

Monte-Carlo neutrino transport in neutron star merger simulations

FRANCOIS FOUCART,¹ MATTHEW D. DUEZ,² FRANCOIS HEBERT,³ LAWRENCE E. KIDDER,⁴ HARALD P. PFEIFFER⁵
AND MARK A. SCHEEL³

¹*Department of Physics & Astronomy, University of New Hampshire, Durham, New Hampshire 03824, USA*

²*Department of Physics & Astronomy, Washington State University, Pullman, Washington 99164, USA*

³*TAPIR, Walter Burke Institute for Theoretical Physics, MC 350-17, California Institute of Technology, Pasadena, California 91125, USA*

⁴*Center for Radiophysics and Space Research, Cornell University, Ithaca, New York, 14853, USA*

⁵*Max-Planck-Institut für Gravitationsphysik, Albert-Einstein-Institut, D-14476 Potsdam, Germany*

ABSTRACT

Gravitational waves and electromagnetic signals from merging neutron star binaries provide valuable information about the the properties of dense matter, the formation of heavy elements, and high-energy astrophysics. To fully leverage observations of these systems, we need numerical simulations that provide reliable predictions for the properties of the matter unbound in these mergers. An important limitation of current simulations is the use of approximate methods for neutrino transport that do not converge to a solution of the transport equations as numerical resolution increases, and thus have errors that are impossible to quantify. Here, we report on a first simulation of a binary neutron star merger that uses Monte-Carlo techniques to directly solve the transport equations in low-density regions. In high-density regions, we use approximations inspired by implicit Monte-Carlo to greatly reduce the cost of simulations, while only introducing errors quantifiable through more expensive convergence studies. We simulate an unequal mass neutron star binary merger up to 5 ms past merger, and report on the properties of the matter and neutrino outflows. Finally, we compare our results to the output of our best approximate ‘M1’ transport scheme, demonstrating that an M1 scheme that carefully approximates the neutrino energy spectrum only leads to $\sim 10\%$ uncertainty in the composition and velocity of the ejecta, and $\sim 20\%$ uncertainty in the ν_e and $\bar{\nu}_e$ luminosities and energies. The most significant disagreement found between M1 and Monte-Carlo results is a factor of ~ 2 difference in the luminosity of heavy-lepton neutrinos.

Keywords: neutron stars, r-process, neutrinos, computational methods

1. INTRODUCTION

Gravitational wave (GW) and electromagnetic (EM) observations of neutron star mergers provide us with important information about the properties of dense matter, the synthesis of heavy nuclei, and high-energy astrophysics, as demonstrated by the first detection of GWs from a neutron star merger (Abbott et al. 2017) and associated EM observations (e.g. Abbott et al. (2017); Kasliwal et al. (2017); Chornock et al. (2017); Smartt et al. (2017); Soares-Santos et al. (2017); Cowperthwaite et al. (2017)). The UV/optical/infrared signal powered by r-process nucleosynthesis in the matter unbound by the merger (kilonova) (Li & Paczynski 1998; Roberts et al. 2011) is of particular interest for nuclear astrophysics for the information that it provides about nucleosynthesis in mergers, and about the equation of state of neutron stars. To extract information from observed kilonovae, however, reliable theoretical models are required. These in turn rely on a good understanding of the properties of the matter ejected during and after merger (Barnes & Kasen 2013), and of nuclear physics in the neutron-rich ejecta (Barnes et al. 2016).

Simulations are our main source of information about merger outflows. However, they suffer from important limitations: they do not capture the growth of magnetic fields from realistic initial strengths (Kiuchi et al. 2015), and use approximate methods for neutrino transport (Foucart et al. 2016; Foucart et al. 2018). As a result, they can miss important physical processes: magnetic fields heat the remnant, drive angular momentum transport in the system, and produce most post-merger outflows, while neutrinos cool the remnant and drive the evolution of its composition.

For neutrino transport, the main issue is the high dimensionality of the problem. Ideally, one would evolve the neutrino distribution function using Boltzmann’s equations of radiation transport. Unfortunately, this is a function of time, position, neutrino energy and momentum, making this a 7-dimensional problem for each neutrino species. The problem is further complicated by the existence of stiff coupling terms between neutrinos and nucleons in dense and hot regions. Merger simulations first included neutrino effects through leakage schemes (Sekiguchi 2010; Deaton et al. 2013) that account for the local cooling effects of neutrinos at an order-of-magnitude level. More recently, grey two-moment schemes (‘M1’ schemes) that evolve the neutrino energy and momentum density but use approximate analytical closures for the neutrino pressure and energy spectrum have been implemented (Shibata et al. 2011; Wanajo et al. 2014; Foucart et al. 2015), as well as a mixed leakage-one moment scheme (Radice et al. 2016). Simulations using M1 schemes have clearly demonstrated that leakage is insufficient to capture the composition of matter outflows in mergers (Wanajo et al. 2014), the most important parameter to determine the outcome of nucleosynthesis in merger outflows. Yet M1 schemes themselves show that outflow composition and neutrino luminosities have non-negligible dependencies on the exact choice of analytical closures (Foucart et al. 2016; Foucart et al. 2018), and that standard closures lead to numerical artifacts in simulations (e.g. neutrino shocks in polar regions). Most importantly, while M1 schemes may be sufficient for many purposes, there is no way to test their accuracy without comparison with a solution to the transport equations, as they do not converge to the correct solution when increasing resolution.

In recent years, general relativistic Monte-Carlo (MC) algorithms have risen as a tempting alternative to provide low-cost neutrino transport in merger and post-merger simulations (Richers et al. 2015; Ryan et al. 2015; Miller et al. 2019). Building on our implementation of a MC algorithm as a closure to a M1 code (Foucart 2018), we present here a first MC transport algorithm for fully general relativistic merger simulations, as well as a first simulation of merging neutron stars with MC transport. The aim of this code is to provide cheap yet reasonably accurate solutions to the transport problem, within a framework that converges to the correct physical solution as more computational resources become available. This code can be used for direct simulations of neutron star mergers, as well as to test the accuracy of existing M1 and leakage results, and can thus greatly improve our understanding of the merger outflows. We note that low-cost MC transport is possible in neutron star mergers because neutrinos impact the evolution of the merger remnant on time scales that are long compared to the simulation time step; e.g. the cooling time scale of a post-merger accretion disk is $\gtrsim (10 - 100)$ ms (Deaton et al. 2013). Additionally, neutrino-driven winds have mass outflows of only $\dot{M} \sim 0.1 M_{\odot} s^{-1}$ even immediately after merger (Foucart et al. 2016), and remain subdominant with respect to viscous/magnetically-driven winds at later times (Just et al. 2015). MC sampling errors thus have a limited impact on the dynamics of the remnant.

2. METHODS

We perform general relativistic radiation hydrodynamics simulations of merging neutron stars with masses $M_1 = 1.27 M_{\odot}$, $M_2 = 1.58 M_{\odot}$, using the ‘DD2’ equation of state from Hempel et al. (2012). The neutron stars have radii $R \sim 13.2$ km and zero spin. We generate initial data with the *Spells* code (Pfeiffer et al. 2003; Foucart et al. 2008) four orbits before merger, and end the simulations ~ 5 ms after merger. This is sufficient to observe tidal ejection and the production of a neutrino-driven wind. Evolution over longer timescales would require us to model angular momentum transport and heating due to turbulence, through magnetic field evolution or the use of a viscous model. We evolve neutrinos using either our new MC transport scheme or approximate ‘M1’ transport. In this section, we briefly discuss our numerical methods. Readers interested in our results and their implications may skip to Section 3.

Simulations are performed with the SpEC code¹. SpEC evolves Einstein’s equations in the Generalized Harmonics formalism (Lindblom et al. 2006) using pseudospectral methods with adaptive mesh refinement (Szilágyi 2014). The fluid equations are evolved on a Cartesian grid, using high-order finite volume shock-capturing methods. We use the methods of Duez et al. (2008); Foucart et al. (2013), except that we allow the time step on the pseudospectral grid to be smaller than the time step on the finite volume grid. At the resolution used in this manuscript, time stepping

¹ <https://www.black-holes.org/code/SpEC.html>

errors remain small compared to other sources of errors. The finite volume grid has a spacing $\Delta x_{\text{FV}} = 188$ m at the beginning of the evolution, and $\Delta x_{\text{FV}} = 200$ m after merger. We use fixed mesh refinement after merger, doubling the grid spacing at each level. Each of our 4 refinement levels has $200 \times 200 \times 176$ points.

We perform three simulations. The first uses the two-moment ('M1') transport scheme from [Foucart et al. \(2015, 2016\)](#), evolving the neutrino energy density, momentum density, and number density. It uses approximate analytical closures to estimate the pressure tensor and energy spectrum. The others use MC transport, with different numbers of packets in order to estimate sampling errors in the simulations. The core of our MC algorithm is described in [Foucart \(2018\)](#). Here, we summarize the main components of the algorithm, and new features needed to obtain stable and accurate evolution in merging neutron stars.

All simulations assume that neutrino-matter interactions can be described by an emissivity η , absorption opacity κ_a , and elastic scattering opacity κ_s . We use tabulated values produced with the NuLib library ([O'Connor 2015](#)). The table includes reaction rates for the charged current reactions

$$p + e^- \leftrightarrow n + \nu_e \quad n + e^+ \leftrightarrow p + \bar{\nu}_e; \quad (1)$$

scattering of neutrinos on protons, neutrons, α -particles and heavy nuclei; and, for the muon and tau (anti)neutrinos only, $e^+e^- \leftrightarrow \nu\bar{\nu}$ and Bremsstrahlung. Under these assumptions, muon and tau (anti)neutrinos all behave in the same manner, and we treat them as a single species (called ν_x). The table is logarithmically spaced in neutrino energies (16 groups up to $E = 528$ MeV), density (86 points in $[1e6, 3.2e15]$ g/cm³) and temperature (65 points in $[0.05, 150]$ MeV), and linearly spaced in the electron fraction Y_e (51 points in $[0.01, 0.6]$).

The basic idea of the MC method is to sample, for each neutrinos species, the distribution function $f_{(\nu)}(t, x^i, p^\mu)$ using a discrete number P of packets that each represent a number N of neutrinos. More precisely,

$$f_{(\nu)}(t, x^i, p_\mu) \approx \sum_{k=1}^P N_k \delta^3(x^i - x_k^i) \delta^3(p_i - p_i^k) \quad (2)$$

with (t, x^i) the coordinate time and position, p^μ the 4-momentum, (x_k^i, p_i^k) the position and spatial components of the momentum one-form of packet k at time t , and N_k the number of neutrinos that this packet represents. Packets are created from an isotropic distribution in the fluid frame, propagated along null geodesics, and scattered/absorbed with probabilities set by κ_a, κ_s (see [Foucart \(2018\)](#)). We use a split operator method where the fluid and metric are evolved first, and neutrino packets second. We also allow the MC code to take time steps covering multiple fluid steps when possible: the MC code aims to take steps with $c\Delta t = (0.5 - 1)\Delta x_{\text{FV}}$. Neutrinos deposit/remove energy and momentum from the fluid at the end of the MC step. For the interactions considered here, changes in the fluid variables due to neutrino-matter interactions follow the equations for conservation of energy, momentum, and lepton number:

$$\nabla_\mu T_{\text{fl}}^{\mu\nu} = -\eta u^\nu + \sum_k (\kappa_a J_k u^\mu + [\kappa_a + \kappa_s] H_k^\mu); \quad \nabla_\mu (\rho_0 Y_e u^\mu) = - \sum_s s_s \eta_{N,s} + \sum_k s_k \kappa_a \frac{J_k}{\nu_k} \quad (3)$$

where $T_{\text{fl}}^{\mu\nu}$ is the stress-energy tensor of the fluid, η the total neutrino energy emissivity, $\eta_{N,s}$ the number emissivity of neutrinos of species s , u^μ the fluid 4-velocity, J_k, H_k^μ the energy density and momentum density of neutrinos in packet k , and ν_k their average energy, in the fluid frame. The electron fraction $Y_e = n_p/(n_p + n_n)$ (with $n_{p,n}$ the number density of protons and neutrons) parametrizes the composition of the fluid. $s = 1$ for electron neutrinos, -1 for electron antineutrinos, and 0 otherwise. We discuss below how these terms are estimated.

The main components of this code were used in [Foucart et al. \(2018\)](#) to estimate errors in the M1 method. However, at the time we could not use MC methods in the densest, hottest regions of the merger. Hot regions are problematic for MC algorithms as large emissivities and opacities cause rapid creation and destruction of packets, and numerical instabilities. To avoid this, we adapt to the merger problems the ideas of implicit MC, partially following the work of [Fleck & Cummings \(1971\)](#). Once the absorption opacity becomes too large, the emissivities and absorption coefficients are modified according to

$$\kappa'_a = (1 - \alpha)\kappa_a; \quad \kappa'_s = \kappa_s + \alpha\kappa_a; \quad \eta' = (1 - \alpha)\eta \quad (4)$$

for some constant α , thus reducing (η, κ_a) without modifying the neutrino diffusion rate or equilibrium energy density. As opposed to [Fleck & Cummings \(1971\)](#), we choose α separately for each species and energy bin. We require that

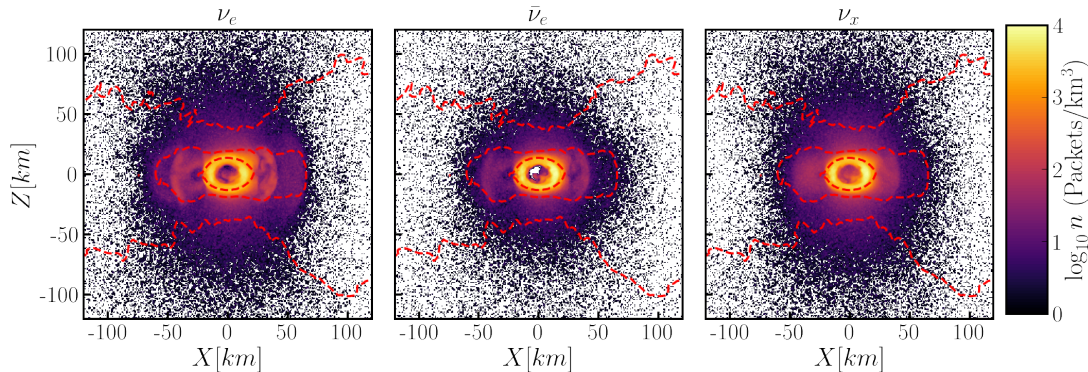


Figure 1. Number density of MC packets for each species, in a vertical slice at the end of the high-resolution simulation. The dashed red lines are density contours at $10^{9,11,13}$ g/cm³. Fig. 3 shows the same snapshot. Most packets are in the hot region close to the stellar surface. Few packets are needed far away from the remnant, or in cold regions of the neutron star.

$\kappa'_a \Delta t < 0.5$, effectively guaranteeing that the equilibration time scale is always at least a few time steps. We also require

$$\alpha > \frac{\beta \Delta t \kappa_a}{1 + \beta \Delta t \kappa_a}; \quad \beta = \min \left(\frac{du_{\text{rad}}}{du_{\text{fl}}} \Big|_{(\rho, Y_e)}, \frac{dn_{\text{rad}}}{dn_{\text{fl}}} \Big|_{(\rho, T)} \right) \quad (5)$$

with $u_{\text{fl,rad}}$ and $n_{\text{fl,rad}}$ the energy density and lepton number density of the fluid and of neutrinos in equilibrium with that fluid. This aims to prevent instabilities in the joint evolution of the fluid and neutrino radiation. While implicit MC in Fleck & Cummings (1971) was specifically designed to match a given time discretization of the original transport equations, our scheme only does so in the limit $\Delta t \rightarrow 0$. This is because we couple neutrinos to both the composition and internal energy of the fluid (as opposed to the energy only for photons), and because Fleck & Cummings (1971) used the same α for all energy groups, and an effective *inelastic* scattering cross-section such that neutrinos have an equilibrium energy spectrum post-scattering.

As we fix the number of MC packets emitted per unit time (see below), the effective number of packets in optically thick cells is $\propto (\kappa'_a)^{-1}$. Decreasing κ'_a allows us to reduce statistical errors without increasing the cost of simulations, while introducing a small error that converges away with resolution. To test this method in circumstances reasonably similar to neutron star merger conditions, we consider the evolution of a post-bounce supernova remnant, performed with much coarser resolution than our merger simulations (Abdikamalov et al. 2012; Foucart 2018). We find that the ν_e and $\bar{\nu}_e$ luminosities are very accurate, while the ν_x luminosity is impacted at the $\sim 20\%$ level. This gives us confidence that even in an unfavorable configuration for this approximation, the method provides reasonably accurate results. Nevertheless, improving this part of the algorithm without drastically increasing the cost of simulations or causing instabilities may be useful in the future. We note that transforming the absorption opacity into a scattering opacity would not provide much of a gain if we always treated scattering events explicitly. However, we developed in Foucart (2018) a cost-effective method to approximate many scatterings as a diffusion process. A different implementation of that idea is also used by Richers et al. (2015).

To increase stability, we also compute η, κ_a, κ_s using *predicted values* for the temperature and composition of the fluid. These are computed by solving implicitly the energy and lepton number conservation equations (Eq.3), using the expectation values of the source terms and neglecting changes in the momentum of the fluid.

Coupling terms between neutrinos and the fluid are evaluated by integrating the right-hand sides of Eq. 3 over a MC time step. This requires integrals, for each packet, of J_k and H_k^μ over the worldline of a packet. In the MC formalism, and for any time interval Δt between neutrino-matter interactions, these are simply (see Foucart (2018))

$$\int dt J_k = \Delta t N_k \frac{\nu_k^2}{p_k^t \sqrt{-g} \Delta V}; \quad \int dt H_k^\mu = \Delta t N_k \frac{\nu_k}{p_k^t \Delta V \sqrt{-g}} p_\nu^k (g^{\mu\nu} + u^\mu u^\nu), \quad (6)$$

with $\delta^3(x_k^i - x^i) = (\Delta V)^{-1}$ in the finite volume discretization, and ΔV the coordinate volume of a cell.

Finally, we need to choose the number of neutrinos N_k represented by a packet. We could take the total energy $E_k = N_k \nu_k$ of a packet to be constant, or fix the total number of packets. However, none of these choices proved sufficient to limit noise in our low-cost simulations. Instead, we choose a *minimum energy* (E_{min}), a *desired number*

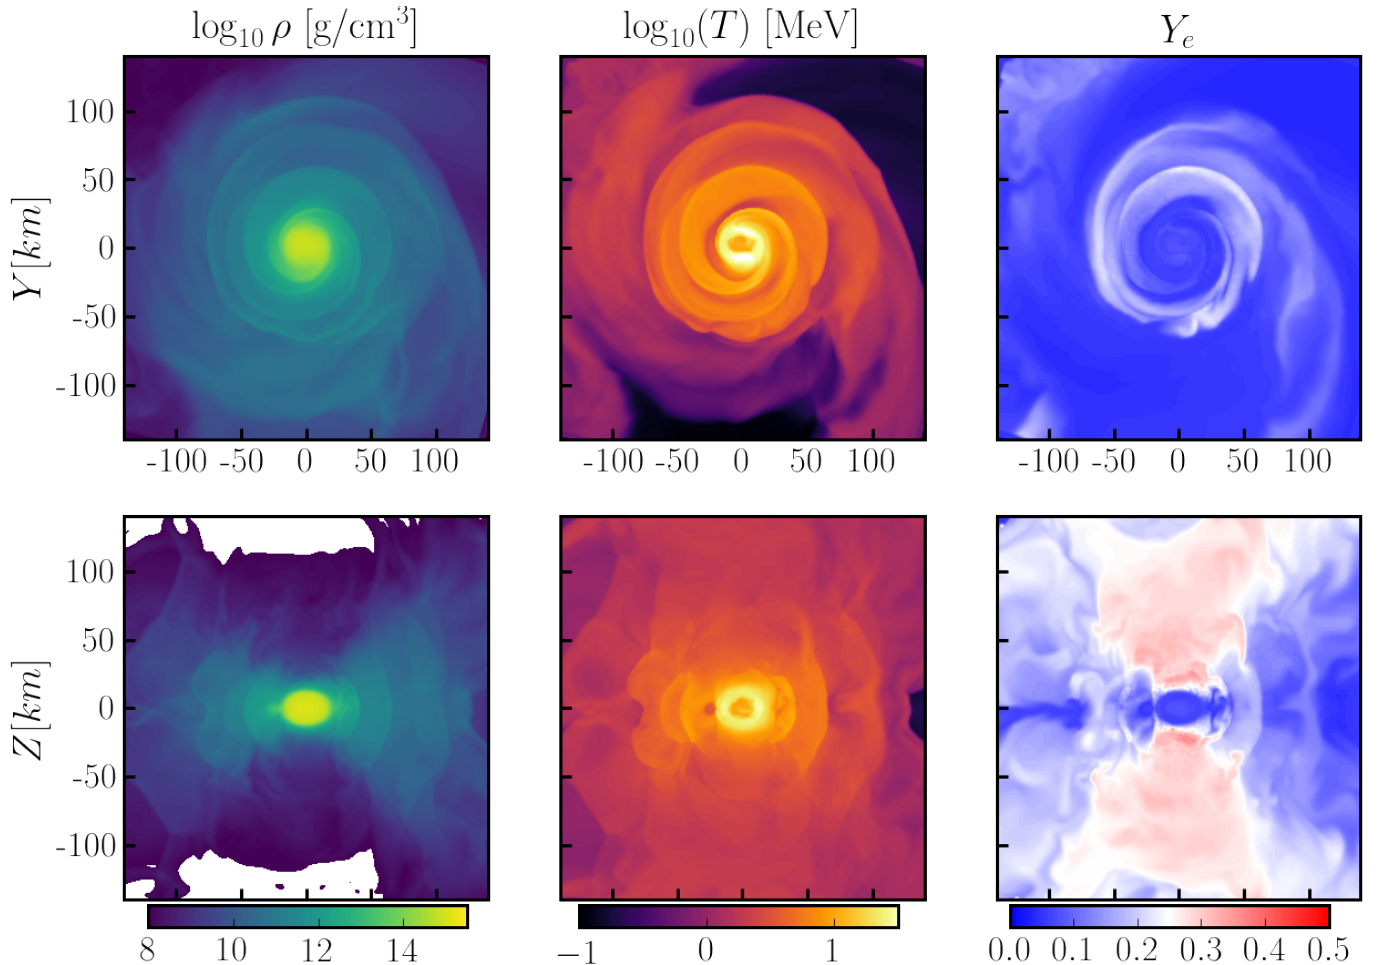


Figure 2. Snapshot of the MC simulation with the largest number of packets, 5 ms after merger. We show the density (*Left*), temperature (*Center*), and electron fraction (*Right*) in the equatorial plane (*Top*) and a vertical slice passing through the center of the remnant (*Bottom*). We can see the post-merger hypermassive neutron star at the center of the figures, surrounded by a dense torus. Low-density outflows are launched mostly along the edges of the torus.

of packets to be emitted within a grid cell per light-crossing time of the cell ($N_{\text{em,target}}$), and a *maximum number of packets* per species over the whole domain ($N_{\text{max,tot}}$). We first set E_k to get the desired $N_{\text{em,target}}$, and reset it to E_{min} if $E_k < E_{\text{min}}$. Whenever the number of packets of a given species grows above $N_{\text{max,tot}}$, we increase E_{min} for that species by 10%, and resample the existing packets: 10% of the packets are randomly destroyed, while the surviving packets multiply their N_k by $(1.1)^{-1}$. This guarantees a minimum number of packets per cell in hot regions, without reducing too much the number of packets elsewhere. We use $N_{\text{max,tot}} = (12, 3) \times 10^7$, $E_{\text{min}} = (1, 4) \times 10^{-14} M_{\odot} c^2$, and $N_{\text{em,target}} = (100, 25)$ for our two MC simulations. The factor of 4 between simulations halves the expected sampling noise. Even with that method, a majority of packets are concentrated in the hottest regions (see Fig. 1). Accordingly, instead of evolving MC packets on the processor responsible for the fluid cell in which they are located (as in Foucart et al. (2018)), our code can offload the evolution of all MC packets within a given finite volume cell to another processor, as needed for load-balancing.

3. SIMULATION RESULTS

The qualitative features of the evolution are similar in all simulations. The neutron stars perform ~ 4 orbits before merging. Due to the mass asymmetry, a significant amount of mass is unbound in a cold, neutron-rich tidal tail, while hotter, less neutron-rich matter is ejected following the collision of the neutron star cores. Within the following ~ 5 ms, a massive torus forms around the dense merger remnant, while a neutrino-driven wind develops along the edges of the torus and in the polar regions. Figure 2 shows the main properties of the remnant 5 ms after merger:

Table 1. Matter outflows in our 3 simulations. We provide the total ejected mass, average Y_e of the ejecta, and average asymptotic velocity of the ejecta. We also provide these same quantities for the polar ejecta, defined as unbound material with a velocity vector inclined by less than 45° with respect to the rotation axis.

Sim	$M_{\text{ej}} (10^{-3} M_\odot)$	$\langle Y_e \rangle_{\text{ej}}$	$\langle v \rangle_{\text{ej}}$	$M_{\text{ej,pol}} (10^{-3} M_\odot)$	$\langle Y_e \rangle_{\text{ej,pol}}$	$\langle v \rangle_{\text{ej,pol}}$
MC-low	11.56	0.135	0.214c	0.53	0.228	0.192c
MC-high	8.25	0.130	0.206c	0.49	0.234	0.191c
M1	8.31	0.129	0.201c	0.34	0.259	0.184c

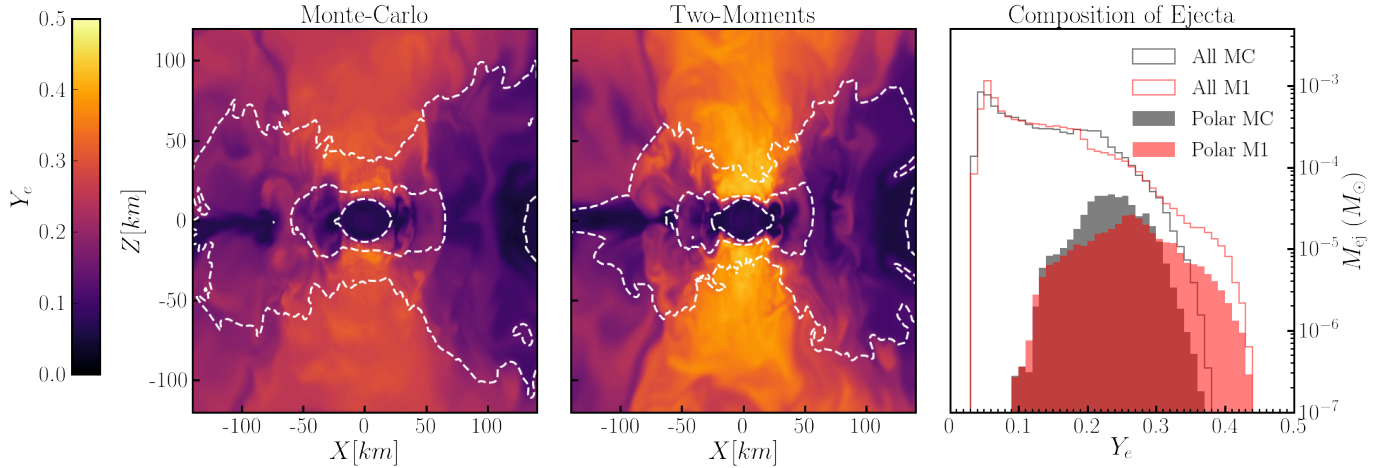


Figure 3. *Left:* Electron fraction of the remnant 5 ms after merger in the MC simulation with the highest packet count. Dashed white lines are density contours at $10^{9,11,13} \text{ g/cm}^3$. *Center:* The same, for the M1 simulation. Regions just outside of the neutron star and disk are slightly denser in the MC simulations (possibly due to the higher ν_x luminosity), and matter is more neutron rich in the outflow regions. *Right:* Composition of the unbound material. Filled histograms show the polar outflows only.

the density, temperature, and electron fraction. The compact remnant reaches temperatures up to $\sim 60 \text{ MeV}$, and remains neutron rich ($Y_e \lesssim 0.1$). The surrounding torus has temperature of a few MeVs and higher Y_e (0.15-0.25), while the polar outflows are even less neutron rich. The main roles of neutrinos are to cool the system and, for ν_e and $\bar{\nu}_e$, drive changes in Y_e . Over longer time scales, we expect the remnant to form a uniformly rotating star surrounded by a more axisymmetric disk. A large fraction of that disk ($\gtrsim 30\%$) is unbound by magnetically-driven winds, viscous angular momentum transport, and neutrino-driven winds, with neutrinos playing a major role in setting Y_e in the outflows (Metzger & Fernández 2014; Fujibayashi et al. 2020).

The mass, composition, velocity, and geometry of the outflows are of particular interest, as they set the main observable properties of kilonovae (Barnes & Kasen 2013). High- Y_e outflows produce optical kilonovae evolving on timescales of days and do not produce the heaviest r-process nuclei, while low- Y_e outflows produce infrared, week-long kilonovae that mostly produce heavier nuclei. The (approximate) boundary between these two outcomes is $Y_e \sim 0.25$ (Lippuner & Roberts 2015). The mass and velocity of the outflows also impact the brightness and duration of kilonovae. Figs. 3-4 and Table 1 summarize the main properties of the outflows. While there are some meaningful differences between simulations, these are generally smaller than when comparing our 'best' M1 scheme (used here) to a simpler version with an approximate treatment of neutrino energies (Foucart et al. 2016), and dwarfed by differences between neutrino transport and leakage schemes (Wanajo et al. 2014; Foucart et al. 2016). This is an encouraging sign for both M1 and MC schemes. The main differences that cannot be explained by MC sampling errors are a $\sim 10\%$ lower value of Y_e in the polar outflows ($< 45^\circ$ from the rotation axis), and a lower cutoff for the maximum value of Y_e reached by the outflows. This difference grows in a $\sim 5 \text{ km}$ region along the polar axis, around and outside of the neutrinospheres of (10–20) MeV electron (anti)neutrinos. In that region, Y_e is mostly constant in the MC simulation, but increases in the M1 simulation. One possible explanation is that this region is farther from the remnant in the MC simulation (Fig. 3), and that the M1 closure increases the density of neutrinos at the poles – so that even though the neutrino luminosities are higher in the MC simulation, neutrino fluxes are smaller. However, as this is a region

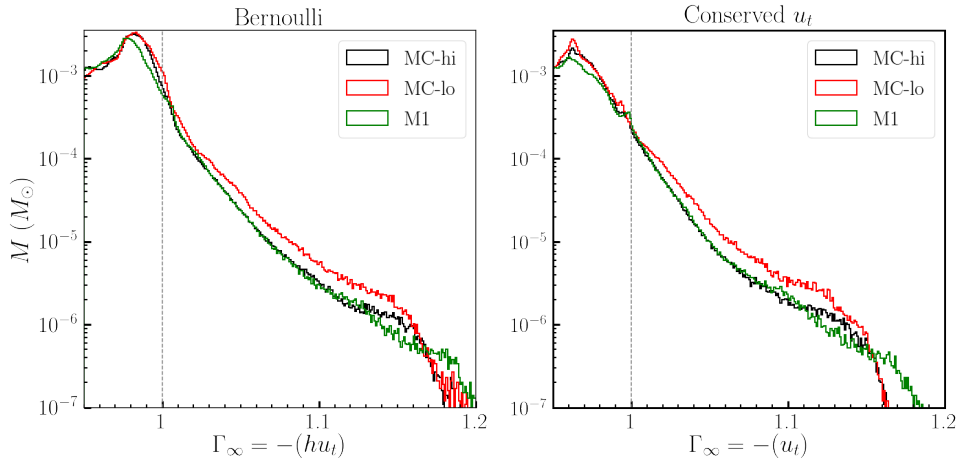


Figure 4. Estimated Lorentz factor of the ejecta at infinity. We show results using the Bernoulli criteria (*Left*), which slightly overestimates ejected masses, and assuming that u_t is constant (*Right*), which typically underestimates the ejected mass. $\Gamma_\infty = 1$ (dashed grey line) is the estimated boundary between bound and unbound material.

with rapid variations of the fluid properties, significant neutrino pressure, and non-negligible impacts of the M1 energy closure, other potential sources of errors cannot be ignored.

The total mass of unbound material (Table 1) and estimated velocity distribution of the outflows (Fig. 4) show broad agreement between simulations, with just two noticeable differences. First, the MC simulation with a lower number of packets overproduce outflows by $\sim 30\%$. This is less than e.g. the factor of 2 uncertainty associated with the definition of the ejecta itself (see Fig. 4), but it is the most noticeable effect of MC sampling errors. The excess ejecta has $Y_e \sim 0.2$, is close to the orbital plane ($\theta > 60^\circ$), and is observed (1–2) ms post-merger and in the final snapshot. It is thus most likely ejected during core bounces, from the hottest regions of the remnant, where neutrino pressure is the most significant. This would most naturally explain the properties of the excess ejecta, and the fact that the neutrino luminosity surprisingly varies less with MC sampling rate than the ejected mass (see below). Fig. 4 also shows why it is so difficult to accurately estimate the amount of unbound material in a simulation: the steepness of the Lorentz factor distributions implies that a very small error in the estimated location of the boundary between bound and unbound material has large effects on the predicted mass of unbound material. The second, more minor difference is that the cutoff of the velocity distribution is slightly lower in both MC simulations ($v_{\max} \sim 0.5$) than with M1 ($v_{\max} \sim 0.55$). We can compare these results to Sekiguchi et al. (2016), who performed a longer simulation with M1 transport, the same equation of state, and slightly more symmetric masses. They find $0.005M_\odot$ of ejecta with $\langle Y_e \rangle = 0.2$ and $\langle v \rangle = 0.19$, and a Y_e distribution fairly similar to those obtained with our M1 code – results that are broadly consistent if the higher mass asymmetry led to the ejection of an additional $0.003M_\odot$ of neutron-rich material.

The properties of neutrinos escaping the system are summarized in Fig. 5. After the first 2 ms, the average energy of neutrinos and the ν_e and $\bar{\nu}_e$ luminosities in the M1 and MC simulations agree within $\sim 20\%$. The ν_x luminosity, on the other hand, is nearly twice as large in the MC simulation. Differences between the two MC simulations are negligible compared to differences between MC and M1. The luminosity of heavy-lepton neutrinos is one of the most uncertain observable in our M1 scheme, because a large region close to the neutron star surface has negligible κ_a but large κ_s . In that region, neutrinos are trapped, but out of thermal equilibrium with the fluid. As the energy closure is the most ad-hoc part of our M1 algorithm, and the diffusion rate of neutrinos strongly depends on the choice of energy spectrum, this is a particularly difficult situation. The MC scheme has no particular reason to perform poorly in that regime: while it corrects large absorption opacities, its treatment of high-scattering regions is better motivated than that of the grey M1 scheme. Nevertheless, it may be premature to assume that the MC scheme provides the better answer, as the two schemes may be impacted in different ways by the grid resolution. In Sekiguchi et al. (2016), neutrino luminosities reach values similar to those measured here within ~ 5 ms, before decreasing on the cooling time scale of the remnant.

The angular distribution of neutrinos shows clearer differences between simulations (Fig. 5), and is much better captured by the MC algorithm than by the M1 algorithm. In polar regions, there is a nearly 50% excess of neutrinos in the M1 scheme, due to artificial radiation shocks caused by the analytical closure (Foucart et al. 2018). This mildly

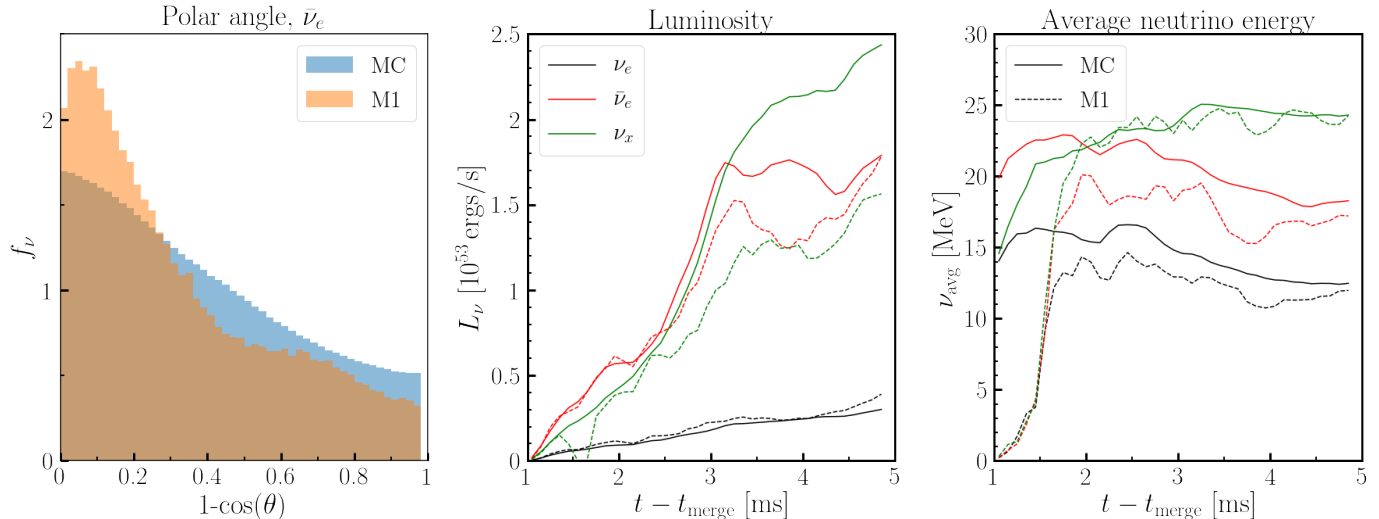


Figure 5. *Left:* Angular distribution of the electron antineutrinos leaving the grid (3 – 5) ms after merger; θ is the angle with respect to the angular momentum vector, i.e. polar neutrinos are on the left, equatorial neutrinos on the right, and each bin covers the same surface area. *Center:* Neutrino luminosity in the high resolution MC simulation and the M1 simulation. *Right:* Average (number-weighted) energy of neutrinos leaving the computational grid in the same simulations.

impacts the composition of the winds, and would lead to large errors in the calculation of the energy deposited by $\nu\bar{\nu}$ annihilation in polar regions.

Finally, we report on simulation costs: the two MC simulations cost 230k and 310k CPU-hrs on the Frontera cluster, at 30k and 40k CPU-hrs per millisecond at the end of the evolution (the merger phase is costlier). The M1 simulation is not directly comparable at early times, as it was performed on the Comet cluster, but it evolved at 35k CPU-hrs per millisecond by the end of the evolution on Frontera. The new MC code is thus competitive with our best M1 code, and cheap enough to be used for at least small parameter space surveys of neutron star mergers on (10 – 20) ms time scales, or for a small number of longer evolutions. This is because a small number of MC packets is sufficient to capture the most important neutrino effects: our simulations use only $\sim (1, 4)$ packets per finite volume cell.

4. IMPLICATIONS FOR NUMERICAL RELATIVITY AND ASTROPHYSICS

The first general relativistic simulations of binary neutron stars with MC radiation transport reported here demonstrates that cheap, yet reasonably accurate evolution of the transport equations with MC methods is possible. The only noticeable effect of MC sampling noise in the most important observables of our simulations is a $\sim 30\%$ increase in the outflow mass in our lower accuracy MC simulation. This is comparable to other sources of error in current simulations, even for the very low number of neutrino packets used in this study. The composition and velocity of the outflows, and the neutrino luminosities and energies are largely unaffected by sampling noise. For the specific binary studied in this manuscript, we find $\sim 0.008M_{\odot}$ of neutron-rich equatorial ejecta, mostly in a tidal tail, as well as an incipient neutrino-driven wind with higher electron fraction (Y_e) ~ 0.23 . Both ejecta components have average velocity $\sim 0.2c$. Detailed information about the properties of the ejected nuclear matter and the neutrino outflows are available upon request. Our MC simulations also provide full snapshots of the neutrino and matter distribution every 0.5 ms that can be used for studies of neutrino physics in mergers.

A first comparison of MC results with our approximate M1 code provides generally reassuring results for both methods, with agreement at the $\sim 20\%$ level for the neutrino luminosities (except ν_x) and energies, and very good agreement in the outflow masses and velocity between the M1 simulation and our most accurate MC simulation. The M1 simulation overestimates the average and maximum electron fraction of the outflows by $\sim 10\%$, and the two methods disagree by a factor of 2 on the luminosity of heavy-lepton neutrinos, but other observables are in remarkable agreement given the resolution of the simulation and the number of MC packets used. This indicates that earlier studies comparing our latest M1 code with more approximate M1 (Foucart et al. 2016) and leakage methods (Foucart et al. 2016), that found more significant differences, provide good estimates of the uncertainties of simpler neutrino treatments. Differences in ν_x luminosities may however have more of an impact if ν_x oscillate into $\nu_e, \bar{\nu}_e$ (e.g. due to neutrino-matter resonances [Vlasenko & McLaughlin (2018)]), and thus affect Y_e . Our new MC algorithm has a

low computational cost, will allow us to add new neutrino physics to simulations (e.g. pair annihilation, inelastic scatterings,...), and converges to a solution of the transport equations. On the other hand, the fact that an advanced M1 code is shown to already provide reasonably accurate results is reassuring for current outflow and kilonova models.

ACKNOWLEDGMENTS

We are grateful to Sherwood Richers, Ernazar Abdikamalov, Ben Ryan, Charles Gammie, Dan Kasen, Jonah Miller and Josh Dolence for useful discussions regarding neutrino transport and Monte-Carlo methods. We also thank Aurore Bertranhandy for her help with the NuLib library. We are grateful to the Yukawa Institute for Theoretical Physics the Center for Computational Astrophysics for hosting some of these discussions. FF gratefully acknowledges support from the DOE through Early Career award DE-SC0020435, from the NSF through grant PHY-1806278, and from NASA through grant 80NSSC18K0565. M.D gratefully acknowledges support from the NSF through grant PHY-1806207. F.H. and M.S. acknowledge funding from NSF Grants PHY-1708212 and PHY-1708213. L.K. acknowledges support from NSF grant PHY-1606654 and PHY-1912081. F.H., L.K. and M.S. also thank the Sherman Fairchild Foundation for their support.

REFERENCES

- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017, *Astrophys. J. Lett.*, 848, L13
- Abbott, B. P., et al. 2017, *Phys. Rev. Lett.*, 119, 161101
- Abdikamalov, E., Burrows, A., Ott, C. D., et al. 2012, *Astrophys. J.*, 755, 111
- Barnes, J., & Kasen, D. 2013, *Astrophys. J.*, 775, 18
- Barnes, J., Kasen, D., Wu, M.-R., & Mart'inez-Pinedo, G. 2016, ArXiv e-prints, arXiv:1605.07218
- Chornock, R., Berger, E., Kasen, D., et al. 2017, "Astrophys. J. Lett.", 848, L19
- Cowperthwaite, P. S., et al. 2017, *Astrophys. J.*, 848, L17
- Deaton, M. B., Duez, M. D., Foucart, F., et al. 2013, *Astrophys. J.*, 776, 47
- Duez, M. D., Foucart, F., Kidder, L. E., et al. 2008, *Phys. Rev. D*, 78, 104015
- Fleck, J. A., J., & Cummings, J. D. 1971, *Journal of Computational Physics*, 8, 313
- Foucart, F. 2018, *Mon. Not. Roy. Astr. Soc.*, 475, 4186
- Foucart, F., Duez, M. D., Kidder, L. E., et al. 2018, *Phys. Rev.*, D98, 063007
- Foucart, F., Kidder, L. E., Pfeiffer, H. P., & Teukolsky, S. A. 2008, *Phys. Rev. D*, 77, 124051
- Foucart, F., O'Connor, E., Roberts, L., et al. 2016, *Phys. Rev.*, D94, 123016
- Foucart, F., Deaton, M. B., Duez, M. D., et al. 2013, *Phys. Rev. D*, 87, 084006
- Foucart, F., O'Connor, E., Roberts, L., et al. 2015, *Phys. Rev. D*, 91, 124021
- Foucart, F., Haas, R., Duez, M. D., et al. 2016, *Phys. Rev.*, D93, 044019
- Fujibayashi, S., Kiuchi, K., Nishimura, N., Sekiguchi, Y., & Shibata, M. 2018, *Astrophys. J.*, 860, 64
- Fujibayashi, S., Wanajo, S., Kiuchi, K., et al. 2020, arXiv:2007.00474
- Hempel, M., Fischer, T., Schaffner-Bielich, J., & Liebendörfer, M. 2012, *Astrophys. J.*, 748, 70
- Just, O., Bauswein, A., Ardevol Pulpillo, R., Goriely, S., & Janka, H.-T. 2015, *Mon. Not. Roy. Astr. Soc.*, 448, 541
- Kasliwal, M. M., Nakar, E., Singer, L. P., et al. 2017, *Science*, 358, 1559
- Kiuchi, K., Cerdá-Durán, P., Kyutoku, K., Sekiguchi, Y., & Shibata, M. 2015, *Phys. Rev. D*, 92, 124034
- Li, L.-X., & Paczynski, B. 1998, *Astrophys. J.*, 507, L59
- Lindblom, L., Scheel, M. A., Kidder, L. E., Owen, R., & Rinne, O. 2006, *Class. Quantum Grav.*, 23, S447
- Lippuner, J., & Roberts, L. F. 2015, *Astrophys. J.*, 815, 82
- Metzger, B. D., & Fernández, R. 2014, *Mon. Not. Roy. Astr. Soc.*, 441, 3444
- Miller, J. M., Ryan, B. R., Dolence, J. C., et al. 2019, *Phys. Rev. D*, 100, 023008
- O'Connor, E. 2015, *ApJ Suppl.*, 219, 24
- Pfeiffer, H. P., Kidder, L. E., Scheel, M. A., & Teukolsky, S. A. 2003, *Comput. Phys. Commun.*, 152, 253
- Radice, D. 2017, *Astrophys. J. Lett.*, 838, L2
- Radice, D., Galeazzi, F., Lippuner, J., et al. 2016, *Mon. Not. Roy. Astr. Soc.*, 460, 3255
- Richers, S., Kasen, D., O'Connor, E., Fernandez, R., & Ott, C. D. 2015, *Astrophys. J.*, 813, 38
- Roberts, L. F., Kasen, D., Lee, W. H., & Ramirez-Ruiz, E. 2011, *Astrophys. J. Lett.*, 736, L21
- Ryan, B. R., Dolence, J. C., & Gammie, C. F. 2015, *Astroph. J.*, 807, 31
- Sekiguchi, Y. 2010, *Class. Quantum Grav.*, 27, 114107

- Sekiguchi, Y., Kiuchi, K., Kyutoku, K., Shibata, M., & Taniguchi, K. 2016, *Phys. Rev.*, D93, 124046
- Shibata, M., Kiuchi, K., Sekiguchi, Y., & Suwa, Y. 2011, *Progress of Theoretical Physics*, 125, 1255
- Smartt, S. J., Chen, T.-W., Jerkstrand, A., et al. 2017, "Nature", 551, 75
- Soares-Santos, M., Holz, D. E., Annis, J., et al. 2017, "Astrophys. J. Lett.", 848, L16
- Szilágyi, B. 2014, *Int. J. Mod. Phys. D*, 23, 1430014
- Vlasenko, A., & McLaughlin, G. 2018, *Phys. Rev. D*, 97, 083011
- Wanajo, S., Sekiguchi, Y., Nishimura, N., et al. 2014, *Astrophys.J.Lett.*, 789, L39