Robustness of neutron star merger simulations to changes in neutrino transport and neutrino-matter interactions

Francois Foucart,[1] Patrick Chi-Kit Cheong,[2, 1] Matthew D. Duez,[3]
Lawrence E. Kidder,[4] Harald P. Pfeiffer,[5] and Mark A. Scheel[6]

[1]*Department of Physics & Astronomy, University of New Hampshire, 9 Library Way, Durham NH 03824, USA*
[2]*Department of Physics, University of California, Berkeley, Berkeley, CA 94720, USA*
[3]*Department of Physics & Astronomy, Washington State University, Pullman, Washington 99164, USA*
[4]*Cornell Center for Astrophysics and Planetary Science, Cornell University, Ithaca, New York, 14853, USA*
[5]*Max Planck Institute for Gravitational Physics (Albert Einstein Institute), Am Mühlenberg 1, 14476 Potsdam, Germany*
[6]*TAPIR, Walter Burke Institute for Theoretical Physics, MC 350-17,*
California Institute of Technology, Pasadena, California 91125, USA

Binary neutron star mergers play an important role in nuclear astrophysics: their gravitational wave and electromagnetic signals carry information about the equation of state of cold matter above nuclear saturation density, and they may be one of the main sources of r-process elements in the Universe. Neutrino-matter interactions during and after merger impact the properties of these electromagnetic signals, and the relative abundances of the produced r-process elements. Existing merger simulations are however limited in their ability to realistically model neutrino transport and neutrino-matter interactions. Here, we perform a comparison of the impact of the use of state-of-the art two-moment or Monte-Carlo transport schemes on the outcome of merger simulations, for a single binary neutron star system with a short-lived neutron star remnant ($(5-10)$ ms). We also investigate the use of different reaction rates in the simulations. While the best transport schemes generally agree well on the qualitative impact of neutrinos on the system, differences in the behavior of the high-density regions can significantly impact the collapse time and the properties of the hot tidal arms in this metastable merger remnant. The chosen interaction rates, transport algorithm, as well as recent improvements by Radice *et al* to the two-moment algorithms can all contribute to changes at the $(10-30)\%$ level in the global properties of the merger remnant and outflows. The limitations of previous moment schemes fixed by Radice *et al* also appear sufficient to explain the large difference that we observed in the production of heavy-lepton neutrinos in a previous comparison of Monte-Carlo and moment schemes in the context of a low mass binary neutron star system.

I. INTRODUCTION

The collisions of two neutron stars have proven themselves over the last decade to be particularly interesting events for the study of nuclear physics. Gravitational waves, gamma-ray bursts, and kilonovae all carry information about the properties of cold neutron rich matter above nuclear saturation density, thus providing us with a rare opportunity to study the high-density state of quantum chromodynamics. This was most clearly demonstrated by the multimessenger detection of GW170817 [1–14], though joint detections of gamma-ray bursts and kilonovae [15–21] could also provide useful information if our ability to model these signals was not so limited by uncertainties in numerical simulations and nuclear physics inputs. Neutron star mergers are also probably the systems with the strongest observational evidence for the production of r-process elements in the Universe [22]; yet how much these mergers practically contribute to nucleosynthesis remains very uncertain, as is the abundance patterns of the elements that they produce. These two open questions have in fact clearly overlapping uncertainties, as the radioactive decay of the ashes of the r-process is what powers kilonovae. If we could gain a proper understanding of the connection between kilonova signals and the properties of the matter ejected by mergers, and from there to the properties of the merging compact objects, we would improve our ability to understand both the equation of state of dense matter and the role of mergers in astrophysical nucleosynthesis.

Besides observational uncertainties, our ability to understand kilonova signals and r-process nucleosynthesis from mergers is limited by a number of theoretical unknowns in the properties of neutron-rich nuclei (reaction rates, opacities,...) [23, 24], as well as uncertainties in the simulation of neutron star mergers and their matter outflows. For merger simulations, the most important issues are likely the poorly resolved impact of magnetic fields on the generation of matter outflows [25], and the impact of neutrino transport and neutrino-matter interactions on the composition of these outflows [26, 27]. Here, we focus on the latter issue, which directly impacts the outcome of nucleosynthesis [28] and the color and duration of kilonovae [29].

Numerical algorithms for general relativistic neutrino transport in merger simulations have made very significant progresses over the last decade, moving from existing order-of-magnitude accurate leakage schemes [30–33] to hybridization of the leakage scheme with an approximate two-moment [26, 34] and one-moment [35] transport algorithm, standalone two-moment transport [36, 37],

two-moment transport complemented with the evolution of the neutrino number density [38], significant improvements to the methods used to take implicit timesteps in high-density regions [39], low-order Monte-Carlo transport [27], higher-order time-stepping with Monte-Carlo [40], and the potential use of Monte-Carlo transport as a closure for two-moment schemes [41, 42]. Even the best of these schemes either remain approximation to the equations of radiation transport (two-moment schemes), or can only be used at very low accuracy (Monte-Carlo). Additionally, there has been only limited attention paid to the role of the chosen interaction rates in simulations using the most modern transport scheme.

A number of comparisons of neutrino transport schemes have already been performed through direct simulations of binary neutron star mergers. In [43], we used the SpEC code to compare neutrino leakage with an early version of our two-moment transport scheme in a low-mass binary neutron star system, finding sigificant differences in the composition of the remnant and outflows between the two schemes (average Y_e rising from 0.1 to 0.2 in that system) – a result in agreement with the interpretation of the neutron-poor outflows previously seen in the first two-moment simulations of Wanajo *et al* [26]. The same binary neutron star configuration was also evolved with leakage in [44] and two-moment schemes in [45], for a system collapsing to a black hole on relatively short time scales. The collapse time itself was significantly impacted by the choice of transport scheme, with the two-moment simulation collapsing after 12 ms and the leakage simulation after 22 ms. In [38] we showed for the same system that the ad-hoc choices made for the energy spectrum of neutrinos when computing neutrino-matter interactions similarly impacted outflow composition, though mainly in the polar regions. Radice *et al* [39] used the THC code to perform comparisons of moment simulations including their one-moment scheme [35], two-moment scheme with Eddington closure, and two-moment scheme with Minerbo closure for a low-mass (non-collapsing) and a high-mass (collapsing) systems. The one-moment scheme is found to differ from the two-moment scheme by a factor of two in neutrino luminosities (still closer to the two-moment scheme than a leakage scheme likely would be), and lower electron fraction of the outflows, while the different two-moment schemes are in significantly better agreement. In [27], we compared Monte-Carlo transport with the SpEC gray two-moment scheme in a low-mass system, finding $O(10\%)$ differences in the neutrino luminosities and outflow composition (and larger errors for the heavy-lepton luminosity). We note of course that these results focused on direct comparisons of neutrino transport scheme are just a fraction of the community effort going to the study of neutron star mergers (see e.g. [46–48] and references therein).

Here, we continue this effort to address the capabilities of various neutrino transport scheme by studying a binary neutron star merger with a collapse timescale of ~ 10 ms, a case already studied with a hybrid two-moment+leakage scheme by Sekiguchi *et al* [49]. We used the same physical system in [50] to study the impact of residual eccentricity and equation of state implementation on merger results. We perform a direct comparison of the improvements made to two-moment schemes by Radice *et al* [39] with results obtained using the pre-existing SpEC two-moment algorithm [38], as well as a comparison of Monte-Carlo and two-moment transport in a collapsing binary neutron star system. We additionally consider the impact of the inclusion of different reactions in the simulations by considering simulations using NuLib rates in which we ignore pair processes for electron type neutrinos [51], and simulations using the energy-integrated rates of [31] including an approximate treatment of pair processes.

In Sec. II, we describe our initial conditions and numerical methods, with particular emphasis on the implementation within SpEC of the improvements to the two-moment scheme proposed by Radice *et al* [39]. Sec. III discuss the outcome of the simulations in terms of the merger remnant, early matter ejection, and neutrino distribution. We focus on the early evolution, up to ~ 3 ms post-merger. This allows us to study the dynamical ejecta and formation of the compact remnant. We also follow two of the simulations through collapse to a black hole.

II. METHODS

A. Initial Conditions

In all simulations presented here, we consider the merger of neutron stars with gravitational masses $1.4M_\odot$ and $1.3M_\odot$. Both stars are initially non-rotating, and start their evolution about 6 orbits before merger. We start from the exact same initial conditions as in the low eccentricity simulations of [50] ($e \sim 0.003 - 0.004$), using either the Compose [52, 53] or StellarCollapse [54] version of the SFHo equation of state. These initial conditions were generated with the Spells initial data solver [55, 56]. In [50], we showed that in the absence of neutrinos and magnetic fields, simulations using the two versions of this equation of state cannot be differentiated within estimated numerical errors. We note that the use of this equation of state results in neutron stars with circumferential radius $R = 11.9$ km for the two masses simulated. This case uses the same binary parameters system as simulation SFHo-130-140 of Sekiguchi *et al* [49].

B. GR-Hydro methods

The evolution of the metric and fluid variables follow the exact same methods as in our neutrinoless simulations of the same system [50]. These methods are described in more details in [57, 58]. To summarize, we

evolve Einstein's equations in the generalized harmonic formalism [59] on a pseudospectral grid using p-adaptive mesh refinement [60], and the fluid equations in the flux-conservative Valencia formalism [61] using high-order shock capturing methods (WENO5 reconstruction [62–64] and HLL approximate Riemann solver [65]) on a finite volume grid with fixed mesh refinement. The grid resolution on the finest grid, which covers both neutron stars, is $\Delta x \sim (160 - 200)\,\mathrm{m}$ during inspiral[1], and $\Delta x \sim 160\,\mathrm{m}$ during and after merger. During and after merger, each refinement level has 256^3 cells, with the grid spacing increasing by a factor of 2 between refinement level. The finest level thus covers a cube of side $40\,\mathrm{km}$. We use a total of 4 refinement levels in these simulations. In [50], we also presented a higher resolution simulation, which provides an estimate of numerical errors due to the coupled evolution of the metric and fluid.

C.   Neutrino transport

The main objective of this manuscript is to study the impact of neutrino transport and neutrino-matter interaction rates on the merger. We consider the following variations:

Base transport algorithm: We consider either a two-moment scheme with Minerbo closure [66, 67], or the Monte-Carlo transport scheme recently implemented in SpEC [68]. The two-moment scheme is generally expected to behave better in high-density regions (with caveats discussed below for some implementations of this scheme), but is limited by the lack of spectral information in the decoupling region and the use of approximate analytical closure for the pressure tensor in the optically thin region. The Monte-Carlo scheme, on the other hand, is expected to converge to the correct solution of Boltzmann's equations, but very slowly and with significant shot noice at the number of packets affordable in our simulations. In addition, the use of implicit Monte-Carlo in high-opacity regions should further slow down convergence of the scheme in those regions [68].

Detailed methods for two-moment schemes: Radice *et al* [39] recently developed improved methods for the use of two-moment schemes in merger simulations. These include changes to the treatment of numerical fluxes in high-density regions, as well as an improved treatment of the implicit solve required to evolve the moment equations in time. The preexisting SpEC methods, on the other hand, used an approximate implicit solve, linearizing the problem around $E = F^i = 0$ (with E, F^i the neutrino energy and flux density in the laboratory frame)

and ignoring terms quadratic and higher in the fluid velocity. In this manuscript, we use simulations implementing both the modified fluxes in high-density regions, and the improved treatment of the implicit solve from [39]. We will refer to these methods as "SpEC" and "Radice", although all simulations described in this manuscript are performed with the SpEC code. Our implementation of the methods of [39] in SpEC, which differ in minor ways from the original, are described in more details in Sec. II E. We note that an intermediate model accounting for the fluid velocity at all orders was implemented in the BAM code by Schianchi *et al* [69]; that method is not tested in the simulations presented here.

Neutrino-matter interactions: The simulations using Monte-Carlo methods use a NuLib [51] table with 16 energy groups, logarithmically spaced up to $528\,\mathrm{MeV}$. The table includes ν_e absorption on neutrons, $\bar{\nu}_e$ absorption on protons, as well as production of $\nu_\mu \bar{\nu}_\mu$ and $\nu_\tau \bar{\nu}_\tau$ pairs from $e^+ e^-$ annihilation and nucleon-nucleon Bremsstrahlung. The pair production channels are not included for electron type neutrinos. All inverse reactions are also included, in such a way that the emissivity η and absorption opacity κ_a satisfy Kirchoff's law $\eta/\kappa_a = B_\nu$, with B_ν the black-body distribution function of neutrinos in equilibrium with the fluid, integrated over the relevant energy bin. We also include in the table elastic scattering on protons, neutrons, $^4_2\mathrm{He}$, and heavy nuclei. For the two-moment method, we first consider a gray version of the NuLib table, i.e. by taking a weighted average of the opacities that assumes an equilibrium distribution of neutrinos, then correcting for the estimated average energy of the neutrinos as described below. We also consider approximate gray emission and absorption rates calculated on the fly, using the rates of [31]. In this latter version of the interaction rates, we use the same reactions as for the NuLib table, but now adding the contribution from pair production to the electron type neutrinos and additionally considering plasmon decay $\gamma\gamma \leftrightarrow \nu\bar{\nu}$. The two sets of reaction rates also make different approximations for the spectrum of neutrinos and blocking factors, even for reactions included in both sets of interaction rates.

D.   Simulations performed

The specific combinations of methods used in this manuscript include:

1. Monte-Carlo transport with an energy-dependent NuLib table and a maximum of 10^8 packets per species (MC-HR)

2. Monte-Carlo transport with an energy-dependent NuLib table and a maximum of 4×10^7 packets per species (MC-LR; MC-LR and MC-HR start from the same pre-merger evolution, as before merger we do not need many neutrino packets to capture the very small impact of neutrinos on the evolution)

---

[1] The grid spacing changes as our grid contracts to follow the motion of the neutron stars; we interpolate the solution on a new grid with $\Delta x \sim 200\,\mathrm{m}$ every time the grid contracts by a factor of 0.8.

ID	Transport	Fluxes	Implicit Solve	Rates
MC-HR	MC (HR)	NA	NA	NuLib
MC-LR	MC (LR)	NA	NA	NuLib
M1-Radice	M1	Radice	Radice	On-the-fly
M1-SpEC	M1	SpEC	SpEC	On-the-fly
M1-NuLib	M1	SpEC	SpEC	GrayNuLib
Hydro	None	NA	NA	NA

TABLE I. List of simulations performed for this manuscript. We list the transport scheme used (M1 or MC), the method used to calculate interaction rates, and, for M1 simulations, the method used to calculate numerical fluxes and to perform the implicit solve. The "Hydro" simulation is performed without neutrino transport (simulation "SC-MR" from [50])

3. Two-moment transport using Radice's fluxes and implicit solve, and on-the-fly calculation of interaction rates (M1-Radice).

4. Two-moment transport with SpEC's fluxes and implicit solve, as well as on-the-fly calculation of interaction rates (M1-SpEC).

5. Two-moment transport using SpEC's fluxes and implicit solve, and the gray NuLib table (M1-NuLib)

6. No neutrinos at all. This is simulation "SC-MR" from [50], renamed "Hydro" here, and thus not a new simulation in this manuscript.

These simulations are also listed on Table I.

All simulations are performed to at least $3\,\mathrm{ms}$ post-merger. The MC-HR simulation as well as the M1-NuLib simulation are followed up to the collapse of the remnant to a black hole.

E. Implementation details for the two-moment scheme

Most of the simulations presented in this manuscript use some variations of the two-moment scheme in which we evolve the energy density, momentum density, and number density of neutrinos in the laboratory frame. This formalism was first developed in [66, 67], and its SpEC implementation is described in [36, 38]. Improvements to the two-moment methods, and especially to the treatment of the neutrino spectrum and implicit time-stepping, were presented in [39]. We follow the methods of [38] for the 'SpEC' simulations and of [39] for the 'Radice' simulations; in this section, we will present a brief overview of the moment methods, then focus on implementation details of the 'Radice' algorithm within the SpEC code.

The main equations evolved in the two-moment scheme

are

$$\partial_t \tilde{E} + \partial_i \left(\alpha \tilde{F}^i - \beta^i \tilde{E} \right) = \alpha \tilde{P}^{ij} K_{ij} - \tilde{F}^j \partial_j \alpha$$
$$- \alpha \tilde{S}^\mu n_\mu \qquad (1)$$
$$\partial_t \tilde{F}_j + \partial_i \left(\alpha \tilde{P}^i_j - \beta^i \tilde{F}_j \right) = -\tilde{E} \partial_j \alpha + \tilde{F}_k \partial_j \beta^k$$
$$+ \frac{\alpha}{2} \tilde{P}^{ik} \partial_j \gamma_{ik} + \alpha \tilde{S}^\mu \gamma_{j\mu} \quad (2)$$
$$\partial_t \tilde{N} + \partial_j \left(\alpha \sqrt{\gamma} \mathcal{F}^j - \beta^j \tilde{N} \right) = \alpha \sqrt{\gamma} C_{(0)} \qquad (3)$$

with E the neutrino energy density, F^i the neutrino momentum density, N the neutrino number density, and tilde quantities denoting densitized variables (i.e. $\tilde{E} = \sqrt{\gamma} E$, with γ the determinant of the spatial 3-metric γ_{ij}). The metric is decomposed according to

$$ds^2 = -\alpha^2 dt^2 + \gamma_{ij} \left(\beta^i dt + dx^i \right) \left(\beta^j dt + dx^j \right) \qquad (4)$$

with α the lapse and β^i the shift vector. The term S^α includes coupling between the neutrinos and the fluid. Its exact form will depend on the interactions included in the simulations. In our current simulations, we account for an isotropic emissivity in the fluid frame η, an absorption opacity κ_A, and an isotropic elastic scattering opacity κ_S. Then

$$S^\alpha = \eta u^\mu - \kappa_A J u^\mu - (\kappa_A + \kappa_S) H^\mu \qquad (5)$$

with u^μ the 4-velocity of the fluid, J the neutrino energy density in the fluid frame, and H^μ the neutrino momentum density in the fluid frame. Finally, $C_{(0)}$ is given by

$$C_{(0)} = \eta_N - \kappa_N \frac{J}{\langle \nu \rangle} \qquad (6)$$

with η_N, κ_N the number emissivity and absorption opacity, and ν the average neutrino energy. We use

$$\langle \nu \rangle = \frac{W(E - F_i v^i)}{N} \qquad (7)$$

with v^i the spatial 3-velocity of the fluid. The main approximations of a gray two-moment scheme comes in the analytical form used to estimate the pressure tensor P^{ij} and the number flux $\mathcal{F}^j$, as well as in the assumptions made about the neutrino spectrum that go into the calculations of the energy-averaged $\eta, \eta_N, \kappa_A, \kappa_S, \kappa_N$. We refer the reader to [38] for more details about our standard choices for these calculations and a discussion of the impact of these approximations on the outcome of the radiation transport problem in neutron star mergers.

Numerical fluxes: The first important change when going from the standard SpEC algorithm to the methods of [39] is the way the flux terms on the left-hand-side of the evolution equations are computed, i.e. all the divergence terms of the form $\partial_i f^i$. In low-density regions, these fluxes can be treated with the same high-order shock-capturing methods as the similar terms coming into the evolution of the fluid variables. In high

optical-depth regions, however, the numerical diffusion that is naturally introduced by these methods can become larger than the physical diffusion rates of neutrinos through dense matter. Then, using these methods can lead to a large overestimate of the neutrino diffusion rate. In previous SpEC simulations, we used a complex procedure described in [36, 70] to drive the flux in the evolution of $\tilde{E}$ to the desired solution of the diffusion equation when $\kappa_A \Delta x \gg 1$. Radice *et al* [39] however pointed out that a simpler and more reliable methods is to use a weighted average of a 'low-order' flux taken from shock capturing methods (Lax-Friedrich flux), and a 'high-order' flux that does not include diffusive shock-capturing methods (second order finite differencing). The detailed calculations of the fluxes in SpEC exactly matches that of [39].

Time stepping method: The second modification to our algorithm is in the method used to capture the neutrino-matter coupling (i.e. the source terms S^α). In the two moment equations, most terms in the evolution of $(\tilde{E}, \tilde{F}^i)$ are treated explicitly (in SpEC, using a second-order Runge-Kutta algorithm). The source term $C_{(0)}$ is treated implicitly, but the evolution of $\tilde{N}$ can be trivially solved as a post-processing steps after we obtain the values of $(\tilde{E}, \tilde{F}^i)$ at the end of a time step. The terms involving $\tilde{S}^\alpha$ are more complex to handle. In high-opacity regions, they need to be treated implicitly, due to the very short timescale for neutrino-matter equilibration. We do this neutrino species by neutrino species, given a guess for the fluid properties at the end of a time step (see below for the method used to guess at the fluid temperature and electron fraction; the other properties of the fluid are assumed constant during a time step in this calculation). Practically, this requires solving an equation of the type

$$U^{n+1} = U^n + S^n_{\exp}dt + S^{n+1}_{\imp}dt \tag{8}$$

with $U = (\tilde{E}, \tilde{F}^i)$ the evolved variables, $S_{\exp}$ the explicit source terms and fluxes, and $S_{\imp}$ the implicit source terms. The superscripts n and $n+1$ refer to values at the beginning and end of a time step (or half step for the intermediate step of the Runge-Kutta algorithm). In previous simulations, we used a linearization of S^{n+1} around $E = F^i = 0$ to solve this equation, neglecting terms quadratic in the fluid velocity. Radice *et al* [39] instead derived the full jacobian matrix $\partial(S_{\imp})/\partial U$ for generic U. In simulation M1-Radice, we solve the implicit problem using the dog-legged method and the jacobian matrix of [39]. We note that this does not exactly match the method of [39]. This is in part because we use a different time stepping algorithm, and in part because we make a slightly different approximation in regions of high optical depth. In [39], the authors note that in these regions the solver can fail, and that this issue can be resolved by calculating the jacobian under the assumption of an 'optically thick' closure for the pressure tensor. In our simulation, we do not wait for a failure and automatically assume the optically thick closure if the equilibration timescale $\tau_{\eq} = 1/\sqrt{\kappa_A(\kappa_A + \kappa_S)} \leq 0.1\alpha dt$, with dt

the time step. We also treat all source terms explicitly if $\tau_{\eq} \geq 10\alpha dt$. Finally, we use as initial guess for the solution

$$U^{n+1}_{\guess} = U^{*,n}(1 - f) + U_{\eq}f \tag{9}$$

with $U_{\eq}$ representing the values of the evolved variables for neutrinos in equilibrium with the fluid, discussed below, $U^{*,n}$ the value of the evolved variables *after application of the explicit source terms*, and

$$f = \exp\left(-\tau_{\eq}/(\alpha dt)\right). \tag{10}$$

Guess for temperature and electron fraction in high-opacity regions: in dense, hot regions, the temperature T and electron fraction Y_e of the fluid may be stiffly coupled to the neutrinos. Calculating interaction rates using the fluid variables at the beginning of a time step is then occasionally unstable. In simulation M1-Radice, we updated our methods to guess at the fluid properties at the end of a step following methods similar to Radice al [39]. Specifically, we calculate the total (fluid plus neutrinos) specific internal energy $\epsilon_{\tot}$ and total (fluid and neutrinos) lepton number fraction $Y_{e,\tot}$ before performing an implicit solve (but after adding the explicit components of the right hand side to the neutrino moments). We then solve for the values of T, Y_e such that

$$\epsilon_{\tot} = \epsilon_{\fluid}(T, Y_e) + \sum_i \frac{\eta_{\nu_i}(T, Y_e)}{\rho \kappa_{\nu_i}(T, Y_e)} \tag{11}$$

$$Y_{e,\tot} = Y_e + \frac{\eta^N_{\nu_e}(T, Y_e)}{\kappa^N_{\nu_e}(T, Y_e)} \frac{m_b}{\rho} - \frac{\eta^N_{\bar{\nu}_e}(T, Y_e)}{\kappa^N_{\bar{\nu}_e}(T, Y_e)} \frac{m_b}{\rho} \tag{12}$$

with m_b the mass of a baryon, ρ the baryon number density, and the sum over i being over all neutrino species. Practically, this comes down to requiring that neutrinos reach equilibrium with the fluid, and that we satisfy conservation of energy and lepton number. Let us now call $(T_0, Y_{e,0})$ the temperature and electron fraction before the implicit step, and $(T_{\eq}, Y_{e,\eq})$ the solutions of the system of equations above. When calculating neutrino-matter interactions, we then use as our guess for T and Y_e

$$T = fT_0 + (1 - f)T_{\eq}; \; Y_e = fY_{e,0} + (1 - f)Y_{e,\eq} \tag{13}$$

and

$$f = \max\left(0, \min\left[1, 2(\alpha dt - \max_i\left(\frac{1}{\sqrt{\kappa_{\nu_i,a}(\kappa_{\nu_i,a} + \kappa_{\nu_i,s})}}\right)\right]\right) \tag{14}$$

with α the lapse and dt the time step. Practically, $f \approx 1$ when the equilibration timescale between neutrinos and the fluid is long compared to the time step, and $f \approx 0$ if that timescale is short.

Correction of interaction rates in optically thin regions: In optically thin regions, it is quite common for the average energy of neutrinos to be sigificantly higher than the average energy that neutrinos would have if they were in

equilibrium with the fluid. Given the strong dependence of interaction rates on neutrino energy, this causes uncorrected gray schemes to underestimate interaction rates in these regions. In SpEC, we calculate an estimated neutrino temperature T_ν as [38]

$$T_\nu = \frac{F_2(\eta_\nu)}{F_3(\eta_\nu)} W \frac{E - F_i v^i}{N} \tag{15}$$

with W the Lorentz factor, N the neutrino number density (in the laboratory frame),

$$F_k(\eta) = \int_0^\infty dx \frac{x^k}{1 + exp(x - \eta)} \tag{16}$$

the Fermi integrals, $\eta = \mu/T$, T the fluid temperature and μ the chemical potential. We then correct opacities by a factor of $(T_\nu/T)^2$, following the scaling of opacities with neutrino temperature for $T_\nu \approx T$. With the more advanced implicit solver of Radice, this is occationally unstable and we apply a few corrections. First, we do not correct interaction rates when $T_\nu < T$. Second, we take advantage of the fact that neutrinos quickly equilibrate with the fluid when $\kappa_a dt \gtrsim 1$ to ignore corrections in high-opacity regions. Practically, we take

$$a = \frac{T_\nu^2}{T^2}(1 - f) + f; \ f = \frac{\alpha dt \kappa_a}{\alpha dt \kappa_a + 1} \tag{17}$$

with a the multiplicative factor applied to opacities. We note that $a \to 1$ when $\kappa_a \alpha dt \gg 1$. Finally, we follow Radice $et\ al$ [39] and apply the correction to both emissivities and opacities, rather than solely to opacities, in order to still drive the neutrino distribution function to the correct equilibrium energy density. While in this manuscript the changes are only applied to simulation M1-Radice, we believe that they would likely be beneficial to simulations using our old algorithm as well. We also note that a more advanced algorithm to solve for T_ν without explicitly relying on the approximate scaling $\kappa_a \propto T_\nu^2$ has since been developed in Andresen $et\ al$ [71], but is not currently implemented in SpEC.

We note that these are the only modifications made to the SpEC code; in particular, we use the same methods to estimate the shape of the neutrino energy spectrum and number flux density $\mathcal{F}^j$ as in [38], even though these methods differ from the default choices made in THC.

F. Diffusion test

To illustrate the impact of the changes made to the ability of the SpEC code to handle regions with high scattering opacity, we reproduce the diffusion test of Radice $et\ al$ [39]. In this test, a pulse of radiation $E = e^{-9x^2}$ is evolved in a fluid moving at velocity $v^x = 0.5$ and with scattering opacity $\kappa_s = 10^3$. We evolve the pulse over a time $\Delta t = 2$, in slab geometry, with grid spacing $\Delta x = 0.0098$ and Courant factor of 0.25. Each grid
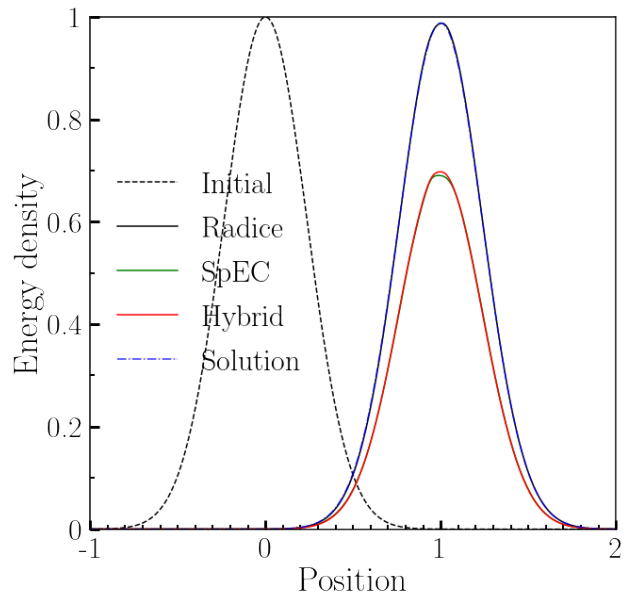

FIG. 1. Diffusion of a pulse of radiation through a medium of high scattering opacity $\kappa_s = 10^3$. We show the initial pulse, as well as results at $t = 2$ for the analytical solution of the diffusion equation, the solution obtained with the 'Radice' methods (which lies exactly underneath the analytical results), the solution obtained with the 'SpEC' methods, and a hybrid algorithm using the 'Radice' fluxes and the 'SpEC' implicit solve. The last two provide nearly identical (but erroneous) results.

cell thus has an optical depth $\tau = 9.8$. The results of the test are shown in Fig. 1. We see that the M1-Radice code correctly reproduces the expected solution. The two simulations using the linearization of the sources around $E = F^i = 0$ and only including $O(v/c)$ terms in the Jacobian do not, regardless of the method used to calculate fluxes in high-opacity regions. As opposed to Radice et al [39], we find that the energy of the pulse decreases when using the approximate implicit solve, though the results are otherwise equally problematic. In Schianchi $et\ al$ [69], the same test was performed using all velocity terms in the linearization of the implicit problems, in which case they find much better agreement with the analytical solution. It is thus likely that the approximations made in the linearization of the implicit problem in the original SpEC code are responsible for its issues with this diffusion test.

III. SIMULATION RESULTS

A. Merger dynamics

Before contact, the various simulations presented here behave nearly identically, and follow largely the same evolution as the simulations without neutrino transport described in [50]. Neutrino emission mostly comes from
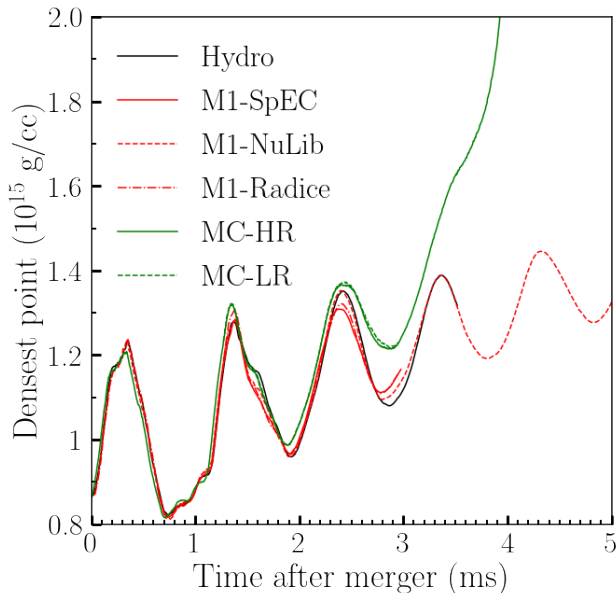
FIG. 2. Maximum value of the baryon density on the grid as a function of time, for all 6 simulations.

Sim	M_{disk}	$\langle T_{disk}\rangle$	$\langle Y_{e,disk}\rangle$	$\langle T_{rem}\rangle$	$\langle Y_{e,rem}\rangle$
	$M_\odot$	MeV		MeV	
MC-HR	0.054	10.8	0.126	28.2	0.066
MC-LR	0.056	11.0	0.131	29.0	0.068
M1-Radice	0.055	10.9	0.133	23.9	0.062
M1-SpEC	0.050	10.1	0.118	25.6	0.053
M1-NuLib	0.045	10.2	0.145	24.8	0.064
Hydro	0.048	11.9	0.046	26.3	0.057

TABLE II. Remnant properties 2.5 ms post-merger. We show the mass below 10^{13} g/cc (*disk mass*), as well as the average temperature and electron fraction of that low-density matter, and the average temperature and electron fraction of the entire remnant.

the surface of the star, which is heated by numerical dissipation. One millisecond before merger, we observe a luminosity of $\sim 2 \times 10^{51}$ erg/s for each species of neutrinos, with only $\sim 20\%$ variations between simulations and neutrino species. As a result, the neutron stars are very slightly cooler and more neutron rich in simulations with neutrino transport than in simulations without transport: the mass weighted average temperature is $\sim 1\%$ lower in most of our simulations than in [50], and the mass weighted electron fraction $\sim 0.2\%$ lower. The exceptions are simulation M1-Radice, for which the temperature is nearly identical to that of the simulation without transport, and the MC simulation, for which both temperature and electron fraction are nearly identical to the fluid-only evolution (despite, in both cases, similar neutrino luminosities). As these effects are small, resolution dependent, and do not impact the dynamics of the binary, we do not investigate these differences further.

There is no clear, local definition of the merger time for binary neutron star mergers. In this manuscript, as in [50], we define as *merger* the time at which the maximum baryon density on the grid increases to 3% above its initial value. This happens soon after contact, as the two neutron star cores collide. We follow all simulations until at least 3.0 ms post-merger, and additionally continue MC-HR and M1-NuLib until collapse of the remnant to a black hole.

The most visible difference between the various simulations is the evolution of the maximum baryon density and eventual collapse of the remnant to a black hole, shown on Fig. 2. The Monte-Carlo simulations exhibit higher maximum baryon density, with the differences remaining small up to ~ 2.5 ms post-merger. As the remnant is close to collapse, this is however enough to lead

to a divergence of the evolution once the MC simulation reaches high enough densities to trigger collapse. As a result, we observe noticeable difference in the collapse time. The MC-HR simulation collapses after 4 ms. The two-moment and pure hydrodynamics simulations are longer lived. To estimate the difference in collapse time, we followed the M1-NuLib simulation up to black hole formation, and find that collapse happens 8.5 ms post-merger for that simulation. That simulation was chosen as the most likely to collapse quickly among the two-moment simulations (it has slightly higher maximum density than the others, as shown on Fig. 2); this is thus likely to be a lower bound on the difference in collapse time between Monte-Carlo and the two-moment methods. The maximum density of the MC-LR simulation is close enough to that of MC-HR that the number of packets in the MC simulation does not appear to be a driving factor for these differences. We note however that the collapse time in such a system is very sensitive to small variations in the evolution (and would also be impacted e.g. by magnetic fields, which are not included here, and by numerical resolution). These stark differences in collapse time are due to very small changes in the evolution of the system in the few milliseconds following merger, and likely particularly important for the binary evolved here, as it sits right at the threshold of collapse at merger. This result is quite similar to the different collapse times observed between leakage and two-moment schemes by Curtis *et al* [45]. We note that the system evolved in this work was alsoe simulated by Sekiguchi *et al* [49] using their mixed two-moment/leakage transport scheme. They find a collapse time of ~ 10 ms, fairly close to our two-moment results[2]

Whether the MC or M1 scheme is most accurate here is a difficult question to answer. On the one hand, the MC scheme properly takes into account the full energy spectrum of neutrinos in dense region, which could lead it to

_____

[2] The exact collapse time is not provided for this configuration in Sekiguchi *et al* [49], it is ~ 10, ms for a range of simulations with the same total mass and SFHo equation of state, but different mass ratios.

better capture transport effects in the remnant (specifically, to better capture the behavior of low-energy neutrinos). On the other hand, for monochromatic radiation, the M1 schemes (and particularly M1-Radice) are more accurate in their treatment of high-opacity regions, where the MC scheme has to rely on a low-accuracy implementation of implicit Monte-Carlo methods and noisy estimates of the radiation pressure. What is clear is that the choice of transport scheme – leakage, two-moment, or Monte-Carlo – can have a significant impact on the collapse time of such metastable remnants.

The overall properties of the remnant between merger and collapse (specifically, $2.5\,\mathrm{ms}$ post-merger) are summarized on Table II, as well as through visualizations of horizontal and vertical slices through the remnant in Figs 3-6. The bulk of the remnant has an average temperature of $(24-29)\,\mathrm{MeV}$ and an average electron fraction $(0.053-0.068)$. The main differences in temperature are between the rapidly collapsing MC simulations and the other configurations, which makes sense if the temperature of the high-density core of the remnant increases as it contracts. For the electron fraction, the main outliers are the M1-SpEC simulation, and the Hydro simulation, which obviously ignores weak interactions. The M1-SpEC simulation has clear issues with the evolution of Y_e in the densest region, as the star neutronizes despite the neutrino luminosity being dominated by $\bar{\nu}_e$. This issue is not shared by any other simulation, including the M1-NuLib simulation which uses the same numerical methods in high-density regions, or previous simulations using the M1-SpEC methods. We will return to this issue in our discussion of the neutrinos within the remnant, but note here that this due to a failure of the M1-SpEC algorithm to satisfy lepton number conservation for a short period of time around merger. The M1-Radice and MC simulations, which should be considered as the 'best' models for radiation transport in this work, are in much better agreement with $\langle Y_e \rangle = (0.062-0.068)$. We note that both the transport algorithms (M1-SpEC, M1-Radice, MC) and the chosen interaction rates (M1-NuLib vs M1-SpEC) appear to have an impact on our results here. Disentangling the impact of any specific interaction rate on these global quantities is however difficult in this nonlinear system in which only the sum of many energy-integrated interaction rates are used in the simulations. Here, we limit ourselves to assess the level at which these changes impact the results; a more detailed study of the impact of each assumption in the reaction rates would require more systematically modifying individual contributions to the interaction rates.

We can also look more specifically at the low-density matter that may eventually form a disk. Here, we use the arbitrary cutoff $\rho \leq 10^{13}\,\mathrm{g/cc}$ to define the "disk" material, noting that the "disk" material remains difficult at this point to distinguish from the surface of the remnant. We see in Table II that the amount of matter in the low-density regions is $(0.045-0.056)M_\odot$, with excellent agreement between the M1-Radice and MC methods

which find $(0.054-0.056)M_\odot$. The same is true for the temperature and electron fraction of the low-density matter, with all simulations showing a range of temperature $(10.1-11.9)\,\mathrm{MeV}$ but the M1-Radice and MC simulations a smaller range $(10.8-11.0)\,\mathrm{MeV}$; and all simulations with neutrino transport showing a range of electron fraction $(0.118-0.145)$ but the M1-Radice and MC simulations showing only a range $(0.126-0.133)$. Unsurprisingly, the Hydro simulation remains at a much lower electron fraction $Y_e = 0.046$ in the low-density regions. It is thus clear that all transport schemes used here capture the dominant effect of the neutrinos on the low-density matter, but with much smaller variations between our best radiation schemes than when we compared those best schemes to the more approximate M1 schemes. The chosen reaction rates additionally have a clear impact on both the mass and composition of the matter below $10^{13}\,\mathrm{g/cc}$ (comparing M1-SpEC and M1-NuLib in the table), with a stronger protonization of the low-density matter in the M1-NuLib simulation.

At even lower densities ($\rho \sim 10^{10-11}\,\mathrm{g/cc}$), most of the matter is in shocked tidal arms (see on Figs. 3-4). The tidal arms will persist at least until collapse to a black hole. Fig. 3 shows a horizontal slice through the remnant $2.5\,\mathrm{ms}$ post-merger for the M1 and Hydro simulations, while Fig. 4 compares our best models M1-Radice and MC-HR, as well as the MC-LR simulation. We note that slight differences in the exact time at which snapshots were available leads to differences in the phase of the rotation of the remnant, which should be ignored in comparisons. At a qualitative level, we see broad agreements in the density and temperature of the remnants. The electron fraction is first impacted by the inclusion or not of neutrinos, with noticeable differences in the electron fraction of lower-density regions between different M1 simulations. These differences are particularly notable in the hot tidal arms, where they are $O(0.1)$. They are present in the temperature as well, but harder to see given the range of temperature that these plots must cover. We can for example look at the dominant tidal arm (above and right of the remnant on the figures) as it crosses the positive x-axis. There, it has temperature $(3-4)\,\mathrm{MeV}$ in the M1-Radice simulation, $(4-5)\,\mathrm{MeV}$ in the M1-SpEC and M1-NuLib simulation, $(6-7)\,\mathrm{MeV}$ in the MC simulations, and $(6-8)\,\mathrm{MeV}$ in the Hydro simulation – much more significant variations than in higher-density regions.

We can also take advantage of Fig. 4 to note that while noise due to the use of an MC algorithm is clearly visible in the electron fraction, the average properties of the matter in the MC-HR and MC-LR simulations are in much better agreement than between the M1-Radice and MC-HR simulations, indicating that MC noise does not have a significant impact on our final results.

Taking both figures together, and focusing on the hot tidal arm region where differences are most notable, we see that in that region there is good agreement between the M1-NuLib and MC simulations, as well as between
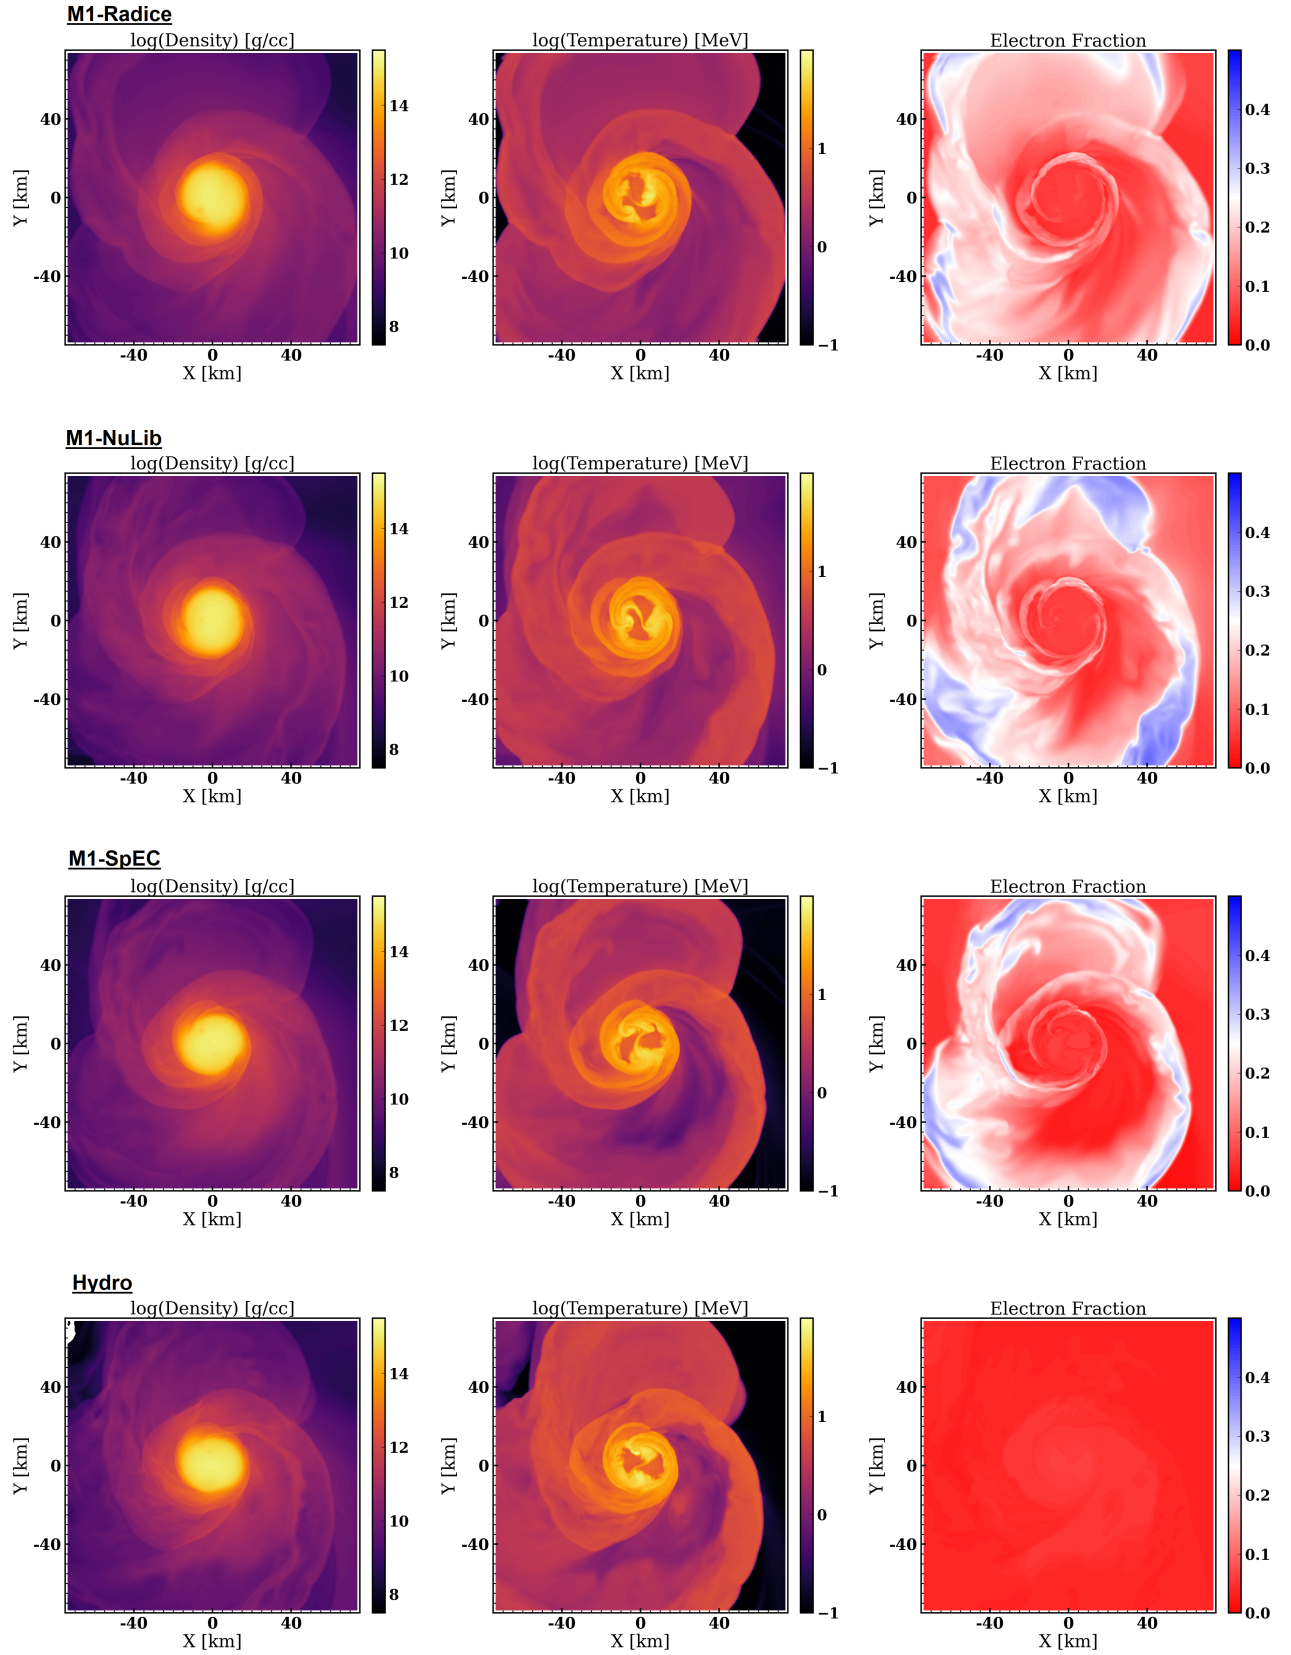
FIG. 3. Horizontal slice through the merger remant 2.5 ms post-merger. We show the baryon density (left), temperature (middle) and electron fraction (right) for simulations M1-Radice (top), M1-NuLib (2nd row), M1-SpEC (3rd row) and Hydro (bottom).
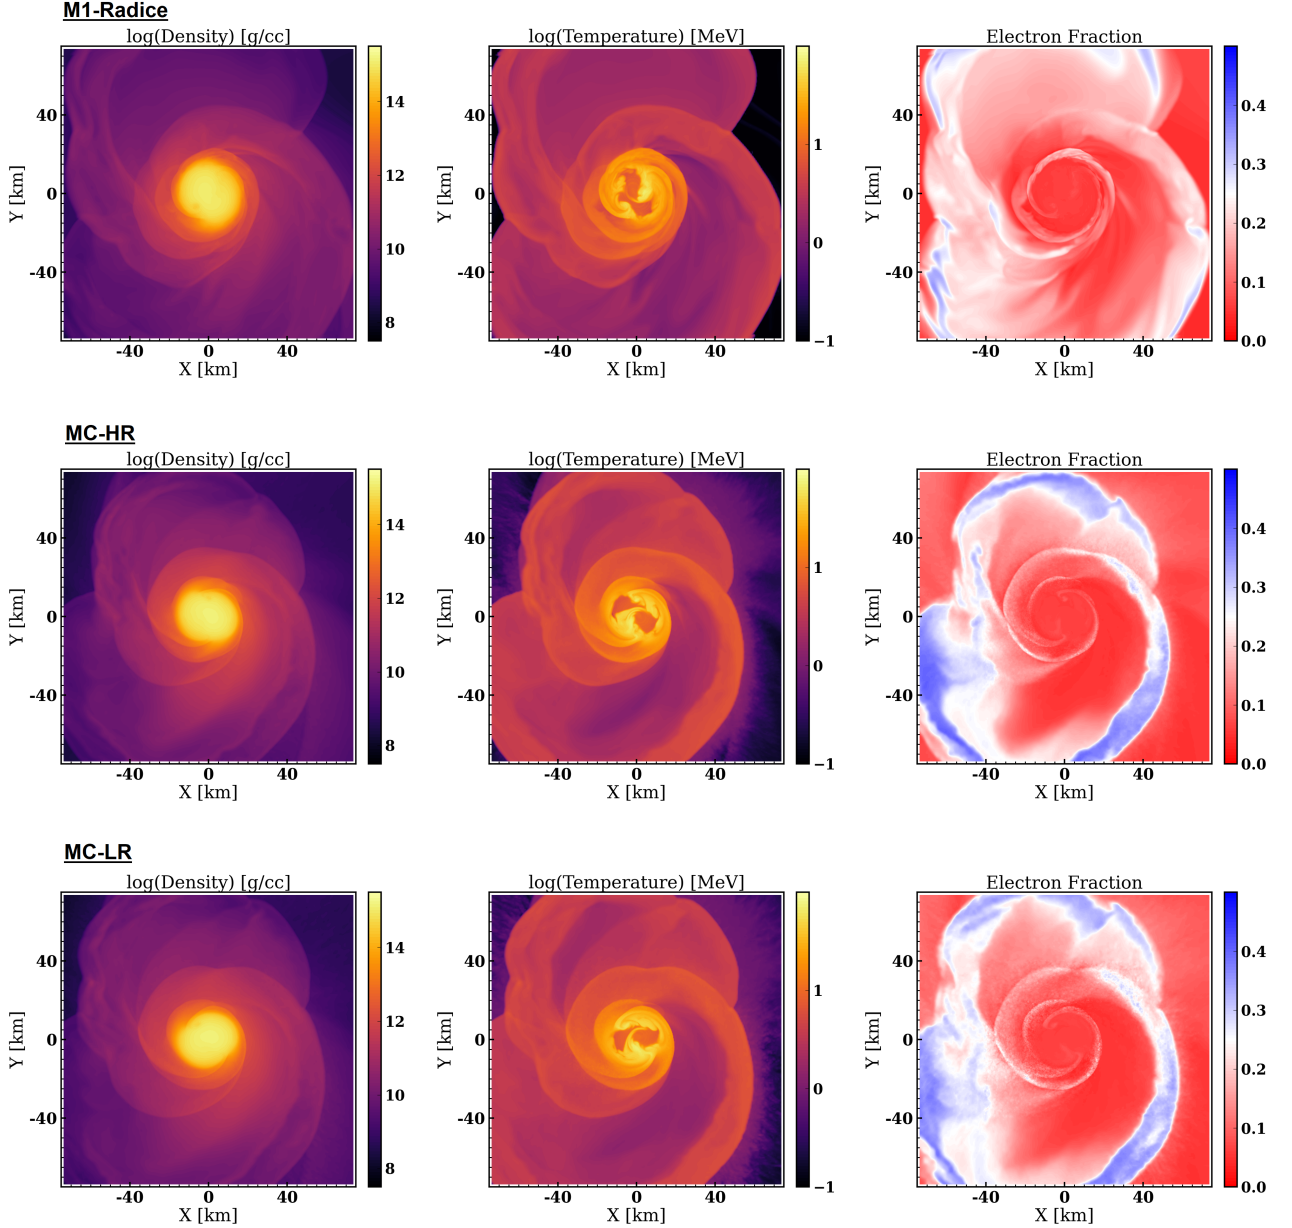
FIG. 4. Same as Fig. 3 but for simulations M1-Radice (top), MC-HR (middle) and MC-LR (bottom).

the M1-Radice and M1-SpEC simulations. This hints that the included reaction rates may have a more significant impact on the evolution of the electron fraction in those regions than the exact transport algorithm used. For simulations with neutrinos, there is a correlation between higher temperature and higher electron fraction. This is expected, as the fluid is driven towards a higher equilibrium Y_e at higher temperatures. What drives differences in temperature is harder to ascertain, though it would be reasonable to guess that both the oscillations of the central compact object, which drive the production of tidal arms, and neutrino emission, which cools the tidal arms, play a role in setting the temperature of these

tidal arms.

Figs. 5-6, which show vertical slices through the post-merger remant for the same snapshots, show similar features. The phase of the evolution of the post-merger remnant is a confusing factor here when comparing density and temperature, yet we see broadly similar features in all simulations. The electron fraction, once again, shows the most significant differences – especially in the tidal arms and polar regions. As for the horizontal slices, the two MC simulations are in good agreement with each other outside of the visible sampling noise, supporting the idea that the number of MC packets is not a critical source of error here. The MC simulations are once more
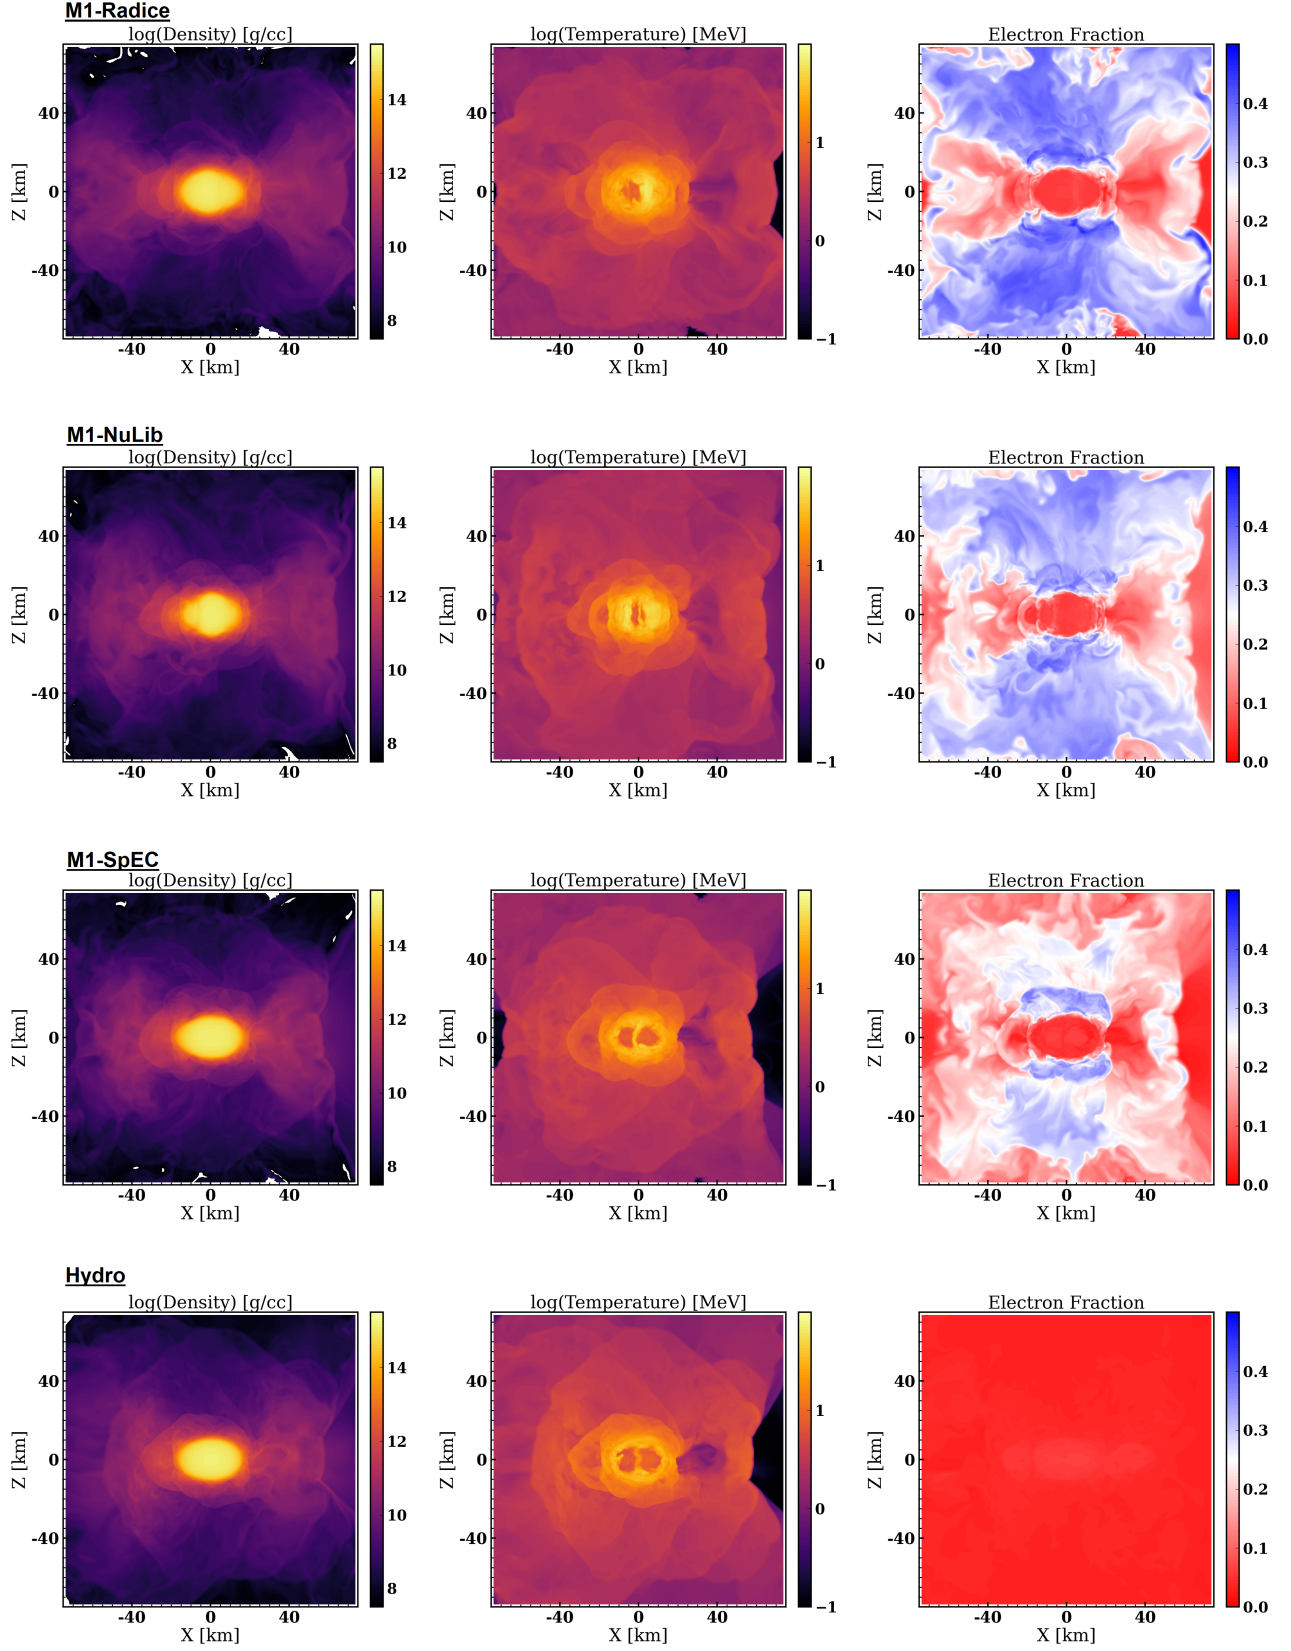
FIG. 5. Vertical slice through the merger remant 2.5 ms post-merger. We show the baryon density (left), temperature (middle) and electron fraction (right) for simulations M1-Radice (top), M1-NuLib (2nd row), M1-SpEC (3rd row) and Hydro (bottom).
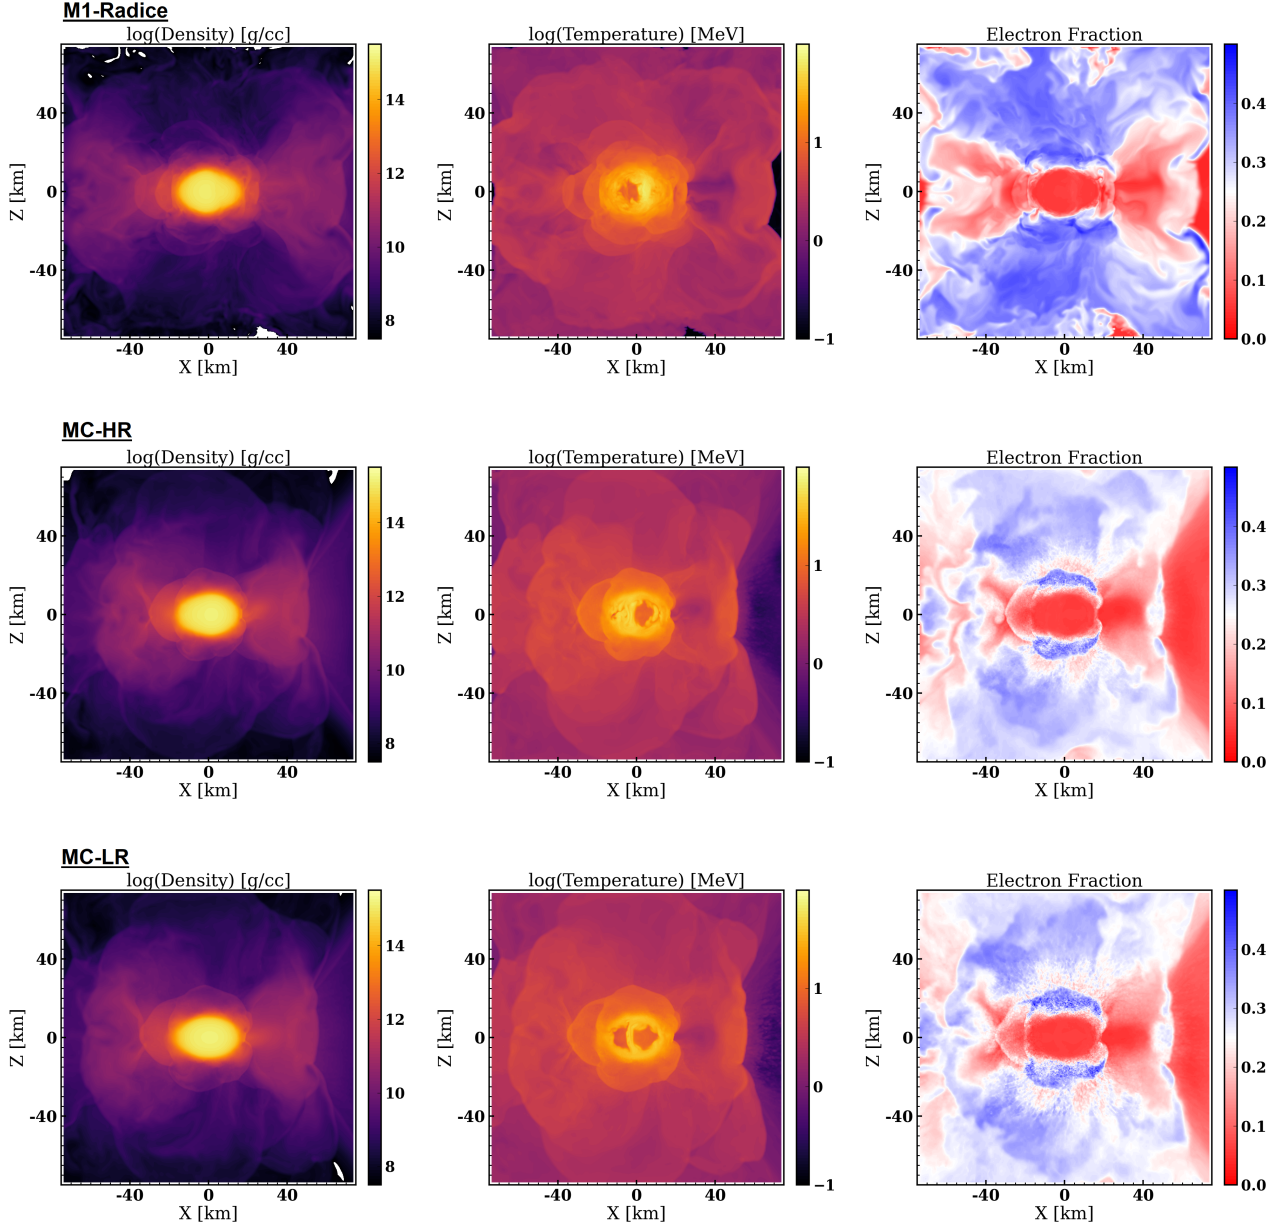
FIG. 6. Same as Fig. 5 but for simulations M1-Radice (top), MC-HR (middle) and MC-LR (bottom).

closest to the M1-NuLib simulation, indicating again that in low-density regions the chosen interaction rates have a more significant impact on the electron fraction than the transport algorithm.

Overall, we thus find that very early in the evolution of the post-merger remnant, the chosen transport algorithm (MC vs M1-NuLib; M1-Radice vs M1-SpEC) causes $\sim (20-30)\%$ relative changes in the average properties of the compact remnant (temperature, electron fraction). MC and M1 simulations in addition differ in the evolution of the densest region towards collapse. In the lower density regions close to the remnant, which are probed by our measurements of average temperature and

electron fraction in the "disk", there is great agreement between the two best transport algorithms (M1-Radice, MC) – as might be expected from regions where the decoupling of the neutrinos from the fluid is the driving factor for the evolution of the matter. In the tidal arms and polar regions, where neutrinos are farther from equilibrium, the interaction rates used in the simulations are the most important driver of differences in composition between simulations, and significant differences (factor of 2) are observed in the temperature of the tidal arms. The number of packets in the MC simulations does not seem to be a significant source of uncertainty, though we caution that this does not guarantee that MC errors are

Sim	$M_{\mathrm{ej}}(rp)$	$\langle v_{\mathrm{ej}}(rp) \rangle$	$\langle Y_e \rangle$
	$M_\odot$	c	
MC-HR	0.0041	0.235	0.215
MC-LR	0.0047	0.239	0.212
M1-Radice	0.0037	0.204	0.209
M1-SpEC	0.0036	0.219	0.187
M1-NuLib	0.0031	0.214	0.210
Hydro	0.0044	0.243	0.036

TABLE III. Unbound mass M_{ej}, as measured on the computational grid 2.5 ms after merger, as well as average velocity and electron fraction of the ejecta.

negligible in the hottest, densest region, where we rely on a low-accuracy implicit MC algorithm. The M1-SpEC simulations, for its part, has significant issues with the evolution of the electron fraction in dense regions, which erroneously decreases over the period considered here.

B. Mass ejection

The mass, composition, and geometry of the outflows is naturally one of the most important prediction of merger simulations. Over the first ~ 3 ms of evolution, the system considered here first ejects a small amount of cold, neutron-rich tidal ejecta, then slighly larger amounts of hotter outflows whose ejection is tied to the core-bounce oscillations of the forming remnant. The former is largely confined to the equatorial plane, while the latter covers a broader range of polar angles.

We can first compare the global properties of the outflows, summarized in Table III. The table considers as outflows all matter marked as unbound 2.5 ms post-merger, according to the criteria of [72] which accounts for r-process heating, neutrino cooling, and the finite time available for heating before particles reach apoastron. We note that material will continue to be ejected over longer timescales, both due to continued core-bounces and the evolution of the remnant on longer timescales. This early time data however allows us to compare the different transport methods without the complications of different collapse times or the large errors associated on longer timescales with the absence of magnetic fields in our simulations. At 2.5 ms post-merger, the simulations flag $(0.0031 - 0.0047) M_\odot$ of material as unbound – a larger relative uncertainty than for many of the remnant properties considered in the previous section, but not far outside of expected numerical error. Indeed, the finite resolution of the grid leads to $\sim 10\%$ relative uncertainty in the ejected mass for the M1 and Hydro simulations [50], while the finite number of packets in the MC simulation leads to $\sim 25\%$ relative uncertainty in the ejected mass of the MC-HR simulation (assuming errors scaling as $\sqrt{N}$, with N the number of packets, so that we estimate the total error in MC-HR as 1.7 times the difference between MC-HR and MC-LR). The composition and velocity of the outflows are in better agreement

($[10 - 20]\%$ relative differences), with the easily explained exception of the composition of the outflows in the neutrinoless simulation. Otherwise, the higher velocity of the outflows in the MC simulations and the lower electron fraction of the M1-SpEC simulation are the main observable differences outside of expected numerical errors.

Fig. 7 provides more fine-grained information about the outflows, for simulations with neutrinos. The corner plots show the distribution of the outflows in the three-dimensional parameter space of polar angle θ, electron fraction Y_e, and predicted asymptotic velocity v_∞. There is in general a fairly striking agreement between the four different neutrino transport schemes. The most notable differences are the underproduction of low-Y_e tidal ejecta and high-Y_e core-bounce ejecta in the M1-SpEC simulation, and a smoother angular distribution of the outflows in the MC simulation than in all M1 simulations. At the qualitative level however, all simulations predict that most of the outflows are within $\sim 60°$ of the equator, with electron fraction peaking just above $Y_e \sim 0.2$ and extending to $Y_e \sim 0.4$, and velocity distribution peaking at $v_\infty \sim 0.2c$ and extending to $v_\infty \sim 0.6c$.

C. Neutrinos in the remnant

Some of the most important differences between the transport schemes used in this manuscript are mostly relevant in high-density, optically thick regions. In particular, one region of parameter space where the M1-Radice method is expected to perform better than the M1-SpEC is regions of high scattering opacity and low absorption opacity [39]. The implicit MC scheme used in our simulations is for its part active in any region where the absorption timescale is shorter than the simulation time step. Accordingly, we will now investigate in more details the distribution of neutrinos throughout the dense regions of the remnant.

To begin, we consider the fluid-frame energy density of neutrinos on horizontal and vertical slices on Fig. 8. Here, we use the M1-Radice simulation as reference considering that, a priori, it is the most accurate of the M1 methods (we will see below that getting clean information from the MC simulations is more difficult). For reference, the maximum energy density of the fluid at the beginning of the simulation (including rest mass energy density), at the center of the most massive neutron star, is $J_{\mathrm{fl}} \sim 8 \times 10^{36}$ erg/cc. Neutrinos are mainly created in the dense, hot regions of the remnant. The $\bar{\nu}_e$ species dominates, with energy density rising to $O(10^{-3})$ of the fluid energy density. Heavy-lepton neutrinos have slightly lower energy density, while ν_e production is relatively suppressed. The same hierarchy is visible in shocked regions of the tidal tail, and in the free-streaming regions. A comparison with Figs. 3-6 shows that the neutrino energy follows the fluid temperature, as expected. In particular, at this early time in the post-merger evolu-
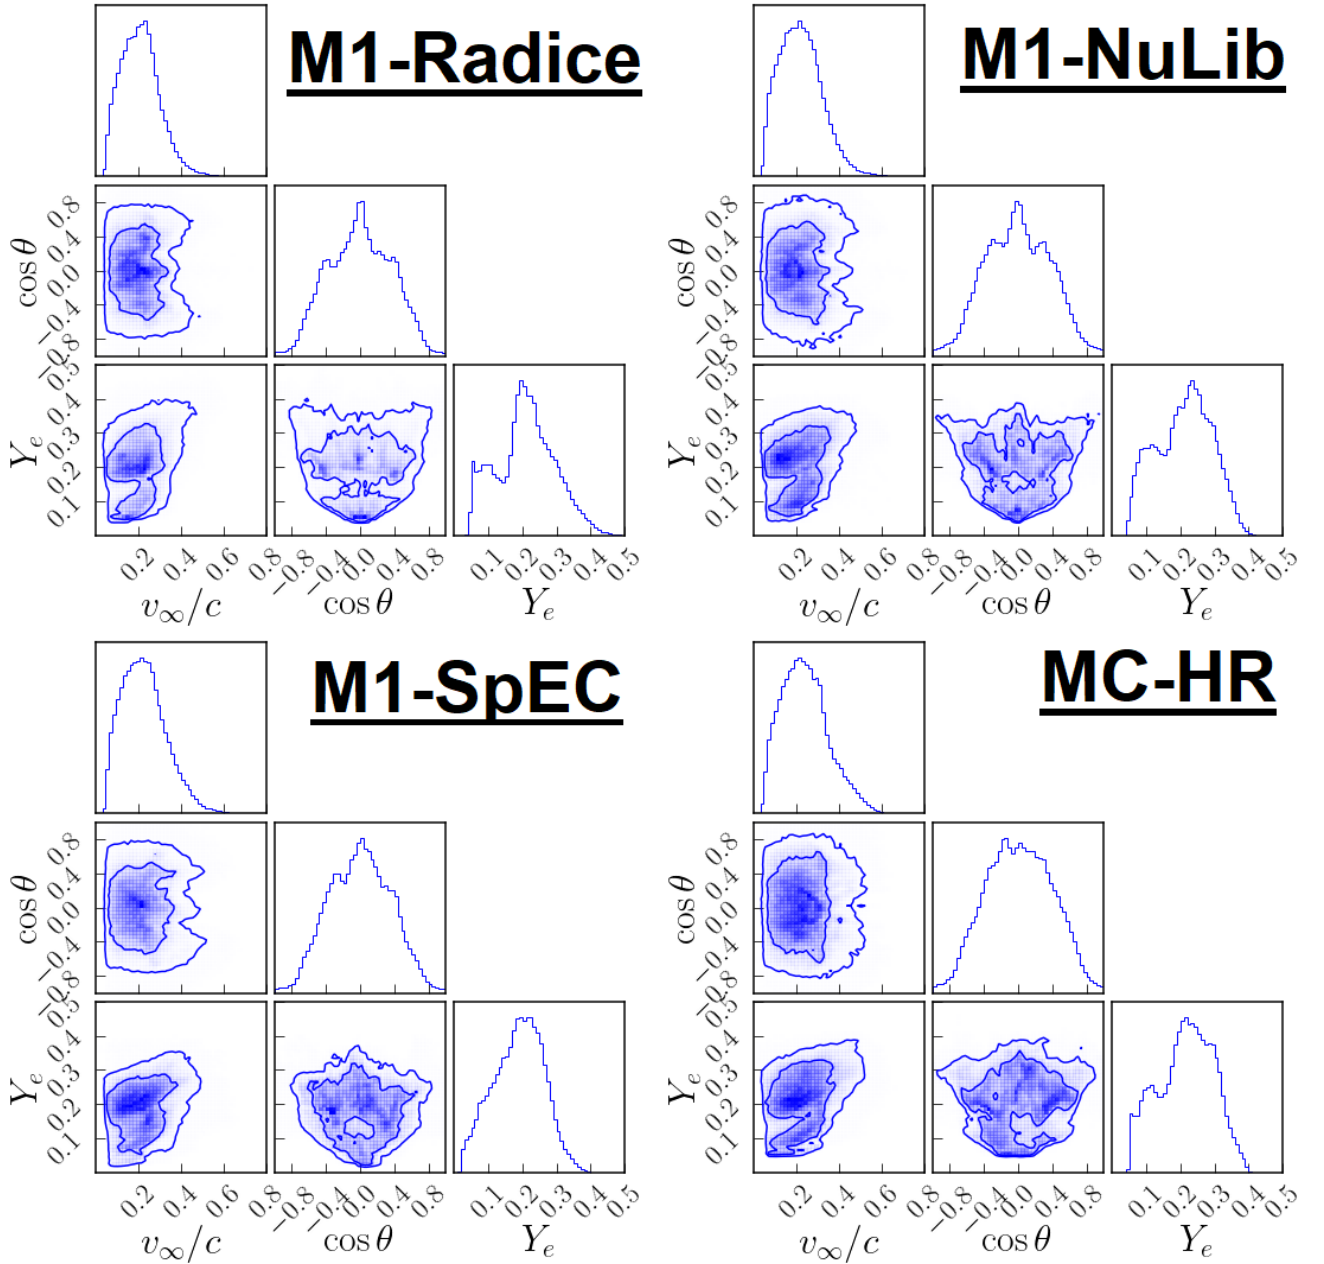
FIG. 7. Distribution of the asymptotic velocity v_∞, azimuthal angle θ, and electron fraction Y_e of the ejected matter 2.5 ms post-merger. We show simulations M1-Radice (top left), M1-NuLib (top right), M1-SpEC (bottom left), and MC-HR (bottom right). The MC-LR is, for these distributions, very similar to MC-HR while the neutrinoless simulation has obviously very different composition.

tion, cold dense regions remain within the remnant where very few neutrinos are created or able to diffuse into.

To compare different radiation transport methods, we find it easier to consider the ratio of the neutrino energy density in the fluid frame to the expected energy density in equilibrium. Indeed, this tends to show more clearly regions where neutrinos are in equilibrium ($J \approx J_{eq}$), regions irradiated by hotter parts of the remnants ($J > J_{eq}$), and (semi-)transparent regions from which neutrinos are partially free-streaming away ($J < J_{eq}$).

These appear, respectively, as white, red, and blue regions on Figs 9-10. We note however that $J \approx J_{eq}$ is not a sufficient condition to guarantee that the neutrinos are trapped and in thermal equilibrium with the fluid; it is only a necessary condition.

In Fig. 9, we compare the M1-Radice and M1-SpEC simulations. The M1-NuLib simulation is quite similar to M1-SpEC when considering the neutrino energy density, despite the use of different reaction rates. We will see below however that the M1-SpEC simulation has prob-
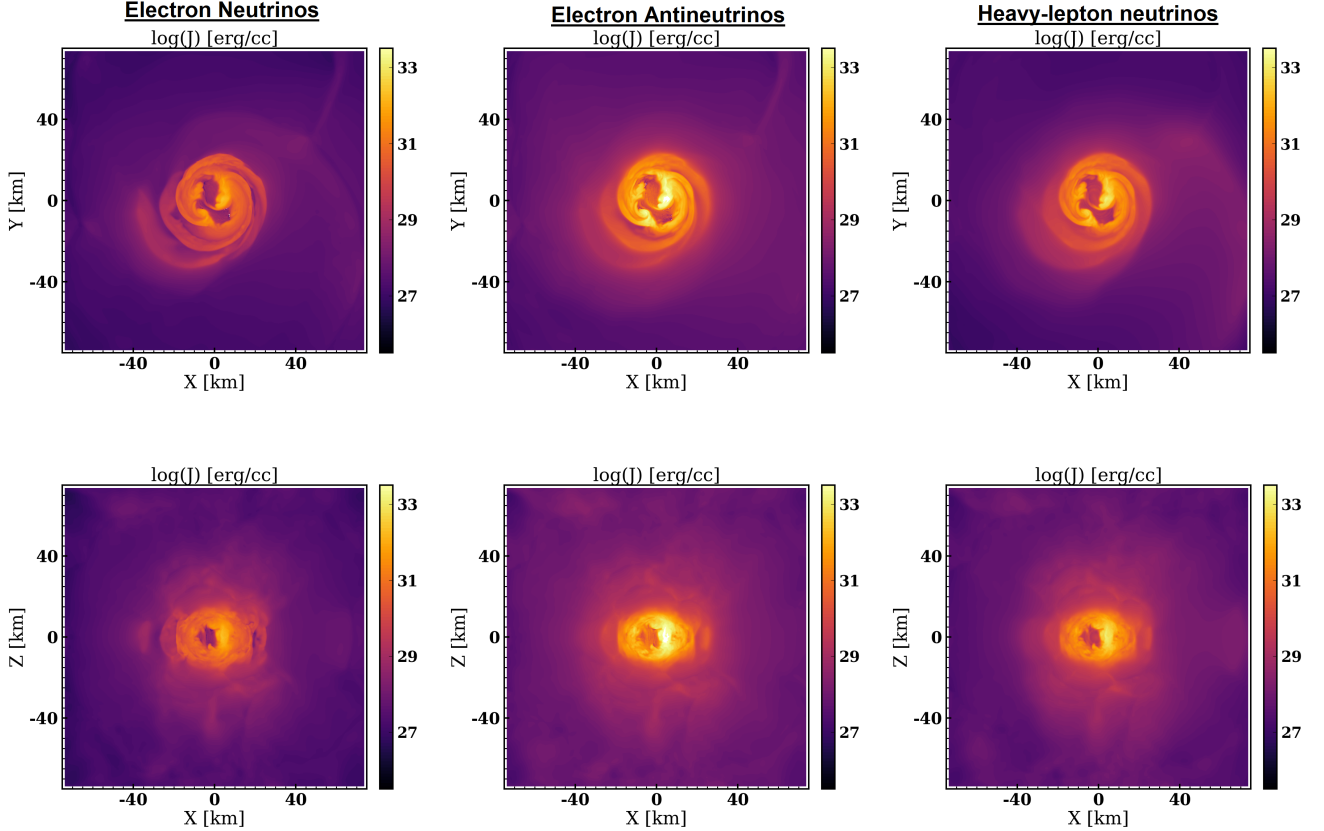
FIG. 8. Energy density of neutrinos in the fluid frame for ν_e (left), $\bar{\nu}_e$ (middle), and each individual species of heavy-lepton neutrinos (right). We show results for the M1-Radice simulation 2.5 ms after merger and on a horizontal (top) and vertical (bottom) slice.

lematic behavior in the evolution of the number density at the time of merger. We note a few important regions where the M1-Radice and M1-SpEC schemes disagree. First, the neutrinos are generally closer to equilibrium in the hot tidal tails in the M1-Radice simulation than in the M1-SpEC simulation. This is visible as blue regions in the tidal tail for M1-SpEC, especially for the heavy-lepton neutrinos and electron antineutrinos. Second, dense regions remain in equilibrium down to lower density in the M1-Radice simulations. Third, cold regions inside the neutron star are more strongly irradiated in the M1-SpEC simulation. The last two could naturally be understood as the inability of the M1-SpEC simulation to properly advect neutrinos along with the fluid in dense regions. The first difference is more puzzling. We can estimate the length scale over which neutrinos are expected to equilibrate in the tidal tail as $L_{\text{eq}} \approx \sqrt{\kappa_a(\kappa_a + \kappa_s)}$ with $\kappa_{a,s}$ averaged over the equilibrium energy spectrum of neutrinos. For the conditions in the tidal tail of the M1-Radice simulations ($\rho \approx 4 \times 10^{10}$ g/cc, $T \approx 3$ MeV, $Y_e \approx 0.15$ at the points where the tidal tail crosses the positive x-axis), we have $L_{\text{eq}} = (20, 120, 2700)$ km for ($\nu_e, \bar{\nu}_e, \nu_x$) respectively. Considering that the tidal tail has width ~ 20 km and height ~ 50 km, we would expect

that in the tidal tail the ν_e would be partially trapped, the $\bar{\nu}_e$ would be in the semi-transparent regime, and the ν_x would be nearly free-streaming. The numerical results are reasonably consistent with these predictions for ν_e and $\bar{\nu}_e$, but clearly not for ν_x. In the M1-SpEC simulations, and for different conditions in the tidal tail ($\rho \approx 5.5 \times 10^{10}$ g/cc, $T \approx 6$ MeV, $Y_e \approx 0.21$), we estimate $L_{\text{eq}} = (4, 17, 180)$ km instead. We observe trapped ν_e, partially trapped $\bar{\nu}_e$, and free streaming ν_x, in better agreement with expectations. One possible explanation, considering that this issue is only observed in ν_x in the tidal tail, is that the equilibrium energy density in the M1-Radice tail is a factor of ~ 15 lower than in the M1-SpEC tail, and thus that energy density may be more easily reached through irradiation from other regions of the simulation – i.e. this might be an example of the fact that $J \sim J_{\text{eq}}$ can be true even if neutrinos are not trapped in a given region.

Finally, we turn to Fig. 10, which performs the same comparison for the M1-Radice and MC-HR simulations. We only show ν_e and ν_x here (the $\bar{\nu}_e$ show a level of agreement between simulations comparable to the ν_e). We note that getting an estimate of the energy density of neutrinos in the MC simulations is more difficult than
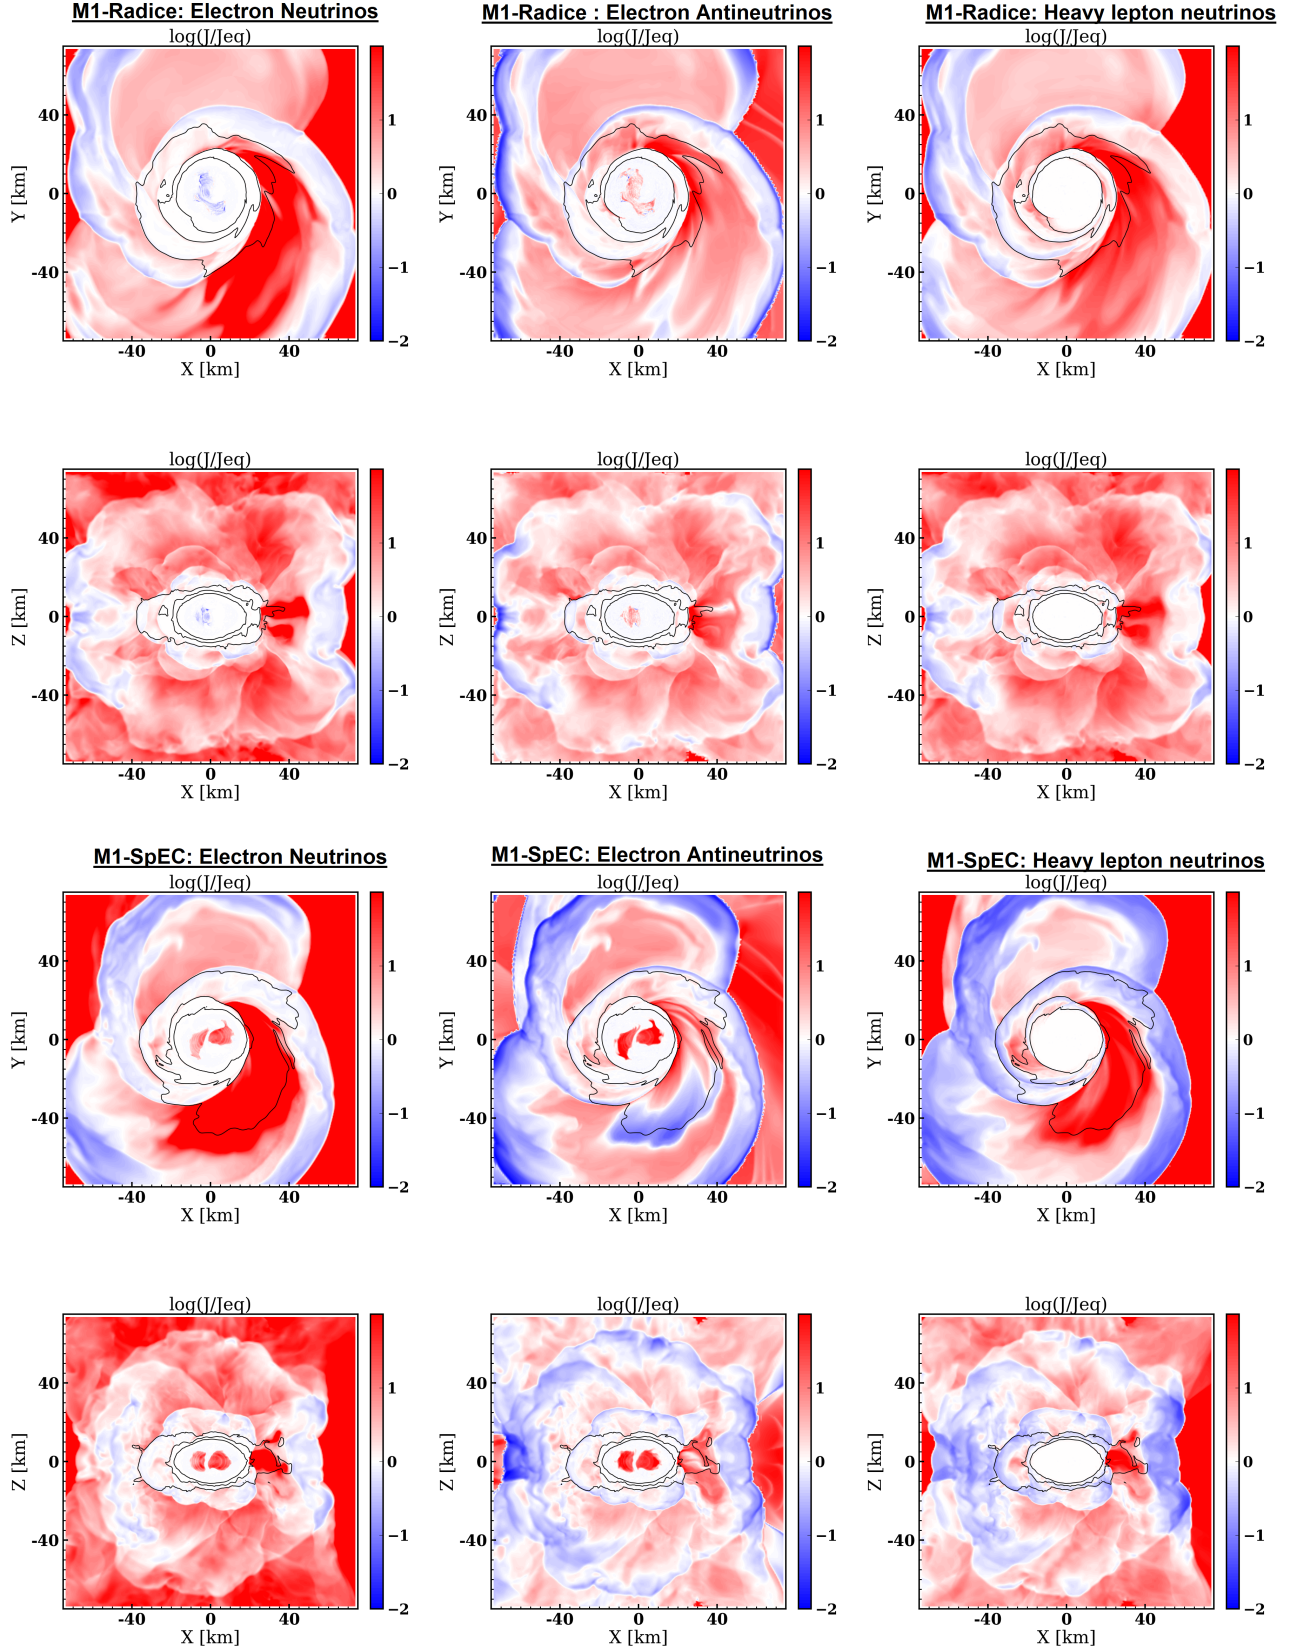
FIG. 9. Logarithm of the ratio of the energy density of neutrinos in the fluid frame to the energy density that neutrinos would have if they were in equilibrium with the fluid. We show results for ν_e (left), $\bar{\nu}_e$ (center) and ν_x (right), in a horizontal (top) and vertical (bottom) slice of the remnant and for the M1-Radice simulation (first two rows) and M1-SpEC simulation (last two rows). Black contours correspond to densities of $10^{11,12,13}$ g/cc. In white regions, neutrinos are close to being in equilibrium with the fluid. Blue regions are semi-transparent emission region, while red regions are irradiated by other regions of the simulations.
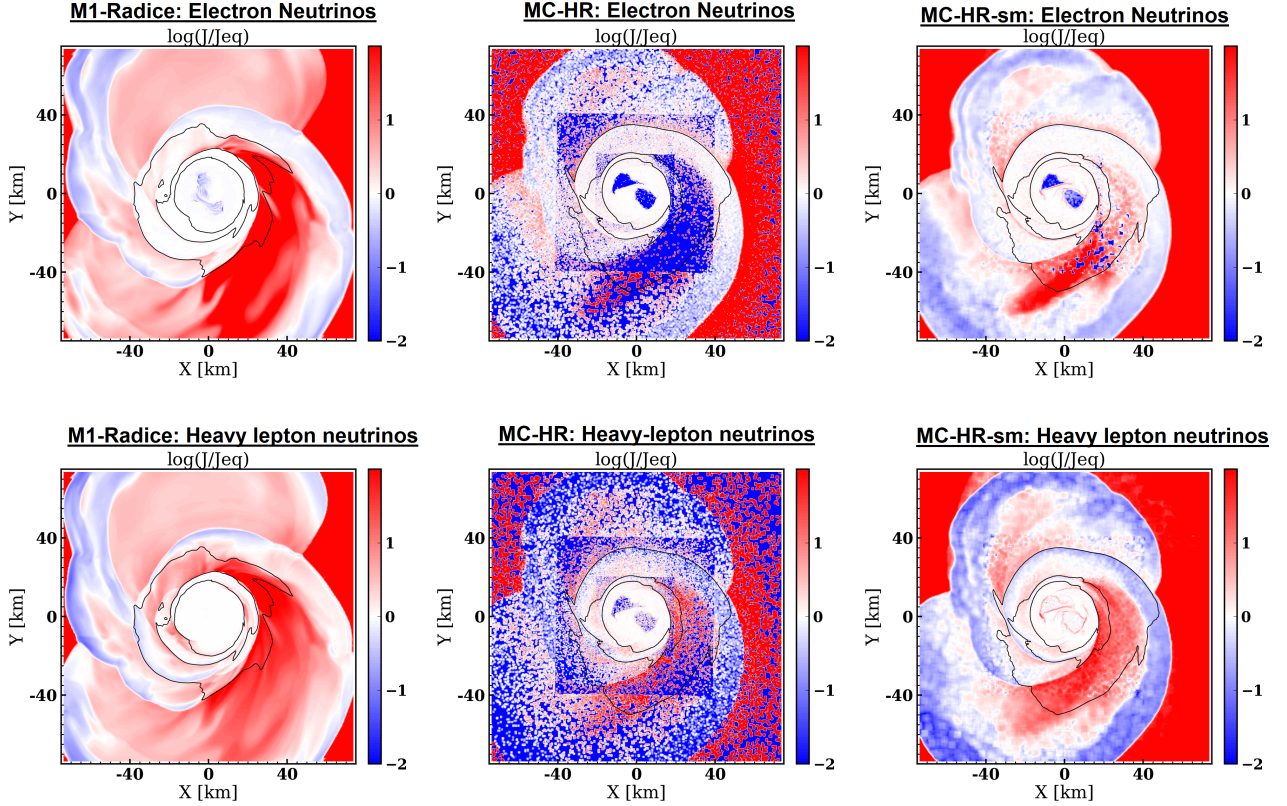
FIG. 10. Same as Fig. 9 but for simulation M1-Radice (left) and MC-HR (center). Note that for the MC simulation, we calculate the energy density from the packets existing during a single step $dt \sim 25\,$ns in a single grid cell, which introduces large sampling noise. In particular, saturated colors (dark red, dark blue) are typically associated with cells with $0-1$ packets, making them entirely unreliable. The right panel shows the data from the central panel smoothed over 5^3 cells, to limit the impact of that noise. The top row shows results for ν_x, and the bottom row for ν_x. The boundary between refinement level is clearly visible on the MC plots because coarser cells have lower sampling noise than immediately neighboring finer cells.

for the M1 simulations. When computing moments in MC simulations, we currently integrate the contribution of all packets over a 4-volume corresponding to a grid cell of proper spatial volume ΔV evolved over a time $\Delta\tau$, then divide the result by $\Delta\tau$ to estimate the energy density (using the proper time of an inertial observer, $\Delta\tau = \alpha\Delta t$, with Δt the elapsed simulation time). We keep track of moments of the distribution function of neutrinos integrated over the latest simulation time step, which at 2.5 ms post-merger leads to $\Delta t \sim 25\,$ns. Without additional averaging, this leads to significant shot noise everywhere but in the densest and hottest regions of the remnant (middle panels of Fig. 10; we note in particular that the grid structure is clearly visible on these plots, because on coarser grids we average over a larger volume, thus lowering the MC noise). To improve on this issue, we average the MC results over a $5 \times 5 \times 5$ group of cells (right panels of Fig. 10). This improves the accuracy of our prediction in lower-density regions, at the cost of inaccuracies in regions where the energy density varies rapidly; e.g. at the boundary between hot and cold regions inside the remnant. We see that the MC simula-

tion agrees quite well with the M1-Radice simulation in dense regions. In particular, the region where neutrinos are in equilibrium with the fluid extends to similar densities, and the cold regions of the remnant are devoid of ν_e (i.e. blue on the figure). This tends to confirm that the differences observed in these regions between M1-Radice and M1-SpEC are indeed due to the limitations of the M1-SpEC simulations in dense regions. In the hot tidal tail, the MC simulations behave closer to the M1-SpEC simulation. As discussed in the previous paragraph, the behavior of the MC simulations is also what we might intuitively expect from the calculation of the equilibration length $L_{\rm eq}$ of neutrinos. As the MC simulation also has a hotter tidal tail than the M1-Radice simulation, this does not however provide new information on the origin of the apparently-trapped ν_x neutrinos in M1-Radice.

Fig. 10 also provides a good illustration of some of the strength and weaknesses of the MC code. The fact that, in the absence of smoothing, the shot noise in the calculation of the energy density is so large in many regions of the simulation is a good reminder that directly estimating the distribution function of neutrinos (or its mo-

ments) from an MC simulation can be very inaccurate. This can be problematic for example when attempting to calculate $\nu\bar{\nu}$ annihilation rate in polar regions, where very few packets are located, or in any region where the neutrino pressure may become dynamically important. The impact of these errors when using MC as a closure for M1 schemes [42] at realistic number of packets remains an interesting open question for hybrid methods as well. On the other hand, the very simple smoothing performed on this figure illustrates why MC methods with relatively low packet numbers work quite well in binary merger simulations. Indeed, for the output of simulations to be minimally impacted by shot noise we only need the distribution function (or its relevant moments) to be well-modeled by the neutrino packets averaged over a timescale $\tau_{\rm coupling}$ over which the fluid variables are modified by neutrino-matter interactions. Going from the middle to the right panels of Fig. 10 is similar to estimating the neutrino energy density by averaging over a time $125\Delta t \approx 3\,\mu$s, which is still much shorter than $\tau_{\rm coupling}$ ($\sim$ ms or more) except in regions where neutrinos are trapped, where the MC sampling noise is generally lower even without smoothing. Averaging the neutrino energy density over 1 ms would in theory reduce shot noise by an additional factor of ~ 20, assuming that the noise scales as $N^{-1/2}$ when using N packets. The same scaling also explains why getting ten percent-level errors from MC simulations is cost-effective, but reducing that error by multiple orders of magnitude would be prohibitively expensive.

Finally, on Fig. 11 we report the luminosity of neutrinos in the $(1-3)$ ms post-merger for which data is available for all simulations. The most noticeable feature here is obviously the large burst of $\bar{\nu}_e$ emission from simulation M1-SpEC over a 0.5 ms period following the merger. We already noted that the high-density regions of the remnant erroneously neutronize in this simulation; this is another indication that the early post-merger composition evolution of the high-density regions is not well captured in M1-SpEC. A more detailed analysis of the neutrino distribution in post-processing shows that during a ~ 0.5 ms period around merger, the M1-SpEC simulation created large numbers of low-energy $\bar{\nu}_e$ in the dense, cold regions of the merging neutron star cores, and that lepton number conservation was not satisfied during that time period (as opposed to what is observed in other simulations using the various M1 schemes). These low-energy $\bar{\nu}_e$ are the source of the extra emission of $\bar{\nu}_e$ at early times, but also of the decrease in the electron fraction observed in dense regions, as the conservation of lepton number was satisfied during their partial reabsorption into the remnant. The exact source of this violation of lepton number conservation is unclear, but likely due to a combination of the inaccuracies of the M1-SpEC scheme in dense cold regions and of corrections applied to the evolution of the number density to guarantee that the average energy of neutrinos remain between 0.01 MeV and 1000 MeV. This is of course a clear indication that the results of the other
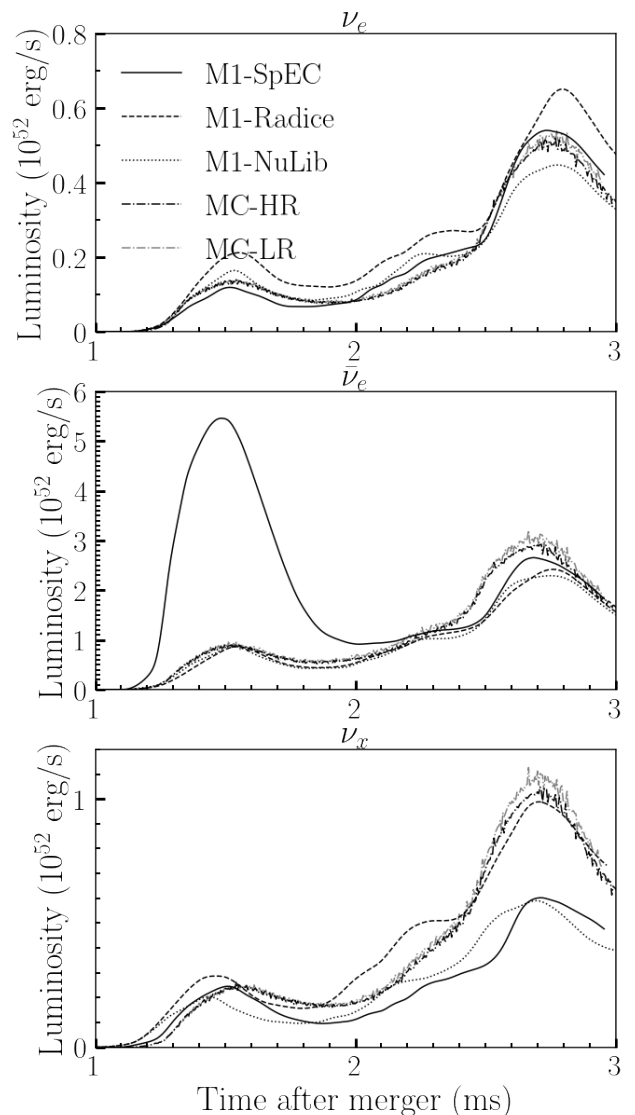

FIG. 11. Luminosity of neutrinos for all simulations, for ν_e (top panel), $\bar{\nu}_e$ (middle panel), and each species of ν_x individually (bottom panel), measured at the outer boundary of our computational domain (a box of side ~ 320 km). Note the different scales for each species.

simulations should be trusted over the M1-SpEC simulation in any region impacted by this problem.

The second main differences between simulations is that both M1-NuLib and M1-SpEC underproduces ν_x by the end of the evolution. Considering that these two simulations use the same inaccurate methods for the implicit solve in the two-moment equation, that these inaccuracies are expected to be particularly important for regions of high scattering opacities but low absorption opacities, and that ν_x neutrinos specifically go through such regions as they exit the remnant, these differences can be reasonably attributed to the limitations of that approximate implicit solve. We note that a similar difference in the luminosity of ν_x had also been noted as the main

difference between M1 and MC in our earlier comparison of transport schemes in a low mass binary neutron star system in [27]. Our results here indicate that this was likely an effect of the approximations made in the implicit solver within the M1 algorithm, rather than a limitation of the M1 formalism itself.

The three simulations that we would consider most reliable in the neutrino sector (MC-HR, MC-LR, M1-Radice) are in fairly good agreement. The M1-Radice simulation produces slightly less $\bar{\nu}_e$ and slightly more ν_e than the MC simulations, with $\sim 20\%$ relative differences, while the two MC simulations have nearly identical neutrino luminosities. This is in this case consistent with the M1-Radice protonizing less in the dense regions (less $\bar{\nu}_e$ emission) and protonizing more in the outflows (more ν_e absorption).

D. Cost of simulations

In order to determine how convenient any of these methods is to perform surveys of the binary parameter space, it is also useful to compare computational costs. These simulations were performed on the Plasma cluster at UNH and on the Wheeler cluster at Caltech, which have similar performance for our simulations (at the $\sim 20\%$ level). We report costs from the beginning of the simulation to $3\,\mathrm{ms}$ post-merger. The cheapest simulations are the pure hydrodynamics simulation and the MC-LR simulation, with costs of $194k$ and $211k$ CPU-hrs respectively. The simulations using the cheaper but less accurate implicit time stepping for the two moment scheme (M1-SpEC, M1-NuLib) as well as the high packet count MC simulation MC-HR are $\sim 50\%$ more costly ($295k$, $309k$, and $288k$ respectively). This is consistent with our previous findings that our old two-momenst scheme and our MC scheme with $\sim 10^8$ packets per species have comparable computational costs [27]. Finally, the two-moments simulation using the more accurate but more expensive implicit time stepping of [39] (M1-Radice) costs about twice as much as the pure hydrodynamics simulation ($412k$). We note of course that the slow convergence of Monte-Carlo methods means that significantly improving the accuracy of the Monte-Carlo simulations would inevitably require simulations much more expensive than the two-moments simulations, and that the cost of the M1-Radice might be impacted by details of the implementation of the implicit solver that have not been fully optimized in SpEC so far. On the other hand, while the two-moment scheme converges faster than the Monte-Carlo scheme, it does not converge to the correct solution to the transport equations.

IV. DISCUSSION

In this manuscript, we performed simulations of a single binary neutron star merger system that forms a compact remnant with expected lifetime of $O(10\,\mathrm{ms})$, using a range of state-of-the art neutrino transport schemes as well as two sets of neutrino-matter interactions. We find that for this metastable system, small differences in the evolution of dense regions right after merger lead to significant difference in the time required for the system to collapse to a black hole. Simulations using an energy-dependent, low-resolution Monte-Carlo transport collapse after $\sim 4\,\mathrm{ms}$, while the simulation using a gray approximate two-moment scheme with similar reactions rates collapses after $\sim 8.5\,\mathrm{ms}$ – in reasonable agreement with preexisting simulations of a similar system with a two-moment scheme by Sekiguchi et al [49]. These differences can be contrasted with the good agreement of the evolution of dense regions between MC and two-moment algorithms for lower mass remnants found in [27]. We note however that the collapse time to a black hole for these metastable systems is also known to be quite sensitive to grid resolution and magnetic field evolution (see e.g. [44] for similar effects in a system with a $\sim (20-30)\,\mathrm{ms}$ collapse time), and may be seen mainly as a confirmation of the high sensitivity of such systems to small changes in the evolution of dense regions.

Besides this very visible difference between MC and two-moment simulations, we find that all neutrino transport used here capture the qualitative effects of neutrino-matter interactions on important observables, with quantitative differences between simulations at the level of $(10-30)\%$ relative uncertainties for the global properties of the remnant and outflows and significant difference in the temperature and composition of lower-density regions, especially the polar regions and hot tidal arms. In the tidal arms, the temperature varies by about a factor of 2 between simulations. These differences are seen when varying the transport method (MC vs M1), the detailed implementation of the moment scheme (M1-Radice vs M1-SpEC), and the interaction rates used in the simulation (M1-SpEC vs M1-NuLib). These results provide us with a useful starting point to determine whether the details of neutrino transport in simulations are important for any given use of simulation results; e.g. predicting kilonova lightcurves or r-process yields at a desired accuracy level.

Our simulations also allow us to assess more directly the impact of improvements to the M1 scheme recently proposed by Radice et al [39]. We find indications that these improvements do indeed modify the behavior of neutrinos in regions with high scattering optical depths, as predicted in [39], and generally lead to differences with our preexisting M1 code that are often comparable to or larger than other sources of uncertainties in the neutrino transport algorithm. We note in addition that the M1-SpEC simulation has a peculiar behavior leading to the emission of low-energy $\bar{\nu}_e$ during merger, violating conservation of electron lepton number. This effect has not been seen in other M1 simulations, yet considering that it occurs in the same high density, cold regions where the M1-SpEC is expected to be least accurate, it is likely a

result of the inaccuracies of that M1 scheme.

The M1-Radice and MC simulations show $\sim 20\%$ level agreements in early-time neutrino luminosities, while more significant differences are observed between those three simulations and those using the approximate implicit solve previously implemented in SpEC. Interestingly, this difference between M1-Radice and M1-SpEC also appears sufficient to explain the largest difference between M1 and MC results observed in our previous comparison of low-mass binary neutron star systems, i.e. the fact that the M1 scheme was underproducing ν_x neutrinos by up to $\sim 50\%$.

Finally, we perform simulations with Monte-Carlo transport for two different number of packets. We find that at the number of packets currently in use in our simulations, Monte-Carlo sampling noise is an important contribution to the uncertainty in the total amount of matter ejected by the merger ($\sim 25\%$ relative error here). For all other observables, packet count is a negligible contribution to the error budget. However, we caution that a more detailed testing of the impact of the use of implicit Monte-Carlo in the dense regions is likely warranted before concluding that these simulations are converged in other observables. In particular, it is at this point unclear whether differences between M1 and Monte-Carlo in dense regions are due to Monte-Carlo shot noise and the use of implicit Monte-Carlo in dense regions, or to the energy and pressure closure used by the M1 code in dense regions where neutrinos are not in equilibrium with the fluid.

[1] B. P. Abbott *et al.* (Virgo, LIGO Scientific), GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral, Phys. Rev. Lett. **119**, 161101 (2017), arXiv:1710.05832 [gr-qc].

[2] B. P. Abbott *et al.* (GROND, SALT Group, OzGrav, DFN, INTEGRAL, Virgo, Insight-Hxmt, MAXI Team, Fermi-LAT, J-GEM, RATIR, IceCube, CAASTRO, LWA, ePESSTO, GRAWITA, RIMAS, SKA South Africa/MeerKAT, H.E.S.S., 1M2H Team, IKI-GW Follow-up, Fermi GBM, Pi of Sky, DWF (Deeper Wider Faster Program), Dark Energy Survey, MASTER, AstroSat Cadmium Zinc Telluride Imager Team, Swift, Pierre Auger, ASKAP, VINROUGE, JAGWAR, Chandra Team at McGill University, TTU-NRAO, GROWTH, AGILE Team, MWA, ATCA, AST3, TOROS, Pan-STARRS, NuSTAR, ATLAS Telescopes, BOOTES, CaltechNRAO, LIGO Scientific, High Time Resolution Universe Survey, Nordic Optical Telescope, Las Cumbres Observatory Group, TZAC Consortium, LOFAR, IPN, DLT40, Texas Tech University, HAWC, ANTARES, KU, Dark Energy Camera GW-EM, CALET, Euro VLBI Team, ALMA), Multi-messenger Observations of a Binary Neutron Star Merger, Astrophys. J. **848**, L12 (2017), arXiv:1710.05833 [astro-ph.HE].

[3] B. P. Abbott *et al.* (Virgo, Fermi-GBM, INTEGRAL, LIGO Scientific), Gravitational Waves and Gamma-rays from a Binary Neutron Star Merger: GW170817 and GRB 170817A, Astrophys. J. **848**, L13 (2017), arXiv:1710.05834 [astro-ph.HE].

[4] M. M. Kasliwal, E. Nakar, L. P. Singer, D. L. Kaplan, D. O. Cook, A. Van Sistine, R. M. Lau, C. Fremling, O. Gottlieb, J. E. Jencson, S. M. Adams, U. Feindt, K. Hotokezaka, S. Ghosh, D. A. Perley, P.-C. Yu, T. Piran, J. R. Allison, G. C. Anupama, A. Balasubramanian, K. W. Bannister, J. Bally, J. Barnes, S. Barway, E. Bellm, V. Bhalerao, D. Bhattacharya, N. Blagorodnova, J. S. Bloom, P. R. Brady, C. Cannella, D. Chatterjee, S. B. Cenko, B. E. Cobb, C. Copperwheat, A. Corsi, K. De, D. Dobie, S. W. K. Emery, P. A. Evans, O. D. Fox, D. A. Frail, C. Frohmaier, A. Goobar, G. Hallinan, F. Harrison, G. Helou, T. Hinderer, A. Y. Q. Ho, A. Horesh, W.-H. Ip, R. Itoh, D. Kasen, H. Kim, N. P. M. Kuin, T. Kupfer, C. Lynch, K. Madsen, P. A. Mazzali, A. A. Miller, K. Mooley, T. Murphy, C.-C. Ngeow, D. Nichols, S. Nissanke, P. Nugent, E. O. Ofek, H. Qi, R. M. Quimby, S. Rosswog, F. Rusu, E. M. Sadler, P. Schmidt, J. Sollerman, I. Steele, A. R. Williamson, Y. Xu, L. Yan, Y. Yatsu, C. Zhang, and W. Zhao, Illuminating gravitational waves: A concordant picture of photons from a neutron star merger, Science **358**, 1559

(2017), arXiv:1710.05436 [astro-ph.HE].

[5] R. Chornock, E. Berger, D. Kasen, P. S. Cowperthwaite, M. Nicholl, V. A. Villar, K. D. Alexander, P. K. Blanchard, T. Eftekhari, W. Fong, R. Margutti, P. K. G. Williams, J. Annis, D. Brout, D. A. Brown, H.-Y. Chen, M. R. Drout, B. Farr, R. J. Foley, J. A. Frieman, C. L. Fryer, K. Herner, D. E. Holz, R. Kessler, T. Matheson, B. D. Metzger, E. Quataert, A. Rest, M. Sako, D. M. Scolnic, N. Smith, and M. Soares-Santos, The Electromagnetic Counterpart of the Binary Neutron Star Merger LIGO/Virgo GW170817. IV. Detection of Near-infrared Signatures of r-process Nucleosynthesis with Gemini-South, "Astrophys. J. Lett." **848**, L19 (2017), arXiv:1710.05454 [astro-ph.HE].

[6] P. S. Cowperthwaite *et al.*, The Electromagnetic Counterpart of the Binary Neutron Star Merger LIGO/Virgo GW170817. II. UV, Optical, and Near-infrared Light Curves and Comparison to Kilonova Models, Astrophys. J. **848**, L17 (2017), arXiv:1710.05840 [astro-ph.HE].

[7] C. D. Kilpatrick, R. J. Foley, D. Kasen, A. Murguia-Berthier, E. Ramirez-Ruiz, D. A. Coulter, M. R. Drout, A. L. Piro, B. J. Shappee, K. Boutsia, C. Contreras, F. Di Mille, B. F. Madore, N. Morrell, Y.-C. Pan, J. X. Prochaska, A. Rest, C. Rojas-Bravo, M. R. Siebert, J. D. Simon, and N. Ulloa, Electromagnetic evidence that SSS17a is the result of a binary neutron star merger, "Science" **358**, 1583 (2017), arXiv:1710.05434 [astro-ph.HE].

[8] C. McCully, D. Hiramatsu, D. A. Howell, G. Hosseinzadeh, I. Arcavi, D. Kasen, J. Barnes, M. M. Shara, T. B. Williams, P. Väisänen, S. B. Potter, E. Romero-Colmenero, S. M. Crawford, D. A. H. Buckley, J. Cooke, I. Andreoni, T. A. Pritchard, J. Mao, M. Gromadzki, and J. Burke, The Rapid Reddening and Featureless Optical Spectra of the Optical Counterpart of GW170817, AT 2017gfo, during the First Four Days, "Astrophys. J. Lett." **848**, L32 (2017), arXiv:1710.05853 [astro-ph.HE].

[9] M. Nicholl, E. Berger, D. Kasen, B. D. Metzger, J. Elias, C. Briceño, K. D. Alexander, P. K. Blanchard, R. Chornock, P. S. Cowperthwaite, T. Eftekhari, W. Fong, R. Margutti, V. A. Villar, P. K. G. Williams, W. Brown, J. Annis, A. Bahramian, D. Brout, D. A. Brown, H.-Y. Chen, J. C. Clemens, E. Dennihy, B. Dunlap, D. E. Holz, E. Marchesini, F. Massaro, N. Moskowitz, I. Pelisoli, A. Rest, F. Ricci, M. Sako, M. Soares-Santos, and J. Strader, The Electromagnetic Counterpart of the Binary Neutron Star Merger LIGO/Virgo GW170817. III. Optical and UV Spectra of a Blue Kilonova from Fast Polar Ejecta, "Astrophys. J. Lett." **848**, L18 (2017), arXiv:1710.05456 [astro-ph.HE].

[10] S. J. Smartt, T.-W. Chen, A. Jerkstrand, M. Coughlin, E. Kankare, S. A. Sim, M. Fraser, C. Inserra, K. Maguire, K. C. Chambers, M. E. Huber, T. Krühler, G. Leloudas, M. Magee, L. J. Shingles, K. W. Smith, D. R. Young, J. Tonry, R. Kotak, A. Gal-Yam, J. D. Lyman, D. S. Homan, C. Agliozzo, J. P. Anderson, C. R. Angus, C. Ashall, C. Barbarino, F. E. Bauer, M. Berton, M. T. Botticella, M. Bulla, J. Bulger, G. Cannizzaro, Z. Cano, R. Cartier, A. Cikota, P. Clark, A. De Cia, M. Della Valle, L. Denneau, M. Dennefeld, L. Dessart, G. Dimitriadis, N. Elias-Rosa, R. E. Firth, H. Flewelling, A. Flörs, A. Franckowiak, C. Frohmaier, L. Galbany, S. González-Gaitán, J. Greiner, M. Gromadzki, A. N. Guelbenzu, C. P. Gutiérrez, A. Hamanowicz, L. Hanlon, J. Harmanen, K. E. Heintz, A. Heinze, M.-S. Hernandez, S. T. Hodgkin, I. M. Hook, L. Izzo, P. A. James, P. G. Jonker, W. E. Kerzendorf, S. Klose, Z. Kostrzewa-Rutkowska, M. Kowalski, M. Kromer, H. Kuncarayakti, A. Lawrence, T. B. Lowe, E. A. Magnier, I. Manulis, A. Martin-Carrillo, S. Mattila, O. McBrien, A. Müller, J. Nordin, D. O'Neill, F. Onori, J. T. Palmerio, A. Pastorello, F. Patat, G. Pignata, P. Podsiadlowski, M. L. Pumo, S. J. Prentice, A. Rau, A. Razza, A. Rest, T. Reynolds, R. Roy, A. J. Ruiter, K. A. Rybicki, L. Salmon, P. Schady, A. S. B. Schultz, T. Schweyer, I. R. Seitenzahl, M. Smith, J. Sollerman, B. Stalder, C. W. Stubbs, M. Sullivan, H. Szegedi, F. Taddia, S. Taubenberger, G. Terreran, B. van Soelen, J. Vos, R. J. Wainscoat, N. A. Walton, C. Waters, H. Weiland, M. Willman, P. Wiseman, D. E. Wright, Ł. Wyrzykowski, and O. Yaron, A kilonova as the electromagnetic counterpart to a gravitational-wave source, "Nature" **551**, 75 (2017), arXiv:1710.05841 [astro-ph.HE].

[11] M. Soares-Santos, D. E. Holz, J. Annis, R. Chornock, K. Herner, E. Berger, D. Brout, H.-Y. Chen, R. Kessler, M. Sako, S. Allam, D. L. Tucker, R. E. Butler, A. Palmese, Z. Doctor, H. T. Diehl, J. Frieman, B. Yanny, H. Lin, D. Scolnic, P. Cowperthwaite, E. Neilsen, J. Marriner, N. Kuropatkin, W. G. Hartley, F. Paz-Chinchón, K. D. Alexander, E. Balbinot, P. Blanchard, D. A. Brown, J. L. Carlin, C. Conselice, E. R. Cook, A. Drlica-Wagner, M. R. Drout, F. Durret, T. Eftekhari, B. Farr, D. A. Finley, R. J. Foley, W. Fong, C. L. Fryer, J. García-Bellido, M. S. S. Gill, R. A. Gruendl, C. Hanna, D. Kasen, T. S. Li, P. A. A. Lopes, A. C. C. Lourenço, R. Margutti, J. L. Marshall, T. Matheson, G. E. Medina, B. D. Metzger, R. R. Muñoz, J. Muir, M. Nicholl, E. Quataert, A. Rest, M. Sauseda, D. J. Schlegel, L. F. Secco, F. Sobreira, A. Stebbins, V. A. Villar, K. Vivas, A. R. Walker, W. Wester, P. K. G. Williams, A. Zenteno, Y. Zhang, T. M. C. Abbott, F. B. Abdalla, M. Banerji, K. Bechtol, A. Benoit-Lévy, E. Bertin, D. Brooks, E. Buckley-Geer, D. L. Burke, A. Carnero Rosell, M. Carrasco Kind, J. Carretero, F. J. Castander, M. Crocce, C. E. Cunha, C. B. D'Andrea, L. N. da Costa, C. Davis, S. Desai, J. P. Dietrich, P. Doel, T. F. Eifler, E. Fernandez, B. Flaugher, P. Fosalba, E. Gaztanaga, D. W. Gerdes, T. Giannantonio, D. A. Goldstein, D. Gruen, J. Gschwend, G. Gutierrez, K. Honscheid, B. Jain, D. J. James, T. Jeltema, M. W. G. Johnson, M. D. Johnson, S. Kent, E. Krause, R. Kron, K. Kuehn, S. Kuhlmann, O. Lahav, M. Lima, M. A. G. Maia, M. March, R. G. McMahon, F. Menanteau, R. Miquel, J. J. Mohr, R. C. Nichol, B. Nord, R. L. C. Ogando, D. Petravick, A. A. Plazas, A. K. Romer, A. Roodman, E. S. Rykoff, E. Sanchez, V. Scarpine, M. Schubnell, I. Sevilla-Noarbe, M. Smith, R. C. Smith, E. Suchyta, M. E. C. Swanson, G. Tarle, D. Thomas, R. C. Thomas, M. A. Troxel, V. Vikram, R. H. Wechsler, J. Weller, Dark Energy Survey, and Dark Energy Camera GW-EM Collaboration, The Electromagnetic Counterpart of the Binary Neutron Star Merger LIGO/Virgo GW170817. I. Discovery of the Optical Counterpart Using the Dark Energy Camera, "Astrophys. J. Lett." **848**, L16 (2017), arXiv:1710.05459 [astro-ph.HE].

[12] N. R. Tanvir, A. J. Levan, C. González-Fernández, O. Korobkin, I. Mandel, S. Rosswog, J. Hjorth, P. D'Avanzo, A. S. Fruchter, C. L. Fryer, T. Kan-

gas, B. Milvang-Jensen, S. Rosetti, D. Steeghs, R. T. Wollaeger, Z. Cano, C. M. Copperwheat, S. Covino, V. D'Elia, A. de Ugarte Postigo, P. A. Evans, W. P. Even, S. Fairhurst, R. Figuera Jaimes, C. J. Fontes, Y. I. Fujii, J. P. U. Fynbo, B. P. Gompertz, J. Greiner, G. Hodosan, M. J. Irwin, P. Jakobsson, U. G. Jørgensen, D. A. Kann, J. D. Lyman, D. Malesani, R. G. McMahon, A. Melandri, P. T. O'Brien, J. P. Osborne, E. Palazzi, D. A. Perley, E. Pian, S. Piranomonte, M. Rabus, E. Rol, A. Rowlinson, S. Schulze, P. Sutton, C. C. Thöne, K. Ulaczyk, D. Watson, K. Wiersema, and R. A. M. J. Wijers, The Emergence of a Lanthanide-rich Kilonova Following the Merger of Two Neutron Stars, "Astrophys. J. Lett." **848**, L27 (2017), arXiv:1710.05455 [astro-ph.HE].

[13] P. A. Evans, S. B. Cenko, J. A. Kennea, S. W. K. Emery, N. P. M. Kuin, O. Korobkin, R. T. Wollaeger, C. L. Fryer, K. K. Madsen, F. A. Harrison, Y. Xu, E. Nakar, K. Hotokezaka, A. Lien, S. Campana, S. R. Oates, E. Troja, A. A. Breeveld, F. E. Marshall, S. D. Barthelmy, A. P. Beardmore, D. N. Burrows, G. Cusumano, A. DAì, P. DAvanzo, V. D Elia, M. de Pasquale, W. P. Even, C. J. Fontes, K. Forster, J. Garcia, P. Giommi, B. Grefenstette, C. Gronwall, D. H. Hartmann, M. Heida, A. L. Hungerford, M. M. Kasliwal, H. A. Krimm, A. J. Levan, D. Malesani, A. Melandri, H. Miyasaka, J. A. Nousek, P. T. OBrien, J. P. Osborne, C. Pagani, K. L. Page, D. M. Palmer, M. Perri, S. Pike, J. L. Racusin, S. Rosswog, M. H. Siegel, T. Sakamoto, B. Sbarufatti, G. Tagliaferri, N. R. Tanvir, and A. Tohuvavohu, Swift and NuSTAR observations of GW170817: Detection of a blue kilonova, "Science" **358**, 1565 (2017), arXiv:1710.05437 [astro-ph.HE].

[14] K. P. Mooley, A. T. Deller, O. Gottlieb, E. Nakar, G. Hallinan, S. Bourke, D. A. Frail, A. Horesh, A. Corsi, and K. Hotokezaka, Superluminal motion of a relativistic jet in the neutron-star merger GW170817, Nature **561**, 355 (2018), arXiv:1806.09693 [astro-ph.HE].

[15] N. R. Tanvir, A. J. Levan, A. S. Fruchter, J. Hjorth, R. A. Hounsell, K. Wiersema, and R. L. Tunnicliffe, A 'kilonova' associated with the short-duration γ-ray burst GRB130603B, Nature **500**, 547 (2013), arXiv:1306.4971 [astro-ph.HE].

[16] E. Berger, W. Fong, and R. Chornock, An r-process Kilonova Associated with the Short-hard GRB 130603B, Ap. J. Lett. **774**, L23 (2013), arXiv:1306.3960 [astro-ph.HE].

[17] Z.-P. Jin, X. Li, Z. Cano, S. Covino, Y.-Z. Fan, and D.-M. Wei, The Light Curve of the Macronova Associated with the Long-Short Burst GRB 060614, Astrophys. J. Lett. **811**, L22 (2015), arXiv:1507.07206 [astro-ph.HE].

[18] B. Yang, Z.-P. Jin, X. Li, S. Covino, X.-Z. Zheng, K. Hotokezaka, Y.-Z. Fan, T. Piran, and D.-M. Wei, A possible Macronova in the late afterglow of the 'long-short' burst GRB 060614, Nature Commun. **6**, 7323 (2015), arXiv:1503.07761 [astro-ph.HE].

[19] J. C. Rastinejad *et al.*, A kilonova following a long-duration gamma-ray burst at 350 Mpc, Nature **612**, 223 (2022), arXiv:2204.10864 [astro-ph.HE].

[20] E. Troja *et al.*, A nearby long gamma-ray burst from a merger of compact objects, Nature **612**, 228 (2022), arXiv:2209.03363 [astro-ph.HE].

[21] J. Yang, S. Ai, B.-B. Zhang, B. Zhang, Z.-K. Liu, X. I. Wang, Y.-H. Yang, Y.-H. Yin, Y. Li, and H.-J. Lü, A long-duration gamma-ray burst with a peculiar origin, Nature **612**, 232 (2022), arXiv:2204.12771 [astro-ph.HE].

[22] D. Kasen, B. Metzger, J. Barnes, E. Quataert, and E. Ramirez-Ruiz, Origin of the heavy elements in binary neutron-star mergers from a gravitational-wave event, Nature **551**, 80 (2017), arXiv:1710.05463 [astro-ph.HE].

[23] J. Barnes, Y. L. Zhu, K. A. Lund, T. M. Sprouse, N. Vassh, G. C. McLaughlin, M. R. Mumpower, and R. Surman, Kilonovae Across the Nuclear Physics Landscape: The Impact of Nuclear Physics Uncertainties on r-process-powered Emission, Astrophys. J. **918**, 44 (2021), arXiv:2010.11182 [astro-ph.HE].

[24] Y. L. Zhu, K. Lund, J. Barnes, T. M. Sprouse, N. Vassh, G. C. McLaughlin, M. R. Mumpower, and R. Surman, Modeling Kilonova Light Curves: Dependence on Nuclear Inputs, Astrophys. J. **906**, 94 (2021), arXiv:2010.03668 [astro-ph.HE].

[25] K. Kiuchi, K. Kyutoku, Y. Sekiguchi, M. Shibata, and T. Wada, High resolution numerical relativity simulations for the merger of binary magnetized neutron stars, Phys. Rev. D **90**, 041502 (2014), arXiv:1407.2660 [astro-ph.HE].

[26] S. Wanajo, Y. Sekiguchi, N. Nishimura, K. Kiuchi, K. Kyutoku, and M. Shibata, Production of All the r-process Nuclides in the Dynamical Ejecta of Neutron Star Mergers, Astrophys.J.Lett. **789**, L39 (2014), arXiv:1402.7317 [astro-ph.SR].

[27] F. Foucart, M. D. Duez, F. Hebert, L. E. Kidder, H. P. Pfeiffer, and M. A. Scheel, Monte-Carlo neutrino transport in neutron star merger simulations, Astrophys. J. Lett. **902**, L27 (2020), arXiv:2008.08089 [astro-ph.HE].

[28] J. Lippuner and L. F. Roberts, r-process Lanthanide Production and Heating Rates in Kilonovae, Astrophys. J. **815**, 82 (2015), arXiv:1508.03133 [astro-ph.HE].

[29] J. Barnes and D. Kasen, Effect of a High Opacity on the Light Curves of Radioactively Powered Transients from Compact Object Mergers, Astrophys. J. **775**, 18 (2013), arXiv:1303.5787 [astro-ph.HE].

[30] M. Ruffert, H.-T. Janka, K. Takahashi, and G. Schaefer, Coalescing neutron stars - a step towards physical models. II. Neutrino emission, neutron tori, and gamma-ray bursts., Astron. Astrophys. **319**, 122 (1997), astro-ph/9606181.

[31] S. Rosswog and M. Liebendörfer, High resolution calculations of merging neutron stars. 2: Neutrino emission, Mon. Not. Roy. Astr. Soc. **342**, 673 (2003), arXiv:astro-ph/0302301 [astro-ph].

[32] M. B. Deaton, M. D. Duez, F. Foucart, E. O'Connor, C. D. Ott, L. E. Kidder, C. D. Muhlberger, M. A. Scheel, and B. Szilágyi, Black Hole-Neutron Star Mergers with a Hot Nuclear Equation of State: Outflow and Neutrino-Cooled Disk for a Low-Mass, High-Spin Case, Astrophys. J. **776**, 47 (2013), arXiv:1304.3384 [astro-ph.HE].

[33] D. Neilsen, S. L. Liebling, M. Anderson, L. Lehner, E. O'Connor, and C. Palenzuela, Magnetized Neutron Stars With Realistic Equations of State and Neutrino Cooling, Phys. Rev. D **89**, 104029 (2014), arXiv:1403.3680 [gr-qc].

[34] Y. Sekiguchi, K. Kiuchi, K. Kyutoku, and M. Shibata, Dynamical mass ejection from binary neutron star mergers: Radiation-hydrodynamics study in general relativity, Phys. Rev. D **91**, 064059 (2015), arXiv:1502.06660 [astro-ph.HE].

[35] D. Radice, F. Galeazzi, J. Lippuner, L. F. Roberts, C. D. Ott, and L. Rezzolla, Dynamical Mass Ejection from Binary Neutron Star Mergers, Mon. Not. Roy. Astr. Soc.

460, 3255 (2016), arXiv:1601.02426 [astro-ph.HE].

[36] F. Foucart, E. O'Connor, L. Roberts, M. D. Duez, R. Haas, L. E. Kidder, C. D. Ott, H. P. Pfeiffer, M. A. Scheel, and B. Szilágyi, Post-merger evolution of a neutron star-black hole binary with neutrino transport, Phys. Rev. D **91**, 124021 (2015), arXiv:1502.04146 [astro-ph.HE].

[37] L. Sun, M. Ruiz, S. L. Shapiro, and A. Tsokaros, Jet launching from binary neutron star mergers: Incorporating neutrino transport and magnetic fields, Phys. Rev. D **105**, 104028 (2022), arXiv:2202.12901 [astro-ph.HE].

[38] F. Foucart, E. O'Connor, L. Roberts, L. E. Kidder, H. P. Pfeiffer, and M. A. Scheel, Impact of an improved neutrino energy estimate on outflows in neutron star merger simulations, Phys. Rev. **D94**, 123016 (2016), arXiv:1607.07450 [astro-ph.HE].

[39] D. Radice, S. Bernuzzi, A. Perego, and R. Haas, A new moment-based general-relativistic neutrino-radiation transport code: Methods and first applications to neutron star mergers, Mon. Not. Roy. Astron. Soc. **512**, 1499 (2022), arXiv:2111.14858 [astro-ph.HE].

[40] K. Kawaguchi, S. Fujibayashi, and M. Shibata, Monte Carlo–based relativistic radiation hydrodynamics code with a higher-order scheme, Phys. Rev. D **107**, 023026 (2023), arXiv:2209.12472 [astro-ph.HE].

[41] F. Foucart, Monte Carlo closure for moment-based transport schemes in general relativistic radiation hydrodynamic simulations, Mon. Not. Roy. Astr. Soc. **475**, 4186 (2018), arXiv:1708.08452 [astro-ph.HE].

[42] M. R. Izquierdo, J. F. Abalos, and C. Palenzuela, Guided moments formalism: A new efficient full-neutrino treatment for astrophysical simulations, Phys. Rev. D **109**, 043044 (2024), arXiv:2312.09275 [astro-ph.HE].

[43] F. Foucart, R. Haas, M. D. Duez, E. O'Connor, C. D. Ott, L. Roberts, L. E. Kidder, J. Lippuner, H. P. Pfeiffer, and M. A. Scheel, Low mass binary neutron star mergers: Gravitational waves and neutrino emission, Phys. Rev. D **93**, 044019 (2016), arXiv:1510.06398 [astro-ph.HE].

[44] P. Mösta, D. Radice, R. Haas, E. Schnetter, and S. Bernuzzi, A magnetar engine for short GRBs and kilonovae, Astrophys. J. Lett. **901**, L37 (2020), arXiv:2003.06043 [astro-ph.HE].

[45] S. Curtis, P. Bosch, P. Mösta, D. Radice, S. Bernuzzi, A. Perego, R. Haas, and E. Schnetter, Magnetized Outflows from Short-lived Neutron Star Merger Remnants Can Produce a Blue Kilonova, Astrophys. J. Lett. **961**, L26 (2024), arXiv:2305.07738 [astro-ph.HE].

[46] L. Baiotti and L. Rezzolla, Binary neutron-star mergers: a review of Einstein's richest laboratory (2016), arXiv:1607.03540 [gr-qc].

[47] D. Radice, S. Bernuzzi, and A. Perego, The Dynamics of Binary Neutron Star Mergers and GW170817, Ann. Rev. Nucl. Part. Sci. **70**, 95 (2020), arXiv:2002.03863 [astro-ph.HE].

[48] K. Kiuchi, General relativistic magnetohydrodynamics simulations for binary neutron star mergers, arXiv (2024), arXiv:2405.10081 [astro-ph.HE].

[49] Y. Sekiguchi, K. Kiuchi, K. Kyutoku, M. Shibata, and K. Taniguchi, Dynamical mass ejection from the merger of asymmetric binary neutron stars: Radiation-hydrodynamics study in general relativity, Phys. Rev. **D93**, 124046 (2016), arXiv:1603.01918 [astro-ph.HE].

[50] s. F. Foucart, M. D. Duez, L. E. Kidder, H. P. Pfeiffer, and M. A. Scheel, Dynamical ejecta from binary neutron star mergers: Impact of residual eccentricity and equation of state implementation, arXiv (2024), arXiv:2404.18674 [astro-ph.HE].

[51] E. O'Connor, An Open-source Neutrino Radiation Hydrodynamics Code for Core-collapse Supernovae, ApJ Suppl. **219**, 24 (2015), arXiv:1411.7058 [astro-ph.HE].

[52] S. Typel, M. Oertel, and T. Klähn, CompOSE CompStar online supernova equations of state harmonising the concert of nuclear physics and astrophysics compose.obspm.fr, Phys. Part. Nucl. **46**, 633 (2015), arXiv:1307.5715 [astro-ph.SR].

[53] M. Oertel, M. Hempel, T. Klähn, and S. Typel, Equations of state for supernovae and compact stars, Rev. Mod. Phys. **89**, 015007 (2017), arXiv:1610.03361 [astro-ph.HE].

[54] E. O'Connor and C. D. Ott, A New Open-Source Code for Spherically-Symmetric Stellar Collapse to Neutron Stars and Black Holes, Class. Quantum Grav. **27**, 114103 (2010), arXiv:0912.2393 [astro-ph.HE].

[55] H. P. Pfeiffer, L. E. Kidder, M. A. Scheel, and S. A. Teukolsky, A multidomain spectral method for solving elliptic equations, Comput. Phys. Commun. **152**, 253 (2003), gr-qc/0202096.

[56] F. Foucart, L. E. Kidder, H. P. Pfeiffer, and S. A. Teukolsky, Initial data for black hole-neutron star binaries: A Flexible, high-accuracy spectral method, Phys. Rev. D **77**, 124051 (2008), arXiv:0804.3787 [gr-qc].

[57] M. D. Duez, F. Foucart, L. E. Kidder, H. P. Pfeiffer, M. A. Scheel, and S. A. Teukolsky, Evolving black hole-neutron star binaries in general relativity using pseudospectral and finite difference methods, Phys. Rev. D **78**, 104015 (2008), arXiv:0809.0002 [gr-qc].

[58] F. Foucart, M. B. Deaton, M. D. Duez, L. E. Kidder, I. MacDonald, C. D. Ott, H. P. Pfeiffer, M. A. Scheel, B. Szilágyi, and S. A. Teukolsky, Black hole-neutron star mergers at realistic mass ratios: Equation of state and spin orientation effects, Phys. Rev. D **87**, 084006 (2013), arXiv:1212.4810 [gr-qc].

[59] L. Lindblom, M. A. Scheel, L. E. Kidder, R. Owen, and O. Rinne, A new generalized harmonic evolution system, Class. Quantum Grav. **23**, S447 (2006), gr-qc/0512093.

[60] B. Szilágyi, L. Lindblom, and M. A. Scheel, Simulations of binary black hole mergers using spectral methods, Phys. Rev. D **80**, 124010 (2009), arXiv:0909.3557 [gr-qc].

[61] F. Banyuls, J. A. Font, J. M. Ibáñez, J. M. Martí, and J. A. Miralles, Numerical {3 + 1} General Relativistic Hydrodynamics: A Local Characteristic Approach, Astrophys. J. **476**, 221 (1997).

[62] X.-D. Liu, S. Osher, and T. Chan, Weighted essentially non-oscillatory schemes, J. Comput. Phys. **115**, 200 (1994).

[63] G.-S. Jiang and C.-W. Shu, Efficient implementation of weighted eno schemes, J. Comput. Phys. **126**, 202 (1996).

[64] R. Borges, M. Carmona, B. Costa, and W. S. Don, An improved weighted essentially non-oscillatory scheme for hyperbolic conservation laws, J. Comput. Phys. **227**, 3191 (2008).

[65] A. Harten, P. D. Lax, and B. van Leer, On Upstream Differencing and Godunov-Type Schemes for Hyperbolic Conservation Laws, SIAM Review **25**, 35 (1983).

[66] K. S. Thorne, Relativistic radiative transfer - Moment formalisms, Mon. Not. Roy. Astr. Soc. **194**, 439 (1981).

[67] M. Shibata, K. Kiuchi, Y. Sekiguchi, and Y. Suwa, Trun-

cated Moment Formalism for Radiation Hydrodynamics in Numerical Relativity, Progress of Theoretical Physics **125**, 1255 (2011), arXiv:1104.3937 [astro-ph.HE].

[68] F. Foucart, M. D. Duez, F. Hebert, L. E. Kidder, P. Kovarik, H. P. Pfeiffer, and M. A. Scheel, Implementation of Monte Carlo Transport in the General Relativistic SpEC Code, Astrophys. J. **920**, 82 (2021), arXiv:2103.16588 [astro-ph.HE].

[69] F. Schianchi, H. Gieg, V. Nedora, A. Neuweiler, M. Ujevic, M. Bulla, and T. Dietrich, M1 neutrino transport within the numerical-relativistic code BAM with application to low mass binary neutron star mergers, Phys. Rev. D **109**, 044012 (2024), arXiv:2307.04572 [gr-qc].

[70] E. Audit, P. Charrier, J. . Chièze, and B. Dubroca, A radiation-hydrodynamics scheme valid from the transport to the diffusion limit, ArXiv Astrophysics e-prints (2002), astro-ph/0206281.

[71] H. Andresen, E. P. O'Connor, O. E. Andersen, and S. M. Couch, Grey Two-moment Neutrino Transport: Comprehensive Tests and Improvements for Supernova Simulations, arXiv (2024), arXiv:2402.18303 [astro-ph.HE].

[72] F. Foucart, P. Moesta, T. Ramirez, A. J. Wright, S. Darbha, and D. Kasen, Estimating outflow masses and velocities in merger simulations: Impact of r-process heating and neutrino cooling, Phys. Rev. D **104**, 123010 (2021), arXiv:2109.00565 [astro-ph.HE].