Lawrence Berkeley National Laboratory

Recent Work

Title

The impact of r-process heating on the dynamics of neutron star merger accretion disc winds and their electromagnetic radiation

Permalink

<https://escholarship.org/uc/item/76x7j6hk>

Journal

Monthly Notices of the Royal Astronomical Society, 510(2)

ISSN

0035-8711

Authors

Klion, H
Tchekhovskoy, A
Kasen, D
[et al.](#)

Publication Date

2022-02-01

DOI

10.1093/mnras/stab3583

Peer reviewed

The impact of r-process heating on the dynamics of neutron star merger accretion disc winds and their electromagnetic radiation

Hannah Klion^{1,2*}, Alexander Tchekhovskoy³, Daniel Kasen^{1,2,4}, Adithan Kathirgamaraju¹, Eliot Quataert⁵, and Rodrigo Fernández⁶

¹*Astronomy Department and Theoretical Astrophysics Center, University of California, Berkeley, Berkeley, CA 94720, USA*

²*Physics Department, University of California, Berkeley, Berkeley, CA 94720, USA*

³*Center for Interdisciplinary Exploration & Research in Astrophysics (CIERA), Physics & Astronomy, Northwestern University, Evanston, IL 60208, USA*

⁴*Nuclear Science Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA*

⁵*Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, USA*

⁶*Department of Physics, University of Alberta, Edmonton, AB T6G 2E1, Canada*

Accepted XXX. Received YYY; in original form ZZZ

ABSTRACT

Neutron star merger accretion discs can launch neutron-rich winds of $> 10^{-2}M_{\odot}$. This ejecta is a prime site for r-process nucleosynthesis, which will produce a range of radioactive heavy nuclei. The decay of these nuclei releases enough energy to accelerate portions of the wind by $\sim 0.1c$. Here, we investigate the effect of r-process heating on the dynamical evolution of disc winds. We extract the wind from a 3D general relativistic magnetohydrodynamic simulation of a disc from a post-merger system. This is used to create inner boundary conditions for 2D hydrodynamic simulations that continue the original 3D simulation. We perform two such simulations: one that includes the r-process heating, and another one that does not. We follow the hydrodynamic simulations until the winds reach homology (60 seconds). Using time-dependent multi-frequency multi-dimensional Monte Carlo radiation transport simulations, we then calculate the kilonova light curves from the winds with and without dynamical r-process heating. We find that the r-process heating can substantially alter the velocity distribution of the wind, shifting the mass-weighted median velocity from $0.06c$ to $0.12c$. The inclusion of the dynamical r-process heating makes the light curve brighter and bluer at ~ 1 d post-merger. However, the high-velocity tail of the ejecta distribution and the early light curves are largely unaffected.

Key words: neutron star mergers – radiative transfer

1 INTRODUCTION

GW170817 was the first gravitational wave detection of a binary neutron star (NS) merger (Abbott et al. 2017a). It was followed by a short gamma ray burst (sGRB), confirming the connection between compact object mergers and sGRBs (Abbott et al. 2017c; Goldstein et al. 2017; Savchenko et al. 2017). An ultraviolet, optical, and infrared counterpart, AT 2017gfo (also SSS17a), was detected hours later, consistent with a $v \sim 0.1c$ outflow of $> 0.01M_{\odot}$ of radioactive material with an average opacity of $1 - 3 \text{ cm}^2\text{g}^{-1}$ (Abbott et al. 2017b; Arcavi et al. 2017; Kasliwal et al. 2017; Drout et al. 2017; Villar et al. 2017; Coulter et al. 2017; Cowperthwaite et al. 2017; Chornock et al. 2017; Soares-Santos et al. 2017). This was consistent with predictions of a “kilonova”, a radioactively-powered, red, and rapidly evolving counterpart of a NS merger (Metzger et al. 2010; Metzger & Berger 2012; Barnes & Kasen 2013; Tanaka & Hotokezaka 2013).

Outflows from NS mergers are expected via multiple channels. As an NS binary merges, mass ejection occurs on dynamical ($\sim$ ms) time-scales due to hydrodynamical and tidal forces. Numerical simulations predict that dynamical ejecta consist of $10^{-4} - 10^{-2}M_{\odot}$

of material with escape velocities of $\sim 0.1 - 0.3c$ (Bauswein et al. 2013; Hotokezaka et al. 2013a; Lehner et al. 2016; Bovard et al. 2017; Dietrich et al. 2017; Radice et al. 2018). Some of the disrupted material is still gravitationally bound and can form an accretion disc of up to $0.3M_{\odot}$ that evolves on longer timescales ($100 \text{ ms} - 10 \text{ s}$) (Oechslin et al. 2007; Hotokezaka et al. 2013b). Initially, neutrino cooling dominates because the disc is hot and dense (Popham et al. 1999; Narayan et al. 2001). As the disc cools and spreads, neutrino cooling becomes inefficient, and the disc becomes fully advective. Weak interactions freeze out, which can lead to a strong neutron-rich wind (Metzger et al. 2009).

Regardless of how it is launched, the ejected material undergoes rapid decompression from nuclear densities. Once the material cools below $\sim 10 \text{ GK}$ and leaves nuclear statistical equilibrium (NSE), neutrons capture onto light seed nuclei faster than the nuclei can undergo beta decays. This rapid- (r-)process nucleosynthesis produces a range of radioactive neutron-rich nuclei that beta decay to stability over weeks (Lattimer et al. 1977; Eichler et al. 1989; Metzger et al. 2010). The bulk of the heat released from r-process decays is deposited in the first few seconds after the material leaves NSE. It then falls off approximately as a power law, powering the kilonova (Li & Paczyński 1998; Metzger et al. 2010).

The distribution of nuclei produced by the r-process primarily depends on the electron fraction Y_e of the nuclear material. When

* E-mail: hklion@berkeley.edu

$Y_e \lesssim 0.25$, the r-process produces material enriched with Lanthanides and Actinides, which have uniquely high opacities. Their partially-filled f orbitals produce a high density of moderately strong atomic lines, which lead to a high quasi-continuum opacity ($\kappa \sim 10 \text{ cm}^2 \text{ g}^{-1}$), especially in the blue and ultraviolet (Kasen et al. 2013; Fontes et al. 2015; Tanaka et al. 2020). Low Y_e ejecta would therefore produce a red, dim, slowly-evolving transient (Barnes & Kasen 2013; Wollaeger et al. 2018). By contrast, higher- Y_e material gives rise to lighter r-process elements, which primarily occupy the second row of the d -block. These elements have lower average opacities $\kappa \sim 0.5 - 1 \text{ cm}^2 \text{ g}^{-1}$ that are again highly wavelength-dependent. The light curve of AT 2017gfo can be modeled by (at least) two distinct components with opacities that roughly correspond to light and heavy r-process products (Cowperthwaite et al. 2017; Kasen et al. 2017; Perego et al. 2017; Villar et al. 2017).

Many past studies of post-merger accretion discs have relied on axisymmetric hydrodynamic simulations, where an imposed shear viscosity transports angular momentum. (e.g. Fernández & Metzger 2013; Just et al. 2015; Fujibayashi et al. 2018). Siegel & Metzger (2017, 2018) presented the first 3D general relativistic magnetohydrodynamic (GRMHD) simulations of the post-merger accretion disc evolution. They track the system out to 400 ms, by which point about half of the wind has been launched. Fernández et al. (2019, hereafter F19) followed the evolution of a post-merger disc system to 9 seconds, allowing all of the disc material to be accreted or ejected. Miller et al. (2019) evolve the wind to 130 ms with concurrent GR Monte Carlo neutrino transport, and with a full nuclear reaction network. They find that the resulting outflow primarily consists of light r-process elements, consistent with a blue disc wind. The wind is also sensitive to the initial magnetic field configuration within the disc; toroidal or weaker fields lead to less massive and slower outflows (Christie et al. 2019).

The very long-term evolution of disc winds remains uncertain. Most disc wind simulations are only evolved to $\lesssim 10$ s, while the wind needs to expand for several times that in order to reach homology, which is required to generate light curves and spectra with most existing photon radiative transfer codes. Longer-term evolution of hydrodynamic winds has been studied by extracting wind properties at a large radius, and using that to set an inner boundary condition on a larger grid (e.g. Kasen et al. 2015; Kawaguchi et al. 2021). This technique has not yet been applied to a GRMHD simulation of a disc wind.

R-process heating provides around 1–3 MeV per nucleon, most of which is deposited within 1 s. If this were completely converted to kinetic energy, it would boost a particle at rest to a velocity of $0.1 - 0.2c$. Prior work has found that r-process heating can have a particularly strong effect on the tidal component of the dynamical ejecta, causing it to inflate and smooth out inhomogeneities (Rosswog et al. 2014; Grossman et al. 2014; Fernández et al. 2015; Darbha et al. 2021).

In this paper, we are interested in understanding how r-process heating affects the structure of a 3D MHD-driven disc wind and, consequently, the resulting kilonova light curves. We evolve a kilonova disc wind to homology both with and without the r-process heating. We construct an inner boundary condition for 2D GR hydrodynamics (GRHD) simulations based on the wind formed in the 3D GRMHD simulation of F19, approximately accounting for r-process heating that occurs before injection. We present our formalism for doing so in Section 2.1, and discuss the details of our 2D GRHD simulations in Section 2.2. After 60 s of evolution, we pass the resulting density and temperature structures to SEDONA, a multi-dimensional, multi-frequency radiation transport code (Sections 2.3 and 2.4). In

Section 3, we compare the results of our GRHD simulations with and without r-process heating to assess the effect of r-process heating on the dynamical evolution of kilonova disc winds. Our predictions for disc wind light curves, with and without dynamic r-process heating, are found in Section 4. We conclude and discuss future research directions in Section 5.

2 METHODS

Our full calculation consists of the following steps:

(i) Simulate a post-merger black hole accretion disc in 3D GRMHD to 10 s. This model is the one presented in F19, and does not include r-process heating.

(ii) Extract wind properties from F19 at a radius of $r_b = 2 \times 10^4$ km, and use them to set an inner boundary condition for a 2D GRHD simulation. There are two versions of these conditions: one that is simply axisymmetrized, and another that is axisymmetrized and then modified to approximately account for r-process heating that occurred within the domain of the 3D simulation, prior to the gas reaching r_b .

(iii) Evolve two 2D GRHD simulations, each with one of the sets of boundary conditions described above. The simulation ‘with r-process’ uses the r-process-adjusted inner boundary conditions and self-consistently includes an r-process heating source term during evolution. Meanwhile, the simulation ‘without r-process’ does not include an r-process heating term, and uses the axisymmetrized boundary conditions (not modified to include r-process heating). We evolve both of these models until 60 s, at which point they are mostly in free expansion.

(iv) Calculate the optical emission from both of the 2D GRHD simulations using SEDONA, a Monte Carlo radiation transport code. In this phase of the calculation, both models include r-process heating, but the evolution is assumed to be homologous. The velocity structure is fixed, so the heating only affects the light curve. This allows us to isolate the effects of r-process heating on the dynamics of the ejecta.

2.1 2D GR hydrodynamic initial conditions

We continue the 3D GRMHD simulation of F19 in 2D GRHD. Unless otherwise stated, we take $G = c = 1$. F19 initialize a torus of mass $0.033M_\odot$ ¹ with a strong poloidal magnetic field around a black hole of mass $3M_\odot$ and spin parameter $a = 0.8$. We extract the time-series of the primitive variables at a radius of $r_b \equiv 2 \times 10^4 \text{ km} = 4452 r_g$ in the 3D simulations: rest mass density ρ , energy density ε , four velocity in Kerr-Schild coordinates $\{u^t, u^r, u^\theta, u^\phi\}$ and composition (electron fraction Y_e and abundances of protons, neutrons, and alpha particles). This radius is small enough that the entire wind crosses through the surface before the end of the 3D simulation, but large enough that no bound material falls back through.

We make the simplifying assumption that the magnetic fields are zero. While they are critical for the jet launch and wind ejection, both of these occur interior to the inner boundary of our simulation. We are primarily interested in the baryon-rich wind, and do not expect the exclusion of magnetic fields to significantly affect our results.

¹ The simulations of F19 were performed before the announcement of GW170817/AT 2017gfo. Subsequent modeling has inferred that the initial torus mass was larger by a factor of ~ 3 (Shibata et al. 2017).

We axisymmetrize the density by averaging in ϕ :

$$\langle \rho \rangle = \frac{\int \rho \sqrt{-g} \, d\phi}{\int \sqrt{-g} \, d\phi}, \quad (1)$$

and calculating mass-weighted averages for the other variables, X :

$$\langle X \rangle = \frac{\int X \rho \sqrt{-g} \, d\phi}{\int \rho \sqrt{-g} \, d\phi}. \quad (2)$$

The determinant of the metric is g . These axisymmetrized values are then used as a time-dependent inner boundary condition in 2D GRHD simulations. We perform two simulations: one with and one without *r*-process heating. When we do not include *r*-process heating, we directly use the axisymmetrized primitives at r_b as the boundary condition.

The majority of the thermal energy input by the *r*-process will be converted to kinetic energy via adiabatic expansion of the hot gas. The decay of *r*-process elements deposits $\sim 4 \times 10^{18} \text{ erg g}^{-1}$ on a timescale of $\sim 1 \text{ s}$, which corresponds to a velocity increase of $\sim 0.1c$ for a particle at rest. Of the mass crossing r_b in F19, $8.7 \times 10^{-3} M_\odot$ (68 per cent) has a velocity $< 0.1c$, so we expect acceleration from this burst of energy to be significant. The effects of Y_e on heating rates are generally modest (Lippuner & Roberts 2015), so we assume a spatially uniform, time-dependent heating rate. We adopt the $Y_e = 0.1$ *r*-process heating rate from Metzger et al. (2010). We approximate the time-dependent heating rate per mass $\eta(t)$ as a broken power law:

$$\eta(t) = \begin{cases} 6.7 \times 10^{17} \text{ erg/g/s} & t_1 < 0.01, \\ 1.1 \times 10^{19} t_1^{0.6} \text{ erg/g/s} & 0.01 \leq t_1 < 0.5, \\ 7.0 \times 10^{17} t_1^{-3.33} \text{ erg/g/s} & 0.5 \leq t_1 < 4, \\ 4.2 \times 10^{16} t_1^{-1.3} \text{ erg/g/s} & t_1 > 4. \end{cases} \quad (3)$$

where $t_1 = t/1 \text{ s}$. This rate assumes that all of the nuclear energy emitted by radioactive decays is converted to thermal energy. In the first several seconds, neutrinos are the only decay product that can easily escape, while the remainder are thermalized. Neutrinos likely carry away $\sim 20 - 30$ per cent of the total energy (Metzger et al. 2010; Barnes et al. 2016). This adjustment would be less than the uncertainty in the *r*-process heating rate ($\sim$ factor of a few). We therefore assume complete thermalization during the hydrodynamic phase of the calculation. We take $t = 0$ for the *r*-process to be at the start of the simulation of F19. Physically, the *r*-process begins once the material leaves NSE, which will occur when the temperature drops below $\sim 10 \text{ GK}$. These temperatures are only achieved within a radius of $50 r_g = 225 \text{ km} \ll r_b$, and at early times $t \lesssim 300 \text{ ms}$. By contrast, the median r_b -crossing time is 1.7 s . By the time the wind enters our domain, it has been out of NSE for some time, and it is a reasonable approximation that it left NSE at $t = 0$.

Before we axisymmetrize, we adjust the velocities and internal energy to include the effects of *r*-process heating. To calculate the modified boundary condition at time t_b , we consider a fluid parcel of mass m that is crossing r_b at time t_b . We assume that it is expanding homologously and has been doing so since the start of the 3D simulation. This is an approximation since the parcel will accelerate, so its evolution is not actually homologous. Also, it would have been ejected from the disc at some time after the start of the 3D calculation. None the less, this should give a reasonable estimate of the distribution of *r*-process energy into thermal and kinetic. Improving on this approximation would require the fluid trajectories, which we do not have. The parcel has volume $V(t)$ at time t . Unless

otherwise indicated, we work in the fluid frame. The gas has been heated radioactively since $t_0 = 0 \text{ s}$ by $\eta(t)$. Its thermal energy due to radioactive heating e_{rad} will evolve as

$$\frac{de_{\text{rad}}}{dt} = \frac{-e_{\text{rad}}(t)}{t} + \eta(t)m. \quad (4)$$

The first term on the right-hand side accounts for thermal losses due to adiabatic expansion. Solving the above differential equation and dividing by the volume $V(t)$ gives the internal energy density due to radioactive decays ε_{rad} ,

$$\varepsilon_{\text{rad}}(t) \equiv \frac{e_{\text{rad}}(t)}{V(t)} = \frac{\rho(t)}{t} \int_{t_0}^t \eta(t') t' \, dt'. \quad (5)$$

We add $\varepsilon_{\text{rad}}(t_b)$ to the internal energy density at r_b . We assume that the remainder of the total *r*-process energy is converted to kinetic energy, giving the parcel a boost Δe_{kin} :

$$\Delta e_{\text{kin}}(t_b) = e_{\text{rad,tot}}(t_b) - e_{\text{rad}}(t_b), \quad (6)$$

where

$$e_{\text{rad,tot}}(t) = m \int_{t_0}^t \eta(t') \, dt' \quad (7)$$

is the total energy deposited by *r*-process heating up to time t . The kinetic energy boost $\Delta e_{\text{kin}} \geq 0$ because the thermal energy attributed to radioactive decays can never exceed the total heating by radioactivity. The conversion of non-radioactive thermal energy to kinetic energy is already handled in the underlying 3D simulation.

The energies and masses in this calculation are in the fluid frame. Ideally, we would use the fluid proper time to calculate the heating rate. Calculating the proper time for each parcel of fluid is infeasible since we do not have information about the trajectory of individual fluid elements. For the sake of consistency, we will use the observer time throughout this paper when evaluating $\eta(t)$. We do not expect this to have a substantial effect on our results. The vast majority of the mass has $v < 0.2c$, so proper and observer time will generally be equal. Additionally, the bulk of the heating occurs in the first few seconds, so the total energy deposited over the $\sim 10^2 \text{ s}$ of the hydrodynamic simulation will be the same to within a few per cent. Effects on the dynamical evolution of the wind should therefore be negligible.

The combined rest mass and kinetic energy of the parcel is γm , where γ is the Lorentz factor. If we add Δe_{kin} in kinetic energy, the new energy of the parcel becomes

$$\gamma_{\text{new}} m = \gamma_{\text{old}} m + \Delta e_{\text{kin}}(t_b), \quad (8)$$

and its new Lorentz factor is

$$\gamma_{\text{new}} = \gamma_{\text{old}} + \Delta e_{\text{kin}}(t_b)/m. \quad (9)$$

The magnitude of the new 3-velocity is v_{new} , which can be calculated directly from γ_{new} . We scale the components of the 3-velocity by $v_{\text{new}}/v_{\text{old}}$. While we set our boundary conditions in terms of densities, e.g. ρ and ε , the physically significant quantity is the flux of mass and energy onto the grid. Accordingly, we scale the boundary ρ and ε by the ratio of the radial 4-velocities: $u_{\text{old}}^r/u_{\text{new}}^r$. This preserves the homologous estimate of the kinetic-thermal energy distribution.

Fig. 1 shows the time-dependent boundary condition used in our axisymmetric GRHD simulations for four representative angles. We show rest mass density, radial four-velocity (u^r) and temperature. We calculate the temperature from the internal energy density ε assuming that the gas is radiation-dominated, $T = (\varepsilon/a)^{1/4}$, where a is the radiation constant. Due to the jet, the polar regions have high velocity, particularly in the first $\sim 0.5 \text{ s}$, where $u^r > 1$. The high

Figure 1. Inner boundary condition at $r_b = 2 \times 10^4$ km in the 2D GRHD simulations of long-term disc wind evolution (see Section 2.1), derived from the 3D GRMHD simulation of F19. The left (right) column shows the conditions used for the simulation without (with) r-process heating. The conditions on the right include approximate effects of r-process heating that occurs before injection. We show rest mass density (top panel), radial four-velocity (middle), and radiation temperature (bottom) as a function of time for a representative selection of angles. Temperature is calculated from internal energy density ε assuming radiation pressure dominates, $T = (\varepsilon/a)^{1/4}$, where a is the radiation constant. The conditions are largely symmetric across the equator, so we only show angles in the Southern hemisphere; angles are measured from the South pole. The effects of r-process heating are more pronounced toward the higher-density, lower-velocity equatorial region.

velocity in the polar regions is largely unaffected by the addition of r-process kinetic energy. However, the slower material in the bulk of the wind is accelerated to a velocity of $\sim 0.1c$. We hold the mass flux constant, so the density of the wind is correspondingly lower in the case with r-process heating.

The r-process boosts the temperature of the early wind. At later times, most of the energy has been converted to kinetic, and the thermal energy increases due to r-process heating are modest. While the thermal energy flux is larger in the model with r-process heating, the energy density, and therefore temperature, is lower due to the much higher velocity of the material.

By construction, the mass injection rate $\dot{M}$ is the same in both GRHD simulations. $\dot{M}$ and its integral are shown in Fig. 2. The bulk of the mass enters the grid in the first 2 – 3 seconds.

We use the boundary data calculated from F19 for the duration of their GRMHD simulation (~ 9 s). We interpolate the boundary conditions linearly in time and meridional angle. After 9 s have elapsed, we linearly taper the velocity at the boundary to zero, while setting the other variables equal to the floor values.

2.2 GRHD simulations in 2D

We use HARMPI² (Tchekhovskoy 2019), a parallel version of the code HARM (Gammie et al. 2003; Noble et al. 2006), to perform 2D axisymmetric GRHD simulations of the long-term disc wind evolution. We use the same Kerr metric as in F19, though the spacetime is approximately flat at the radii of interest. We work in modified spherical-polar Kerr-Schild coordinates. Our domain extends from $r_b = 4552r_g$ to $r_{\text{out}} = 10^6 r_b$, and is discretized into 1024 radial points. The first 895 are spaced logarithmically between r_b and $r_t = 10^4 r_b$. The remaining 129 are spaced progressively more sparsely between r_t and r_{out} . This would allow us to run our simulation until $t \approx r_t/c = 680$ s before encountering any artifacts due to the grid boundary. The meridional grid is the same as in F19, but with double the number of cells. It covers the interval $[0, \pi]$, and consists of 512 cells, with resolution concentrated at the poles and near the equator. We employ an outflow boundary condition at r_{out} , and a reflective boundary condition in the meridional direction. The inner radial boundary condition is a time-dependent condition, as described in the previous section. In this phase, we perform two

² Available at <https://github.com/atckkho/harmmpi>

Figure 2. Mass injection rate $\dot{M}$ (top panel) and total injected mass (bottom panel) in our GRHD simulations as a function of time. The majority of the mass injection occurs within the first few seconds of the simulation, with a median injection time of 1.6 s. By construction, the mass injection rate is the same in the HARM runs with and without *r*-process heating.

simulations, one that uses the axisymmetrized boundary conditions, and another that uses the *r*-process adjusted boundary conditions.

Throughout, we use an ideal gas equation of state with fixed adiabatic index $\gamma_{\text{ad}} = 4/3$. Our simulations include the (anti)neutrino emission and alpha particle recombination models detailed in F19. However, the temperatures in our simulation are below the thresholds where these processes are significant, so the composition of the flow is a passive scalar and does not affect our results.

In the simulation that uses the *r*-process adjusted boundary conditions, we include an *r*-process heating term. During each time step Δt ,

$$\Delta \varepsilon_{\text{rad}} = \rho \eta(t) \frac{\Delta t}{u^t}, \quad (10)$$

is added to the fluid-frame internal energy in between updates to the conserved quantities. While we evaluate the heating rate per mass $\eta(t)$ (equation 3) in the lab frame, we still compute the added energy in the fluid frame. The factor of u^t arises because the fluid-frame time step is $\Delta t/u^t$.

We evolve our simulations to at least $t_H = 60$ s, when the flow has largely reached free expansion, and the structure has stopped evolving. At this point, regions of the ejecta with high Mach number suffer from numerical errors that artificially increase the internal energy. These are most prominent in the simulation with no *r*-process heating, whose thermal structure is not of interest. When constructing the SEDONA input model from the 2D HARM simulation with no *r*-process heating, we re-set the thermal energy density (but not the kinetic energy) to be that expected from radioactive decays plus adiabatic degradation,

$$\varepsilon(t_H) = \frac{\rho(t_H)}{t_H} \int_{t_0}^{t_H} \eta(t') t' dt'. \quad (11)$$

2.3 Radiation Transport Initial Conditions

When the HARM models reach homology, the ejecta are still very optically thick. It is impractical to start radiation transport calculations on timescales of ~ 60 s due to the computational expense. Because we focus on accurately modeling the emission on timescales of > 1 h, it is unnecessary to start radiation transport calculations before $t_S = 300$ s. Between t_H and t_S , evolution is well-modeled by homologous expansion. At the start of the SEDONA calculation, the density will be

$$\rho(t_S) = \rho(t_H) \left(\frac{t_H}{t_S} \right)^3. \quad (12)$$

R-process heating and adiabatic degradation of internal energy will continue as well, giving

$$\varepsilon(t_S) = \varepsilon(t_H) \frac{t_H}{t_S} + \frac{\rho(t_S)}{t_S} \int_{t_H}^{t_S} \eta(t') t' dt'. \quad (13)$$

When constructing the input models for SEDONA, we exclude the high Mach-number regions that are susceptible to numerical errors in the internal energy. We apply density and radial cuts to the ejecta, only including cells where

$$\rho > \rho_{\text{cut}} = 6.9 \times 10^{-7} \left(\frac{t}{60 \text{ s}} \right)^3 \text{ g cm}^{-3} \quad (14)$$

and

$$\frac{r}{t} < 0.8c. \quad (15)$$

This contour is overplotted on density maps in Fig. 4. From 30 seconds onwards, this contour includes approximately 95 per cent of the mass on the grid in both simulations. Neither of these cuts affect the shapes of the light curves. Expanding the r/t cut to $0.9c$ does not affect the light curves. Lowering the value of ρ_{cut} by a factor of 10 increases the mass enclosed to 98 per cent and causes the overall luminosity of the early light curve to increase slightly (by 10 per cent).

2.4 Radiation Transport Simulations

To predict bolometric and broad-band light curves from these disc winds, we use SEDONA, a time-dependent, multi-dimensional, multi-wavelength Monte Carlo radiation transport code (Kasen et al. 2006; Roth & Kasen 2015). SEDONA tracks the emission and propagation of packets of radiation (‘photons’) through the freely-expanding ejecta. Interactions between the particles and fluid are calculated in the fluid frame, which naturally accounts for adiabatic losses as well as most relativistic effects. The one relativistic effect we neglect is in the evaluation of $\eta(t)$, where we do not distinguish between proper and lab-frame time (see Section 2.1). When particles leave the grid, they are tallied according to their wavelength, direction, and observer arrival time, providing time- and viewing-angle-dependent light curves and spectra.

The ejecta density and temperature are discretized on a cylindrical velocity grid. Since SEDONA requires homologous expansion, we reset the velocity from the HARM models to be purely radial and have magnitude $v = r/t$. The regions within our radial and density cuts are already mostly homologous so this should not significantly affect our results. Within the SEDONA run, density evolves as t^{-3} . The temperature is re-calculated at each time-step by equating the energy from photon absorption and radioactive heating with thermal emission. Adiabatic losses arise from the frame-shifting in the particle-fluid interactions.

The initial radiation field is represented by 10^7 particles. At each time step, 3×10^5 photons are created, representing the emission from r-process decays. We use a parametrized, time-dependent r-process heating rate that assumes an initial entropy of $32 k_B \text{ baryon}^{-1}$, an expansion time-scale of 0.84 ms, and electron fraction of 0.13 (Lipuner & Roberts 2015). This rate is adjusted by a time-dependent thermalization fraction that ranges between $\sim 50 - 75$ per cent (Barnes et al. 2016). The resultant net heating rate roughly agrees with equation (3) between t_H and t_S . As in the HARM calculation, we use the lab frame time to evaluate η . This may have a larger effect in the light curve calculation since the instantaneous luminosity is more important than the total energy budget. SEDONA uses a flat spacetime, so the time dilation factor is γ . For a power-law heating rate $\propto t^{-a}$, we underestimate the heating rate by a factor of γ^a , and the total deposited energy by γ^{a-1} . The fastest material on the SEDONA grid has $\gamma = 1.7$. After t_S , $a \approx 1.3$, so for the fastest material, we may underestimate the r-process luminosity by a factor of 2, and the total energy deposited within a time-step by 20 per cent. We expect the overall magnitude of this effect to be small since 98 per cent of the mass in the SEDONA simulation has $v < 0.5c$, for which luminosity is underestimated by less than 20 per cent.

We adopt a uniform grey (frequency-independent) opacity $\kappa = 1 \text{ cm}^2 \text{ g}^{-1}$. This is approximately the Planck mean opacity of a mixture of first-peak r-process elements at 1 day post-merger (Kasen et al. 2013; Tanaka et al. 2020). On a similar time-scale, a mixture of Lanthanides has a higher average opacity, $\sim 10 \text{ cm}^2 \text{ g}^{-1}$. We are primarily interested in the relative differences between ejecta with and without dynamic r-process, so the choice of a particular grey opacity will not affect the comparison, since the models including and excluding dynamic r-process heating will be affected in the same way.

3 HYDRODYNAMIC RESULTS

We follow our 2D GRHD simulations to 60 seconds, by which point the flow is mostly in homology ($v \propto r$). Fig. 3 compares velocity profiles for selected angles against the profile that SEDONA imposes on the ejecta ($v = r/t$). Deviations from homology are small, especially in the lower-velocity regions that contain most of the mass. Further evolution in GRHD brings the polar regions in line with the velocity structure assumed by SEDONA, but allows for more numerical errors to accumulate in the internal energy in high Mach-number regions. After 60 seconds, changes to the light curves due to these errors exceed those from further evolution of the density structure.

Density snapshots from our 2D GRHD simulations are shown in Fig. 4. In the first few seconds, there are minimal differences between the two simulations. The r-process energy has not yet caused the gas to expand or accelerate. By 10 seconds, the effects become noticeable, as the heating blurs some of the small-scale structure in the wind. The broad structure stabilizes around 30 to 60 seconds. R-process heating accelerates the slowest material to $\sim 0.1c$, which in this case, evacuates a sphere of radius $0.1ct$ at the center of the domain.

The initial and final dM/dv_r and $dM/d\theta$ distributions are shown in Fig. 5. In both models, the vast majority of the mass is at relatively low velocity $\lesssim 0.2c$. Without r-process heating, the velocity distribution of the wind is mostly unchanged by further hydrodynamic evolution. There is much more of a change between the initial and final distributions with the inclusion of r-process heating. The light teal line in the top panel of Fig. 5 is the inner boundary condition in our GRHD simulation with r-process heating. This shows a distribution that is much more sharply concentrated in velocity than the

Figure 3. Radial profiles of radial velocity at a set of representative angles at 60 s (dark colors) in our long-term 2D GRHD simulations of NS merger disc winds. Angles are measured from the South pole. Different angles are separated by an offset. The light colored lines show homology ($v_r = r/t$) for a given angle, which is generally an excellent approximation to the simulation results; most of the motion is self-similar. The profiles are only shown where $\rho > \rho_{\text{cut}}$ (equation 14). The top (bottom) panel shows results from the simulations without (with) r-process heating.

final configuration after the subsequent r-process heating and hydrodynamic evolution (dark teal line). The input distribution has almost no material with $v > 0.15c$. That is because the higher-velocity material is the first to enter the domain and has not yet a) experienced as much r-process heating, b) converted it to kinetic energy. A check of the formalism of Section 2.1 (equations 6-9) is to evaluate Δe_{kin} at 60 s, and compare the resulting distribution to the results of the 2D GRHD simulation at 60 s. The formalism accurately predicts the modal velocity of the wind, but underestimates the spread in the radial velocity distribution. The difference is especially pronounced for velocities below the mode ($v \lesssim 0.1c$).

It is also possible that our GRHD simulations with r-process heating underestimate the amount of mass with $v < 0.07$. In the underlying 3D GRMHD simulation, there is marginally-bound material that reaches a maximum radius $< r_b$, but would become unbound with the additional boost from r-process heating. The radial cut at r_b excludes this material, in effect applying a total energy floor to the ejecta. This enforces an artificial minimum velocity on the wind, which could in turn suppress the central density in our r-process model. Ideally, r-process heating should be included from the start of the GRMHD simulations of NS merger discs.

As expected, dM/dv_r is unaffected by r-process heating above a velocity of $0.3c$ (Fig. 6). In that region, we find that $dM/dv_r \propto v_r^{-3.1}$. We exclude material where $\rho < \rho_{\text{cut}}$, but it is still possible that the slope of the mass distribution is affected by the density floor of the 3D GRMHD simulation, so this fit may not be that accurate. Below $0.3c$, r-process heating causes the slope to steepen, giving a rough power law fit $dM/dv_r \propto v_r^{-3.4}$, as compared to the distribution $dM/dv_r \propto v_r^{-1.5}$ found when fitting the distribution in

Figure 4. Snapshots of mass density after 2, 10, 30, and 60 seconds of evolution in the long-term 2D GRHD simulations discussed in section 2.2. The top (bottom) panels show the runs without (with) *r*-process heating. The spatial bounds are scaled by the snapshot time. The black contour bounds the region included in the SEDONA radiation transport simulations ($r/t < 0.8c$ and $\rho > \rho_{\text{cut}}$). The primary consequence of including *r*-process heating is that the low-velocity portion of the wind is accelerated to $\sim 0.1c$. Heating also smooths out the finer density structure. The effects also become more prominent as time progresses, and the gas thermal energy has time to be converted to kinetic energy.

the simulation without *r*-process heating. Fits are calculated over the interval $0.1c \leq v_r \leq 0.25c$ for the model with *r*-process heating, and $0.03c \leq v_r \leq 0.25c$ for the model without. The values of the slopes are unlikely to be that robust, but the steepening of the slope is a general prediction of *r*-process heating.

The angular distributions are also similar between the GRHD simulations with and without *r*-process. In both cases, the mass distribution is more equatorially-concentrated at the end of the 2D GRHD simulations than it is at injection. This is due to the meridional component of the velocity at injection. The inclusion of *r*-process heating smooths some of the features in angular structure, and slightly broadens the structure. The poles remain relatively evacuated, which may be artificial, since 87 per cent of the energy from *r*-process is

Figure 5. Distribution in v_r (top panel) and meridional angle θ (bottom panel) of the material at the end of the 2D GRHD simulations ($t = 60$ s; solid lines and darker colors). We compare the simulations with (teal) and without (purple) *r*-process heating. Top panel: The radial velocity distributions of material injected at the inner boundaries of the 2D simulations are shown in lighter colors and dashed lines. The dotted line shows the mass distribution obtained by applying the formalism of Section 2.1 but evaluating Δe_{kin} at 60 s. Bottom panel: The angular distribution of the input is the same in both simulations and is shown in grey in the bottom panel. The distributions remain similar at 60 s, though the inclusion of *r*-process heating smooths small-scale structure in the distribution. The material near the pole is almost all high-velocity ($v_r > 0.3c$, dashed line in bottom panel), which is not affected by *r*-process heating.

deposited off of the grid, i.e. prior to injection into the 2D GRHD simulation. The large fraction of the energy injected off-grid may lead our models to underestimate the extent to which the lower-velocity material will fill in the jet cavity. We account for the conversion of thermal to kinetic energy, but only apply this in the direction of the existing velocity vector. Our formalism does not account for meridional expansion that could occur from the early heating prior to injection.

That said, almost all of the polar material is moving faster than $0.3c$ (Fig. 5), which is a portion of the wind whose kinematics seem mostly unaffected by *r*-process heating. We also neglect magnetic fields, which will likely resist the gas’s expansion into the cavity at the pole.

Figure 6. Same as the top panel of Fig. 5, but showing the high-velocity tail of the mass distribution, and only showing the resulting distributions at 60 s. The slopes of power-law fits to dM/dv_r for both models are shown in black. We separately fit low-velocity $\lesssim 0.3c$ material and high-velocity matter, since the latter is dynamically unaffected by r-process heating.

4 RADIATION TRANSPORT RESULTS

Using the Monte Carlo radiation transport code SEDONA, we calculate viewing angle-dependent bolometric and band light curves for our disc wind models with and without dynamical r-process heating. We include r-process heating in both SEDONA simulations, since instantaneous r-process heating at time $\sim t$ sets the luminosity at time $\sim t$. The code assumes homologous expansion, which is a good approximation at this stage; the heating does not affect the dynamics and only powers the kilonovae.

Our results are shown in Figs 7 and 8. In both models, polar viewing angles are brighter in the first 2–3 h. Afterwards, the equator brightens relative to the polar regions. While there is not much high-velocity mass, it is concentrated around $\sim 15^\circ$ from the poles, and is optically thick in the first few hours of evolution. Photons from a fast-moving photosphere near the pole are Doppler shifted towards the observer, resulting in particularly bright early polar emission. Due to the low density directly on the poles, polar observers are also able to see ‘deeper’ into the ejecta, which can lead them to see a hotter and brighter photosphere on the poles, even if said photosphere is at a lower velocity (Klion et al. 2021).

The slope of the high-velocity mass distribution, especially at the pole, may be sensitive to detailed physics not included in our calculations (e.g. equation of state, neutrino transport, and the exact magnetic field structure in the remnant disc). However, the general correlation between high-velocity material and bright, early light curves is plausibly robust. At later times, $\gtrsim 1$ d, the polar light curves continue to fall more sharply, while the equatorial light curves remain flat. Between 4 hours and 1 day, the equatorial light curves are brighter than on the poles due to the prolate structure of the wind. The greater projected surface area on the equator causes equatorial light curves to be somewhat brighter (Darbha & Kasen 2020; Korobkin et al. 2021).

The effects of dynamical r-process heating are most clearly seen when comparing light curves from the same viewing angles (Fig. 8). We also compare light curves in the Northern and Southern hemispheres, which gives an estimate of the minimum uncertainty in the light curves. The setup of the 3D GRHMD simulation is equatorially symmetric, so on average, the light curves from the North and South

Figure 7. Viewing angle-dependent isotropic equivalent bolometric light curves for models without (top) and with (bottom) r-process heating during the hydrodynamic calculations. Both models include r-process heating in the light curve calculation. The light curves shown are averages for observers within the given angular ranges. Polar angles are shown in lighter colors, and equatorial angles in darker ones. For clarity, only Southern hemisphere viewing angles are shown; angles are measured from the South pole. Polar angles are brighter in the first few hours, after which equatorial viewing angles are brighter by a factor of ~ 2 . The latter is likely due to the prolate wind structure.

should be the same. However, in our calculations the South pole is substantially brighter than the North when there is no r-process heating in the axisymmetrized GRHD run. This is likely due to a numerical artefact on the North pole of the 3D GRMHD simulation that disrupts the polar density structure. The South pole light curves are likely somewhat more reliable than those on the North pole. However, the details of the polar structure of the ejecta are, like the high-velocity distribution, sensitive to physics details that we do not explore here. The difference between the hemispheres goes away once r-process heating is included, which suggests that r-process heating may homogenize the wind structure.

The early light curve is unaffected by r-process heating, consistent with the interpretation that it primarily arises from high-velocity material that is too fast to be substantially affected by r-process heating. The primary effect of dynamical r-process heating is to make the light curves brighter but fade more quickly. This is consistent with the scaling between light curve time-scale and ejecta velocity, $t_{\text{peak}} \propto v^{-1/2}$ (Metzger 2019). On the South pole, the light curves from the models with and without dynamic r-process heating are quite similar from $\sim 6 - 30$ h. This may be a coincidence.

Our model with dynamic r-process heating has a bolometric lu-

Figure 8. Isotropic equivalent bolometric luminosity light curves, as in Fig. 7, but each panel shows a different viewing angle range. Light curves from the model with (without) dynamic *r*-process are shown in teal (purple). Southern hemisphere light curves ($\theta_S = 180^\circ - \theta_N$) are shown as solid lines; Northern hemisphere light curves are dashed. The main effect of dynamic *r*-process heating is to brighten and flatten the light curve from 0.3 – 3 d. The light curve also declines somewhat more quickly, most noticeably near the poles. This is due to the change in characteristic velocity from $\sim 0.03c$ to $\sim 0.1c$. The high-velocity material is minimally affected by *r*-process heating, so the early time portion of the light curve is nearly the same. The South pole is brighter than the North when there is no dynamical *r*-process heating, likely due to a non-general feature in the 3D GRMHD simulation. Dynamical *r*-process heating eliminates this difference.

minosity of $\sim 3 \times 10^{41}$ erg/s at 1 day. We do not directly compare our models with observations of AT 2017gfo since our wind has a mass of only $0.013M_\odot$, differing by at least a factor of 3 from the $> 0.04M_\odot$ of AT 2017gfo (e.g. Villar et al. 2017). The peak luminosity for a kilonova is expected to scale as $L_{\text{peak}} \propto M^{0.35}$, while the peak time (and therefore time-scale) $t_{\text{peak}} \propto M^{1/2}$ (Metzger 2019). Applying this correction brings our light curves to the rough energy and time-scale of AT 2017gfo (4×10^{41} erg/s at 1.7 days), though the light curves have a different shape before $\lesssim 1$ d. There is somewhat better agreement between the observations and the model with *r*-process heating. That is likely because the characteristic velocity

of our *r*-process wind is higher, more consistent with the inferred velocity of AT 2017gfo.

The broadband light curves follow similar trends to the bolometric. We show *Swift* UVOT W2 (average wavelength 193 nm), Johnson B (442 nm), and Cousins R-band (635 nm) light curves in Fig. 9. Across all bands, the polar angles are brighter than equatorial ones in the first few hours. Subsequently, the equatorial light curves brighten, while the emission near the pole remains constant or starts to fade. On time-scales of a day, equatorial emission is both brighter and bluer than the poles. Similar to the bolometric light curves, the dynamical effects of *r*-process heating are only apparent at later times when the photosphere reaches the slower (and therefore more affected) velocities.

Figure 9. *Swift* UVOT W2 (solid), B (dashed), and Cousins R-band (dashed-dotted) light curves for Southern hemisphere viewing angle ranges (measured from the South pole). Viewing angles $> 47^\circ$ are similar and are therefore shown together. Models without (with) dynamic r-process heating are in purple (teal). As with the bolometric light curves, dynamic r-process only affects the light curves after 8 hours. When dynamic r-process is included, the light curves are brighter and fade more quickly. Equatorial B and R-band light curves peak later than polar viewing angles, which have flatter light curves in the redder bands.

The difference in the evolution time-scales between models with and without r-process heating is more apparent in the band light curves, especially near the pole. The colors predicted here could also change noticeably with the inclusion of more realistic r-process opacities. Line blanketing, particularly in the blue and ultraviolet, is expected. This would make the UVOT W2 and *B*-band light curves dimmer, while causing more re-processed emission in the red and infrared. However, opacities at early ($\lesssim 1$ d) phases of kilonova evolution may be lower than they are later (Banerjee et al. 2020). The effects of realistic opacities on the colours of these models will be an interesting direction for future work.

If we spherically average the ejecta before the SEDONA calculation, we find that the light curve obtained is similar to the average of all viewing angles in the 2D calculation. We make the comparison between the spherical, averaged, and individual 2D light curves in Fig. 10. For most viewing angles, the light curves have a similar shape, mostly unaffected by the asymmetry of the ejecta. Off of the poles, the light curve is well-modeled by a spherical average. The asymmetry only matters for the polar viewing angles.

5 DISCUSSION & CONCLUSION

The kinematic effects of r-process heating on a kilonova disc wind can be substantial, accelerating the ejecta from a mass-weighted median velocity of $0.06c$ to $0.12c$ (Fig. 5). The faster wind leads to a transient that evolves more quickly and is brighter and bluer on $\sim$ day

time-scales (Figs 7 and 9). A factor-of-two increase in ejecta velocity will lead to a transient that evolves 1.4 times as quickly. Quantitative predictions of kilonova light curves from disc wind simulations should therefore account for the kinematic effects of r-process heating. The wind velocities from these simulations are sensitive to the initial magnetic field configuration. A toroidal configuration and/or a weaker field than the one used in F19 would produce a slower wind, which would experience a larger change in velocity when heated by the r-process.

The early ($\lesssim 4$ h) kilonova light curve is primarily produced by the high-velocity ($\gtrsim 0.3c$) ejecta, which is concentrated near the poles. The poles are especially bright early on because the photosphere is at high velocity, and the emission is Doppler shifted towards the polar observers (Kasliwal et al. 2017). The high-velocity portion of the ejecta is mostly unaffected by r-process heating, so the early light curves are similar to one another in the models with and without dynamic r-process heating. This portion of the light curve is likely sensitive to the physics in the underlying 3D GRMHD simulation (e.g. neutrino interactions, realistic magnetic field structure, and equation of state). Additionally, the high-velocity tail is likely to be disrupted as it runs into the dynamical ejecta from the merger. Joint modeling of the dynamical and secular ejecta is likely necessary to predict the early light curves.

After the first few hours, equatorial viewing angles become brighter than polar ones due to the prolate structure of the ejecta, which has a larger projected surface area at the equator. As the photosphere recedes into the slower portions of the ejecta, the kinematic ef-

Figure 10. Comparison of spherically averaged light curves for the model with dynamic *r*-process heating. The black line shows the light curve obtained by spherically averaging the SEDONA model before the radiation transport. At a few hours, this is 20 per cent dimmer than the average of the 2D light curves shown in Fig. 7. Light curves from individual viewing angles are shown in light teal, while the average is in dark teal.

fects of *r*-process heating become apparent. The light curves from the model with dynamical *r*-process heating are flatter at ~ 0.5 – 1.5 days, and also fade earlier. The faster evolution is particularly apparent near the poles.

The difference we observe in the low-velocity distributions with and without *r*-process heating contrasts with the results of Kawaguchi et al. (2021). They find that the effects of including *r*-process heating during the hydrodynamic phase of the calculation are very small. As they discuss, this is likely because they do not account for heating that occurs before the wind enters their simulation grid. In their simulations, most material enters at $\gtrsim 2$ s, which is longer than the 1 s time-scale on which most of the *r*-process heating occurs. In our calculations, we modify the boundary conditions of the GRHD simulation to account for *r*-process heating that occurred before injection, thus capturing in an approximate way the dominant dynamical effects of *r*-process heating.

Our formalism for pre-injection heating roughly captures the conversion of *r*-process heating to kinetic energy. However, we are unable to include all of the effects of the early heating. In particular, we do not account for the meridional expansion of the gas, which may be significant, particularly at the poles. The propagation of the jet leaves a polar cavity that may be filled by hot material from the wind. In our simulation with *r*-process heating, the evacuated region survives. It is possible that the self-consistent inclusion of *r*-process heating at all times would cause the gas to expand more and fill in the polar region. There is also the competing effect of magnetic fields, which resist polar expansion of the ejecta. A 3D GRMHD accretion disc simulation that includes *r*-process heating is necessary to understand the interaction between magnetic fields and hot gas in the polar region.

The low-velocity distribution in the 2D GRHD simulation with *r*-process heating is truncated at $0.07c$, which may not be physical. In the simulation of F19, there is material that is marginally bound and reaches a maximum radius $< r_b = 2 \times 10^4$ km. With *r*-process heating, some of this material would have become unbound and crossed r_b . Because we apply a cut at r_b , this material is excluded, even though it would have formed part of the wind. This may artificially

truncate the velocity distribution, making the late-time light curve fall off more quickly than it would otherwise.

The degree to which the low-velocity distribution is truncated will also have implications for kilonova spectra. Kasen et al. (2015) found that slowly-moving ($\sim 0.03c$) ejecta have resolved absorption lines that could possibly be used to identify elements in kilonova ejecta. Faster (0.1 – $0.3c$) ejecta have broadened lines that make such identification more difficult (Kasen et al. 2013).

This underscores the importance of self-consistently including *r*-process heating within the accretion disc simulations themselves. Accurately capturing the first few seconds of heating and their kinematic effects is critical for predicting both the radial velocity and angular distributions of the wind.

In our radiation transport calculations, we adopt a grey opacity. Kilonova ejecta are composed of dozens of elements with large, highly frequency-dependent opacities. Our qualitative light curve results are likely robust to changes to the opacity (e.g. faster ejecta leading to faster light curve evolution). The effects of more realistic opacities are none the less an important direction for future investigations. The majority of *r*-process products have particularly high opacities in the ultraviolet and blue portions of the spectrum, so our grey prescription may lead us to overestimate the UV and blue emission. However, at very early times (< 0.5 day), the opacity may be suppressed due to the high temperatures and ionization fractions (Banerjee et al. 2020; Klion et al. 2021).

The composition of disc winds is expected to vary, depending on a number of factors including the life time of a central (hyper-massive) neutron star (Lippuner et al. 2017). A single event can produce material with a wide range of Y_e , and the distribution can be both radially and meridionally-stratified (e.g. Just et al. 2015; Siegel & Metzger 2017; Fernández et al. 2019). The overlay of distinct ejecta components with different opacities may have interesting consequences for viewing angle-dependent kilonova light curves (Korobkin et al. 2021).

ACKNOWLEDGEMENTS

This research was funded by the Gordon and Betty Moore Foundation through grant GBMF5076. AT was supported by NASA 80NSSC18K0565 and NSF AST-1815304 grants. EQ was supported in part by a Simons Investigator award from the Simons Foundation. RF acknowledges support from the Natural Sciences and Engineering Research Council (NSERC) of Canada through Discovery Grant RGPIN-2017-04286, and from the Faculty of Science at the University of Alberta. The simulations presented here were carried out and processed using the National Energy Research Scientific Computing Center, a U.S. Department of Energy Office of Science User Facility, and the Savio computational cluster resource provided by the Berkeley Research Computing program at the University of California, Berkeley (supported by the UC Berkeley Chancellor, Vice Chancellor of Research, and Office of the CIO).

Software: HARM (Gammie et al. 2003; Noble et al. 2006), SEDONA (Kasen et al. 2006; Roth & Kasen 2015), matplotlib (Hunter 2007), numpy (Harris et al. 2020), scipy (Virtanen et al. 2020), FSPS filter files (Conroy et al. 2009; Conroy & Gunn 2010).

DATA AVAILABILITY

The data underlying this article will be shared on reasonable request to the corresponding author.

REFERENCES

- Abbott B. P., et al., 2017a, *Phys. Rev. Lett.*, **119**, 161101
- Abbott B. P., et al., 2017b, *ApJ*, **848**, L12
- Abbott B. P., et al., 2017c, *ApJ*, **848**, L13
- Arcavi I., et al., 2017, *Nature*, **551**, 64
- Banerjee S., Tanaka M., Kawaguchi K., Kato D., Gaigalas G., 2020, *ApJ*, **901**, 29
- Barnes J., Kasen D., 2013, *ApJ*, **775**, 18
- Barnes J., Kasen D., Wu M.-R., Martínez-Pinedo G., 2016, *ApJ*, **829**, 110
- Bauswein A., Goriely S., Janka H. T., 2013, *ApJ*, **773**, 78
- Bovard L., Martin D., Guercilena F., Arcones A., Rezzolla L., Korobkin O., 2017, *Phys. Rev. D*, **96**, 124005
- Chornock R., et al., 2017, *ApJ*, **848**, L19
- Christie I. M., Lalakos A., Tchekhovskoy A., Fernández R., Foucart F., Quataert E., Kasen D., 2019, *MNRAS*, **490**, 4811
- Conroy C., Gunn J. E., 2010, *ApJ*, **712**, 833
- Conroy C., Gunn J. E., White M., 2009, *ApJ*, **699**, 486
- Coulter D. A., et al., 2017, *Science*, **358**, 1556
- Cowperthwaite P. S., et al., 2017, *ApJ*, **848**, L17
- Darbha S., Kasen D., 2020, *ApJ*, **897**, 150
- Darbha S., Kasen D., Foucart F., Price D. J., 2021, *ApJ*, **915**, 69
- Dietrich T., Ujevic M., Tichy W., Bernuzzi S., Brüggmann B., 2017, *Phys. Rev. D*, **95**, 024029
- Drout M. R., et al., 2017, *Science*, **358**, 1570
- Eichler D., Livio M., Piran T., Schramm D. N., 1989, *Nature*, **340**, 126
- Fernández R., Metzger B. D., 2013, *MNRAS*, **435**, 502
- Fernández R., Quataert E., Schwab J., Kasen D., Rosswog S., 2015, *MNRAS*, **449**, 390
- Fernández R., Tchekhovskoy A., Quataert E., Foucart F., Kasen D., 2019, *MNRAS*, **482**, 3373
- Fontes C. J., Fryer C. L., Hungerford A. L., Hakel P., Colgan J., Kilcrease D. P., Sherrill M. E., 2015, *High Energy Density Physics*, **16**, 53
- Fujibayashi S., Kiuchi K., Nishimura N., Sekiguchi Y., Shibata M., 2018, *ApJ*, **860**, 64
- Gammie C. F., McKinney J. C., Tóth G., 2003, *ApJ*, **589**, 444
- Goldstein A., et al., 2017, *ApJ*, **848**, L14
- Grossman D., Korobkin O., Rosswog S., Piran T., 2014, *MNRAS*, **439**, 757
- Harris C. R., et al., 2020, *Nature*, **585**, 357
- Hotokezaka K., Kiuchi K., Kyutoku K., Okawa H., Sekiguchi Y.-i., Shibata M., Taniguchi K., 2013a, *Phys. Rev. D*, **87**, 024001
- Hotokezaka K., Kiuchi K., Kyutoku K., Muranushi T., Sekiguchi Y.-i., Shibata M., Taniguchi K., 2013b, *Phys. Rev. D*, **88**, 044026
- Hunter J. D., 2007, *Computing in Science & Engineering*, **9**, 90
- Just O., Bauswein A., Ardevol Pulpillo R., Goriely S., Janka H. T., 2015, *MNRAS*, **448**, 541
- Kasen D., Thomas R. C., Nugent P., 2006, *ApJ*, **651**, 366
- Kasen D., Badnell N. R., Barnes J., 2013, *ApJ*, **774**, 25
- Kasen D., Fernández R., Metzger B. D., 2015, *MNRAS*, **450**, 1777
- Kasen D., Metzger B., Barnes J., Quataert E., Ramirez-Ruiz E., 2017, *Nature*, **551**, 80
- Kasliwal M. M., et al., 2017, *Science*, **358**, 1559
- Kawaguchi K., Fujibayashi S., Shibata M., Tanaka M., Wanajo S., 2021, *ApJ*, **913**, 100
- Klion H., Duffell P. C., Kasen D., Quataert E., 2021, *MNRAS*, **502**, 865
- Korobkin O., et al., 2021, *ApJ*, **910**, 116
- Lattimer J. M., Mackie F., Ravenhall D. G., Schramm D. N., 1977, *ApJ*, **213**, 225
- Lehner L., Liebling S. L., Palenzuela C., Caballero O. L., O'Connor E., Anderson M., Neilsen D., 2016, *Classical and Quantum Gravity*, **33**, 184002
- Li L.-X., Paczyński B., 1998, *ApJ*, **507**, L59
- Lippuner J., Roberts L. F., 2015, *ApJ*, **815**, 82
- Lippuner J., Fernández R., Roberts L. F., Foucart F., Kasen D., Metzger B. D., Ott C. D., 2017, *MNRAS*, **472**, 904
- Metzger B. D., 2019, *Living Reviews in Relativity*, **23**, 1
- Metzger B. D., Berger E., 2012, *ApJ*, **746**, 48
- Metzger B. D., Piro A. L., Quataert E., 2009, *MNRAS*, **396**, 304
- Metzger B. D., et al., 2010, *MNRAS*, **406**, 2650
- Miller J. M., et al., 2019, *Phys. Rev. D*, **100**, 023008
- Narayan R., Piran T., Kumar P., 2001, *ApJ*, **557**, 949
- Noble S. C., Gammie C. F., McKinney J. C., Del Zanna L., 2006, *ApJ*, **641**, 626
- Oechslin R., Janka H. T., Marek A., 2007, *A&A*, **467**, 395
- Perego A., Radice D., Bernuzzi S., 2017, *ApJ*, **850**, L37
- Popham R., Woosley S. E., Fryer C., 1999, *ApJ*, **518**, 356
- Radice D., Perego A., Hotokezaka K., Fromm S. A., Bernuzzi S., Roberts L. F., 2018, *ApJ*, **869**, 130
- Rosswog S., Korobkin O., Arcones A., Thielemann F. K., Piran T., 2014, *MNRAS*, **439**, 744
- Roth N., Kasen D., 2015, *ApJS*, **217**, 9
- Savchenko V., et al., 2017, *ApJ*, **848**, L15
- Shibata M., Fujibayashi S., Hotokezaka K., Kiuchi K., Kyutoku K., Sekiguchi Y., Tanaka M., 2017, *Phys. Rev. D*, **96**, 123012
- Siegel D. M., Metzger B. D., 2017, *Phys. Rev. Lett.*, **119**, 231102
- Siegel D. M., Metzger B. D., 2018, *ApJ*, **858**, 52
- Soares-Santos M., et al., 2017, *ApJ*, **848**, L16
- Tanaka M., Hotokezaka K., 2013, *ApJ*, **775**, 113
- Tanaka M., Kato D., Gaigalas G., Kawaguchi K., 2020, *MNRAS*, **496**, 1369
- Tchekhovskoy A., 2019, HARMPI: 3D massively parallel general relativistic MHD code (ascl:1912.014)
- Villar V. A., et al., 2017, *ApJ*, **851**, L21
- Virtanen P., et al., 2020, *Nature Methods*, **17**, 261
- Wollaeger R. T., et al., 2018, *MNRAS*, **478**, 3298

This paper has been typeset from a $\text{\LaTeX}$ file prepared by the author.