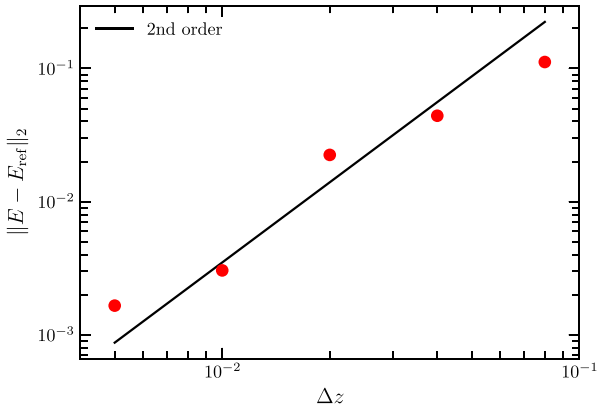

Figure 4. Convergence of the THC_M1 to the reference solution for the diffusion test in a moving medium. We find an approximate second order convergence.

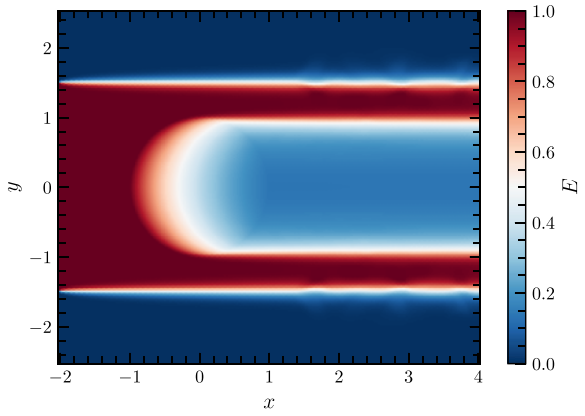

Figure 5. Shadow cast by an absorbing cylinder illuminated by a beam of radiation propagating from the left to the right. THC_M1 correctly captures the formation of the shadow.

We remark that the solution of the full non-linear source term is computationally expensive, however cheaper alternatives fail to capture the correct behaviour of the trapped radiation. A first order in v/c approach obtained by Lorentz transforming the radiation moments to and from the reference frame as discussed in Section 3.2.4 produces a stable evolution, but predicts the wrong advection speed for the radiation energy (Fig. 3). Even worse, this approach predicts different advection speeds for the neutrino number densities (not shown) and the radiation energy density which produce large errors in the average neutrino energies.

The treatment of the optically thick limit of ZelmaniM1 (Roberts et al. 2016), which is similar to the approach used in SpEC (Foucart et al. 2015), is also problematic and affected by two important issues. First, the diffusive fluxes corrected using the (acausal) heat diffusion equation significantly overestimate the rate of diffusion for the radiation, resulting in a significant broadening of the radiation pulse. Minor improvements in the diffusion rate can be obtained by implementing a better treatment using modified HLLE fluxes following Skinner et al. (2019). Secondly, because of the approximation in the source terms, the ZelmaniM1 solution violates energy conservation and the radiation energy density increases with time (Fig. 3). The violation of energy conservation is exacerbated in this problem, because there is no back reaction of the radiation on to the matter. In a more realistic setting, ZelmaniM1 would enforce energy conservation, so the increase in the radiation energy density would come at the expense of the fluid kinetic energy. That is matter would experience an unphysical drag force driving it to rest in the grid frame.

We perform additional simulations with $\Delta z = 0.16, 0.08, 0.04$, and 0.02 in addition to $\Delta z = 0.01$. The L^2 norm of the difference between the THC_M1 solution with the complete treatment of the radiation-matter source terms and the semi-analytic solution is presented in Fig. 4. Overall, we find second-order convergence for THC_M1 in this test.

4.4 Shadow test

As a first multidimensional test, we consider the problem of a beam of radiation interacting with a semi-transparent cylinder with radius $\varpi = 1$ centred at the origin. The absorption opacity in the cylinder is set to $\kappa_a = 1$ and the density to 1. Absorption is zero elsewhere. The scattering opacity κ_s is set to zero. We initialize the radiation fields to zero and inject a beam of radiation from the left of the domain with $F^x = E = 1$. The grid spacing used in this test is $\Delta x = \Delta y = 0.0125$ and the CFL is set to 0.4.

Fig. 5 shows the radiation energy density at time $t = 10$, when the solution has achieved steady state. We observe some lateral diffusion of radiation and the formation of small unsteady oscillation in the radiation field in the wake of the cylinder. The latter are artefacts caused by the non-linear nature of the Minerbo closure. Nevertheless, THC_M1 correctly captures the overall solution. The attenuation of radiation in the cylinder and the formation of a shadow behind it agree with the analytic solution for this problem.

4.5 Homogeneous sphere

The homogeneous sphere test has been considered by many authors, since it reproduces the typical geometry encountered in astrophysical applications. In this test an homogeneous sphere, which we take to be of radius $r = 1$, emits and absorbs radiation at a constant rate $\eta = \kappa_a = 1$. Scattering is neglected in this problem, so it is possible to compute an exact solution of the radiative transfer equations by numerical quadrature. This is an extremely idealized model of a hot protoneutron star or a neutron star merger remnant emitting neutrinos. We perform this test in full 3D and in Cartesian coordinates. The resolution adopted for this test is $\Delta x = \Delta y = \Delta z = 0.0125$. This corresponds to about 80 points along the radius of the 'star', a typical resolution for production neutron star merger simulations. To reduce the computational costs, we impose reflection symmetry across the xy, xz, and yz planes and only simulate the part of the domain with $x, y, z \geq 0$. The CFL is set to 0.3.

Fig. 6 shows the radiation energy density as a function of radius in the diagonal direction at time $t = 10$, when the solution has reached steady state. THC_M1 does not solve the full radiative transfer equations, so the numerical solution is not expected to converge to the exact solution. Nevertheless, the THC_M1 solution shows excellent agreement with the analytic solution and even compares favourably with the full-Boltzmann FP_N solution presented in Radice et al. (2013) for modest angular resolutions.

4.6 Gravitational light bending

Finally, we present a test validating the implementation of space-time curvature source terms in THC_M1. We study the propagation

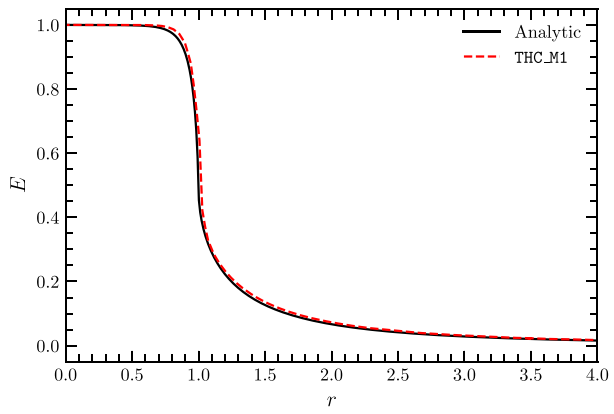

Figure 6. Radiation from an homogeneous absorbing and emitting sphere of radius one. The reference solution is obtained by solving the radiative transfer equation. The THC_M1 solution agrees well with the analytical solution. Small deviations are expected, since the adopted closure is not expected to reproduce the solution to the angle-dependent radiative transfer equations.

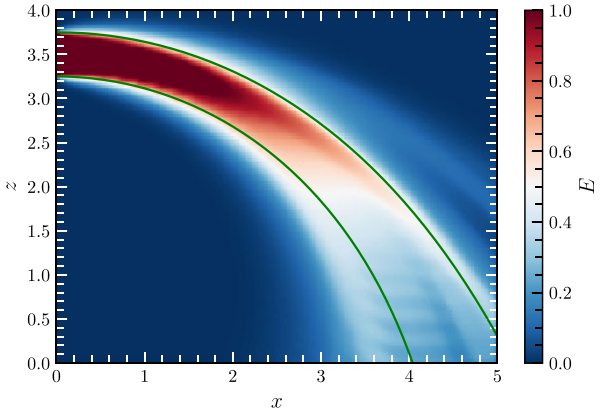

Figure 7. Beam of radiation propagating in the meridional plane of a non-rotating black hole (BH) in Kerr–Schild coordinates. The mass of the BH is set to be one. The green lines show null geodesics. THC_M1 correctly captures the bending of the beam of radiation in the strong gravitational field of the BH.

of a beam of radiation in a black hole (BH) space-time described by the Kerr–Schild metric. The BH spin is set to zero and its mass to one (in geometrized units). The computational grid is centred at the location of the BH. We only consider the region near the meridional plane $y = 0$ and $x, z > 0$. We simulate a beam of radiation injected at the location $x = 0$ and $z = 3.5$ propagating towards the positive x-direction (see Fig. 7). In particular, we set $E = 1$ at the beam injection location. The fluxes F_i are set so that $\alpha F^i - \beta^i E$ is along the x-axis and $F_i F^i = 0.99E^2$. The resolution used for this test is $\Delta x = \Delta z = 0.025$ and the CFL is set to 0.2.

Fig. 7 shows the THC_M1 solution at time $t = 20$, after steady state has been achieved. We also plot two analytically predicted trajectories for photons (null geodesics) in the same metric. THC_M1 correctly captures the bending of radiation due to the BH, indicating that curvature terms have been implemented correctly. The THC_M1 solution shows lateral diffusion of radiation comparable to other GR M1 codes (McKinney et al. 2014; Foucart et al. 2015; Weih et al. 2020a). This later diffusion is a numerical artefact. However, we do not consider this to be as a significant issue for our approach, because isolated beams of radiation are never found in the astrophysical systems we intend to model.

5 NEUTRON STAR MERGERS

As a first application of THC_M1, we consider the late inspiral and merger of a binary of two 1.364 $M_\odot$ NSs. We adopt the SRO(SLy4) EOS (SLy for brevity in the rest of the text; Douchin & Haensel 2001; Schneider, Roberts & Ott 2017). To ease the comparison with previous results, we use the same set of reactions and opacities as in Radice et al. (2018b). We construct initial data with an initial separation of 45 km using the Lorene pseudo-spectral code (Gourgoulhon et al. 2001). We have already considered this initial data in Endrizzi et al. (2020) and Nedora et al. (2021b), to which we refer for more details. The evolution grid employs seven levels of adaptive mesh-refinement (AMR), with the finest grid having finest grid spacing of $h = 0.25\,GM_\odot/c^2$, $0.167\,GM_\odot/c^2$ and $0.125\,GM_\odot/c^2$, respectively denoted as VLR, LR, and SR setups. For this purpose, we use the Carpet AMR driver (Schnetter, Hawley & Hawke 2004; Reisswig et al. 2013) of the Einstein Toolkit (Loffler et al. 2012; Etienne et al. 2021). Carpet implements the Berger-Oilger scheme with refluxing (Berger & Oliger 1984; Berger & Colella 1989). THC can make use of this infrastructure to ensure mass and energy conservation as matter flows between different refinement levels. However, since the current implementation of refluxing in Carpet is memory intensive, we do not employ it for the radiation variables. To have a baseline for comparison, in addition to the simulations with THC_M1, we perform three simulations using the M0+Leakage neutrino scheme (Radice et al. 2016, 2018b). This is the current methodology employed for neutrino transport in production simulations with the THC general-relativistic hydrodynamics code (Radice & Rezzolla 2012; Radice, Rezzolla & Galeazzi 2014b, a, 2015). However, we have updated the M0 scheme to compute neutrino opacities using the approach discussed in Section 3.2.2. Although THC has the ability to model subgrid-scale viscous angular momentum transport using the GRLES formalism (Radice 2017, 2020), we do not employ it in the simulations presented here.

5.1 Merger dynamics

Our simulations span the last ∼4 orbits of the binary prior to merger, the merger, and extend to ∼10 ms after the merger. After the star come into contact, the remnant experiences one centrifugal bounce before collapsing to BH. We use the AHFinderDirect Thornburg (2004) thorn of the Einstein Toolkit to locate an apparent horizon. Both the M0 scheme and THC_M1 excise the region inside the apparent horizon. Both codes handle BH formation well, but, due to the low resolution, the BH experiences an unphysical drift starting from ∼5 to 10 ms after merger. The drift is particularly large for the M0 runs and eventually the code fails when the BH leaves the finest refinement level in the grid. The M1 runs, instead, experience smaller drifts. The M1 LR run fails at ∼12 ms after merger, while the M1 VLR and SR runs remain stable for the entire duration of the simulation. Such drifts are known to be the result of issues in the shift gauge condition, they are often seen in simulations, and some fixes have been proposed (Bruegmann et al. 2008; Most et al. 2021; Shibata et al. 2021). We remark that such drifts are also seen in purely hydrodynamics simulations, so this issue does not appear to be connected with the neutrino treatment. Since we are not interested in evolving the system for long times after BH formation, we do not attempt to address this issue here.

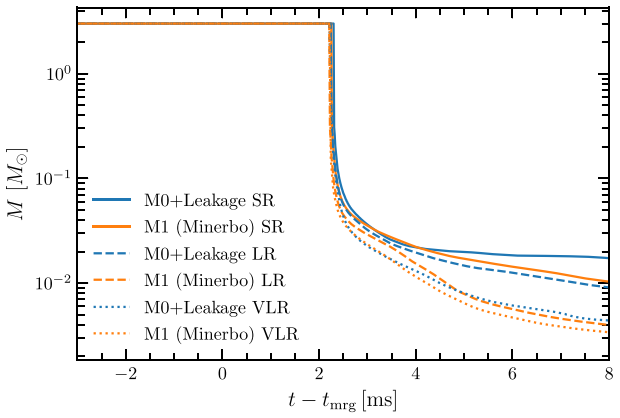

Figure 8. Rest mass outside of the apparent horizon for the SLy4 1.364 $M_\odot$ − 1.364 $M_\odot$ binary as a function of time. The figure shows the results for three resolutions (VLR, LR, and SR) and two radiation transport methods. There is a persistent trend of increasing disc mass with resolution, but we find good agreement between M0 and M1 results.

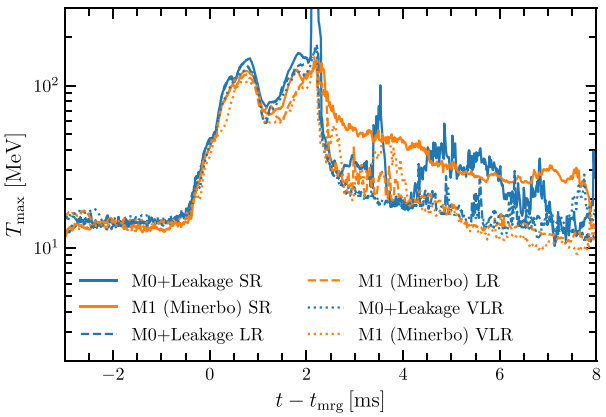

Figure 9. Maximum temperature for the SLy4 1.364 $M_\odot$ − 1.364 $M_\odot$ binary as a function of time. The figure shows the results for three resolutions (VLR, LR, and SR) and two radiation transport methods. We find consistent results among all the simulations, even after BH formation, demonstrating the robustness of our new M1 solver.

Fig. 8 shows the amount of (rest) mass outside of the BH apparent horizon as a function of time. The mass of the accretion disc increases monotonically with resolution and does not appear to have fully converged even at the highest resolution considered in this study. There are differences of up to ∼50 per cent in the disc mass 8 ms after merger between M1 and M0. However, such differences are smaller than the overall uncertainty due to finite resolution effects, suggesting that neutrino transport is not the dominant source of uncertainty in the merger dynamics over these time-scales.

Fig. 9 shows the maximum temperature outside the apparent horizon. During the inspiral, the surface of the stars is artificially heated to temperatures exceeding 10 MeV (Hammond et al. 2021). Equal mass systems with soft EOSs, such as the one considered here, experience the most violent mergers (Radice et al. 2020). Indeed, we observe the temperature to raise to values in excess of 120 MeV. This leads to the production of a dense trapped neutrino gas. This is a very challenging test for a neutrino radiation-hydrodynamics code, since matter is thrown out of weak equilibrium and the radiation-matter coupling becomes very stiff. Our leakage schemes circumvent this

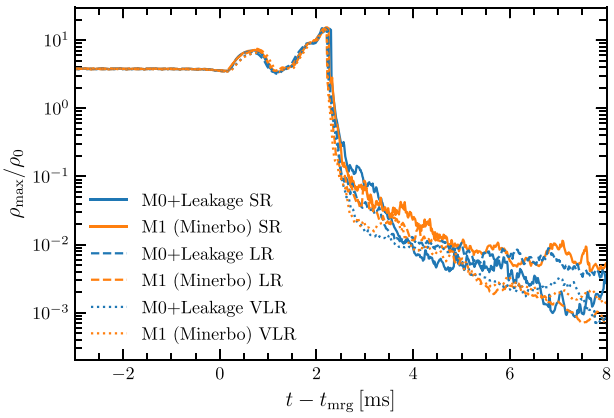

Figure 10. Maximum density for the SLy4 1.364 $M_\odot$ − 1.364 $M_\odot$ binary as a function of time. The figure shows the results for three resolutions and three radiation transport methods. We find consistent results among all the simulations, even after BH formation, demonstrating the robustness of our new M1 solver.

problem by using effective source terms that are not stiff, but does not capture the correct thermodynamic conditions of matter in the remnant (see Perego et al. 2019, for a discussion on the implications). `THC_M1`, instead, captures the correct weak equilibrium of matter inside the star, but at the price of having to solve a stiff set of equations.

After $t - t_{\rm mrg} \simeq 2$ ms, an apparent horizon is found and Fig. 9 shows the maximum temperature in the accretion stream outside of the horizon. Since the highest temperatures are reached close to the horizon, these data are rather sensitive to resolution. It also has large excursions when new grid cells are tagged as being inside the horizon, or when the converse happens. Overall, we find good agreement between the M0+Leakage and the M1 simulations. This test demonstrate that `THC_M1` can handle even the most demanding conditions encountered in NS mergers.

A complementary view on the dynamics of the system can be obtained from Fig. 10 which shows the maximum density outside the apparent horizon. We observe a large oscillation in the maximum density corresponding to the merger and a subsequent centrifugal bounce, followed by the collapse. After $t - t_{\rm mrg} \simeq 2$, the figure shows the maximum density reached in the accretion disc as a function of time. This figure shows that all simulations are in excellent agreement in the description of the bulk motion of matter in the system.

Fig. 11 shows the composition of the remnant accretion torus formed in the highest resolution M1 binary shortly after BH formation. The disc is primarily composed of matter expelled from the inner part of the remnant at the time of merger. The accretion flow is turbulent. The torus has a large $\ell = 2$ deformation, an imprint of the geometry of the remnant shortly after merger (Radice et al. 2018b). We find that the bulk of the torus is very neutron rich, but that its surface layers have higher $Y_e \gtrsim 0.25$ (blue colour in the figure).

5.2 Neutrino luminosities

We compute the emergent neutrino luminosities on a coordinate sphere with radius $r = 300\,GM_\odot/c^2 \simeq 443$ km. The results are shown in Fig. 12. The curves are time shifted to approximately take into account the time of flight of the neutrinos from the remnant to the detection sphere. The neutrino luminosity is artificially large prior to the merger, due to the spurious heating of the stellar surface

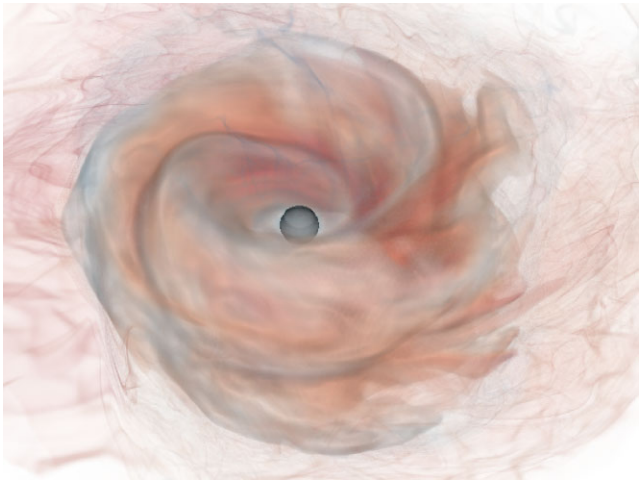

Figure 11. Remnant BH + torus system for the SLy $1.364\,M_\odot - 1.364\,M_\odot$ M1 (Minerbo) SR binary at $t - t_{\rm mrg} \simeq 2.5$ ms. The colour code represents Y_e (red: $Y_e < 0.25$; blue: $Y_e > 0.25$). The inner grey surface shows the approximate location of the apparent horizon ($\alpha = 0.3$; Bernuzzi et al. 2020). The transparency is set to show only matter with density $\rho \gtrsim 5 \times 10^{10}$ g cm^{-3}. The visualization shows the data in a box of diameter 118 km centred at the origin of the coordinate system used in the simulation. We find that the torus is in a turbulent state and is far from axisymmetric.

discussed in the previous subsection. This effect is less severe at higher resolutions. The luminosity peaks shortly after merger and sharply drops following BH formation. As dense neutron rich material is decompressed and heated during merger, it tends to protonize. As a result, the anti-electron neutrino luminosity is the highest among all species, while the electron neutrino luminosity is partially suppressed and is the smallest among all species. Overall, the different resolutions are in good qualitative and quantitative agreement, particularly before BH formation. Discrepancies are found after BH simulation, likely because of the low resolution adopted in this study.

The binary considered here has not yet been simulated by other groups, so detailed comparisons with the literature are not possible. However, the overall neutrino luminosities are in good qualitative agreement with those reported by Foucart et al. (2016b), Vincent et al. (2020), and Foucart et al. (2020) for similar binaries. An important difference is that our luminosities peak at the time of merger and then drop rapidly after BH formation, while the luminosities shown in the aforementioned works increase monotonically, since no BH is formed in those cases. Moreover, those works report the neutrino luminosity only for $t - t_{\rm mrg} > 0$. The luminosities predicted by THC_M1 are in good agreement with the M0 luminosities, but the M0 data (not shown in Fig. 12 to avoid overcrowding the figure) is truncated shortly after BH formation. A more quantitative comparison between M1 and M0 is discussed in Section 6.4. The luminosities predicted by M1 are a factor of several smaller than those predicted by the leakage scheme alone (not taking into account reabsorption; cf. Sekiguchi et al. 2011; Palenzuela et al. 2015; Lehner et al. 2016; Radice et al. 2016). Our luminosities are also a factor of several smaller than those predicted by the M1+Leakage scheme of Sekiguchi et al. (2015, 2016).

The average neutrino energies are also computed on a coordinate sphere of radius $r = 300\,GM_\odot/c^2 \simeq 443$ km and are shown in Fig. 13. With the exception of the average energy anti-electron neutrinos in the LR resolution simulation, we find excellent quantitative agreement between the simulations. The average energies satisfy the expected hierarchy $\langle\epsilon_{\nu_\mu}\rangle > \langle\epsilon_{\bar\nu_e}\rangle > \langle\epsilon_{\nu_e}\rangle$ (Ruffert & Janka 1998; Foucart et al. 2016b; Endrizzi et al. 2020; Cusinato et al. 2021) and are in good quantitative agreement with the Monte Carlo simulations of Foucart et al. (2020), with the caveat that we are not considering the same binary configuration. At $t - t_{\rm mrg} \simeq 2.5$ ms we observe the formation of a shock in the collapsing remnant of the LR simulation, just outside the apparent horizon. This generates a burst of neutrinos that is responsible for the peak in $L_{\bar\nu_e}$. Because the radiation is highly redshifted this results in a dip in $\langle\epsilon_{\bar\nu_e}\rangle$. This feature is absent in the other resolutions.

5.3 Dynamical ejecta

Material is ejected dynamically during the merger by tidal torques and shocks (Shibata & Hotokezaka 2019). We monitor this dynamical ejecta by computing the flux of matter on a coordinate sphere of radius $r = 300\,GM_\odot/c^2 \simeq 443$ km. We consider a fluid element to be unbound if its velocity is larger than the escape velocity from the system ($-u_t > 1$). This is the so-called geodesic criterion (e.g. Kastaun & Galeazzi 2015).

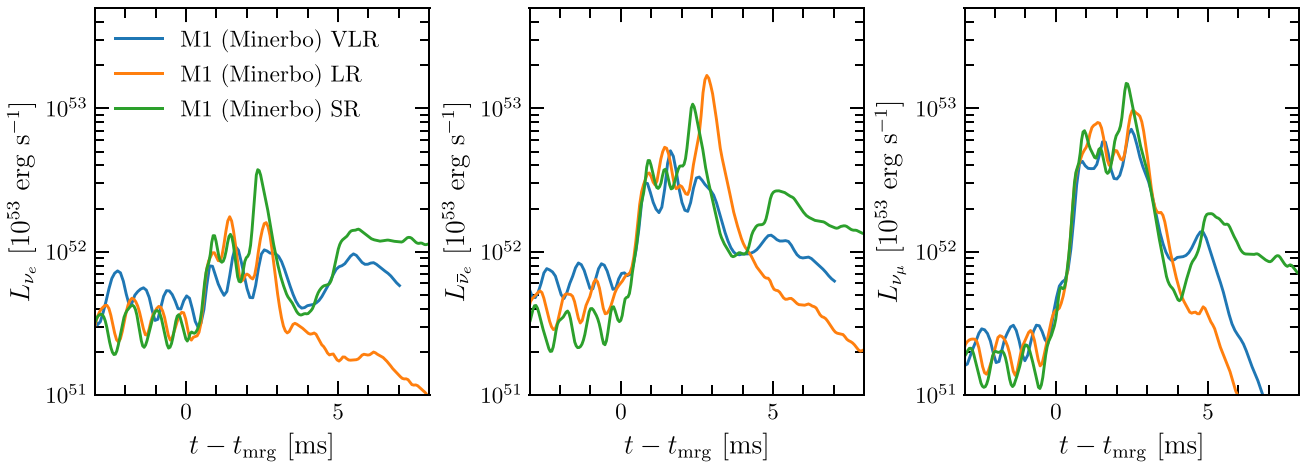

Figure 12. Neutrino luminosity for the SLy $1.364\,M_\odot - 1.364\,M_\odot$ binary computed with THC_M1 at three resolutions. The simulations are in good qualitative agreement at the peak of the neutrino burst, but diverge after BH formation, indicating that the collapse phase is not well resolved in these simulations.

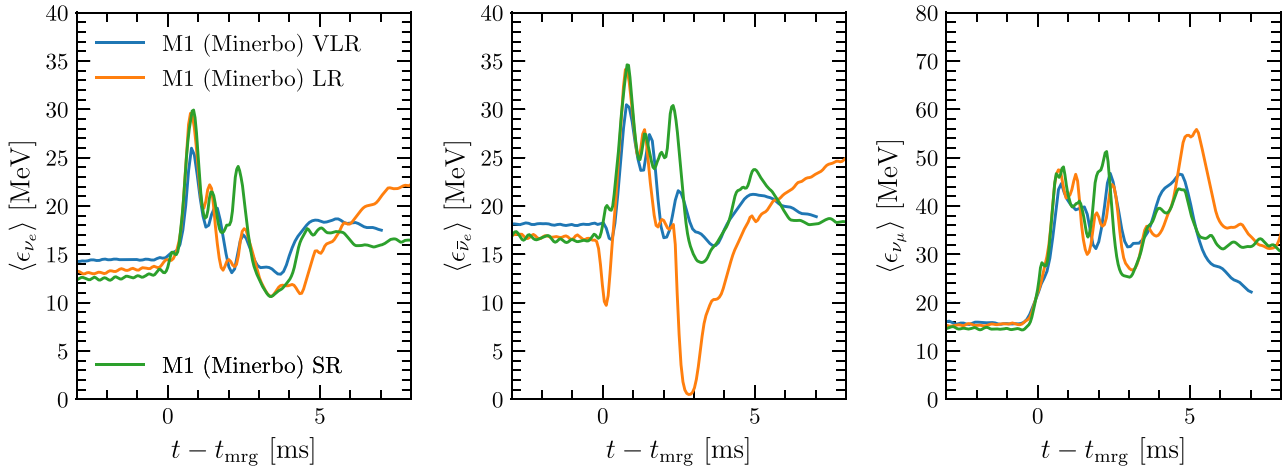

Figure 13. Average neutrino energies for the SLy $1.364\,M_\odot - 1.364\,M_\odot$ binary computed with `THC_M1` at three resolutions. We find good qualitative and quantitative agreement between the three resolutions. The dip in the average $\bar{\nu}_e$ for the LR resolution is due to a burst of highly redshifted radiation originating in the vicinity of the BH.

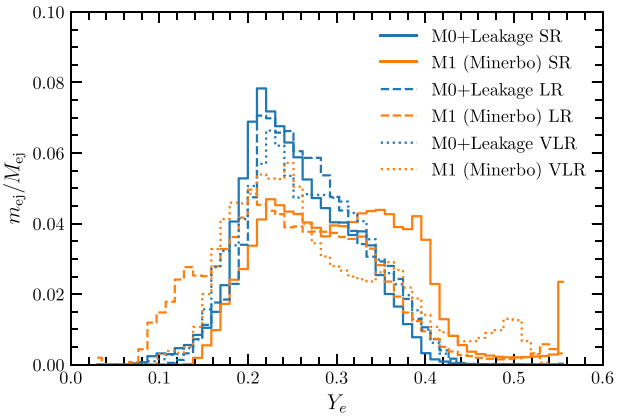

Figure 14. Electron fraction distribution for the dynamical ejecta from the SLy4 $1.364\,M_\odot - 1.364\,M_\odot$ binary. M1 predicts a broad distribution in Y_e extending to $Y_e \simeq 0.4$. The M0 ejecta distribution, instead, clearly peaks at $Y_e \simeq 0.2$. The M1 results appear to be more sensitive to resolution.

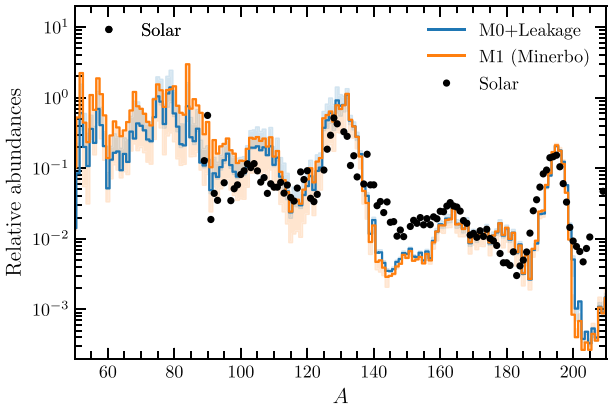

Figure 15. Nucleosynthesis yields for the SLy4 $1.364\,M_\odot - 1.364\,M_\odot$ binary. Compared to the Solar abundance pattern from Arlandini et al. (1999), this binary overproduces r-process elements with $A \simeq 100$, or, equivalently, underproduces second and third peak elements, according to all schemes. Despite the qualitative differences in the Y_e distribution, well resolved M0 and M1 simulations produce similar abundance patterns.

Neutrino irradiation is known to have a strong impact on the composition of the dynamical ejecta from NS mergers (Sekiguchi et al. 2015; Foucart et al. 2016a, b; Radice et al. 2016; Perego et al. 2017b; Foucart et al. 2020), which, in turn, has a strong impact on their nucleosynthesis yields (Lippuner & Roberts 2015; Thielemann et al. 2017; Cowan et al. 2021; Perego, Thielemann & Cescutti 2021). Not surprisingly, we find that the composition of the dynamical ejecta, shown in Fig. 14, is sensitive to the adopted neutrino transport scheme. In particular, the M0+Leakage simulations show a characteristic peak in the Y_e distribution at $Y_e \simeq 0.2$, while the SR M1 shows a broader distribution extending to $Y_e \simeq 0.4$. It also predict the presence of a proton-rich component of the ejecta with $0.55 < Y_e < 0.6$. This component is lumped in the highest Y_e bin in our analysis and is responsible for the bump in the histogram at $Y_e \simeq 0.55$. That said, while the Y_e distribution of the M0 runs is consistent across all resolutions, the Y_e distribution for M1 vary significantly with resolution. The VLR results are in better agreement with the M0 calculations, apart from the presence of a high-Y_e peak for $Y_e \simeq 0.5$. The LR M1 simulations, instead, predict a lower Y_e than the M0 simulations.

The differences in composition are reflected in the final abundances after r-process nucleosynthesis, shown in Fig. 15. The abundances are obtained using a grid of pre-computed trajectories with `SkyNet` (Lippuner & Roberts 2017), as discussed in detail in Radice et al. (2018b). We normalize the relative abundances by fixing the height of the third r-process peak ($A \simeq 190$). We also report Solar r-process abundances from Arlandini et al. (1999) in the same figure. However, we emphasize that even if NS mergers were the sole contributor of r-process elements, there is no reason to expect that every merger should produce ejecta with relative abundances close to Solar. Indeed, variability between the yields of different mergers is required to explain observed abundances in metal-poor stars (Holmbeck et al. 2019). Overall, the simulations span a factor $\sim$2 in the ratio of $A \simeq 100$ to third r-process peak. However, the difference between the M0 and M1 at the SR resolution, which is the resolution we use for production simulation, are modest compared to the systematic uncertainties from the unknown NS EOS and to the variability due to the binary mass ratio (Radice et al.

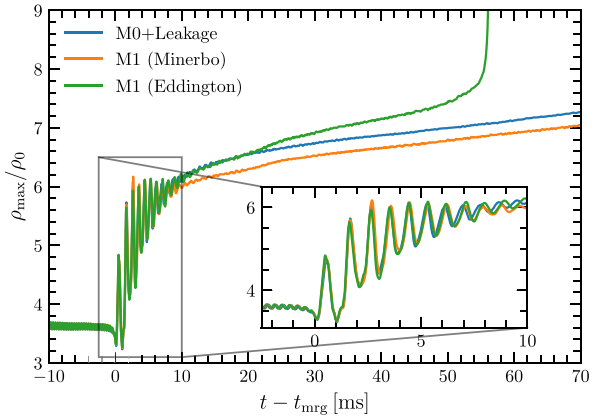

Figure 16. Maximum density as a function of time in millisecond from the merger for the SLy $1.3\,M_\odot - 1.3\,M_\odot$ binary. Small differences in the evolution of the merger remnant are seen starting from ∼10 ms after merger.

2018b; Nedora et al. 2021b). Clearly, strong conclusions cannot be drawn from this limited study alone, but our simulations suggest that the uncertainties in the yields from mergers arising from neutrino radiation treatment are modest. This is also supported by the results of Foucart et al. (2020). They compared M1 and Monte Carlo neutrino transport in the context of NS mergers and reported only a modest ∼10 per cent difference in the Y_e of the ejecta between the two schemes. Interestingly, they reported that M1 systematically overestimates the Y_e of the ejecta, so we cannot exclude that the M0+Leakage results are actually more accurate than the results obtained with THC_M1. That said, it is important to emphasize that this comparisons has only been made for the dynamical ejecta and not for the secular ejecta, which we discuss in Section 6.5.

6 LONG-TERM POST-MERGER EVOLUTIONS

The main application we envision for THC_M1 is to simulate the diffusion of neutrinos out of the merger remnant and the production of winds on secular time-scales after merger. These winds are currently thought to constitute the bulk of the outflow from binary mergers (Shibata & Hotokezaka 2019; Siegel 2019; Radice et al. 2020; Nedora et al. 2021b). In this section, we demonstrate the viability of this approach by performing long-term post-merger simulations for a binary producing a long-lived remnant. In particular, we consider the merger of two identical $1.3\,M_\odot$ NSs simulated with the SLy EOS. Initial data produced with the Lorene code are prepared at an initial separation of 45 km, and have already been considered in Breschi et al. (2019). We perform simulations with THC_M1 with both the Eddington and Minerbo closures. Additionally, we perform a simulation with the M0+Leakage scheme used in production simulations with THC. The M1 (Eddington) simulation is discontinued shortly after BH formation ($t - t_{\rm mrg} \simeq 55$ ms), while the M0+Leakage and the M1 (Minerbo) simulations are carried out until $t - t_{\rm mrg} \simeq 77$ ms. The simulation setup is the same as in that of the calculations presented in the previous section. However, due to the large computational costs, we only present results with the VLR grid spacing.

6.1 Qualitative dynamics

Fig. 16 shows the maximum density for the three simulations. These are in good agreement, especially during the first 10 ms after the merger. Systematic differences appear at later times. In particular, the M1 simulation with Eddington closure collapses to BH at $t - t_{\rm mrg} \simeq 55$ ms, while the other remnants remain stable for the entire simulation time. That said, we caution the reader that the collapse time of the remnant NS is known to be sensitive to resolution and small perturbations, so these differences might not be related to the different neutrino treatment. A detailed investigation of possible neutrino effects on the evolution of the remnant would require many more simulations at higher resolution, so it is outside of the scope of this work.

It has been proposed that out-of-weak-equilibrium effects in the post-merger could give raise to an effective bulk viscosity (Alford et al. 2018, 2020; Hammond et al. 2021; Most et al. 2022). Such effect cannot be captured with leakage schemes, but can be captured with THC_M1, since our code does not assume thermodynamic equilibrium between matter and neutrinos. Our M1 simulations do not show evidences of enhanced damping of the radial oscillations of the remnant compared to the M0 runs. This suggests that the impact of bulk viscosity cannot be too large. That said, higher resolution simulations with a variety of possible EOSs would be required to draw firm conclusions. We also leave this to future work.

The dynamics of the binary is imprinted in the GW signal. We show the dominant $\ell = 2$, $m = 2$ component of the strain in Fig. 17. As for the maximum density, we find that the strain from the three simulations agree both qualitatively and quantitatively. There is a small dephasing between the three waveforms in the post-merger, as can be observed in the figure inset. However, this dephasing is well within the estimated uncertainties in the post-merger signal at this resolution (Radice et al. 2017; Breschi et al. 2019). The most substantial difference between the waveforms is that the M1 (Eddington) GW emission abruptly shuts off at the time of BH formation. Overall, our results show that leakage simulations are adequate to study the GW emission and the early evolution of binary NS remnants. This is not surprising, given the typical neutrino cooling time-scale for the remnant is of a few seconds (Sekiguchi et al. 2011), while most of the GW energy is radiated within ∼20 ms of the merger (Bernuzzi et al. 2016; Zappa et al. 2018).

6.2 Dynamical ejecta

Fig. 18 shows the electron fraction of the ejecta in the meridional plane of the binary about 12 ms after the merger. Overall, we find that the M0+Leakage scheme tends to underestimate the proton fraction in the ejecta, when compared to the M1 scheme. This is consistent with our findings in Section 5, but the 2D plot reveals two interesting systematic differences.

First, the M1 simulations find pockets of moderate Y_e material also in the equatorial regions. This is material that is shock heated and irradiated as the tidal tail and the shocked ejecta collide. The M0 simulations also exhibits an interaction between the tidal tail and the shocked ejecta, however the material remains very neutron rich $Y_e \lesssim 0.2$. A possible explanation for this difference is that the irradiating neutrinos are not propagating radially, so they are not correctly treated by the M0 scheme. This is suggested by the fact that there is a strong density and temperature gradient in the ejecta along the azimuthal direction. This effect is more prominent in the M1 simulation with the Minerbo closure, likely because the Eddington closure limits the propagation velocity of free streaming neutrinos to $c/\sqrt{3}$. This implies that neutrinos interact with the ejecta at larger radii, where they are more diluted.

Secondly, the M1 simulations predict the formation of a tenuous, but rapidly expanding neutrino driven wind with $Y_e \simeq 0.5$ starting few milliseconds after the merger. A similar wind also develops in

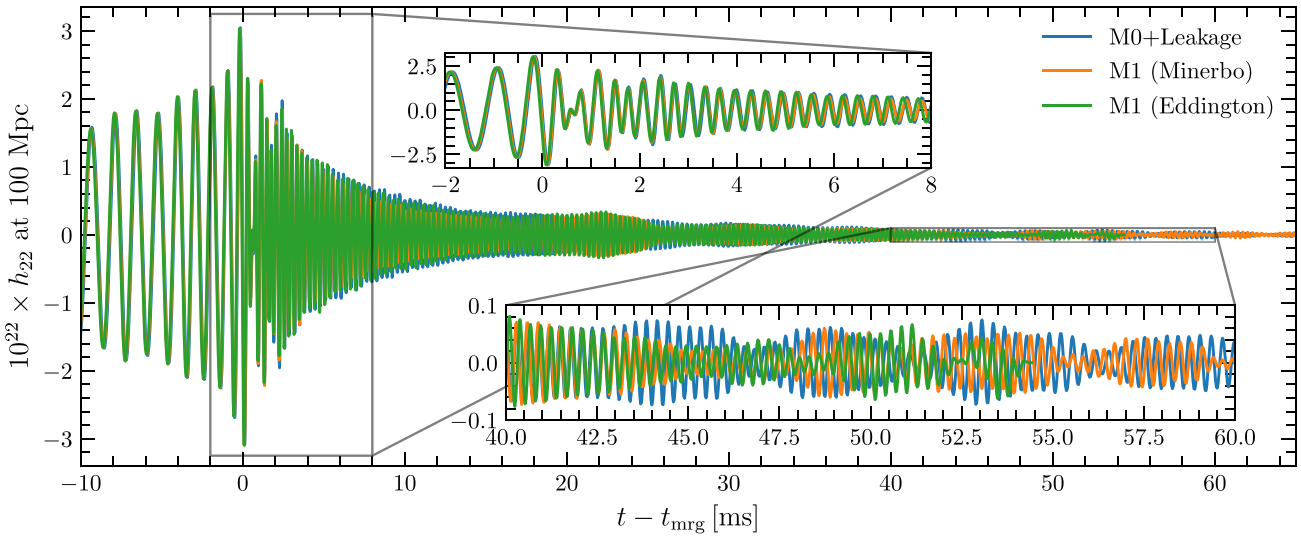

Figure 17. Real part of the $\ell = 2$, $m = 2$ dominant GW strain mode for the SLy $1.3\,\mathrm{M}_\odot - 1.3\,\mathrm{M}_\odot$ binary. Small differences in the evolution of the merger remnant are seen starting from ∼10 ms after merger.

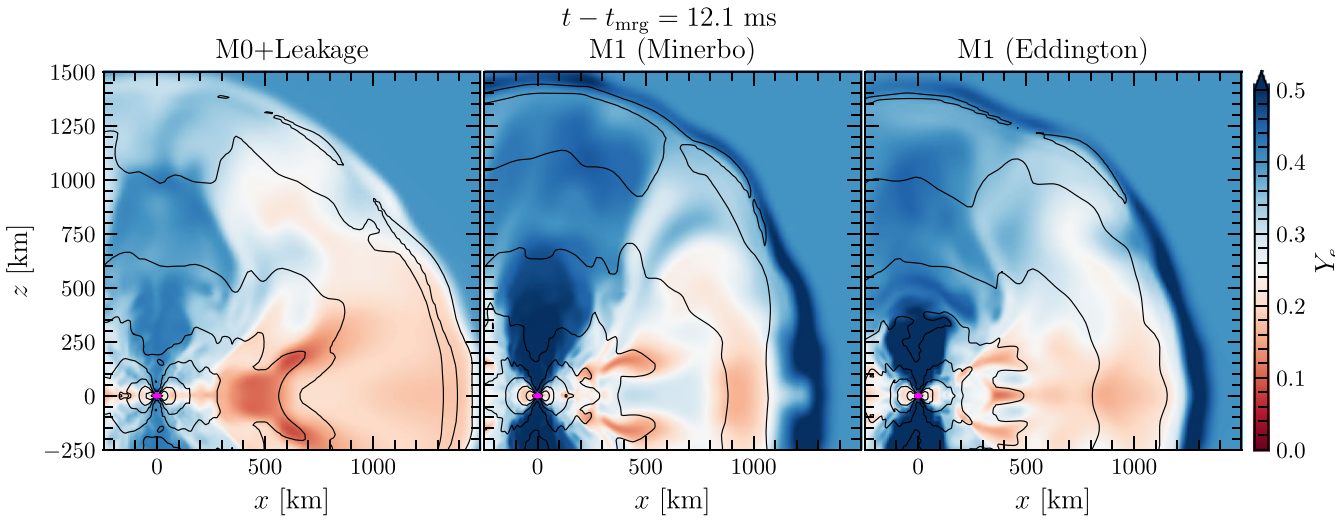

Figure 18. Electron fraction (colour) of the dynamical ejecta cloud formed for the SLy $1.3\,\mathrm{M}_\odot - 1.3\,\mathrm{M}_\odot$ binary. The black lines are isodensity contours of $\rho = 10^5$, 10^6, 10^7, 10^8, 10^9, 10^{10}, 10^{11}, and $10^{12}\,\mathrm{g\,cm^{-3}}$. The purple contour shows corresponds to $\rho = 10^{13}\,\mathrm{g\,cm^{-3}}$ and denotes the approximate location of the surface of the merger remnant. M0 and M1 results are in good qualitative agreement, but M1 predicts higher electron fractions for both the polar and equatorial ejecta.

the M0 case, but with a delay of ∼10–15 ms from the merger. The properties of the neutrino driven winds are discussed in more detail in Section 6.3 and in Nedora et al. (2021b). We speculate that the reason for this discrepancy is that the M0 scheme only models neutrino heating in optically thin regions[2] and might not be able to capture the sharp transition from optically thick to thin conditions along the spin axis of the remnant. As a result, the wind needs to be bootstrapped by the presence of a sufficient amount of low density material ($\rho \lesssim 10^{11}\,\mathrm{g\,cm^{-3}}$) in the polar region of the remnant. This speculation is tentatively confirmed by the fact that the M0 luminosities for electron-flavour neutrinos are larger by a factor of a few compared to the M1 luminosities (see Section 6.4 and Fig. 23), as expected if neutrinos do not entrain baryons in their way out. We remark that the wind is present with both the Minerbo and Eddington closure, so it is not the result of the well known beam-crossing artefact of non-linear M1 closures (Frank et al. 2007).

[2]Absorption is included also at high optical-depth, but is suppressed with a factor $O(e^{-\tau})$, τ being the optical depth, to be consistent with the effective sources of the leakage scheme.

6.3 Remnant structure

Fig. 19 shows the structure and composition of the merger remnant ∼55 ms after the merger. We find good qualitative agreement between the three numerical schemes. In particular, all simulations predict a very neutron rich composition ($Y_e \lesssim 0.2$) for the accretion torus and the presence of a high-Y_e wind at high latitudes. They also predict

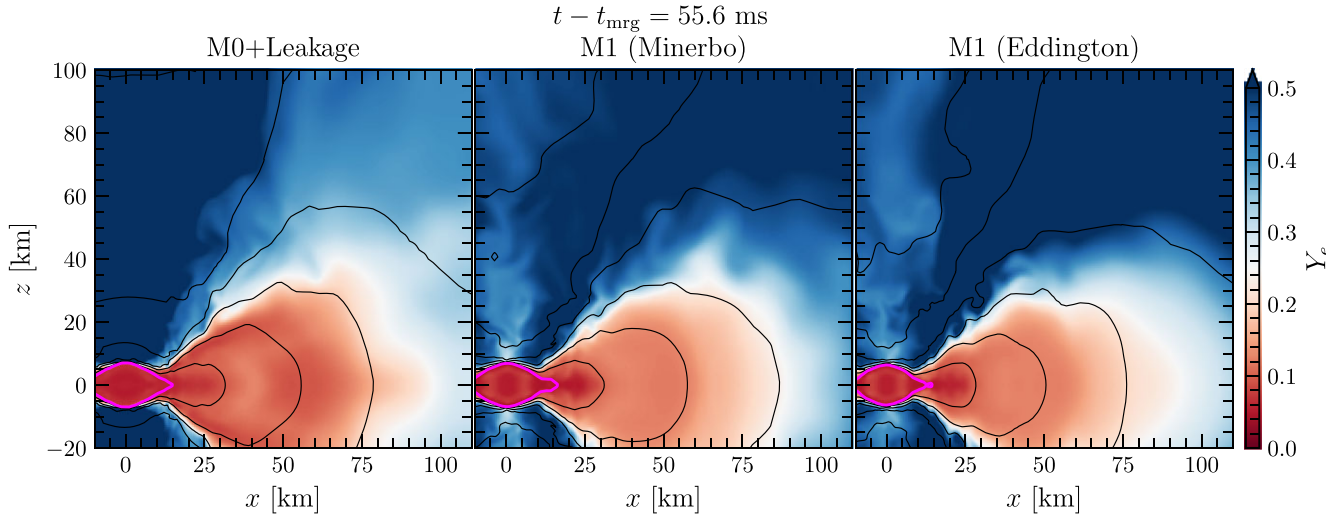

Figure 19. Electron fraction (colour) of the of the SLy $1.3\,M_\odot - 1.3M_\odot$ merger remnant ∼55 ms after merger. The black lines are isodensity contours of $\rho = 10^5$, 10^6, 10^7, 10^8, 10^9, 10^{10}, 10^{11}, and 10^{12} g cm^{-3}. The purple contour shows corresponds to $\rho = 10^{13}$ g cm^{-3} and denotes the approximate location of the surface of the merger remnant. M0 and M1 results are in good qualitative agreement, but M1 predicts higher electron fractions in the disc corona and somewhat smaller electron fraction in the neutrino-driven wind along the rotation axis.

a shift to higher Y_e at densities below 10^{11} g cm^{-3}, where thermal electron-type neutrinos are expected to decouple (Endrizzi et al. 2020). However, there are important quantitative differences.

First, the M0+Leakage scheme systematically underpredicts the Y_e in the corona of the disc. This is because M0 only transports neutrinos radially, so it cannot model the irradiation of the corona by neutrinos emerging from the disc below, while THC_M1 does not have this limitation.

Secondly, there are small, but important differences in the Y_e of the remnant. These differences arise because our leakage scheme does not model the presence of a trapped component of $\bar{\nu}_e$ in the remnant. THC_M1, instead, correctly captures the protonization of the region of the remnant around $\rho = 10^{14}$ g cm^{-3} and the creation of a trapped component of anti-electron neutrinos, in agreement with the predictions of Perego et al. (2019). This trapped neutrino component can impact the pressure at the several percent level (Perego et al. 2019), which might be sufficient to impact the remnant stability (Radice et al. 2018a).

Thirdly, the M1 simulations produce a denser neutrino-driven wind, as can be seen from the isodensity contours in Fig. 19. This wind also entrains material from the outer layers of the central remnant, so it is more neutron rich than that predicted by the M0 simulation. This difference could have been anticipated, because the M0+Leakage scheme only models the transport and reabsorption of free streaming neutrinos, while M1 can also capture the heating of the outer layers of the remnant due to the diffusion of neutrinos along the steep density and temperature gradient along the rotational axis of the binary. In particular, because the opacity in the M0+Leakage scheme is weighted with the optical depth, this scheme systematically underestimates heat deposition for optical depths $\tau \gtrsim 1$.

Fig. 20 shows the neutrino energy density for the M1 (Minerbo) simulation ∼55 ms after the merger. This is a representative time for the neutrino field in the post-merger. However, we emphasize that the neutrino energy density oscillates and shows quasi-periodic bursts, especially shortly after merger. The M1 (Eddington) neutrino radiation energy densities are qualitatively and quantitatively similar. We observe the formation of a trapped component of neutrinos. As previously discussed, $\bar{\nu}_e$ are the dominant neutrino species in the inner part of the remnant. However, we find trapped neutrinos of all flavours in the central part of the remnant and in the accretion disc. Radiation is geometrically focused in the polar direction and most intense ∼10–20 km above the surface of the massive NS. Equatorial outflows are shielded from the intense neutrino radiation from the inner part of the remnant by the torus, but they are instead irradiated by neutrinos produced directly in the disc.

There are effectively two sources of electron-flavour neutrinos. The massive NS at the centre and the disc. Neutrinos from the massive NS have ∼50 per cent higher average energies (see Fig. 21), so their interaction cross-section with matter is ∼3 times larger. However, only material outflowing in the polar direction is directly exposed to these neutrinos. The neutrinos from the disc are less energetic, but fill a significantly larger area (Fig. 21). The net effect is to enhance the differences in the Y_e of polar and equatorial ejecta and to increase the anisotropic character of the resulting kilonova emission (Perego et al. 2017b; Kawaguchi, Shibata & Tanaka 2020; Korobkin et al. 2021).

Fig. 22 shows the average neutrino energy obtained with the Eddington closure. There are small differences with the Minerbo closure in the location of the separatix between the stream of neutrinos emerging from the massive NS and the disc. This is because, on the one hand, the Minerbo closure artificially prevents different neutrino streams from mixing. On the other hand, the Eddington closure tends to smooth out structures in the radiation energy density profile. Most notably, the x-shaped feature present in the M1 (Minerbo) run for both E_{ν_μ} and $\langle \epsilon_{\nu_\mu} \rangle$ close to the massive NS is absent in the Eddington simulations. This suggests that this feature is likely to be an artefact of the Minerbo closure. That said, Minerbo and Eddington closure are broadly consistent with each other, suggesting that the results discussed so far are robust.

6.4 Neutrino emission

We show the angle integrated neutrino luminosities for the SLy $1.3\,M_\odot - 1.3\,M_\odot$ binary in Fig. 23. The neutrino luminosities for

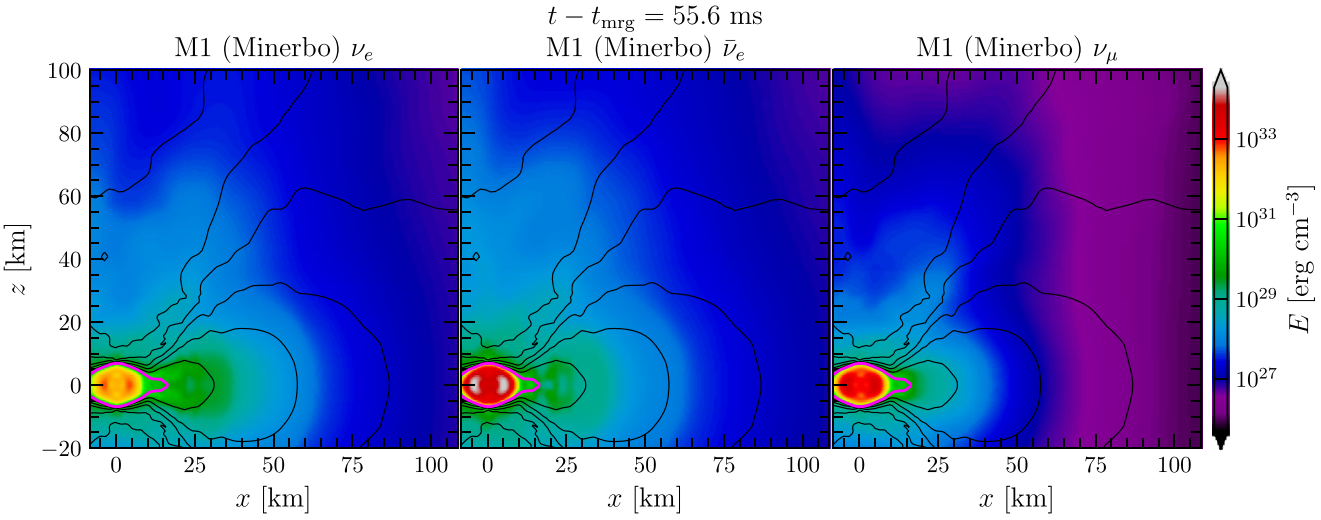

Figure 20. Neutrino radiation energy density (colour) for the SLy $1.3\,\mathrm{M}_\odot - 1.3\mathrm{M}_\odot$ binary ∼55 ms after merger. The black lines are isodensity contours of $\rho = 10^5$, 10^6, 10^7, 10^8, 10^9, 10^{10}, 10^{11}, and $10^{12}\,\mathrm{g\,cm^{-3}}$. The purple contour shows corresponds to $\rho = 10^{13}\,\mathrm{g\,cm^{-3}}$ and denotes the approximate location of the surface of the merger remnant. Due to the geometry of the post-merger, radiation is preferentially focused in the polar regions.

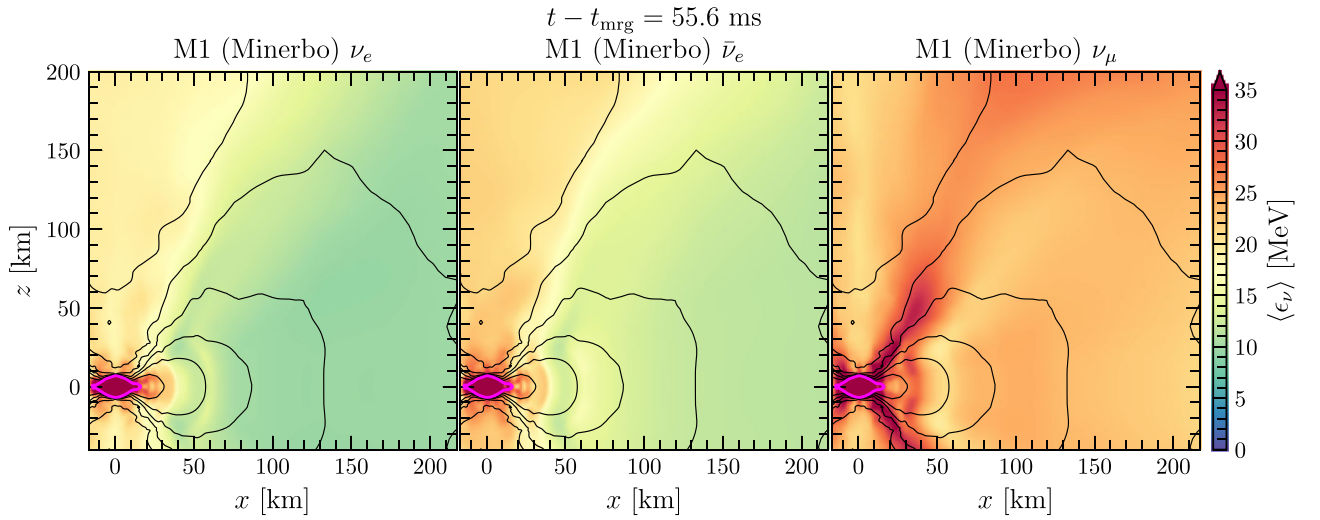

Figure 21. Average neutrino energies (colour) for the SLy $1.3\,\mathrm{M}_\odot - 1.3\mathrm{M}_\odot$ binary ∼55 ms after merger. The black lines are isodensity contours of $\rho = 10^5$, 10^6, 10^7, 10^8, 10^9, 10^{10}, 10^{11}, and $10^{12}\,\mathrm{g\,cm^{-3}}$. The purple contour shows corresponds to $\rho = 10^{13}\,\mathrm{g\,cm^{-3}}$ and denotes the approximate location of the surface of the merger remnant. The average neutrino energy is highly anisotropic, especially for electron-type neutrinos, since the disc is optically thick to high energy neutrinos.

the M1 (Minerbo) and M1 (Eddington) simulations are extracted on a coordinate sphere of radius $r = 300\,GM_\odot/c^2 \simeq 443$ km. The luminosities of the M0+Leakage scheme are computed at the outer boundary of the M0 spherical grid $512\,GM_\odot/c^2 \simeq 756$ km. All data are time shifted to account for the neutrino time of flight. As anticipated, we find that M0+Leakage systematically overestimates the luminosity of electron-flavour neutrinos. Good agreement is found for heavy-lepton neutrons, instead. In all cases, the luminosities peak within a few milliseconds of the merger, in contrast to Vincent et al. (2020), and then decay exponentially. The oscillations in the luminosity are not due to a numerical artefact, but are associated with the oscillations of the massive NS remnant (Cusinato et al. 2021). As was the case for the SLy $1.364\,\mathrm{M}_\odot - 1.364\,\mathrm{M}_\odot$ binary, $L_{\bar{\nu}_e} > L_{\nu_\mu} > L_{\nu_e}$, showing that the remnant is protonizing.

The average neutrino energies for the SLy $1.3\,\mathrm{M}_\odot - 1.3\,\mathrm{M}_\odot$ binary are shown in Fig. 24. We find excellent agreement in the average electron-flavour neutrino energies for all schemes. The M0+Leakage scheme predicts a ∼50 per cent smaller average energy for heavy-lepton neutrinos. Moreover, the M0+Leakage scheme predicts a nearly constant heavy-lepton neutrino energy as a function of time. This is because M0 does not properly diffuse neutrinos through the remnant. Instead, each part of the remnant cools at a rate that depends on its optical depth. In contrast, `THC_M1` models the diffusion of neutrinos to the neutrino spheres and their thermalization. Also in this case, we find that $\langle \epsilon_{\nu_\mu} \rangle > \langle \epsilon_{\bar{\nu}_e} \rangle > \langle \epsilon_{\nu_e} \rangle$, as expected.

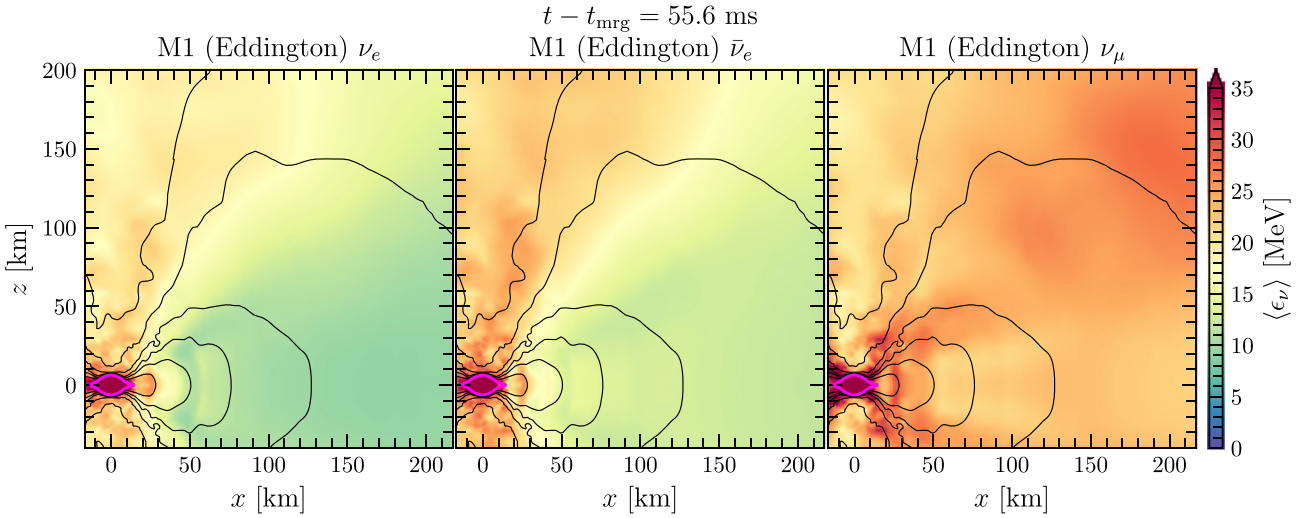

Figure 22. Average neutrino energies (colour) for the SLy $1.3\,M_\odot - 1.3M_\odot$ binary ∼55 ms after merger. The black lines are isodensity contours of $\rho = 10^5$, 10^6, 10^7, 10^8, 10^9, 10^{10}, 10^{11}, and 10^{12} g cm^{-3}. The purple contour shows corresponds to $\rho = 10^{13}$ g cm^{-3} and denotes the approximate location of the surface of the merger remnant. This figure should be contrasted with Fig. 21, which shows the same profiles with obtained with the Minerbo closure.

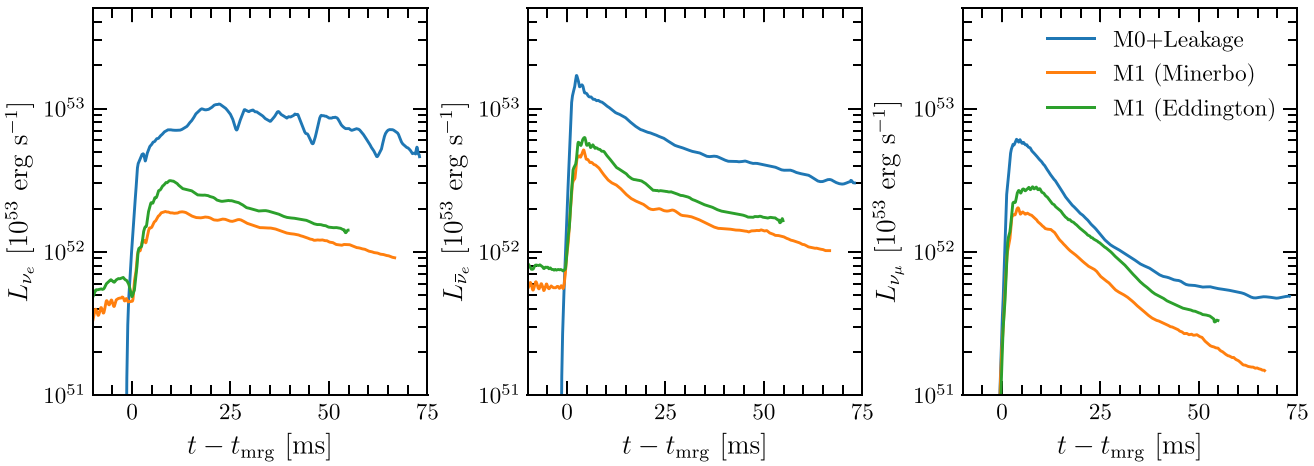

Figure 23. Neutrino luminosity for the SLy $1.3\,M_\odot - 1.3\,M_\odot$ binary. The data are smoothed using a rolling average with width of 1 ms. We find that at this resolution M0 systematically overestimates the ν_e and $\bar{\nu}_e$ luminosities by about a factor of two. The M1 Eddington and Minerbo luminosities are in good agreement.

6.5 Secular ejecta

We observe the emergence of an outflow driven by hydrodynamics torques: the so-called spiral-wave wind (Nedora et al. 2019, 2021b), in addition to the aforementioned neutrino-driven wind. This secular ejecta is extracted at the same extraction radius of the dynamical ejecta ($r = 300\,GM_\odot/c^2 \simeq 443$ km), but we use the Bernoulli criterion ($-hu_t > 1$), which is more appropriate for a steady wind. See Foucart et al. (2021a) for a recent discussion of the issues connected to the discrimination between gravitationally bound and unbound outflows. The time-integrated outflow rate is shown in Fig. 25. We find that the leakage+M0 simulation produces a more robust wind with a larger $\dot{M}$, while the two M1 simulations are in good agreement with each other. However, we warn the reader that, at this resolution, the numerical uncertainties in the outflow is $\gtrsim 50$ per cent (Nedora et al. 2021b), so these differences might not be particularly meaningful. In particular, our previous simulations at higher resolution (Breschi et al. 2019), but with simpler neutrino physics, suggest that this binary might form a BH few tens of milliseconds after merger. Since the spiral wave wind ceases with BH formation (Nedora et al. 2019), the uncertainty in the BH formation time is likely to dominate the overall error budget on the total ejecta mass for this binary.

Fig. 26 shows the composition of the overall ejecta (dynamical + secular) for the SLy $1.3\,M_\odot - 1.3\,M_\odot$ binary. We find that all schemes produce a wide distribution in Y_e. The results are qualitatively consistent with our previously published M0 simulations (Nedora et al. 2019, 2021b). An important quantitative difference is that the M0 scheme predicts a peak in the electron fraction distribution at $Y_e \simeq 0.3$. The outflows in the M1 simulations are, instead, characterized by a peak in their electron fraction at ∼0.5. We attribute this difference to the irradiation of outflows at intermediate latitudes by neutrinos from the disc, an effect that is not captured by the M0 scheme (see Fig. 19). Some differences are also found in the low-Y_e tail of the ejecta, which is primarily of dynamical origin, as anticipated by Fig. 18.

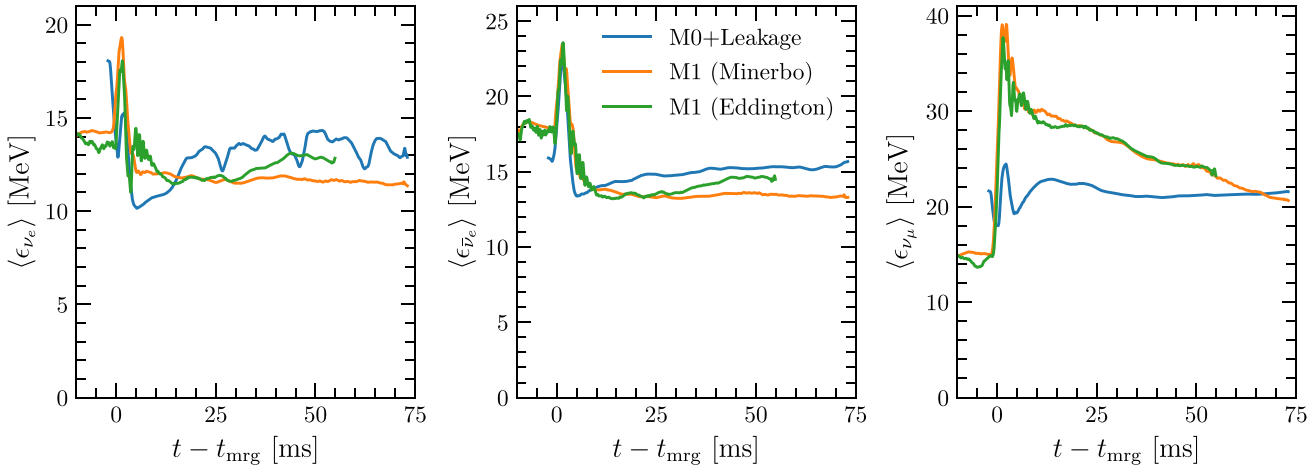

Figure 24. Average neutrino energies for the SLy $1.3\,M_\odot - 1.3\,M_\odot$ binary. The data are smoothed using a rolling average with width of 1 ms. We find excellent agreement in the average neutrino energies for electron type neutrinos. The M0 scheme predicts smaller average energies for heavy-lepton flavour neutrinos.

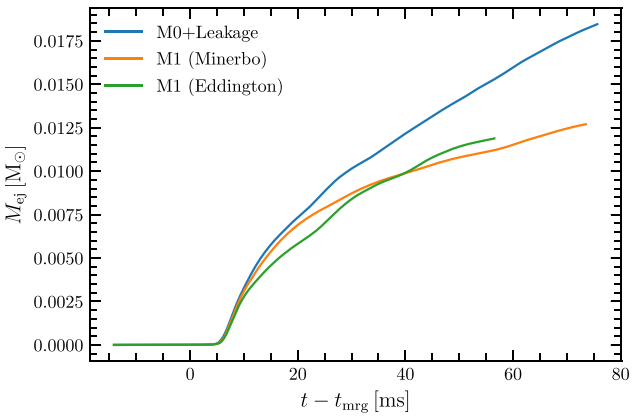

Figure 25. Ejecta mass for the SLy $1.3\,M_\odot - 1.3\,M_\odot$ binary. The M0 simulations slightly overestimate the outflow rate for the post-merger wind compared to M1. The results of the Minerbo and Eddington closures are consistent with each other.

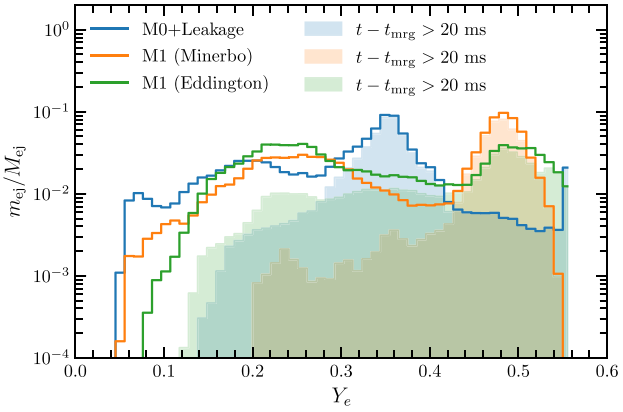

Figure 26. Ejecta Y_e for the SLy $1.3\,M_\odot - 1.3\,M_\odot$ binary. The shaded regions show the contributions to the Y_e histogram due to the material ejected after the first 20 ms of the merger. The ejecta distribution peaks at significantly larger Y_e in the M1 runs. The Eddington and Minerbo results are in good agreement, but the Eddington simulations produce a smaller amount of very neutron rich ejecta ($Y_e \sim 0.1$).

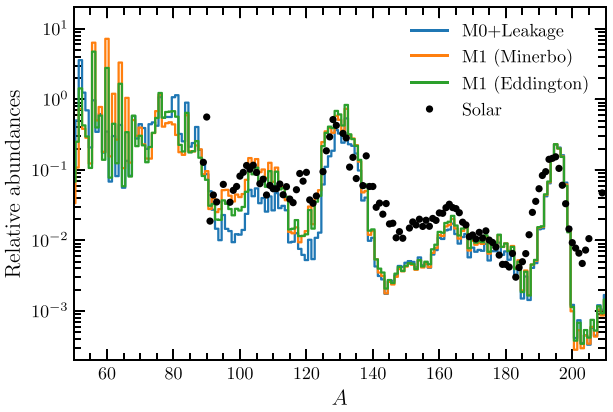

Figure 27. Normalized nucleosynthesis yield for the SLy $1.3\,M_\odot - 1.3\,M_\odot$ binary. The M1 runs predict elemental abundances that are in better agreement with the Solar pattern, while M0 underproduces r-process elements with $A \sim 110$. Overall, however, the differences between M0 and M1 are modest.

These changes in Y_e do not contribute to very large differences in the nucleosynthesis. This is because the main effect of M1 is to shift the peak of the Y_e distribution from 0.3 to 0.5, but both peaks correspond to a regime in which only light r-process elements are produced. The integrated nucleosynthesis yields for the three SLy $1.3\,M_\odot - 1.3\,M_\odot$ simulations are shown in Fig. 27. The relative abundances of light to heavy r-process peak elements differs by about a factor of two between the M1 and the M0+Leakage runs. This is a significant, but not substantial discrepancy, considering the large variabilities of the yields with EOS and mass ratio (Radice et al. 2018b; Nedora et al. 2021b). The differences between the M1 (Minerbo) and M1 (Eddington) simulations are below the level of finite resolution uncertainties (see Section 5.3).

7 CONCLUSIONS

We have presented `THC_M1`, a new moment-based neutrino transport code for numerical relativity simulations of merging NSs. `THC_M1` handles radiation advection using a high-resolution shock capturing scheme that can capture both the free streaming and the diffusive regimes. `THC_M1` simultaneously evolves the frequency-integrated

energy and neutrino number density equations. Ours is one of the first GR radiation transport codes, the first in the merger context, to include velocity-dependent effects at all orders in v/c. We have shown that this full treatment, while technically more complex than that used in other codes, is necessary to correctly capture neutrino trapping in relativistically moving media, such as rotating NSs remnants.

After having validated our new code with a stringent series of tests, we have coupled it with the THC relativistic hydrodynamics code to perform merger simulations of two equal mass binaries: an intermediate mass binary resulting in a short lived remnant that quickly collapses to BH, and a low-mass binary that produces a long-lived remnant. We have studied numerical resolution effects using the first binary, while we have performed long-term evolutions at a fixed resolution for the second binary using two different closure relations, the so-called Eddington and Minerbo closures. To have a baseline for comparison, we have also simulated the same systems using a more approximated M0+Leakage scheme, which is currently used for production simulations with THC.

The intermediate mass binary experiences a violent merger. The remnant is rapidly heated to temperatures exceeding 100 MeV following the collision between the stars. The massive NS formed in the collision undergoes one centrifugal bounce that launches a shock in the merger debris and drives a massive outflow, before collapsing to BH. This is one of the most challenging binary to model due to the high temperatures and the BH formation. We find that THC_M1 is as robust, if not more robust, than our production leakage code. The predicted neutrino luminosities and average energies are consistent with theoretical expectations and other results from the literature.

The remnant of the low mass binary merger also experiences a series of violent oscillations at birth, with maximum density jumping by more than 50 per cent on a dynamical time-scale. However, the remnant eventually settles into a massive, differentially rotating NS evolving on secular time-scales. Even though THC_M1 includes out-of-weak-equilibrium effects which have been suggested to result in an effective bulk viscosity (Alford et al. 2018), we do not find any evidence of additional damping of the remnant oscillations in the M1 runs, compared to simulations that do not model them. That said, simulations with a more comprehensive set of reactions, with more EOSs, and at more resolutions are needed before firm conclusions can be drawn.

We have performed simulations extending for over 70 ms after the merger. For comparison, the longest published simulations performed with a neutrino-transport scheme having comparable sophistication only extended to 10 ms into the post-merger (Vincent et al. 2020). We find that the post-merger GW signal is not sensitive to details in the neutrino transport. However, the inner structure of the massive NS is modified by the presence of a trapped component of anti-electron neutrinos. This could impact the stability of the remnant of higher mass binaries. We find that, due to the geometry of the system, neutrino radiation is most intense along the rotational axis of the system. Matter at lower latitudes is shielded from the direct irradiation from the massive NS by the disc. Instead, it is irradiated by lower energy neutrinos produced in the accretion disc. Because neutrino absorption cross-sections roughly scale with the square of the incoming neutrino energy, this enhances the Y_e difference between polar and equatorial ejecta and has implications for the viewing angle dependence of kilonovae.

We have computed integrated neutrino luminosities and average neutrino energies from our simulations. Consistently with previous studies, we find that anti-electron neutrinos have the highest luminosity and that heavy-lepton neutrinos have the highest average energies. Our M1 data is in good qualitative and quantitative agreement with results published by the SXS collaboration using SpEC. On the other hand, we find that our older M0 neutrino scheme can overestimate electron-flavour neutrino luminosities by as much as a factor two. Discrepancy with the results from leakage calculations, either performed by us or by other groups, are significantly larger and amount to factors of several. We find an excellent agreement between M1 and M0+Leakage in the neutrino average energies, instead.

Neutrino transport impacts the neutron richness of both the dynamical and the secular ejecta in our simulations. In particular, we find that there is a systematic tendency of M0+Leakage to underestimate the electron fraction of the ejecta. This is because the M0 scheme does not model the irradiation of material at intermediate latitudes with neutrinos generated in the remnant accretion disc. However, because the net effect is to reprocess material with $Y_e \simeq 0.2$–0.35 to $Y_e \simeq 0.4$–0.55, this has only a modest impact on the final abundances of the r-process nucleosynthesis.

THC_M1 represents a step forward in the modelling of neutrinos in mergers, particularly over long time-scales over which diffusion of neutrinos from the inner part of the remnant needs to be taken into account. However, this study still has some important limitations to be addressed. Most importantly, our work used a rather crude set of weak reactions and accounted for the energy-dependence of neutrino-matter cross-sections in a simplistic way. We plan to update the set of weak reactions included in our code and to use Planck-averaged opacities that take into account the average incoming neutrino energy. We also plan to perform a larger campaign of simulations spanning a range of binary masses, mass ratios, and EOSs, in order to understand the general features of neutrino-driven winds from NS mergers and the role of non-equilibrium effects in the post-merger. Finally, our work has neglected quantum kinetic effects in the neutrino transport (Zhu, Perego & McLaughlin 2016; Deaton et al. 2018; Richers et al. 2019; George et al. 2020; Li & Siegel 2021; Richers, Willcox & Ford 2021). Future work should quantify the importance of these effects for mergers.

ACKNOWLEDGEMENTS

We wish to thank Luke Roberts for help with ZelmaniM1, and Aviral Prakash and Francesco Zappa for having beta-tested an early version of THC_M1. We also thank Xinyu Li and Patrick Cheong for having spotted some typos in an earlier version of the manuscript. DR acknowledges funding from the U.S. Department of Energy, Office of Science, Division of Nuclear Physics under Award Number(s) DE-SC0021177 and from the National Science Foundation under Grants No. PHY-2011725, PHY-2020275, PHY-2116686, and AST-2108467. SB acknowledges support from the EU H2020 under ERC Starting Grant, no. BinGraSp-714626, and from the Deutsche Forschungsgemeinschaft, DFG, project MEMI number BE 6301/2-1. RH acknowledges funding from the National Science Foundation under Grants No. OAC-2004879, OAC-1550514, and ACI-1238993. NR simulations were performed on ARA cluster at Friedrich Schiller University Jena, SuperMUC-NG at the Leibniz-Rechenzentrum (LRZ) Munich, HPE Apollo Hawk at the High Performance Computing Center Stuttgart (HLRS); Marconi-CINECA (ISCRA-B project HP10BMHFQQ, INF20_teongrav and INF21_teongrav allocation); Bridges, Comet, Stampede2 (NSF XSEDE allocation TG-PHY160025), NSF/NCSA Blue Waters (NSF AWD-1811236), Joliot-Curie at GENCI@CEA (PRACE-ra5202) supercomputers. This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. Computations for this

research were also performed on the Pennsylvania State University's Institute for Computational and Data Sciences' Roar supercomputer. The ARA cluster is funded in part by DFG grants INST 275/334-1 FUGG and INST 275/363-1 FUGG, and ERC Starting Grant, grant agreement no. BinGraSp-714626. The authors acknowledge the Gauss Centre for Supercomputing e.V. (www.gauss-centre.eu) for funding this project by providing computing time on the GCS Supercomputer SuperMUC-NG at LRZ (allocation `pn68wi`).

DATA AVAILABILITY

Data generated for this study will be made available upon reasonable request to the corresponding authors.

REFERENCES

Alford M. G., Bovard L., Hanauske M., Rezzolla L., Schwenzer K., 2018, Phys. Rev. Lett., 120, 041101
Alford M., Harutyunyan A., Sedrakian A., 2020, Particles, 3, 500
Andersson N., Lopez-Monsalvo C. S., 2011, Class. Quantum Gravity, 28, 195023
Anninos P., Fragile P. C., 2020, ApJ, 900, 71
Arcones A., Thielemann F.-K., 2013, J. Phys. G: Nucl. Phys., 40, 13201
Ardevol-Pulpillo R., Janka H.-T., Just O., Bauswein A., 2019, MNRAS, 485, 4754
Arlandini C., Käppeler F., Wisshak K., Gallino R., Lugaro M., Busso M., Straniero O., 1999, ApJ, 525, 886
Audit E., Charrier P., Chieze J. P., Dubroca B., 2002, preprint (arXiv:astro-ph/0206281)
Berger M. J., Colella P., 1989, J. Comput. Phys., 82, 64
Berger M. J., Oliger J., 1984, J. Comput. Phys., 53, 484
Bernuzzi S., Radice D., Ott C. D., Roberts L. F., Mösta P., Galeazzi F., 2016, Phys. Rev. D, 94, 24023
Bernuzzi S. et al., 2020, MNRAS, 497, 1488
Bollig R., Yadav N., Kresse D., Janka H.-T., Müller B., Heger A., 2021, ApJ, 915, 28
Breschi M., Bernuzzi S., Zappa F., Agathos M., Perego A., Radice D., Nagar A., 2019, Phys. Rev. D, 100, 104029
Brügmann B., González J. A., Hannam M., Husa S., Sperhake U., Tichy W., 2008, Phys. Rev. D, 77, 24027
Burrows A., Vartanyan D., 2021, Nature, 589, 29
Burrows A., Radice D., Vartanyan D., Nagakura H., Skinner M. A., Dolence J. C., 2020, MNRAS, 491, 2715
Chan C., Müller B., 2020, MNRAS, 496, 2000
Cowan J. J., Sneden C., Lawler J. E., Aprahamian A., Wiescher M., Langanke K., Martínez-Pinedo G., Thielemann F.-K., 2021, Rev. Mod. Phys., 93, 15002
Cusinato M., Guercilena F. M., Perego A., Logoteta D., Radice D., Bernuzzi S., Ansoldi S., 2021, preprint (arXiv:2111.13005)
Davis S. W., Stone J. M., Jiang Y.-F., 2012, ApJS, 199, 9
Deaton M. B., O'Connor E., Zhu Y. L., Bohn A., Jesse J., Foucart F., Duez M. D., McLaughlin G. C., 2018, Phys. Rev. D, 98, 103014
Dessart L., Ott C., Burrows A., Rosswog S., Livne E., 2009, ApJ, 690, 1681
Douchin F., Haensel P., 2001, A&A, 380, 151
Eichler D., Livio M., Piran T., Schramm D. N., 1989, Nature, 340, 126
Endrizzi A. et al., 2020, Eur. Phys. J. A, 56, 15
Etienne Z. et al., 2021, The Einstein Toolkit. Available at: https://doi.org/10.5281/zenodo.4884780
Fernández R., Tchekhovskoy A., Quataert E., Foucart F., Kasen D., 2019, MNRAS, 482, 3373
Foucart F. et al., 2015, Phys. Rev. D, 91, 124021
Foucart F. et al., 2016a, Phys. Rev. D, 93, 044019
Foucart F., O'Connor E., Roberts L., Kidder L. E., Pfeiffer H. P., Scheel M. A., 2016b, Phys. Rev. D, 94, 123016
Foucart F., Duez M. D., Hebert F., Kidder L. E., Pfeiffer H. P., Scheel M. A., 2020, ApJ, 902, L27
Foucart F., Moesta P., Ramirez T., Wright A. J., Darbha S., Kasen D., 2021a, Phys. Rev. D, 104, 123010
Foucart F., Duez M. D., Hebert F., Kidder L. E., Kovarik P., Pfeiffer H. P., Scheel M. A., 2021b, ApJ, 920, 82
Frank M., Klar A., Larsen E. W., Yasuda S., 2007, J. Comput. Phys., 226, 2289
Fujibayashi S., Sekiguchi Y., Kiuchi K., Shibata M., 2017, ApJ, 846, 114
Fujibayashi S., Kiuchi K., Nishimura N., Sekiguchi Y., Shibata M., 2018, ApJ, 860, 64
Fujibayashi S., Shibata M., Wanajo S., Kiuchi K., Kyutoku K., Sekiguchi Y., 2020a, Phys. Rev. D, 101, 083029
Fujibayashi S., Wanajo S., Kiuchi K., Kyutoku K., Sekiguchi Y., Shibata M., 2020b, ApJ, 901, 122
Galassi M., 2009, GNU Scientific Library: Reference Manual. Network Theory, Bristol
George M., Wu M.-R., Tamborra I., Ardevol-Pulpillo R., Janka H.-T., 2020, Phys. Rev. D, 102, 103015
Gizzi D., Lundman C., O'Connor E., Rosswog S., Perego A., 2021, MNRAS, 505, 2575
Glas R., Just O., Janka H.-T., Obergaulinger M., 2019, ApJ, 873, 45
Gourgoulhon E., Grandclement P., Taniguchi K., Marck J.-A., Bonazzola S., 2001, Phys. Rev. D, 63, 064029
Hammond P., Hawke I., Andersson N., 2021, Phys. Rev. D, 104, 103006
Hiscock W. A., Lindblom L., 1985, Phys. Rev. D, 31, 725
Holmbeck E. M., Frebel A., McLaughlin G. C., Mumpower M. R., Sprouse T. M., Surman R., 2019, ApJ, 881, 5
Hotokezaka K., Kiuchi K., Shibata M., Nakar E., Piran T., 2018, ApJ, 867, 95
Just O., Obergaulinger M., Janka H. T., Bauswein A., Schwarz N., 2016, ApJ, 816, L30
Just O., Goriely S., Janka H.-T., Nagataki S., Bauswein A., 2022, MNRAS, 509, 1377
Kastaun W., Galeazzi F., 2015, Phys. Rev. D, 91, 064027
Kawaguchi K., Shibata M., Tanaka M., 2020, ApJ, 889, 171
Korobkin O. et al., 2021, ApJ, 910, 116
Kullmann I., Goriely S., Just O., Ardevol-Pulpillo R., Bauswein A., Janka H.-T., 2022, MNRAS, 510, 2804
Kuroda T., Takiwaki T., Kotake K., 2016, ApJS, 222, 20
Laiu M. P., Endeve E., Chu R., Harris J. A., Messer O. B., 2021, ApJS, 253, 52
Lehner L., Liebling S. L., Palenzuela C., Caballero O., O'Connor E., Anderson M., Neilsen D., 2016, Class. Quantum Gravity, 33, 184002
Lentz E. J. et al., 2015, ApJ, 807, L31
LeVeque R. J., 1992, Numerical Methods for Conservation Laws, 2nd edn, Lectures in Mathematics. Birkhäuser, Berlin
Li X., Siegel D. M., 2021, Phys. Rev. Lett., 126, 162701
Liebendoerfer M., Whitehouse S., Fischer T., 2009, ApJ, 698, 1174
Lippuner J., Roberts L. F., 2015, ApJ, 815, 82
Lippuner J., Roberts L. F., 2017, ApJS, 233, 18
Loffler F. et al., 2012, Class. Quantum Gravity, 29, 115001
McClarren R. G., Hauck C. D., 2010, J. Comput. Phys., 229, 5597
McClarren R. G., Lowrie R. B., 2008, J. Comput. Phys., 227, 9711
McKinney J. C., Tchekhovskoy A., Sadowski A., Narayan R., 2014, MNRAS, 441, 3177
Melson T., Janka H.-T., Bollig R., Hanke F., Marek A., Müller B., 2015, ApJ, 808, L42
Metzger B. D., Fernández R., 2014, MNRAS, 441, 3444
Mezzacappa A., Endeve E., Messer O. B., Bruenn S. W., 2020, Living Reviews in Computational Astrophysics, 6, 4
Miller J. M. et al., 2019a, Phys. Rev. D, 100, 023008
Miller J. M., Ryan B. R., Dolence J. C., 2019b, ApJS, 241, 30
Most E. R., Papenfort L. J., Tootle S. D., Rezzolla L., 2021, MNRAS, 506, 3511
Most E. R. et al., 2022, MNRAS, 509, 1096
Murguia-Berthier A. et al., 2021, ApJ, 919, 95
Nagakura H., Sumiyoshi K., Yamada S., 2014, ApJS, 214, 16
Nedora V., Bernuzzi S., Radice D., Perego A., Endrizzi A., Ortiz N., 2019, ApJ, 886, L30

Nedora V., Radice D., Bernuzzi S., Perego A., Daszuta B., Endrizzi A., Prakash A., Schianchi F., 2021a, MNRAS, 506, 5908
Nedora V. et al., 2021b, ApJ, 906, 98
Neilsen D., Liebling S. L., Anderson M., Lehner L., O'Connor E., Palenzuela C., 2014, Phys. Rev. D, 89, 104029
O'Connor E., 2015, ApJS, 219, 24
O'Connor E. P., Couch S. M., 2018a, ApJ, 854, 63
O'Connor E. P., Couch S. M., 2018b, ApJ, 865, 81
O'Connor E., Ott C. D., 2010, Class. Quantum Gravity, 27, 114103
Obergaulinger M., Just O., Janka H., Aloy M., Aloy C., 2014, in Pogorelov N. V., Audit E., Zank G. P., eds, ASP Conf. Ser. Vol. 488, 8th International Conference of Numerical Modeling of Space Plasma Flows (ASTRONUM 2013). Astron. Soc. Pac., San Francisco, p. 255
Palenzuela C., Liebling S. L., Neilsen D., Lehner L., Caballero O., O'Connor E., Anderson M., 2015, Phys. Rev. D, 92, 084001
Perego A., Rosswog S., Cabezón R. M., Korobkin O., Käppeli R., Arcones A., Liebendörfer M., 2014, MNRAS, 443, 3134
Perego A., Cabezón R., Käppeli R., 2016, ApJS, 223, 22
Perego A., Yasin H., Arcones A., 2017a, J. Phys. G, 44, 084007
Perego A., Radice D., Bernuzzi S., 2017b, ApJ, 850, L37
Perego A., Bernuzzi S., Radice D., 2019, Eur. Phys. J. A, 55, 124
Perego A., Thielemann F.-K., Cescutti G., 2021, in Bambi C., Katsanevas S., Kokkotas K. D., eds, Handbook of Gravitational Wave Astronomy. Springer, New York, p. 1
Radice D., 2017, ApJ, 838, L2
Radice D., 2020, Symmetry, 12, 1249
Radice D., Rezzolla L., 2012, A&A, 547, A26
Radice D., Abdikamalov E., Rezzolla L., Ott C. D., 2013, J. Comput. Phys., 242, 648
Radice D., Rezzolla L., Galeazzi F., 2014a, Class. Quantum Gravity, 31, 075012
Radice D., Rezzolla L., Galeazzi F., 2014b, MNRAS, 437, L46
Radice D., Rezzolla L., Galeazzi F., 2015, ASP Conf. Ser. Vol. 498, High-Order Numerical-Relativity Simulations of Binary Neutron Stars. Astron. Soc. Pac., San Francisco, p. 121
Radice D., Galeazzi F., Lippuner J., Roberts L. F., Ott C. D., Rezzolla L., 2016, MNRAS, 460, 3255
Radice D., Bernuzzi S., Del Pozzo W., Roberts L. F., Ott C. D., 2017, ApJ, 842, L10
Radice D., Perego A., Bernuzzi S., Zhang B., 2018a, MNRAS, 481, 3670
Radice D., Perego A., Hotokezaka K., Fromm S. A., Bernuzzi S., Roberts L. F., 2018b, ApJ, 869, 130
Radice D., Bernuzzi S., Perego A., 2020, Ann. Rev. Nucl. Part. Sci., 70, 95
Rahman N., Just O., Janka H. T., 2019, MNRAS, 490, 3545
Reisswig C., Haas R., Ott C., Abdikamalov E., Mösta P., Pollney D., Schnetter E., 2013, Phys. Rev. D, 87, 064023
Richers S., 2020, Phys. Rev. D, 102, 083017
Richers S., Nagakura H., Ott C. D., Dolence J., Sumiyoshi K., Yamada S., 2017, ApJ, 847, 133
Richers S. A., McLaughlin G. C., Kneller J. P., Vlasenko A., 2019, Phys. Rev. D, 99, 123014
Richers S., Willcox D., Ford N., 2021, Phys. Rev. D, 104, 103023
Rider W. J., Lowrie R. B., 2002, Int. J. Numer. Methods Fluids, 40, 479
Roberts L. F., Reddy S., 2017, in Alsabti A. W., Murdin P., eds, Chapter of the forthcoming Handbook of Supernovae. Springer International Publishing AG, New York, 1605
Roberts L. F., Ott C. D., Haas R., O'Connor E. P., Diener P., Schnetter E., 2016, ApJ, 831, 98
Rosswog S., Ramirez-Ruiz E., 2002, MNRAS, 336, L7
Ruffert M., Janka H., 1998, A&A, 338, 535
Ruffert M., Janka H., Schaefer G., 1996, A&A, 311, 532
Ryan B. R., Dolence J. C., 2020, ApJ, 891, 118
Schneider A., Roberts L. F., Ott C. D., 2017, Phys. Rev. C, 96, 065802
Schnetter E., Hawley S. H., Hawke I., 2004, Class. Quantum Gravity, 21, 1465
Sekiguchi Y., 2010, Prog. Theor. Phys., 124, 331
Sekiguchi Y., Kiuchi K., Kyutoku K., Shibata M., 2011, Phys. Rev. Lett., 107, 051102
Sekiguchi Y., Kiuchi K., Kyutoku K., Shibata M., 2015, Phys. Rev. D, 91, 064059
Sekiguchi Y., Kiuchi K., Kyutoku K., Shibata M., Taniguchi K., 2016, Phys. Rev. D, 93, 124046
Shibata M., 2005, Phys. Rev. Lett., 94, 201101
Shibata M., Hotokezaka K., 2019, Ann. Rev. Nucl. Part. Sci., 69, 41
Shibata M., Kiuchi K., Sekiguchi Y.-i., Suwa Y., 2011, Prog. Theor. Phys., 125, 1255
Shibata M., Fujibayashi S., Sekiguchi Y., 2021, Phys. Rev. D, 104, 063026
Siegel D. M., 2019, Eur. Phys. J. A, 55, 203
Siegel D. M., Metzger B. D., 2017, Phys. Rev. Lett., 119, 231102
Skinner M. A., Dolence J. C., Burrows A., Radice D., Vartanyan D., 2019, ApJS, 241, 7
Thielemann F. K., Eichler M., Panov I., Wehmeyer B., 2017, Ann. Rev. Nucl. Part. Sci., 67, 253
Thornburg J., 2004, Class. Quantum Gravity, 21, 743
Thorne K. S., 1981, MNRAS, 194, 439
van Riper K., Lattimer J., 1981, ApJ, 249, 270
Vaytet N. M. H., Audit E., Dubroca B., Delahaye F., 2011, J. Quant. Spectrosc. Radiat. Transfer, 112, 1323
Vincent T., Foucart F., Duez M. D., Haas R., Kidder L. E., Pfeiffer H. P., Scheel M. A., 2020, Phys. Rev. D, 101, 044053
Weih L. R., Olivares H., Rezzolla L., 2020a, MNRAS, 495, 2285
Weih L., Gabbana A., Simeoni D., Rezzolla L., Succi S., Tripiccione R., 2020b, MNRAS, 498, 3374
Zalamea I., Beloborodov A. M., 2011, MNRAS, 410, 2302
Zappa F., Bernuzzi S., Radice D., Perego A., Dietrich T., 2018, Phys. Rev. Lett., 120, 111101
Zhu Y.-L., Perego A., McLaughlin G. C., 2016, Phys. Rev. D, 94, 105006

This paper has been typeset from a TEX/LATEX file prepared by the author.