

Two-moment Neutrino Flavor Transformation with Applications to the Fast Flavor Instability in Neutron Star Mergers

Evan Grohs¹ , Sherwood Richers² , Sean M. Couch^{3,4,5} , Francois Foucart⁶ , Julien Froustey^{1,7} , James P. Kneller¹ , and Gail C. McLaughlin¹ 

¹ Department of Physics, North Carolina State University, Raleigh, NC 27695, USA; ebgrohs@ncsu.edu

² Department of Physics and Astronomy, University of Tennessee Knoxville, Knoxville, TN 37996, USA

³ Department of Physics and Astronomy, Michigan State University, East Lansing, MI 48824, USA

⁴ Department of Computational Mathematics, Science, and Engineering, Michigan State University, East Lansing, MI 48824, USA

⁵ Facility for Rare Isotope Beams, Michigan State University, East Lansing, MI 48824, USA

⁶ Department of Physics & Astronomy, University of New Hampshire, 9 Library Way, Durham, NH 03824, USA

⁷ Department of Physics, University of California Berkeley, Berkeley, CA 94720, USA

Received 2023 September 1; revised 2023 December 7; accepted 2023 December 9; published 2024 February 21

Abstract

Multi-messenger astrophysics has produced a wealth of data with much more to come in the future. This enormous data set will reveal new insights into the physics of core-collapse supernovae, neutron star mergers, and many other objects where it is actually possible, if not probable, that new physics is in operation. To tease out different possibilities, we will need to analyze signals from photons, neutrinos, gravitational waves, and chemical elements. This task is made all the more difficult when it is necessary to evolve the neutrino component of the radiation field and associated quantum-mechanical property of flavor in order to model the astrophysical system of interest—a numerical challenge that has not been addressed to this day. In this work, we take a step in this direction by adopting the technique of angular-integrated moments with a truncated tower of dynamical equations and a closure, convolving the flavor-transformation with spatial transport to evolve the neutrino radiation quantum field. We show that moments capture the dynamical features of fast flavor instabilities in a variety of systems, although our technique is by no means a universal blueprint for solving fast flavor transformation. To evaluate the effectiveness of our moment results, we compare to a more precise particle-in-cell method. Based on our results, we propose areas for improvement and application to complementary techniques in the future.

Unified Astronomy Thesaurus concepts: [Neutrino decoupling \(1101\)](#); [Supernova neutrinos \(1666\)](#); [Neutron stars \(1108\)](#)

1. Introduction

Both neutron star mergers and core-collapse supernovae are true multi-messenger events, as they produce neutrinos, photons, gravitational waves, and chemical elements. In the coming decade, there will be a wealth of data from all of these messengers, see, e.g., Lien & Fields (2009), LSST Science Collaboration et al. (2009), Cowperthwaite et al. (2017), Bellm et al. (2018), Tartaglia et al. (2018), Holmbeck et al. (2020), Kalogera et al. (2021), and Baxter et al. (2022). In order to produce the most realistic theoretical predictions to compare with future data, much theoretical development is still needed. One significant area that requires attention is the neutrino physics of hot and dense systems (for a recent review, see Volpe 2023).

Stellar explosions that reach extreme temperatures and densities, such as core-collapse supernovae and neutron star mergers, produce enough neutrinos that they account for a substantial portion of the energy budget (for recent estimates, see Burrows et al. 2019; Bollig et al. 2021; Cusinato et al. 2022; Hayashi et al. 2022; Foucart et al. 2023; Fujibayashi et al. 2023). The majority of these neutrinos are in the energy range of tens of megaelectronvolts. In neutrino-rich regions, the ratio of neutrons to protons is influenced by electron neutrino and electron antineutrino capture reactions. This neutron-to-

proton ratio is a key factor influencing element synthesis (e.g., McLaughlin et al. 1996; Freiburghaus et al. 1999; Surman et al. 2006; Lippuner & Roberts 2015; Curtis et al. 2019, 2023; Miller et al. 2020; Reichert et al. 2021).

Exploratory work has demonstrated the importance of accurately understanding the impact of changes in neutrino flavor as a function of both time and position in the exploding object. For example, the distribution of neutrinos among flavors influences the outcome of the supernova explosions in one-dimensional high-mass CCSNe simulations (Stapleford et al. 2020; Ehring et al. 2023a, 2023b) and the results of element synthesis obtained in neutrino-cooled accretion disks (e.g., Markus et al. 2012; Just et al. 2022), supernovae (e.g., Duan et al. 2011; Mukhopadhyay 2022; Fujimoto & Nagakura 2023), and hypermassive neutron star outflows (e.g., George 2020; Li & Siegel 2021; Fernández et al. 2022). Since the type of elements that are produced in ejecta depend sensitively on the ratio of neutrons to protons, these studies conclude that there is an impact on the elements that are produced in Core-Collapse SuperNovae (CCSNe) and Neutron Star Mergers (NSMs).

The Quantum Kinetic Equations (QKEs), where the terms representing the interactions of the neutrinos are expanded in a series, are often taken as a starting point for calculating the outcome of neutrino transport and propagation. The first term in this series corresponds to evolution through a potential while the second term corresponds to momentum-changing collisions (e.g., Volpe et al. 2013; Vlasenko et al. 2014; Blaschke & Cirigliano 2016; Froustey et al. 2020). Different groups use the



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

phrase QKE in a variety of manners, referring to the specific terms included/excluded in the series expansion. In this work, we will use the phrase QKE to denote any equation that modifies the neutrino-density matrices in time. The starting point for classical neutrino transport can be obtained from the neutrino QKEs under the approximation that neutrino-density matrices are always on-diagonal. In this case, only the second term in the QKE series expansion, the collision term, is relevant.

Modern codes aiming to perform global 3D simulations of CCSNe (Just et al. 2015; Kuroda et al. 2016; Skinner et al. 2019; Bruenn et al. 2020) and/or NSMs and post-merger remnants (Ruffert & Janka 1999; Rosswog & Liebendörfer 2003; Neilsen et al. 2014; Wanajo et al. 2014; Foucart et al. 2016a; Foucart et al. 2021; Perego et al. 2016; Ardevol-Pulillo et al. 2019; Gizzi et al. 2019; Radice et al. 2022) that incorporate neutrinos generally use such classical transport algorithms. Given the difficulty of accurately solving the transport problem in global simulations with sufficient resolution and detailed microphysics, these codes inevitably need to make a range of additional approximations. On the methods side, these might include the use of approximate transport schemes (leakage, truncated moments) or low-resolution Monte Carlo methods, and/or the use of energy-integrated transport or ray-by-ray transport. On the microphysics side, this often includes the use of approximate interaction rates, reduced number of neutrino flavors, or simply ignoring interactions that are too costly to calculate in practice, e.g., pair processes and inelastic scattering. While global neutrino transport algorithms are rapidly improving, they are still having significant difficulties in capturing all important aspects of the classical transport equation (e.g., Nagakura et al. 2014, 2018; Iwakami et al. 2020 for CCSNe; and Miller et al. 2019a, 2019b for NSMs).

Meanwhile, the starting point for flavor transformation in the absence of collisions is often studied by evolving the flavor field using only the first term in the QKE series expansion. The evaluation of this term is done by use of operator splitting of the Hamiltonian (so-called *mean-field*). This set-up has been studied extensively: for example, the part of this Hamiltonian associated with neutrino coherent forward scattering on other neutrinos, in combination with other Hamiltonian terms, gives rise to the phenomenon of bipolar oscillations (for a review see Duan et al. 2006) and matter neutrino resonance transitions (Malkus et al. 2012, 2014; Wu et al. 2016; Frensel et al. 2017; Tian et al. 2017b; Vlasenko & McLaughlin 2018). Additionally, Fast Flavor Conversion (FFC) stems from the combination of specific angular distributions of neutrinos, the mean-field Hamiltonian, and inclusion of neutrino advection, e.g., Sawyer 2005; Dasgupta et al. 2017; Izaguirre et al. 2017. The relevant angular distributions are expected to occur in both supernovae, e.g., Abbar et al. (2019), Nagakura et al. (2021), Nagakura (2023), and neutron star mergers (Wu & Tamborra 2017) at positions close to the central object. A number of works exist that evaluate classically computed angular distributions to determine whether these distributions have a Fast Flavor Instability (FFI), using a variety of techniques (Dasgupta et al. 2018; Johns & Nagakura 2021; Nagakura & Johns 2021; Richers 2022; Abbar 2023). The hallmark of a test for whether an instability will exist is to look for a “crossing” between a curve that represents the number density of neutrinos as a function of angle, and the curve that represents the number density of antineutrinos as a function of angle, e.g., Dasgupta et al. (2009), Abbar & Duan (2018), Dasgupta (2022).

Ideally one wishes to use both of the first two terms in the quantum kinetic equation series and some efforts have been undertaken with the inclusion of both. When including both the first and second term in the QKE series, the collision term most typically produces decoherence of the neutrinos, e.g., Richers et al. (2019), and if that term is sufficiently large, the neutrinos tend to drift into flavor states. However, under the right conditions, the combination of the two terms can also enhance flavor transformation through collisional instabilities (Johns & Xiong 2022; Johns 2023; Xiong et al. 2023a; 2023c).

At present, there are questions about whether the QKEs can ever completely capture the behavior of neutrinos in these astrophysical systems, specifically because of the operator-splitting technique that is used to write down the Hamiltonian. There are ongoing efforts to analyze the evolution of neutrinos due to the forward scattering part of the many-body Hamiltonian under the assumption of continuous temporal interactions of all neutrinos with all other neutrinos (Balantekin & Pehlivan 2007; Pehlivan et al. 2011; Rrapaj 2020; Patwardhan et al. 2021; Cervia et al. 2022; Lacroix et al. 2022; Roggero et al. 2022; Balantekin et al. 2023; Martin et al. 2023a, 2023b; Siwach et al. 2023).

Notwithstanding these questions, efforts have been made to compute the QKEs by capturing the evolution of many neutrino “packets” in many different directions (e.g., Sawyer 2005; Bhattacharyya & Dasgupta 2020; Martin et al. 2020; Richers et al. 2021a; Zaizen & Morinaga 2021; George et al. 2023; Nagakura 2022). These methods provide useful benchmarks but are at present too computationally expensive to use extensively.

An alternative is to use a reformulation of the QKEs in terms of the angular moments (Zhang & Burrows 2013; Richers et al. 2019). This reformulation creates a series of equations describing the time evolution of each moment and one then evolves only a small number of these equations. One then has to choose what to do with the moments that appear in the evolution equations but that are not explicitly evolved. One approach is simply to ignore the evolution of the moments above some order, however, what is found in practice is that one needs to retain a large number of the moment evolution equations making this approach computationally inefficient (Dasgupta et al. 2018; Johns et al. 2020a, 2020b). An alternative solution to the truncation problem is to use a “closure” that links the unevolved moments to the lower order, evolved moments via some relationship. An example of calculations using this closure method can be found in Myers et al. (2022) and, using this approach, moment methods have been able to reproduce fast flavor transformations in neutron star merger-like conditions (Grohs et al. 2023).

In this manuscript, we consider in detail the efficacy of a two-moment implementation of the QKEs neglecting the collision term and using a closure. We illustrate our method using an example quantum closure that is a relatively straightforward extension of the classical maximum entropy closure. In Section 2, we detail the QKE formalism and apply it to moments. We elucidate the angular distributions corresponding to the closure and how they imply lepton number crossings in Section 3. Section 4 gives an exposition of our implementation of neutrino flavor transformation in the framework of FLASH (Fryxell et al. 2000; Dubey et al. 2009) along with the initial and boundary conditions. In Section 5, we compare the results of our moment treatment to a more exact method for several well-studied test problems before turning to the presentation of results for three kinds of neutron star

merger-like conditions. We conclude and discuss the need for further exploration of moment QKE methods in Section 6. With regard to units: we use two conventions. When writing the neutrino flavor-transformation equations, we use natural units where $\hbar = c = 1$. When giving results of numerical calculations, we use cgs units.

2. The Moment Evolution Equations

2.1. General Formalism

We begin from the general QKEs describing the neutrino and anti-neutrino evolution adopting the nomenclature of Sigl & Raffelt (1993), Vlasenko et al. (2014), Froustey et al. (2020), and, in particular, Blaschke & Cirigliano (2016). The evolved variable in the QKEs is a generalized density matrix for neutrinos, $\varrho = \varrho(t, \mathbf{x}, \mathbf{p})$, and corresponding generalized density matrix $\bar{\varrho}$ for antineutrinos, which are functions of time t , spatial location \mathbf{x} , and momentum \mathbf{p} . In the treatment of Blaschke & Cirigliano (2016), the generalized density matrices are one-body reduced density matrices (Volpe et al. 2013; Froustey et al. 2020), and hereafter we will call ϱ and $\bar{\varrho}$ simply “density matrices” for the sake of brevity. If we are describing neutrinos and antineutrinos with two chiral states, ϱ and $\bar{\varrho}$ are $2n_f \times 2n_f$ Hermitian matrices for n_f flavors. However, in this work, we only consider left-chiral neutrinos and right-chiral antineutrinos, and ignore any kind of spin coherence (Cirigliano et al. 2015; Tian et al. 2017a). As a result, the size of the density matrices is reduced to $n_f \times n_f$ for each of the neutrinos and antineutrinos. In the case of three flavors of neutrinos, namely e, μ, τ , we can write the neutrino-density matrix as

$$\varrho = \begin{pmatrix} \varrho_{ee} & \varrho_{e\mu} & \varrho_{e\tau} \\ \varrho_{\mu e} & \varrho_{\mu\mu} & \varrho_{\mu\tau} \\ \varrho_{\tau e} & \varrho_{\tau\mu} & \varrho_{\tau\tau} \end{pmatrix} \quad (1)$$

with a similar expression for antineutrinos. In Equation (1), the diagonal terms indicate the occupation numbers for a given flavor. The off-diagonal terms of Equation (1) encode the quantum coherence between two flavors. The expressions for the number density, energy density, and the number flux vector are obtained by taking appropriate phase-space integrals of ϱ :

$$\mathcal{N}(t, \mathbf{x}) = \frac{1}{(2\pi)^3} \int d^3p \varrho(t, \mathbf{x}, \mathbf{p}), \quad (2)$$

$$\mathcal{E}(t, \mathbf{x}) = \frac{1}{(2\pi)^3} \int d^3p p \varrho(t, \mathbf{x}, \mathbf{p}), \quad (3)$$

$$\mathcal{F}^i(t, \mathbf{x}) = \frac{1}{(2\pi)^3} \int d^3p \frac{p^i}{p} \varrho(t, \mathbf{x}, \mathbf{p}), \quad (4)$$

where the superscript i indicates a component of a three-vector. Note that we approximate neutrinos as ultrarelativistic by setting the neutrino energy equal to the three-momentum magnitude p and that \mathcal{N} , \mathcal{E} , and the components of \mathcal{F} are all $n_f \times n_f$ matrices. The expressions for antineutrinos are analogous.

The QKEs for ϱ and $\bar{\varrho}$ are

$$\frac{\partial \varrho}{\partial t} + \dot{\mathbf{x}} \cdot \frac{\partial \varrho}{\partial \mathbf{x}} + \dot{\mathbf{p}} \cdot \frac{\partial \varrho}{\partial \mathbf{p}} = -\iota [H, \varrho] + C, \quad (5)$$

$$\frac{\partial \bar{\varrho}}{\partial t} + \dot{\mathbf{x}} \cdot \frac{\partial \bar{\varrho}}{\partial \mathbf{x}} + \dot{\mathbf{p}} \cdot \frac{\partial \bar{\varrho}}{\partial \mathbf{p}} = -\iota [\bar{H}, \bar{\varrho}] + \bar{C}, \quad (6)$$

where the single dot over a variable indicates differentiation with respect to time. In Equations (5) and (6), C and \bar{C} are collision terms that can change neutrino flavor, number, or momenta. We shall ignore them throughout this work given the large separation of scales between the fast flavor instability growth rate and the collision rates simulated here. In addition, we will consider systems where the particle three-momenta do not change with time, implying we may exclude the force term on the lhs of Equations (5) and (6). To study flavor transformation, we employ Hamiltonian-like operators in Equations (5) and (6) consistent with mean-field treatments. When working to first order in power counting of the QKEs (Vlasenko et al. 2014), the Hamiltonian operators are a sum of three potentials. Specifically,

$$H = H_V + H_M + H_\nu, \quad (7)$$

$$\bar{H} = H_V - H_M - H_\nu^*, \quad (8)$$

denoting the vacuum (H_V), matter (H_M), and self-interaction (H_ν) terms, and where * denotes complex conjugation. The vacuum term arises from nonzero neutrino rest masses, and we write it as

$$H_V = \frac{1}{2p} U M^2 U^\dagger, \quad (9)$$

where U is the PMNS matrix and $M^2 = \text{diag}(m_1^2, m_2^2, m_3^2)$ is the diagonal matrix of squared neutrino masses. The matter term is linear and familiar in the context of oscillations with solar neutrinos. Electrons and positrons interact weakly with neutrinos in a flavor-dependent manner, which we denote by the following expression in the case the matter fluid has zero velocity

$$H_M = \sqrt{2} G_F n_e I_e, \quad (10)$$

where $G_F \simeq 1.166 \times 10^{-11} \text{ MeV}^{-2}$ is the Fermi coupling constant, n_e is the difference between the number density of electrons and positrons, and I_e is the electron-flavor projection operator, i.e., $I_e = \text{diag}(1, 0, 0)$ for three flavors. We will work in a frame comoving with the matter fluid implying Equation (10) is valid (see Appendix B). Finally, the self-interaction potential is a consequence of neutrinos interacting with the background of other neutrinos

$$H_\nu = \frac{\sqrt{2} G_F}{(2\pi)^3} \int d^3q (1 - \cos \vartheta) [\varrho(t, \mathbf{x}, \mathbf{q}) - \bar{\varrho}^*(t, \mathbf{x}, \mathbf{q})], \quad (11)$$

where ϑ is the angle between the free variable \mathbf{p} and the integration variable \mathbf{q} .

2.2. Moment Quantum Kinetic Equations

In general, the density matrices are seven dimensional, since they depend upon time, space, and momentum. Solving the QKEs for the density matrices with sufficient temporal, spatial, and momentum resolution to ensure numerical convergence will be very computationally expensive. An alternative approach is to recast the QKEs as an infinite set of transport equations for the moments of the density matrices, and then truncate the number of moments that one solves by adopting a closure. Since moments are only five-dimensional quantities, solving their transport equations with sufficient fidelity to ensure convergence is a more feasible, though still difficult,

computational challenge. In this paper, we adopt a two-moment scheme in which we evolve only the “zeroth” and “first” angular-integrated moments. We define the zeroth, first, and second moment of ϱ to be

$$E(t, \mathbf{x}, \mathbf{p}) = \frac{p^3}{(2\pi)^3} \int d\Omega_p \varrho(t, \mathbf{x}, \mathbf{p}), \quad (12)$$

$$F^i(t, \mathbf{x}, \mathbf{p}) = \frac{p^3}{(2\pi)^3} \int d\Omega_p \frac{p^i}{p} \varrho(t, \mathbf{x}, \mathbf{p}), \quad (13)$$

$$P^{ij}(t, \mathbf{x}, \mathbf{p}) = \frac{p^3}{(2\pi)^3} \int d\Omega_p \frac{p^i p^j}{p^2} \varrho(t, \mathbf{x}, \mathbf{p}), \quad (14)$$

where $i, j \in \{x, y, z\}$ are spatial indices. Note that a different convention was chosen with respect to Myers et al. 2022 (no $1/4\pi$ prefactors), for consistency with Grohs et al. (2023). Analogous expressions exist for the antineutrinos. The integrals in these definitions are only over the momentum-space solid angle Ω_p i.e., the propagation directions of the neutrinos at a given spacetime location and comoving-frame neutrino energy, and not the entire phase-space as in Equations (2)–(4). Furthermore, the p^3 in the prefactor of Equations (12)–(14) indicates that these are the differential energy density and differential energy flux. These are the quantities we have chosen to time evolve in the FLASH code since this is the convention used in many instances of classical neutrino moment transport. However we will oftentimes show results using instead the differential number density, which is related to $E(t, \mathbf{x}, \mathbf{p})$ via

$$N(t, \mathbf{x}, \mathbf{p}) = \frac{1}{p} E(t, \mathbf{x}, \mathbf{p}). \quad (15)$$

From the definition of the moments we see we can recover the expressions in Equations (2)–(4) by integrating the moments over the neutrino energy p , i.e.,

$$\mathcal{N}(t, \mathbf{x}) = \int dp N(t, \mathbf{x}, \mathbf{p}), \quad (16)$$

$$\mathcal{E}(t, \mathbf{x}) = \int dp E(t, \mathbf{x}, \mathbf{p}), \quad (17)$$

$$\mathcal{F}^i(t, \mathbf{x}) = \int dp \frac{F^i(t, \mathbf{x}, \mathbf{p})}{p}. \quad (18)$$

In this manuscript, we will only consider monoenergetic neutrinos and, as a result, our expressions for number density and number-density moment differ by a factor of energy-bin width Δp . Note that F is the specific energy flux, but \mathcal{F} is the energy-integrated number flux. For future reference, at this point, we introduce the flux factor vector (actually a vector of matrices) which we define to be

$$\mathbf{f}_{ab} = \frac{\mathbf{F}_{ab}}{E_{ab}}, \quad (19)$$

with the norm for a component of the flavor matrix defined as

$$f_{ab} \equiv |\mathbf{f}_{ab}| = \sqrt{\sum_i (f_{ab}^i)^2}, \quad (20)$$

where i runs over the spatial indices x, y, z .

From comparing the definitions of the moments and Equation (11) for self-interactions, we observe that the self-

interaction term can be written as

$$H_\nu = H_E - \frac{1}{p} \mathbf{p} \cdot \mathbf{H}_F, \quad (21)$$

where the moment self-interaction terms are

$$H_E = \sqrt{2} G_F (\mathcal{N} - \mathcal{N}^*), \quad (22)$$

$$H_F^j = \sqrt{2} G_F (\mathcal{F}^j - \mathcal{F}^{*j}). \quad (23)$$

We will use these moment self-interaction expressions for evolving our dependent variables of E and F^j . Note, however, that when writing Equations (22) and (23), the energy-integrated quantities \mathcal{N} and \mathcal{F}^j appear. Indeed, for the particular physical phenomena we study here, namely the FFI, the ELN crossing depends on the number moments and not the energy ones (\mathcal{E} and intensity). The simulations we present in this work are for monoenergetic neutrino distributions where N and E are equal up to a units factor as shown in Equation (15). Nevertheless, we make the distinction between \mathcal{N} and \mathcal{E} for the FFI under the guise of an eventual incorporation of multi-energy distributions.

Although the physical systems that we model do occur in environments where general relativity has a pronounced effect over large distances, we will do all of our calculations in a local Minkowski reference frame where we may specify the spacetime metric as $g^{\mu\nu} = \text{diag}(-1, 1, 1, 1)$. Therefore, the three-vector contraction in Equation (21) (and all other subsequent three-vector contractions in this work) is equivalent to the 3D dot product of Euclidean space. In addition, we will assume that gradients of the fluid velocity are locally approximately zero, which further simplifies the equations of motion.

With everything defined, we can now write out the evolution equations for E and \mathbf{F} by performing moment integrations of Equations (5) and (6) and scaling by appropriate factors of $p^3/(2\pi)^3$. In Cartesian coordinates, the transport equations for the neutrino moments are

$$\frac{\partial E}{\partial t} + \frac{\partial F^j}{\partial x^j} = -\iota [H_V + H_M + H_E, E] + \iota [H_F^j, F_j], \quad (24)$$

$$\frac{\partial F^j}{\partial t} + \frac{\partial P^{jk}}{\partial x^k} = -\iota [H_V + H_M + H_E, F^j] + \iota [H_F^k, P_k^j], \quad (25)$$

where we have ignored the collision and force terms and assume the background matter to be homogeneous with zero velocity in the tetrad frame. Note that this form would also be applicable if we assume a Minkowski metric with nonzero velocity gradients for energy-integrated moments, but here we explicitly assume zero velocity everywhere. Note also that the second (pressure) moment is not an evolved quantity in a M1 transport scheme and is calculated algebraically as a function of the two time-evolved moments E and \mathbf{F} using a closure relation. The closure relation we use will be discussed in Section 3.1. Finally for completeness, we give the moment evolution equations for the antineutrinos

$$\frac{\partial \bar{E}}{\partial t} + \frac{\partial \bar{F}^j}{\partial x^j} = -\iota [H_V - H_M - H_E^*, \bar{E}] + \iota [H_F^{j*}, \bar{F}_j], \quad (26)$$

$$\frac{\partial \bar{F}^j}{\partial t} + \frac{\partial \bar{P}^{jk}}{\partial x^k} = -\iota [H_V - H_M - H_E^*, \bar{F}^j] + \iota [H_F^{k*}, \bar{P}_k^j]. \quad (27)$$

Equations (24), (25), (26), and (27) give the equations of motion for the neutrino field under study. They comprise a coupled set of matrix equations with spacetime indices $\{j, k\}$ ranging over the 3D space indices $\{1, 2, 3\}$. As an interesting aside, we note that it is possible to write the equations of motion in a 4D spacetime framework. To begin, construct the following neutrino arrays in the laboratory (Euler) frame

$$J_{(\nu)}^\alpha = \begin{pmatrix} \mathcal{N} - \bar{\mathcal{N}}^* \\ \mathcal{F} - \bar{\mathcal{F}}^* \end{pmatrix}, \quad (28)$$

$$T^{\alpha\beta} = \begin{pmatrix} E & \mathbf{F} \\ \mathbf{F}^T & \mathbf{P} \end{pmatrix}, \quad (29)$$

$$H^\alpha = -u^\alpha H_V - \sqrt{2} G_F (J_{(e)}^\alpha + J_{(\nu)}^\alpha), \quad (30)$$

$$\bar{H}^\alpha = -u^\alpha H_V + \sqrt{2} G_F (J_{(e)}^\alpha + J_{(\nu)}^\alpha)^*, \quad (31)$$

where $J_{(e)}^\alpha = u^\alpha (n_e - \bar{n}_e) I_e$. In the above Equation (29), \mathbf{F}^T denotes the transpose of the row vector \mathbf{F} into a column vector. In addition, we define the four-velocity of the reference frame $u^\mu = (1, 0, 0, 0)$ in Equations (30) and (31). With these definitions, we are able to cast the QKEs for neutrinos and antineutrinos as

$$\nabla_\beta T^{\alpha\beta} = -\iota [H_\beta, T^{\alpha\beta}], \quad (32)$$

$$\nabla_\beta \bar{T}^{\alpha\beta} = -\iota [\bar{H}_\beta, \bar{T}^{\alpha\beta}]. \quad (33)$$

In Equations (32) and (33), we adopt the convention where repeated indices are contracted with respect to the metric, i.e., $A^\alpha A_\alpha = g_{\alpha\beta} A^\alpha A^\beta$ with a $(-, +, +, +)$ convention.

3. Moment Closure Relation and Lepton Number Crossing

3.1. The Maximum Entropy Closure

The evolution equations for the fluxes \mathbf{F} and $\bar{\mathbf{F}}$ involve the spatial gradients of the pressure tensors P and \bar{P} ; the evolution of the pressure tensors involve the spatial gradients of the next moment. This pattern continues in perpetuity and results in an infinite tower of equations. This is an unavoidable property of moment decomposition. Nevertheless, in some situations the infinite set of equations can be solved: for example, when the radiation field is strongly interacting, an equation of state will relate the pressure to the energy density under the assumption of Local Thermodynamic Equilibrium (LTE) thereby closing the set of equations for the first two moments. But in general—and neutrinos in CCSNe and NSMs are both such cases—no such equation of state exists that naturally closes the set of evolution equations. The simplest approach is to propose a local, analytic relation to close the tower of equations suited for the individual problem under study that matches analytic results in the trapped and free-streaming limits. This relation is called the closure relation (or “closure” for brevity). We will adopt this same approach when proposing a closure for the quantum moments of the neutrinos that must be able to account for both neutrino advection and the flavor transformation. We begin our explanation of the closure we adopt by ignoring the flavor structure of the moments for the time being, and consider

the Maximum Entropy Closure (MEC) often used in “classical” moment transport.

By definition of the MEC, the neutrinos of a particular species assume an angular distribution in momentum-space such that the angular entropy is extremized (Minerbo 1978; Cernohorsky & Bludman 1994). In other words, the neutrinos are distributed in the momentum-space angles such that an entropy-like function is maximized under the constraints of a net number density and flux. These constraints relate directly to the dynamical variables of interest in Equations (12) and (13). As with any reasonable closure, the MEC exactly represents the radiation field in the limits far from a source (where all radiation is moving in one direction) and when the radiation is in equilibrium. We will utilize the MEC for our flavor-mixing neutrino and anti-neutrino distributions, which we summarize below for completeness.

Under the constraints of number density and flux, the neutrino distribution of a particular species a per unit solid angle of momentum-space, ψ_{aa} , is (Minerbo 1978; Cernohorsky & Bludman 1994)

$$\psi_{aa} = \frac{E_{aa}}{4\pi} \frac{Z_{aa}}{\sinh(Z_{aa})} e^{Z_{aa}\mu}, \quad (34)$$

where E_{aa} is the energy-density moment for species a , and $\mu = \hat{F}_{\text{tet}} \cdot \Omega$ gives the angular dependence. Ω is the direction unit vector and $\hat{F}_{\text{tet}} = \mathbf{F}_{aa}/E_{aa}$. The parameter Z_{aa} follows from the constraint on the magnitude of the flux factor vector f_{aa}

$$f_{aa} = \frac{1}{E_{aa}} \int d\Omega \mu \psi_{aa} = \coth Z_{aa} - \frac{1}{Z_{aa}}. \quad (35)$$

Equation (35) must, in general, be inverted numerically to obtain the value of Z_{aa} corresponding to a given f_{aa} . Once Z_{aa} is obtained, we can construct the angular distribution of the neutrinos in Equation (34). If we adopt initial neutrino distributions from a core-collapse supernova or neutron star merger simulation that uses an MEC, then this is consistent with the assumptions in the original simulations. While the full angular information is assumed in our multi-direction calculations, our two-moment scheme simply uses the MEC to determine the pressure moment.

Borrowing the terminology from classical radiation hydrodynamics, we interpolate the pressure moment between the optically thin and thick limits as

$$P^{ij} = \frac{3\chi - 1}{2} P_{\text{thin}}^{ij} + \frac{3(1 - \chi)}{2} P_{\text{thick}}^{ij}, \quad (36)$$

where the thin and thick limits are

$$P_{\text{thin}}^{ij} = E \frac{F^i F^j}{F^2}, \quad (37)$$

$$P_{\text{thick}}^{ij} = \frac{E}{3} \delta^{ij}, \quad (38)$$

and we have suppressed the flavor indices for ease in notation. Minerbo (1978) demonstrated that these assumptions lead to a simple functional form of the Eddington factor χ :

$$\chi = \frac{1}{3} + \frac{2}{15} f^2 (3 - f + 3f^2), \quad (39)$$

such that Equation (36) becomes consistent with the second angular moment of Equation (34). In Section 4.1, we explore extending this concept to matrix-valued moments necessary for quantum neutrino transport.

3.2. Lepton Number Crossing with the Maximum Entropy Closure

At this point, we give a brief interlude to discuss lepton number crossings in the context of the MEC. Assuming the momentum-space angular distributions of two neutrino species follow an MEC, one can analytically determine whether two distributions cross (Johns & Nagakura 2021; Richers 2022). Such crossings (most straightforwardly between electron neutrino and antineutrino distributions) herald flavor instabilities (Morinaga 2022). Although the conditions for neutrino flavor instability are more general and involve the other-flavor lepton numbers, we shall consider initial conditions where the x -flavor Lepton Number (XLN) is zero and therefore the ELN crossing is the sole source of the instability. We note that although only an ELN crossing is initially present, an XLN crossing can subsequently appear during the evolution.

The intersection of the angular distributions is the boundary of a 2D surface in the 3D momentum-space. Solving for the boundary is, in general, a difficult problem. However, for the purposes of FFC, simply identifying whether or not the intersection exists suffices to determine whether the system is unstable or not. Therefore, we can look at the 2D cross-sectional slice of the 3D distributions in the plane of both flux vectors to determine whether there is an ELN crossing or not.

We will determine whether an ELN crossing exists using energy-density distributions, yet an ELN crossing utilizes number-density distributions by definition. However, we stress that our energy and number variables are simply related by a constant of proportionality for monoenergetic distributions, and as a result, we will continue using the energy quantities below. Let E_{ee} and Z_{ee} define the maximum entropy distributions for electron neutrinos, and similarly \bar{E}_{ee} and \bar{Z}_{ee} for electron antineutrinos (see Equations (34) and (35)). Furthermore, we assume that the flux factors are separated by an angle θ , i.e.,

$$\cos \theta = \frac{f_{ee} \cdot \bar{f}_{ee}}{f_{ee} \bar{f}_{ee}}. \quad (40)$$

The distributions cross if (Richers 2022)

$$\frac{\eta^2}{\alpha^2 + \gamma^2} \leq 1, \quad (41)$$

where

$$\eta = \ln \left[\frac{E_{ee} Z_{ee} \sinh(\bar{Z}_{ee})}{\bar{E}_{ee} \bar{Z}_{ee} \sinh(Z_{ee})} \right], \quad (42)$$

$$\alpha = \bar{Z}_{ee} \sin \theta, \quad (43)$$

$$\gamma = \bar{Z}_{ee} \cos \theta - Z_{ee}. \quad (44)$$

We use the criterion in Equation (41) to indicate the presence of FFI when choosing the locations from NSM simulations to consider in Section 5.2.

4. Methods

We have written four QKEs (one energy density and three components for the flux density) in Equations (24) and (25).

Along with the equations for the antineutrinos, this set of coupled matrix equations comprises 32 evolution variables per energy bin per spatial cell. Our goal will be to integrate these equations under the conditions of FFI to see if this method can capture the behavior of FFC. Before presenting results of test and NSM simulations, we give some more of the pertinent details on the numerical implementation of the moment method into FLASH. In addition, we give a brief exposition on the Particle-In-Cell (PIC) method in EMU and how it was tailored to compare with FLASH.

4.1. FLASH

To study neutrino flavor transformation with moments, we use the FLASH radiation hydrodynamics code (Fryxell et al. 2000; Dubey et al. 2009), further modified by O’Connor & Couch (2018) which includes an M1 moment scheme for classical neutrino transport. It evolves the energy and flux density moments for three species: ν_e , $\bar{\nu}_e$, and ν_{other} for all other neutrinos. We modify the classical code by distinguishing $\bar{\nu}_{\text{other}}$ from ν_{other} and adding the flavor off-diagonal components of the moments. This yields a total of eight effective species that follow from the generalized density matrices of Section 2.1, and specifically a two-flavor version of Equation (1). For example, these eight species for the energy-density moments are E_{ee} , $E_{\mu\mu}$, $\text{Re}(E_{e\mu})$, and $\text{Im}(E_{e\mu})$ and the four charge-conjugate counterparts for the antineutrinos. In reality, neutrinos would oscillate between e and the other two flavors, namely, μ and τ . Our implementation of flavor mixing is two-flavor for simplicity, which artificially assumes that half of the heavy-lepton neutrinos and antineutrinos do not participate in flavor conversion, but the number of flavors does not alter our qualitative conclusions and it will be possible to implement an 18-species framework for three flavors in the future.

We decompose the domain into cells and group cells together into blocks to parallelize the computation over processors. Each block contains 16^3 cells along with ghost cells. The choice of 16^3 cells per block results, in part, from the computational resources we use for this work. For a different platform, we would be free to change the size of the blocks depending on the number of cores and available memory. We use six ghost cells in each dimension so communication occurs only at the end of each full timestep, given a stencil size of two in each direction and three substeps within each full step.⁸ The three-step integrator was originally designed to ensure consistency in the hydrodynamic evolution in FLASH using a general tabulated equation of state. Our calculations do not evolve the hydrodynamics and add an unnecessary computational cost, but we leave the structure in place to ensure future consistency with the full FLASH framework used for ab initio compact object simulations.

We extend the three-species transport subroutines from Appendix B of O’Connor & Couch (2018) to the eight species needed for flavor mixing. We use the same Harten–Lax–van Leer–Einfeldt (HLL) Riemann solver (Harten et al. 1983) to compute fluxes between grid cells for all eight species. We use a first-order method to reconstruct the interface flux and pressure values, instead of the second-order TVD reconstruction employed in O’Connor & Couch (2018). Our advection timestep is set to 0.4 times the grid cell light-crossing time.

⁸ The stencil size is the number of grid cells referenced when evaluating a numerical derivative.

The advection and mixing evolution is done using an operator-split method, where the mixing derivatives are given by the rhs of Equations (24), (25), (26), and (27). To calculate the commutators of 2×2 matrices, we decompose the density matrix into components as detailed above and use the commutation relations of the Pauli matrices. Mixing is only treated locally, with the Hamiltonian-like terms specified at a given x , or equivalently a given cell. Unlike the solver for the advection, we use an adaptive, explicit fifth-order Runge–Kutta Cash–Karp (RKCK) method (Press et al. 1992) in the mixing subroutine closely following the implementation in Grohs et al. (2016). The timestep is determined by requiring that the difference between the embedded fourth- and fifth-order solutions is smaller than one part in 10^6 for each timestep and violations in unitarity (i.e., particle number conservation) are smaller than one part in 10^3 .

Finally, we discuss the MEC as implemented for mixing. The Eddington factor in Equation (39) is a derived result from the assumptions of a classical distribution maximizing angular entropy. Although our evolved quantities are expressed in a particular flavor basis, the physical evolution should not be basis dependent. Naively evaluating flux factors and Eddington factors using Equation (39) would break basis independence. We could diagonalize the energy-density moment such that the off-diagonal components of E are zero. We would also need to apply the same unitary transformation to each vector component of F , but there is no guarantee that E and F^i are simultaneously diagonalizable. In addition, the flux factors of the flavor off-diagonal quantities are in general complex and can be arbitrarily large or small irrespective of whether the radiation is in the trapped or free-streaming regime, making naive flux factors for flavor off-diagonal components a poor choice for interpolating between these regimes.

To ameliorate these issues, we can make the assumed pressure tensor independent of the flavor basis if we calculate a single χ for neutrinos and a single $\bar{\chi}$ for antineutrinos using flavor-traced flux factors. Specifically, those flavor-traced flux factors are defined as

$$f^{(\text{FT})} = \frac{|\text{Tr}[F]|}{\text{Tr}[E]} \quad (45)$$

$$= \frac{|(F_{ee}^i + F_{xx}^i)\hat{x}_i|}{E_{ee} + E_{xx}}, \quad (46)$$

and a similar expression for the antineutrinos and $\bar{f}^{(\text{FT})}$, where \hat{x}_i are the Cartesian unit vectors. These flavor-traced flux factors are substituted into Equation (39) to obtain χ and $\bar{\chi}$, which are in turn used for all flavor components in Equation (36). This also prevents the flavor off-diagonal components from appearing to be in the optically thick regime when the flavor off-diagonal components are in the free-streaming regime. Note, however, the principal direction of the pressure tensor is computed as in Equation (36) separately for *each* flavor component, i.e., ϱ_{ee} , ϱ_{xx} , $\text{Re}[\varrho_{ex}]$, and $\text{Im}[\varrho_{ex}]$. In other words, Equations (36), (37), and (38) all have flavor indices on each quantity, except for χ .

This scheme has the disadvantage that in the limit of no flavor mixing, it does not reduce to the original two-moment transport scheme, since different flavors are no longer allowed to have independent flux factors. However, for many of the cases we study in Section 5, we are in the optically thick limit for both e and x species, and as a result $f_{ee} \sim f_{xx} \sim f^{(\text{FT})}$ and similarly for the antineutrino flux factors. We leave a more

detailed analysis of possible closures to future work (J. Kneller et al. 2024, in preparation).

4.2. EMU

EMU (Richers et al. 2021b) is a three-dimensional particle-in-cell neutrino flavor-transformation code that evolves Equations (5)–(6) individually for a large number of computational particles. To evaluate the self-interaction part of the Hamiltonian (Equation (11)) we collect the contributions of each particle to the background angular moments of the distribution using a second-order shape function, and interpolate the Hamiltonian from the grid to each particle using the same second-order shape function. The advection terms are accounted for by simply translating the position of each computational particle. The flavor-density matrix and positions of each computational particle are evolved with a global fourth-order Runge–Kutta method. The snippets of EMU code that depend on the number of neutrino flavors are automatically generated using `sympy` (Meurer et al. 2017) to carry out symbolic matrix operations, simplify the expressions, and output C++ code. This allows us to run simulations assuming either two or three neutrino flavors. EMU is publicly available at Willcox & Richers (2021).

Whereas FLASH is a moment method and only evolves two angular moments for each flavor component of the neutrino distribution, EMU simulates particles moving in many individual directions. The EMU results we present in this paper were computed to 378 particles per cell corresponding to an angular resolution of roughly 11° , following the resolution tests in Richers et al. (2021a).

4.3. Initial and Boundary Conditions

We assign the flavor-diagonal values to the first two moments at every point in the domain for a FLASH calculation. There are eight neutrino species with four values (one energy-density and three flux components), for a total of 32 initial values per cell which we need to assign. However, in practice we always begin with identical moments for ϱ_{xx} and $\bar{\varrho}_{xx}$, since all heavy-lepton neutrino and antineutrino species are gathered in a single species ν_{other} in Foucart et al. (2016a) due to their very similar evolution in the absence of flavor transformation.

Our calculations of the FFI will only include the self-interacting term for the Hamiltonian-like operator in Equations (24), (25), and the antineutrino counterparts. We set the vacuum and matter potentials to zero so as to focus on the FFI, leaving the interesting physics cases of slow collective modes and matter-neutrino resonances to future work. In an actual astrophysical object, such as a CCSN or NSM, the vacuum potential would act to seed the flavor off-diagonal elements as a function of path length and neutrino energy. Since we simulate a local volume within a larger global system and thus have no information about the advection of perturbed neutrinos into our domain, we choose to take precise manual control over the initial seeds and exclude the vacuum potential. We seed the off-diagonal flavor components with a perturbation of $\mathcal{O}(10^{-6})$ compared to the diagonal components. The scale 10^{-6} is chosen such that the growth in the off-diagonal components begins in the linear regime. We have verified that starting with even smaller perturbations does not change the outcome for FLASH calculations. This is expected, as the growth should be the same in the linear regime, and thus smaller initial perturbations only take longer and use more computing resources. For the pattern of the perturbations, we use random numbers in

Table 1
List of Simulation Parameters and Initial Conditions for the Three 3D Test Simulations

Name	\mathcal{N}_{ee} (10^{32} cm^{-3})	$\overline{\mathcal{N}}_{ee}$ (10^{32} cm^{-3})	$\Sigma \mathcal{N}_{(x)}$ (10^{32} cm^{-3})	\mathbf{f}_{ee}	$\overline{\mathbf{f}}_{ee}$	$\mathbf{f}_{(x)}$	L (cm)	N_{gp}
Fiducial	4.89	4.89	0	(0, 0, 1/3)	(0, 0, -1/3)	(0, 0, 0)	8	128^3
90Degree	4.89	4.89	0	(0, 1/ $\sqrt{18}$, 1/ $\sqrt{18}$)	(0, 1/ $\sqrt{18}$, -1/ $\sqrt{18}$)	(0, 0, 0)	8	128^3
TwoThirds	4.89	3.26	0	(0, 0, 0)	(0, 0, -1/3)	(0, 0, 0)	32	128^3

Note. The first three columns show the number densities of each (anti)neutrino flavor. For clarity, the third column shows the sum of all four heavy lepton (anti)neutrino densities. Three-flavor simulations assume $\mathcal{N}_{\mu\mu} = \overline{\mathcal{N}}_{\mu\mu} = \mathcal{N}_{\tau\tau} = \overline{\mathcal{N}}_{\tau\tau} = \Sigma \mathcal{N}_{(x)}/4$, while two-flavor simulations assume $\mathcal{N}_{xx} = \overline{\mathcal{N}}_{xx} = \Sigma \mathcal{N}_{(x)}/4$, where we assume that the other half of the heavy-lepton neutrinos do not participate in flavor mixing. The next three columns show the flux factor vectors, the norm of which are the flux factors. The seventh column shows the length of each side of the domain and the eighth column the number of grid points for the baseline simulation.

each cell in order to remain agnostic to the scale of the initial perturbations that would be present in nature.

Specifically, we use the following to seed the initial perturbations in the off-diagonal components of the energy densities in FLASH

$$\delta E_{ab}(\mathbf{x}) = 10^{-6} p \max_c \{N_{cc}\} [A_{ab}(\mathbf{x}) + iB_{ab}(\mathbf{x})], \quad (47)$$

where $-1 < A, B < 1$ are uniform random numbers at each location \mathbf{x} , N_{cc} are the initial number-density moments for the diagonal components, and the Hermitian perturbation is only applied for $a \neq b$. For the flux moment, we copy the energy-density moment perturbation into the flux moment and use flux factor vectors to weight the direction

$$\delta \mathbf{F}_{ab}(\mathbf{x}) = \delta E_{ab}(\mathbf{x}) \frac{\Sigma_c N_{cc} \mathbf{f}_{cc}}{\Sigma_c N_{cc}}, \quad (48)$$

implying the initial perturbations for the off-diagonal components of E and \mathbf{F} are correlated. Analogous expressions exists for the antineutrino moments.

In the EMU calculations, each particle is assigned a four-momentum vector, weight, and density matrix. The four-momentum vectors are distributed uniformly in space, but assigned initial weights and density matrices to approximate the maximum entropy distribution (Equation (34)) separately for each flavor. In this way, the zeroth and first moments are reproduced under an appropriate angular integration for each flavor-diagonal element of the density matrix. We impose a random perturbation to the flavor off-diagonal elements of the density matrix at the level of 10^{-6} and adjust the diagonal values accordingly to preserve the length of each polarization vector.

The random numbers are determined at run time, so although the bulk properties of the instability are expected to be related by the similar initial conditions, the exact values in a particular cell have no correspondence between EMU and FLASH calculations. We will make all comparisons in the aggregate between the two sets of calculations. Furthermore, if we calculate the energy density for EMU using Equation (12) (where the angular integration becomes a sum over particle index), we would expect the incoherent sum for δE_{ab} to be reduced by \sqrt{n} for n particles per cell, implying an effectively smaller perturbation on the initial moments. To reiterate, this only impacts the length of the linear phase of the instability, not the growth rate or saturation properties. We give more details on the differences in the initial conditions between FLASH and EMU in Appendix C.

In both FLASH and EMU calculations, we use a 3D cubic box with Cartesian coordinates. The domain sizes and resolutions

of the FLASH production simulations are listed in Tables 1 and 2. We choose the domain size and cell count so that we have the resolution to resolve the fastest growing mode in the FFI, along with enough of a spatial domain to contain a few wavelengths of that fastest growing mode. We do not know the properties of the fastest growing mode a priori, so we perform convergence checks inline with the presentation of the results. The simulation durations are generally longer than the light-crossing time of the domain, implying the initial particles/densities will have advected out of the domain before the end of the simulation. We implement periodic boundary conditions for both sets of calculations implicitly assuming that the initial distribution is reasonably approximated as periodic on scales larger than the simulation domain. We also verify that changing the domain size does not impact the results.

5. Results

Our results comparing the ability of the two-moment method to reproduce the fast flavor instability are split into two parts. First, we consider in Section 5.1 the three test problems in three dimensions that were previously studied with EMU in Richers et al. (2021a). In Section 5.2, we move to consideration of conditions extracted from a dynamical neutron star merger simulation. We use the symbol $\text{Im}(\Omega)$ to denote the growth rate of $|N_{ex}|$ during instability. In addition, we use the symbol $|k|_{\text{max}}$ to denote the fastest growing mode in the discrete Fourier transform of N_{ex} during instability.

5.1. 3D Test Problems

The three 3D test problems we consider are named as Fiducial, 90° (labeled as 90Degree in Table 1 and Figures 1, 2, and 3), and TwoThirds, all of which are described in detail in Richers et al. (2021a, 2021b). None of the three tests have analytic solutions,⁹ so the comparison is based on how well the moment method of FLASH can reproduce the PIC results.

5.1.1. Test Parameters

Table 1 gives the initial conditions for simulation parameters of the three tests. The first three columns of Table 1 give the initial values of the flavor-diagonal number-density moment. All three tests start with nonzero numbers of electron neutrinos and antineutrinos, and zero other-flavor neutrinos. The fourth through sixth columns give the flux factor vectors. Although these particular flux factor vectors need at most two dimensions to be

⁹ The Fiducial calculations of Richers et al. (2021b) assume a slightly different angular distribution.

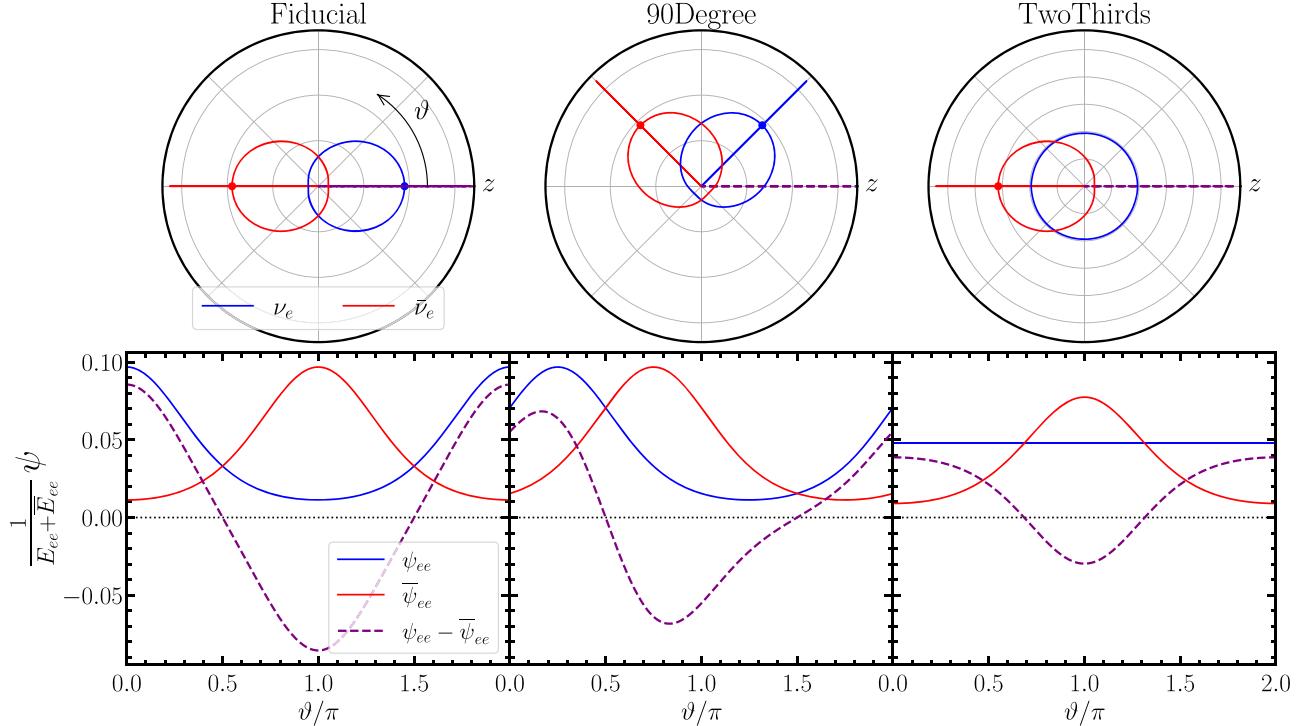


Figure 1. [Top] Polar representations of angular distributions for electron neutrino (blue) and electron antineutrinos (red) for the three tests at the beginning of the simulation. Blue (red) vectors indicate the net flux direction. Purple vector is the difference of blue and red vectors. [Bottom] Curves as given by Equation (34) for ν_e , $\bar{\nu}_e$ and the difference (purple) as a function of polar angle ϑ . Angular distributions for ν_x and $\bar{\nu}_x$ are zero at the beginning of the simulation. Lepton number crossings occur when purple line crosses 0.

Table 2
List of Baseline Simulation Parameters for the FLASH NSM Simulations

Name	\mathcal{N}_{ee} (10^{32} cm^{-3})	$\bar{\mathcal{N}}_{ee}$ (10^{32} cm^{-3})	$\Sigma \mathcal{N}_{(x)}$ (10^{32} cm^{-3})	f_{ee}	\bar{f}_{ee}	$f_{(x)}$	L (cm)	N_{gp}
NSM 1	14.22	19.15	19.65	$\begin{pmatrix} 0.0974 \\ 0.0421 \\ -0.1343 \end{pmatrix}$	$\begin{pmatrix} 0.0723 \\ 0.0313 \\ -0.3446 \end{pmatrix}$	$\begin{pmatrix} -0.0216 \\ 0.0743 \\ -0.5354 \end{pmatrix}$	7.87	128^3
NSM 2	23.29	28.53	60.11	$\begin{pmatrix} 0.0086 \\ -0.0174 \\ -0.1635 \end{pmatrix}$	$\begin{pmatrix} 0.0070 \\ -0.0142 \\ -0.2338 \end{pmatrix}$	$\begin{pmatrix} -0.0476 \\ -0.0231 \\ -0.2679 \end{pmatrix}$	8.27	256^3
NSM 3	28.80	37.42	19.32	$\begin{pmatrix} 0.0004 \\ -0.0033 \\ 0.0044 \end{pmatrix}$	$\begin{pmatrix} 0.0003 \\ -0.0025 \\ -0.1306 \end{pmatrix}$	$\begin{pmatrix} -0.0008 \\ -0.0051 \\ -0.1292 \end{pmatrix}$	5.80	512^3

Note. Column labels are the same as Table 1. Note that all corresponding EMU simulations were run with the same parameters, but using $N_{gp} = 128^3$ grid cells due to the longer wavelength of the fastest growing mode.

fully described, we stress that the calculations are three-dimensional and individual cells will generally develop flux moments where all three spatial components are nonzero. The seventh column gives the side length of the domain, and the eighth column the number of cells. Under the MEC, the initial angular distributions of all three tests in Table 1 exhibit an ELN crossing and are therefore unstable to FFC.

To visualize the geometry of these three tests, Figure 1 shows the neutrino angular distributions [Equation (34)] for the electron neutrinos (blue) and antineutrinos (red). The MEC distributions are 3D as emphasized above and as a result, we plot polar representations of 2D cross-sectional slices in the top row of Figure 1. We measure the polar angle $\vartheta \in [0, 2\pi]$ counter-clockwise from the \hat{z} axis. We take the slices such that the maximum values of the distributions (that is the directions of the fluxes F_{ee} and \bar{F}_{ee}) lie in the same plane. The polar plots

in the upper panels show a more intuitive representation of the magnitude of the distribution in different directions, but the size and depth of the ELN crossings are more apparent in the standard plots on the lower panel. The vectors originating from the origin on the polar plots show the peak direction of the distributions. The difference of the blue and red vectors is shown in dashed purple, for instance, it is coincident with the blue vector on the Fiducial case, and the vector difference is shown vividly in the 90° test. For all three tests, the coordinates are chosen so that the lepton number flux (i.e., the purple vector) lies along the z -axis. We then orient the polar plane so that this axis points in the rightward direction, and indicate this direction as $\vartheta = 0$ in the lower plots. As in the polar plots, the blue and red curves give the electron neutrino and the electron antineutrino distributions. Here, the purple curve gives the ELN distribution, properly normalized by the sum of the

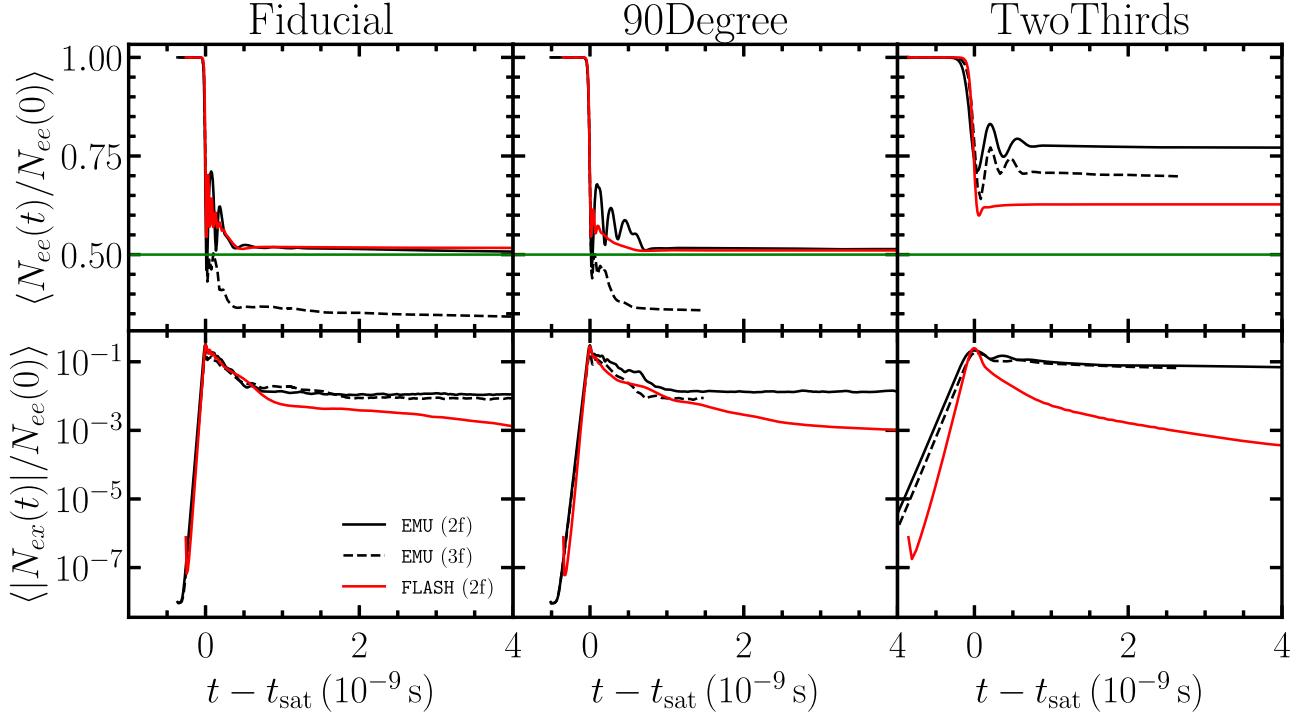


Figure 2. Density matrix elements versus time for the three tests. The horizontal axis is $t - t_{\text{sat}}$ where t_{sat} differs between test case and method of calculation. We use this definition for visualization purposes and stress the calculations are not simultaneous (see Figure 15 for the same plots using directly the simulation time for the horizontal axis). Two-flavor FLASH quantities are plotted in red. Two- (Three-) flavor EMU quantities are plotted in solid (dashed) black. All quantities are averaged over the spatial domain, and in addition over particle number for EMU. [Top] Plotted is the ee component of the number-density moment flavor matrix N (i.e., number density of ν_e) scaled by $N_{ee}(t = 0)$. [Bottom] Plotted is the magnitude of the off-diagonal component of N scaled by $N_{ee}(0)$. For three-flavor EMU calculations, we take the $e\mu$ component of N .

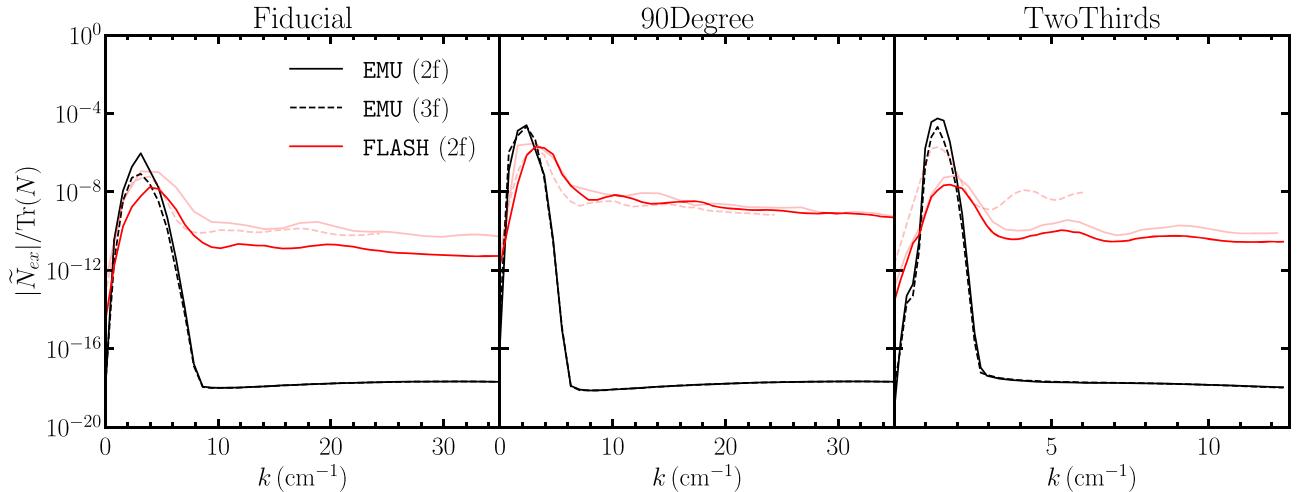


Figure 3. Magnitude of the discrete Fourier transform of N_{ex} for all three tests plotted against wavenumber. The DFTs are calculated at a time prior to saturation. The DFTs for the Fiducial and 90° test cases are taken 0.1 ns before saturation, and TwoThirds 0.35 ns before saturation. Two-flavor (three-flavor) EMU simulations correspond to the black solid (dashed) curve; FLASH simulations to the red solid curve. The two light-red curves correspond to different resolution tests for FLASH. The light-red solid curve has half the box side length and half the number of grid points compared to columns 7 and 8 in Table 1. The light-red dashed curve has the same box side length and half the number of grid points.

energy-density moments. As can clearly be seen in all three tests, the purple curves cross the horizontal axis implying a lepton number crossing.

5.1.2. Time Evolution and FFI

We show the time evolution of the domain-averaged values of $N_{ee}(t)/N_{ee}(0)$ (top) and $|N_{ex}(t)|/N_{ee}(0)$ (bottom) in Figure 2 for FLASH, two-flavor EMU, and three-flavor EMU simulations.

We emphasize that, given the different initial perturbations for FLASH and EMU calculations (see Section 4.3), we do not expect identical time evolution. Subsequently, the saturation time, t_{sat} , when the off-diagonal terms saturate (located at the peak of the $|N_{ex}|/N_{ee}(0)$ curves) depends on the initial conditions, as well as the kind of calculation. To aid in visualization when comparing the growth, saturation, and decoherence phases between FLASH and EMU, we define the horizontal axes in Figure 2 as $t - t_{\text{sat}}$ using a different t_{sat} for

each calculation. We stress that none of the calculations are simultaneous with one another in simulation time; the alignment at $t - t_{\text{sat}} = 0$ is a construct of the plot. Finally, we plot a horizontal green line on the top panel to indicate the expected number of electron neutrinos in a two-flavor calculation if the system were to completely mix flavor.

In all three tests, the FLASH simulations exhibit fast flavor instability with a growth rate very similar to the true value. Considering the red (FLASH) curves in Figure 2, there exists a period of exponential growth in $\langle |N_{ex}| \rangle$, evidencing one of the defining characteristics of the FFI. $\langle |N_{ex}| \rangle$ continues to grow until the off-diagonal magnitude reaches the same order of magnitude as the initial electron-flavor number-density moment. When $|N_{ex}| \lesssim N_{ee}$ at saturation, there are rapid oscillations in the diagonal components, evidencing the other defining characteristic of FFI. Saturation is a nearly instantaneous event with decoherence, i.e., decreasing $|N_{ex}|$, succeeding the rapid oscillations. The decoherence continues as oscillations damp, with an end result of $\langle N_{ee} \rangle$ approaching an asymptote at a value less than the starting condition. In summary, the results presented in Figure 2 are quite remarkable in that, even though the instability criterion in FFI depends upon angular crossings of the ELN, the two-moment method accurately showcases the growth of the FFI without access to crossing information. Of course, crossings are implied by the MEC distribution used to generate the closure relation, but the MEC distribution is nowhere explicitly used in the code.

The growth rate in the FLASH simulations is quantitatively very similar to that in the full EMU simulations, and even the final asymptotic neutrino distributions match well in certain cases. The Fiducial and 90° tests show strong agreement between the two methods. We see nearly identical growth rates for both tests, with FLASH producing a slightly higher $\text{Im}(\Omega)$. Specifically for the Fiducial test case, $\text{Im}(\Omega) = 7.1 \times 10^{10} \text{ s}^{-1}$ in FLASH, as compared to $\text{Im}(\Omega) = 6.3 \times 10^{10} \text{ s}^{-1}$ in EMU, for a difference of $\sim 10\%$. Results are similar for the 90° test, with $\text{Im}(\Omega) = 5.4(4.4) \times 10^{10} \text{ s}^{-1}$ in FLASH (EMU).

Also, the asymptotic values for $\langle N_{ee} \rangle$ are nearly the same, with differences of $\sim 1\%$ for both tests. We nevertheless see differences between these two tests in the saturation and decoherence periods. The oscillations in the top panels of Figure 2 for EMU appear to have larger amplitudes and persist longer than those of FLASH. Note that, during the post-saturation decoherence, there appear to be two phases indicated by different slopes in N_{ex} . In the first, immediately after saturation, N_{ex} decreases rapidly, but then numerical artifacts take over and decrease the decoherence rate (e.g., at around $t - t_{\text{sat}} = 0.75 \text{ ns}$ in the Fiducial case). In the case of the FLASH calculations, this is due to the numerical diffusion from finite grid spacing, and, in the case of EMU, this is due to the finite number of computational particles that achieve a state of random uncorrelated fluctuations, the amplitude of which scale very slowly as $N_p^{-1/2}$. Therefore, the decoherence phase right after saturation is a robust physical prediction, but the late-time values of N_{ex} show numerical artifacts.

The small amplitude oscillations for FLASH are especially evident in the TwoThirds test case. Here, we see a noticeable difference between FLASH and EMU for the asymptotic values of $\langle N_{ee} \rangle$. The growth rate is faster for FLASH by $\sim 40\%$ and the loss of coherence falls off faster. There is a smaller amount of time when $|N_{ex}| \lesssim N_{ee}$, and thus fewer oscillations in the flavor-

diagonal term. The result is an asymptotic value which is $\gtrsim 10\%$ of $N_{ee}(0)$.

For the TwoThirds test, we speculate that the reason the moment calculations do not asymptote at large times to the same value of $\langle N_{ee} \rangle$ as found by EMU is due to our imposition of the MEC. Recall that, for the FLASH calculations, we use the quantum implementation of the MEC at every timestep and substep of the evolution. In contrast, the EMU calculations use Equation (34) only when generating the initial conditions, and the future evolution depends directly on the general distribution. There is no guarantee that the neutrino distributions in EMU follow the MEC at any point except for initialization.

Although we have argued above that the use of the MEC in FLASH necessarily restricts the shape the distributions may take during flavor evolution, there does exist the striking convergence between FLASH and EMU of $\langle N_{ee} \rangle$ in the asymptotic limit for the Fiducial and 90° tests. This is not a coincidence, but rather a result of the symmetry of both of these tests. Initially, the system contains both charge-conjugation parity (CP) and rotational symmetries. The MEC is agnostic to CP but does preserve the rotational invariance for constant flavor-traced flux factors. As the energy-density moment is equal for E_{ee} and \bar{E}_{ee} , and the initial neutrino distributions are rotations of the antineutrino distributions, an ELN crossing is inevitable. The initial conditions and conservation of three-momentum ensures that neutrinos and antineutrinos will never have the same flux factor vectors at any point in the test calculations. As our system of equations is CP symmetric (except for the initial conditions in the flux factors), we expect any flavor transformation for E_{ee} to be accompanied by a commensurate change in \bar{E}_{ee} . Because three-momentum is conserved, the flux factors are invariant and the ELN crossing persists to all times. We have numerically verified that indeed \bar{E}_{ee} mirrors the evolution of E_{ee} and an ELN crossing exists in perpetuity. In other words, the distributions shown in the top panels of Figure 1 only scale in radial coordinate during their evolution. However, the results in Figure 2 clearly show a stable system post saturation. For either the Fiducial or 90° system to become stable, an XLN crossing must develop, canceling the omnipresent ELN crossing (Nagakura & Zaizen 2022; Xiong et al. 2023b; Zaizen & Nagakura 2023). Furthermore, the ν_x and $\bar{\nu}_x$ distributions have the same vector flux factors and use the same flavor-traced flux factor, implying those distributions are identical to the ones in the top panels of Figure 1 except for a difference in the radial coordinate. In the presence of non-trivial ELN and XLN crossings, a zero net lepton number at all angles requires identical energy, flux, and pressure moments for the xx components as compared to their ee counterparts—implying near flavor equilibration. Even if the distributions do not follow Equation (34) and the MEC, the symmetry of the system guarantees that $\langle N_{ee} \rangle$ must converge to 50% of the flavor trace [equivalent to $N_{ee}(0)/2$] in both the Fiducial and 90° tests. This need not be the situation in the TwoThirds test case as the system neither exhibits CP nor rotational symmetry. Here, the MEC is not an accurate representation of the distributions at later times, and as a result, the FLASH and EMU calculations show a stark divergence. Discrepancies between moment and multi-angle methods were also seen in Myers et al. (2022).

5.1.3. Pressure Moment

As discussed above, the MEC need not be a true representation of the distribution even if we find flavor convergence in the asymptotic limit. Figure 4 gives a plot of the zz pressure tensor component for the electron neutrinos in

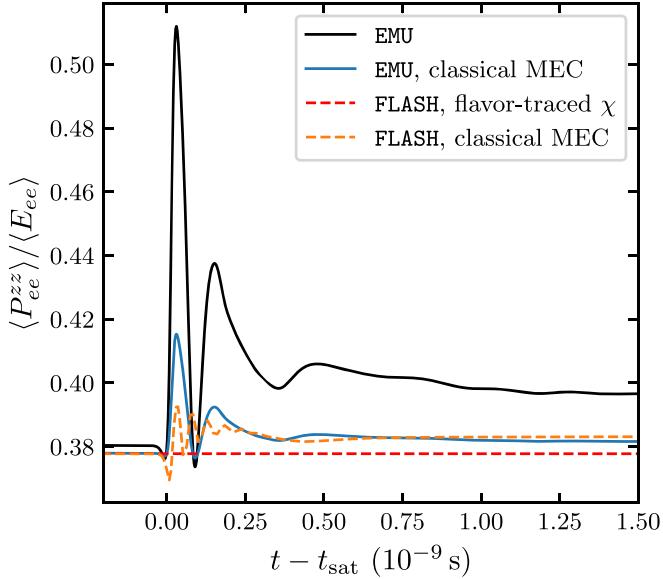


Figure 4. The zz component of the pressure tensor for electron-flavor neutrinos plotted against time for the Fiducial test case. The pressure tensor component is normalized by the energy-density moment. The solid black curve gives the pressure tensor for the EMU simulation, while the dashed red curve gives the same for the FLASH simulation. Also included are diagnostic quantities for FLASH (dashed orange) and EMU (solid blue) using a classical MEC along with the number and flux moments as given by the simulations. All quantities are averaged over the simulation domain.

the Fiducial test case. We pick the zz component for P_{ee} as \hat{z} is the direction of the net neutrino flux. For the geometry of the Fiducial test case, the thin and thick components of the pressure tensor reduce simply to $P_{\text{thin}}^{zz}/E = 1$ and $P_{\text{thick}}^{zz}/E = 1/3$, implying that the interpolated value from Equation (36) is $P^{zz}/E = \chi$. For the purposes of analyzing our moment and PIC simulations, we plot the averaged values of $P_{ee}^{zz}/E_{ee} \sim \chi$ against the time as measured from the saturation peak. We choose this representation of our data as we do not expect qualitative differences for different cells. In contrast, Nagakura & Zaizen (2023) plot the time-averaged values of the pressure against the radial coordinate when comparing multi-angle results to closure approximations in a global CCSN simulation, showing the transition from the optically thick to thin limit. The solid black curve in Figure 4 corresponds to the baseline EMU calculation, i.e., the solid black curve in the upper-left plot of Figure 2. To calculate P_{ee}^{zz} for EMU, we use Equation (14) to sum over the particles in a given cell and obtain the second angular moment of the distribution. We subsequently average over the simulation domain and normalize by the energy-density moment. The dashed red curve gives the same quantity for the FLASH simulation. For FLASH, we first calculate the domain average of the energy and flux density moments. Along with the flavor-traced Eddington factor from Equation (46), we use Equation (36) with the averaged E_{ee} and F_{ee} to obtain P_{ee}^{zz} . Finally, we normalize by $\langle E_{ee} \rangle$. The constant value of the red dashed line shows that FLASH is conserving both neutrino energy density (i.e., particle number) and neutrino flux density (i.e., three-momentum).

For diagnostic purposes, we include two other pressure quantities in Figure 4. The dashed orange curve gives the pressure using the output FLASH energy and flux moments along with the classical MEC prescription (i.e., an Eddington factor calculated using pure diagonal flux factors without a

trace over flavor). The solid blue curve gives the same but for output EMU quantities. By comparing the blue and black curves, we see how much the distribution in the PIC calculation differs from the classical MEC. Note that, at times $t < t_{\text{sat}}$, a finite number of particles causes the black curve to deviate from the blue one (we verified that increasing the number of particles reduces this discrepancy). After saturation, the black curve exhibits a larger amplitude of oscillations as compared to the blue curve. For large values of $\langle P_{ee}^{zz} \rangle$, the actual PIC calculation is more forward peaked than the MEC approximation. The opposite would be true for small values of $\langle P_{ee}^{zz} \rangle$, although it appears that the black and blue curves do not differ much at their minima. This finding is consistent with Nagakura & Zaizen (2023) during periods of significant flavor transformation (see Figures 8 and 9 of the aforementioned work), despite the differences in plotting axes. Nagakura & Zaizen (2023) show that the closure relation cannot always describe the shape of the flavor-transformed distribution during rapid flavor oscillations as the Eddington factor falls outside of the classically allowed range. Although $\langle P_{ee}^{zz} \rangle / \langle E_{ee} \rangle$ always falls within the classically allowed range for our Fiducial case, the difference between the blue and black lines is most acute at the maxima, and implies the MEC does not capture the multi-angle distribution at all times. Last, the black and blue curves oscillate nearly in phase with one another, indicating that the MEC contains the correct scaling of P_{ee}^{zz} with E_{ee} and F_{ee} but not the correct sensitivity.

We notice another difference in sensitivity when comparing the solid blue curve to the dashed orange curve of FLASH. The classical MEC calculation using FLASH data shows a smaller amplitude of oscillation, along with a larger frequency. We attribute the smaller amplitudes to the fact that the MEC underestimates the degree of forward peaking of the distribution. The larger frequency correlates with the smaller time-scales exhibited by FLASH, and observed in all three test cases. Notice that the dashed orange and solid blue curves do asymptote to similar values during the decoherence period, implying that the zeroth and first moments have similar values between the two methods of calculation. Finally, we note that our choice of utilizing the flavor-traced Eddington factor (dashed red curve in Figure 4) over the classical MEC in the FLASH simulation results in a value of $\langle P_{ee}^{zz} \rangle$ differing by $\sim 1\%$ of $\langle E_{ee} \rangle$. As we operate in the optically thick limit at all times for this test case, we do not foresee that adopting the classical MEC prescription for calculating χ would alter the results in Figure 2 by more than a few percent.

5.1.4. Fourier Space Analysis

We have discussed averages of the number density and pressure moments when presenting Figures 2 and 4. In Figure 3, we show information on the structure in the simulation domain by using Discrete Fourier Transforms (DFTs). The horizontal axes give the wavenumber k , and the vertical axis the magnitude of the DFT of the complex flavor-off-diagonal number-density moment N_{ex} , normalized by the flavor trace. We will refer to wavenumber values as “modes.” The solid red, solid black, and dashed black lines all correspond to the same simulations as Figure 2. In lighter shades of red, we have plotted DFTs from two additional FLASH calculations of the same test cases. The light-red solid curve is from a simulation with the same number of grid points per cm but with a box side length of half the original simulation

Table 3
Numerical Results for FLASH and EMU (2f) Calculations for the Three 3D Test Simulations

Name	$\langle N_{ex} /N_{ee}(0) \rangle _{t=t_{sat}}$	$\langle N_{ex}(t_{dec}) \rangle / \langle N_{ex}(t_{sat}) \rangle$	$\langle N_{ee}/N_{ee}(0) \rangle _{t \rightarrow \infty}$	$\text{Im}(\Omega)$ (10^{10} s^{-1})	$ k _{\max}$ (cm^{-1})
Fiducial					
FLASH	0.314	0.292	0.517	7.1	3.9(4)
EMU (2f)	0.333	0.360	0.506	6.3	3.1(4)
90Degree					
FLASH	0.281	0.191	0.510	5.4	3.1(4)
EMU (2f)	0.303	0.333	0.516	4.4	2.4(4)
TwoThirds					
FLASH	0.248	0.216	0.627	2.0	1.8(1)
EMU (2f)	0.214	0.579	0.771	1.2	1.4(1)

Note. First column gives the ratio $\langle |N_{ex}|/N_{ee}(0) \rangle$ when $t = t_{sat}$. Second column gives the ratio of off-diagonal magnitudes at two different times: t_{sat} and $t_{dec} = t_{sat} + 0.2$ ns. Third column gives the asymptotic ratio $\langle N_{ee}/N_{ee}(0) \rangle$ post saturation. Fourth column gives the growth rate $\text{Im}(\Omega)$ when the system is unstable in units of 10^{10} s^{-1} . Last column gives the fastest growing mode in the domain $|k|_{\max}$ in units of cm^{-1} , with an associated uncertainty in parentheses. The values of the first two columns are much more variable with the initial random perturbations than the last three.

compared to the values in columns 7 and 8 of Table 1; the light-red dashed curve is from a calculation with the same domain size, but half the number of grid points per centimeter and a smaller maximum value of k .

The three DFTs in each panel of Figure 3 are all from a time before saturation during the growth period: ~ 0.1 ns before saturation for the Fiducial and 90° cases; ~ 0.35 ns before saturation for the TwoThirds case. While similar, the times of the snapshots used in the DFTs are not exactly equal between different calculations so the values of \widetilde{N}_{ex} cannot be compared across either the simulations or the resolution tests. For this reason, comparisons should be restricted to within an individual calculation, i.e., the relative heights of peaks.

The DFTs show the scales, via wavenumber k , where there exists a sinusoidal pattern in the flavor off-diagonal number-density moment. This superposition of sinusoids need not have growing amplitudes for each mode. A priori, only one mode is necessary to explain the growth phase in Figure 2. However, during the growth phase, all modes in Figure 3 do indeed grow in power until saturation, implying there are many unstable modes in the system.

All three tests show a discernible peak in the dark-red curves of Figure 3. Soon after the simulations begin the DFT exhibits a peak with an associated wavenumber as evidenced in Figure 3. The peak remains at that location in k , although with growing height, until saturation. The DFTs for the resolution tests show similar behavior in the peak position and growth phase, indicating that the dark-red curve for the simulation is indeed spatially resolved. We call the wavenumber at this peak the fastest growing mode $|k|_{\max}$. The growth rate in Figure 2 is linked to the fastest growing mode via a dispersion relation, with details provided in Froustey et al. (2023).

The FLASH and EMU calculations both have discernible peaks with similar fastest growing modes. The wavenumber of the fastest growing modes for FLASH are slightly larger, reflecting a smaller scale. For example, $|k|_{\max} = 3.9(3.1) \text{ cm}^{-1}$ for FLASH (EMU) in the Fiducial test case, and $|k|_{\max} = 3.1(2.4) \text{ cm}^{-1}$ in the 90° case. Also, it appears that the noise floor of the DFT is larger for FLASH, or equivalently, there exists relatively less power in the fastest growing mode. Last, there are a few harmonics visible in FLASH but not present in EMU. This is true for all three test cases, and more pronounced for the TwoThirds case. These harmonics, however, only crest slightly above the noise floor.

In summary, Table 3 gives numerical results of FFC to compare between FLASH and the two-flavor EMU calculations for all three tests. The values in columns 1–4 are deduced from Figure 2 and are the following, respectively: the maximum value of $\langle |N_{ex}|/N_{ee}(0) \rangle$ in the bottom panels; the ratio of $\langle |N_{ex}| \rangle$ at the saturation time and a time $t_{dec} = t_{sat} + 0.2$ ns during the decoherence phase; the asymptote of $\langle N_{ee}/N_{ee}(0) \rangle$ in the top panels; and the slope of the line (in semi-log space) in the bottom panels. The fifth column gives the value of k at the peak of the DFT in Figure 3. We give an uncertainty in parentheses for $|k|_{\max}$ due to the finite box size, namely, $\delta k = \pm \pi/L$. All tests show the FLASH calculations have larger values of $\text{Im}(\Omega)$ and $|k|_{\max}$ compared to EMU. In addition, the rate of decline of $|N_{ex}|$ is larger for the FLASH calculations in all three tests. However, even with the different growth and loss of coherence rates, oscillations occur while the average value of $|N_{ex}|$ exhibits quite similar values for each test. This value, in the first column, is similar within a given test but not uniform across all three tests. Moreover, it varies with the random initial perturbations and should not be taken as a robust prediction for each calculation, contrary to the growth rate, instability length scale, and amount of flavor transformation.

5.2. Neutron Star Merger

Our next set of simulations use initial conditions extracted from the three-dimensional neutron star merger simulation of Foucart et al. (2016a). This simulation is general relativistic, but simulating neutrino oscillations in curved spacetime is beyond the scope of this work. Appendix B describes our procedure on transforming the distributions defined in a general spacetime metric to distributions defined in an orthonormal tetrad comoving with the fluid. In this frame, the construction of a flavor-transformation simulation is more intuitive, since we can treat the spacetime as locally flat.

We analyze and simulate neutrino distributions at a selection of points in the polar slice of a snapshot at 5 ms post merger shown in Figure 5. Green contours give matter densities of $\{10^{11}, 10^{12}, 10^{13}, 10^{14}\} \text{ g cm}^{-3}$ and the inner contours show the position of the central compact object at the center of the domain. The red pixels indicate where an ELN crossing exists according to Equation (41). White pixels indicate that no such ELN crossing exists, although these regions are still subject to flavor transformation via other processes (e.g., the matter-neutrino resonance) or advection of flavor-transformed distributions into those regions of space. We

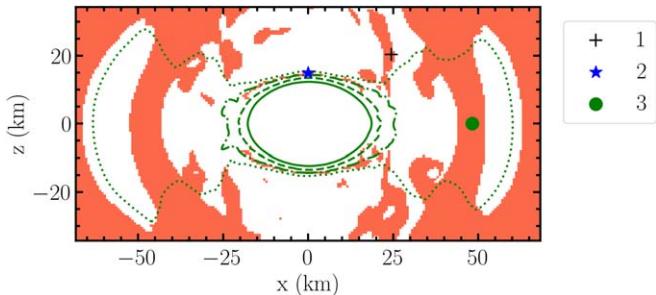


Figure 5. NSM crossing information from Foucart et al. (2016a). The snapshot is taken 5 ms post merger. Red pixels indicate locations where an ELN crossing exists. The three symbols (black cross, blue star, green circle) indicate the locations for the three flavor-transformation simulations we consider in Section 5.2. From outside to inside, the green contours indicate matter densities of $\{10^{11}, 10^{12}, 10^{13}, 10^{14}\} \text{ g cm}^{-3}$.

select three points to simulate in FLASH and EMU, indicated by the black cross, blue star, and green circle, to model regions above the accretion disk, within the disk, and just outside of the compact object, respectively. The black cross is the same point detailed in Grohs et al. (2023).

Table 2 and Figure 6 are the NSM analogs to Table 1 and Figure 1 of Section 5.1, and use the same notation and plotting conventions. Note that the orthogonalization procedure in Appendix B is location dependent and, as a result, the directions of the fluences in Table 2 cannot be compared to one another between points. In other words, for this particular study, our flavor-transformation simulations are restricted to the local area of each point and do not affect one another through advection. However, the neutrinos for points 1 and 2 are generally moving upward, so the leftward direction in the polar plots in Figure 6 roughly correspond to the \hat{z} direction in Figure 5. The leftward direction in the polar plot for point 3 in Figure 6 roughly corresponds to the \hat{x} direction in Figure 5. Note that, while there are healthy ELN crossings in points 1 and 3, the crossings in point 2 are quite tenuous, as would be expected given the very thin band of instability just above the compact object in Figure 5.

5.2.1. Time Evolution and FFI

Figure 7 shows the time evolution of the neutrino number-density moment for all three NSM points. Contrary to the previous test cases, the conditions in the NSM dictate nonzero initial distributions of heavy lepton neutrinos. This results in different two-flavor complete-mixing lines for each simulation, shown in green in Figure 7. The lighter opacity lines are for different resolution tests (see descriptions in caption). For illustrative purposes, we also include a three-flavor EMU simulation for the first NSM point and plot the $e\mu$ component in the bottom panel.

In all three NSM points, we see growth, saturation, and decoherence phases as we did in Section 5.1 and Figure 2. Growth of $|N_{ex}|$ begins soon after the start of the simulation and proceeds until $\langle |N_{ex}| \rangle \sim 0.1 \langle N_{ee} \rangle$ in both FLASH and EMU calculations. Rapid oscillations develop and effect a decrease in $\langle N_{ee} \rangle$ toward complete flavor-mixing. Notice that, for the first NSM point treated by FLASH, the conditions are such that $\langle N_{ee} \rangle$ falls below the complete-mixing green line and asymptotes to a value less than 50% of the flavor trace. This is the case for all three resolution tests, including the light-red dashed curve with half the grid spacing compared to the

baseline simulation. The EMU results also briefly dip below the 50% line, and it seems that the more rapid decoherence in the moment method halts the flavor transformation before it can oscillate back up. In all other calculations (FLASH and EMU), $\langle N_{ee} \rangle$ remains above the green line at all times. Decoherence enters after saturation in much the same manner as the three test calculations in Section 5.1. In the first and second NSM points, the baseline and resolution tests for FLASH begin to lose convergence in the decoherence phase. The divergence occurs well after saturation and at a point where $\langle N_{ee} \rangle$ has reached a steady-state value. The FLASH baseline and resolution tests for the third NSM point maintain convergence longer: a result of this set of calculations having smaller grid spacings compared to the other two points.

We identify some general trends in the FLASH and EMU results. FLASH generally shows faster growth rates, faster decoherence fall-offs, and less oscillations in the N_{ee} moment, similar to the test cases in the previous section. The discrepancies are particularly apparent for the second NSM point. This point is unique in that the distribution described by the MEC is only marginally unstable. This type of condition is expected to lead to slower growth rates, less total flavor transformation, and more dependence on details in the small angular region between the ELN crossing points (e.g., Richers et al. 2021b; Bhattacharyya & Dasgupta 2022). Specifically for FLASH (EMU), $\text{Im}(\Omega) = 5.2(1.1) \times 10^{10} \text{ s}^{-1}$ for this point. The bottom panels of Figure 6 show that ELN crossings are initially present for all three points, but the crossing is most shallow for the second point. The two-moment method plus MEC we are employing in FLASH is not able to capture the FFI behavior as well for this scenario as it is for the more pronounced ELN of the first point, and seems to behave as if there was a more significant instability than present in the detailed angular distribution, although we again emphasize that FLASH still demonstrates a characteristic evolution pattern for a fast-flavor-unstable distribution in general.

In Figure 8, we show 3D volume renderings of the FLASH simulation for the first NSM point at four different simulation times: $t = 0.01 \text{ ns}$; $t = 0.18 \text{ ns}$; $t = 0.21 \text{ ns}$; and $t = 0.59 \text{ ns}$. The contours in each panel are for the phase, ϕ_{ex} , of the complex number N_{ex} . The spatial structure in ϕ_{ex} reflects the phase of the growing mode, and so reflects the 3D structure of the peak of the DFT during the linear phase and a combination of the persisting mode structure and random decoherence after the instability saturates. We plot three contours for the phase: $\phi_{ex} = -\pi/2$ (blue); $\phi_{ex} = 0$ (white); and $\phi_{ex} = \pi/2$ (red). The first panel shows a time close to the start of the simulation, where little flavor transformation has occurred and the contours are close to the initial conditions of the random distributions with no structure [Equation (47)]. We can see some structure in the second panel during the growth phase, where the distance between planar structures reflects the wavelength of the fastest growing mode and the planes are roughly perpendicular to the direction of the net ELN flux. The phases are distorted when the evolution becomes nonlinear as the instability saturates in the third panel. The last panel is during the decoherence phase when the flavor field is no longer unstable, yet there still exists structure in ϕ_{ex} . The pattern seen here in Figure 8 is qualitatively similar to that seen in Figure 2 of Richers et al. (2021a), albeit for a different simulation that employed the PIC method. Nevertheless, the similarity in the growth of structure of ϕ_{ex} again shows that the moment method reproduces many features of the FFI on large and small scales.

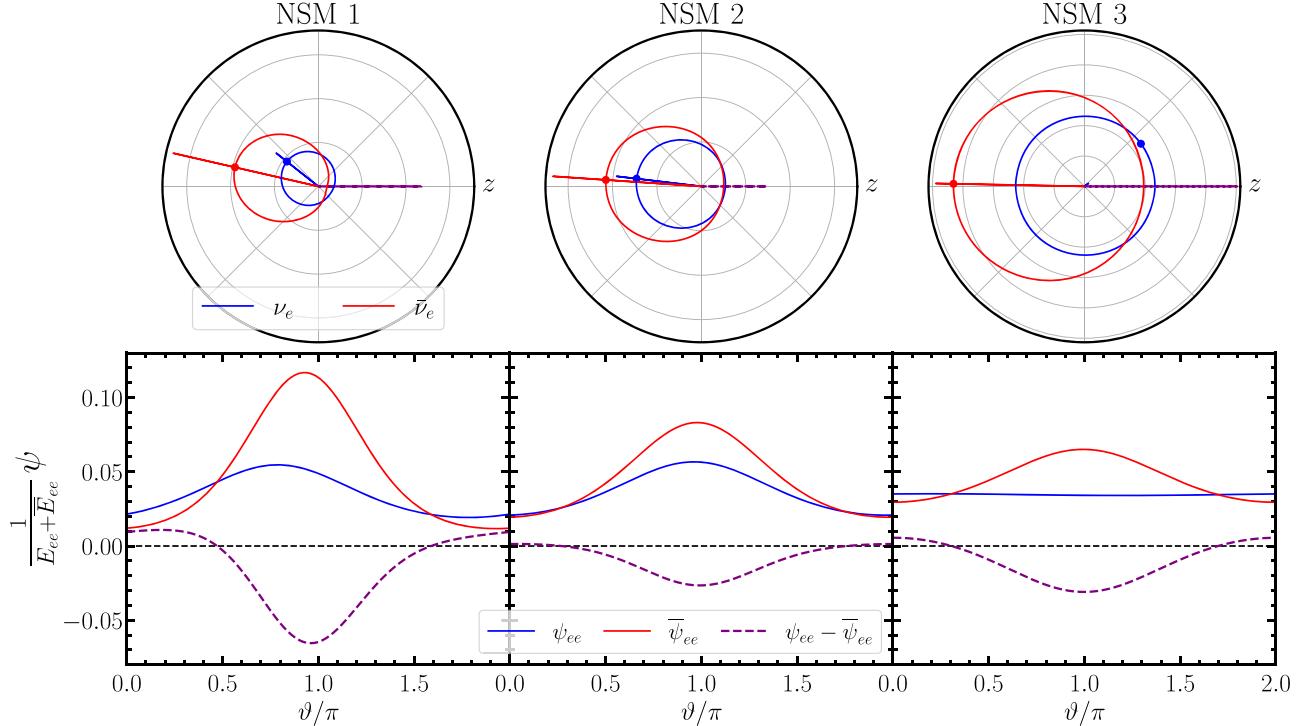


Figure 6. Polar and Cartesian representations of the initial ν_e and $\bar{\nu}_e$ distributions for the NSM points. Plotting conventions are the same as Figure 1. Note that the z -axis, corresponding to $\vartheta = 0$, is a local coordinate chosen differently at each NSM point to be coincident with the lepton number flux direction.

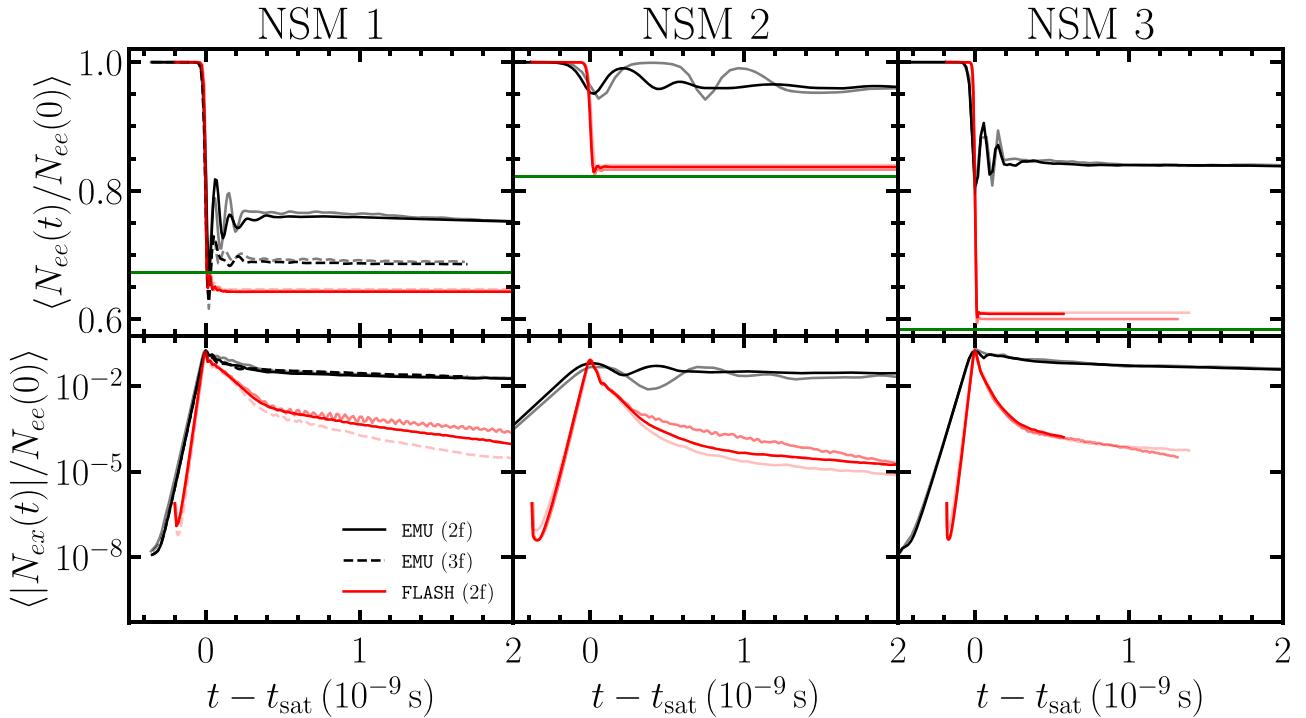


Figure 7. Domain-averaged components for the number-density moment plotted against time measured from the point of saturation. The parameter t_{sat} differs between calculations in the same manner as in Figure 2 (see Figure 16 for the same plots using directly the simulation time for the horizontal axis). Panels and curve conventions are similar to Figure 2 and simulation computational parameters are given in the last two columns of Table 2. Gray and light-red lines give results from resolution tests and are dependent on the NSM point. For NSM 1, the solid (dashed) gray lines are for EMU two-flavor (three-flavor) simulations with half the side length and half the number of grid points per side, for an identical grid spacing. The solid medium-red line also is for a FLASH simulation with half the side length and half the number of grid points. In addition, the dashed light-red line is for a simulation with half the side length but the same number of grid points, for half the grid spacing. For the three-flavor EMU calculation in NSM 1, $ex = e\mu$. NSM 2 and 3 follow identical conventions for resolution testing compared to one another. Gray lines are for two-flavor EMU simulations with half the domain size and half the number of grid points per side. Medium-red FLASH simulations are also half the domain size and half the number of grid points per side. The light-red FLASH calculation is for the same domain side length, but only half the number of grid points resulting in twice the grid spacing of the baseline simulation.

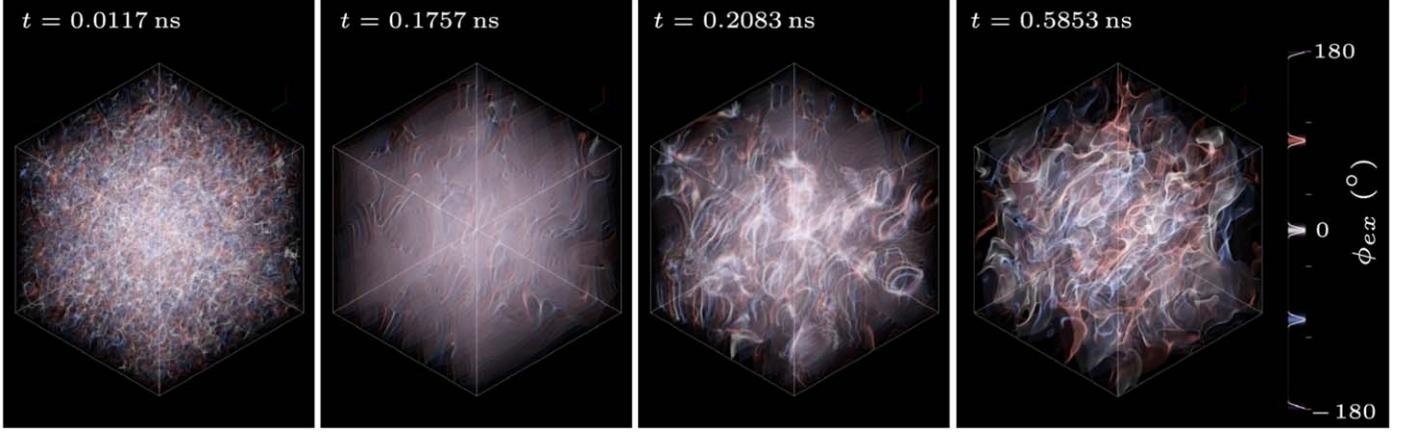


Figure 8. Volume rendering of contours of the phase of N_{ex} , ϕ_{ex} , for NSM 1 (black cross in Figure 5). Blue, white, and red contours correspond to $\phi_{ex} = -\pi/2, 0, \pi/2$, respectively. The four panels are at four different times and roughly correspond to different phases of FFC. From left to right: $t = 0.01$ ns (initial conditions); $t = 0.18$ ns (growth phase); $t = 0.21$ ns (saturation point); and $t = 0.59$ ns (decoherence phase).

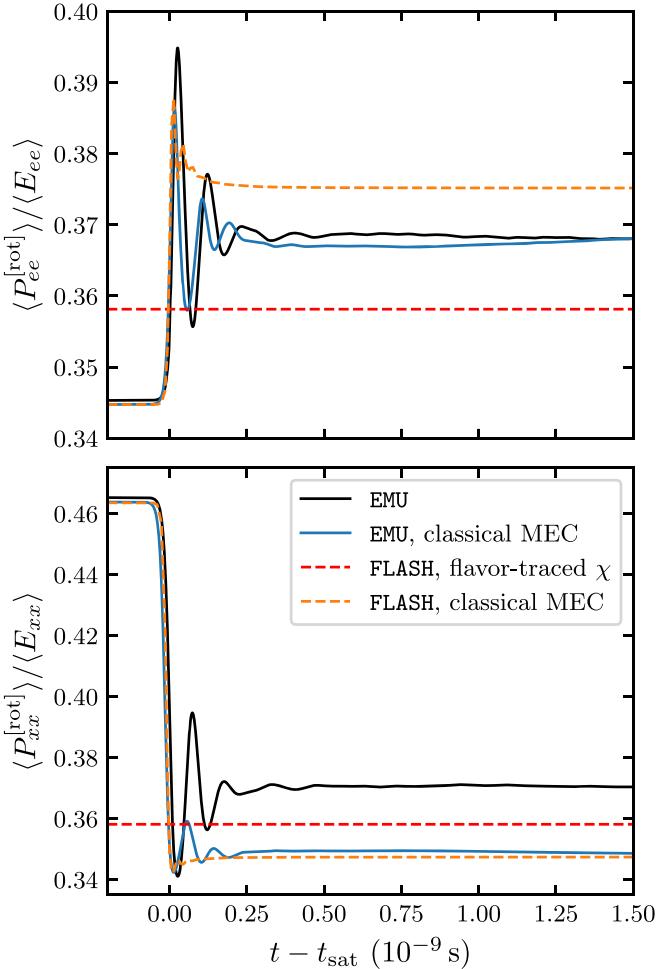


Figure 9. Component parallel to flux of pressure tensor plotted against time for the NSM 1 point. Line and axes conventions are the same as Figure 4. Top panel gives electron-flavor pressure tensor, while bottom panel gives heavy-lepton x flavor. The rotation of the pressure moment is different for each flavor and at each timestep.

5.2.2. Pressure Moment

In analogy to the Fiducial test case in Section 5.1, we compare components of the pressure tensor between FLASH and EMU calculations for the NSM 1 point in Figure 9. For the

Fiducial test case, we chose the pressure component along the symmetry axis \hat{z} when drawing Figure 4. No such symmetry axis exists for the NSM 1 point, so we instead rotate into a primed reference frame where a principal axis is aligned with the flux vector for a given density matrix component. In other words, if we define a basis $(\hat{x}', \hat{y}', \hat{z}')$ such that \hat{z}' is the unit vector in the F_{aa} direction, we compute $P_{aa}^{[rot]} \equiv P_{aa}^{\hat{x}'\hat{z}'}$, which we obtain after an appropriate spatial rotation of the pressure tensor. The rotation is different for each flavor, i.e., $a = e, x$, and computed at each timestep. The solid black curve represents this quantity, averaged over the simulation domain and normalized by $\langle E_{aa} \rangle$, for electron and heavy lepton flavor neutrinos. If the pressure moment is obtained from the closure relation (Equation 36), then by construction, $P_{aa}^{[rot]}/E_{aa} = \chi_{aa}$. This is indeed the case for the dashed red curve in the top and bottom panels of Figure 9, which is equal to the flavor-traced Eddington factor used in FLASH. For diagnostic purposes, we also represent the pressure moment computed using the classical MEC prescription, i.e., the non-flavor-traced Eddington factor obtained from the first two angular moments in EMU (solid blue curve) and FLASH (dashed orange curve).

Similar to the Fiducial test case and Figure 4, the MEC is able to capture some of the features of the underlying distribution. We note that the solid blue curve tracks the black curve quite closely and asymptotes to nearly identical values for the electron-flavor pressure (similar to Figure 4, the initial discrepancy between the black and blue curves is due to the finite number of particles). Similar to the Fiducial test case, the black curve tends to have more extreme maxima and minima, implying that the MEC underestimates the degree of forward peaking (Nagakura & Zaizen 2023). For FLASH, the dashed orange curve follows the solid blue curve for roughly half a period during the onset of rapid flavor oscillations immediately after saturation. These oscillations terminate prematurely for the FLASH simulation and continue for EMU, implying final asymptotic values for $\langle P_{ee}^{[rot]} \rangle$ that differ by few percent. However, we observe that, for the electron-flavor component, there is little variation in $\langle P_{ee}^{[rot]} \rangle$ over the simulation time, and we remain close to the optically thick limit for its entirety. The differences in the solid black, solid blue, and dashed orange lines are only a few percent compared to the flavor-traced Eddington factor displayed in the dashed red curve.

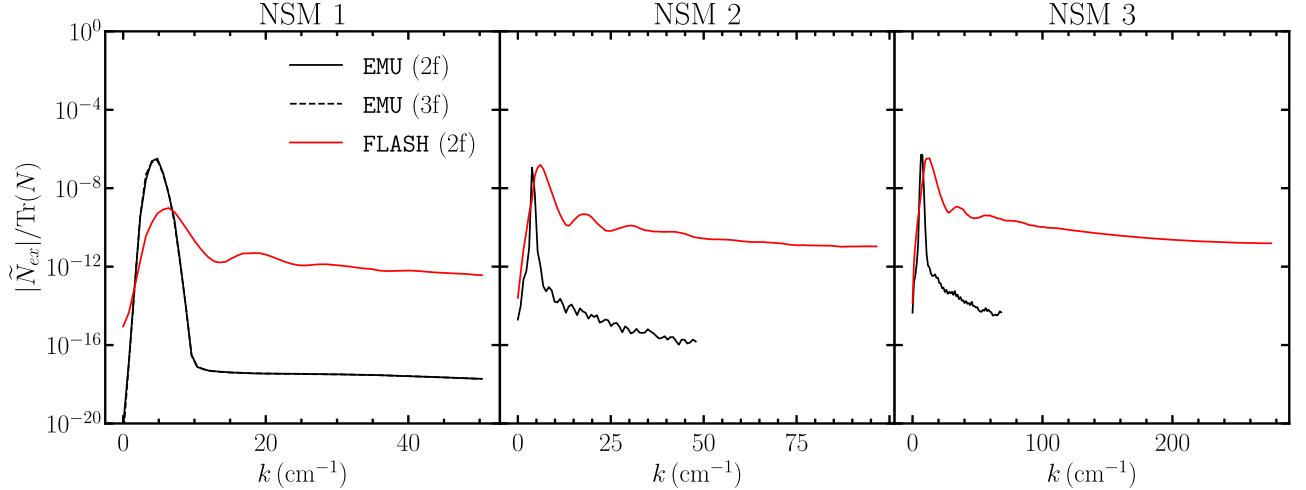


Figure 10. The magnitude of the discrete Fourier transform of N_{ex} in the three NSM simulations plotted against wavenumber. DFTs are at a time prior to saturation. For NSM 1, both FLASH and EMU DFTs are 0.1 ns before saturation. For NSM 2, the DFT for FLASH occurs 0.5 ns before saturation and 0.1 ns for EMU. For NSM 3, the DFT for FLASH occurs 0.15 ns before saturation and 0.05 ns for EMU.

Table 4
Results for FLASH and EMU (2f) Calculations for the Three 3D NSM Simulations

Name	$\langle N_{ex} /N_{ee}(0) \rangle _{t=t_{sat}}$	$\langle N_{ex}(t_{dec}) \rangle / \langle N_{ex}(t_{sat}) \rangle$	$\langle N_{ee}/N_{ee}(0) \rangle _{t \rightarrow \infty}$	$\text{Im}(\Omega) \text{ (} 10^{10} \text{ s}^{-1} \text{)}$	$ k _{\max} \text{ (cm}^{-1}\text{)}$
NSM 1					
FLASH	0.158	0.0766	0.643	8.1	6.4(4)
EMU (2f)	0.178	0.306	0.743	5.6	4.8(4)
NSM 2					
FLASH	0.0845	0.0395	0.837	5.2	6.1(4)
EMU (2f)	0.0640	0.413	0.960	1.1	3.8(4)
NSM 3					
FLASH	0.181	0.00707	0.609	10.7	13.0(5)
EMU (2f)	0.170	0.624	0.831	4.2	6.5(5)

Note. Column labels are the same as Table 3.

In the bottom panel of Figure 9, we again see good agreement between the black, blue, and orange curves prior to saturation. The expression $\langle P_{xx}^{\text{rot}} \rangle$ encompasses a larger range than $\langle P_{ee}^{\text{rot}} \rangle$. Unlike the electron flavor, here the blue and orange curves show the best agreement for the x flavor, indicating that the FLASH simulation captures the first two moments of ϱ_{xx} in accordance with EMU over a large timespan. Furthermore, regardless of the flavor, the orange and blue curves agree closely with the black curve in the growth phase. Although we use a flavor-traced Eddington factor in FLASH, the first two moments remain accurate in that linear phase.

5.2.3. Fourier Space Analysis

Figure 10 gives the DFTs for the three NSM points at 0.1 ns before saturation, roughly corresponding to the second panel in Figure 8. Because the growth rates are different for the second (third) points, we pick a time before saturation of 0.5 ns (0.15 ns) for FLASH, and 0.1 ns (0.05 ns) for EMU, all of which allow us to most clearly capture the fastest growing mode before it begins to nonlinearly couple to other modes. For the first NSM point, we find $|k|_{\max} = 6.4 \text{ cm}^{-1}$, corresponding to a wavelength of 1.0 cm and matching the distance between planar structures in the second panel of Figure 8 (domain size $L \sim 8 \text{ cm}$).

Similar to the results of Section 5.1, the DFTs from the FLASH simulations show the fastest growing modes with smaller wavelengths, higher noise floors, and visible harmonics for all three simulation points. The EMU simulations were run with a larger grid cell size to most optimally resolve the larger unstable wavelengths, resulting in a DFT that is cut off at smaller maximum k . Despite the differences between the FLASH and EMU calculations, we emphasize that all of the FLASH simulations once again reflect characteristic FFI behavior with discernible peaks with reasonable values of the fastest growing mode.

To summarize, Table 4 gives numerical results of FFC to compare between FLASH and the two-flavor EMU calculations for the three NSM simulations. The results in Table 4 are presented in the same way as Table 3. Like the Fiducial, 90°, and TwoThirds tests, the FLASH calculations for the three NSM points have larger values of $\text{Im}(\Omega)$ and $|k|_{\max}$ compared to EMU, but still exhibit reasonable behavior characteristic of the FFI.

6. Conclusions

Core collapse supernovae and merging neutron stars are complex systems that require the melding of many different physics aspects including magnetohydrodynamics, general relativity, equation of state physics, and neutrino physics. When neutrino moment methods are used currently in large-

scale simulations, they consist of classical neutrino physics and typically employ two angular moments with a closure for each neutrino species. In this work, we have extended the moment method framework in the context of the FLASH code to take into account neutrino flavor transformation in a two-moment scheme with a quantum closure. While, at present, we have tested the flavor transformation alone, the success of the moment method in modeling many features of the fast flavor instability lends confidence to incorporating it into large-scale multi-physics simulations in the future.

We found that neutrino transformation behavior is well captured in a number of test problems, including vacuum oscillations, Mikheyev–Smirnov–Wolfenstein (MSW) resonance, and bipolar oscillations. Additionally, we approximately reproduced the results of three multidimensional PIC simulations in the literature at a fraction of the computational cost. For those tests, in the key metric of the final electron neutrino number density, we find very good agreement among the first two ($\sim 1\%$) and more qualitative agreement with the last ($\sim 15\%$). Similarly, the two-moment simulations were able to almost exactly match instability growth rates in the first two although our moment method has somewhat faster growth in the last. The moment method also shows fastest growing modes that peak at a similar, but slightly higher, wavenumber than in the PIC calculations. Further analysis is required to tease out the details of the dispersion relation under the two-moment approximation in regions of instability (Froustey et al. 2023) and compare with numerical simulation.

We then performed multidimensional simulations of the FFI in three separate neutrino angular distributions taken from a full-scale classical simulation of an NSM. In key quantities, we found similar levels of agreement as we found in the TwoThirds test. Specific differences include a lower electron neutrino number density at saturation with the moment method than with EMU, and a faster growth rate with the moment method. While we always found qualitative agreement, the different methods naturally show the largest deviations when the distribution is only marginally unstable. Quantitative agreement is best with deep crossings and tends to worsen in the case of shallow crossings. This shortcoming will be important to improve upon in future advances of the algorithm, since the rapid onset of the FFI is likely to drive ELN crossings to remain shallow in astrophysical environments.

Our two-moment algorithm is based on an extension of the classical MEC relevant to quantum neutrino transport (Richers 2020). Figures 4 and 9 show that the classical MEC by itself cannot fully characterize the underlying distribution, and most certainly leads to discrepant results for the pressure tensor on the order of a few percent. As the closure accounts for missing physics from the unevolved higher-order moments, we anticipate that future efforts to develop quantum closures will improve the agreement between this two-moment method and more exact methods. To go beyond the simple prescription implemented here requires using all of the components of E and F in a basis-independent way as suggested by J. Kneller et al. (2024, in preparation).

Because of the small scales on which flavor transformation occurs, our method is at present still too computationally expensive to directly place in a full-scale CCSN or NSM simulation. However, we anticipate that in the future, methods can be developed to incorporate the very small-scale physics of flavor transformations into large-scale simulations. Nevertheless, by

virtue of only following two moments, this method is substantially computationally cheaper than exact methods that evolve neutrino distributions along hundreds or thousands of directions. Given the successes of capturing many features of the FFI in this work, we believe that moment methods can complement the higher fidelity methods such as PIC and multi-angle. For example, moment methods could be useful in doing faster realistic calculations when searching configuration space (Johns & Nagakura 2021), with follow-up post-processing being done by PIC or multi-angle codes. Machine learning is also a possibility, with moments being used to train the algorithms (Abbar 2023).

Despite the caveats above, we find that our angular-moment implementation of the QKEs reasonably and effectively captures the complex and confounding phenomenon of neutrino flavor transformation in conditions that are plausible in the environments of CCSNe and binary NSMs. Specifically, in anisotropic conditions where the angular neutrino distributions exhibit a lepton number crossing, the two moments of an M1 transport scheme manifest the phases of fast flavor conversion: exponential growth during unstable conditions; peak saturation of the off-diagonal density matrix component; rapid flavor oscillations of the diagonal components; and post-saturation decoherence with subsequent freeze-out.

There are many possible extensions that one can make to the work presented here. Improvements on the closure, the addition of the collision matrix, and an extension to three flavors are a few. Finally, while we have included advection in our simulations, the inclusion of both advection and flavor mixing in the context of a large-scale simulation will likely alter the angular distributions of the neutrinos (Padilla-Gay et al. 2021; Nagakura 2023). These improvements will widen the conditions for which this method can be used and allow us to probe other predicted phenomena, such as collisional instabilities, bipolar oscillations, and matter-neutrino resonance transitions.

Including flavor transformation in 3D general-relativistic-magnetohydrodynamics astrophysics simulations is a major computational challenge for multi-messenger astrophysics theory. Our moment-method flavor calculations offer a contribution to this important field of study.

Acknowledgments

We thank George Fuller and MacKenzie Warren for useful conversations. E.G., J.P.K., and G.C.M. are supported by the Department of Energy Office of Nuclear Physics award DE-FG02-02ER41216. E.G. acknowledges support by the National Science Foundation grant No. PHY-1430152 (Joint Institute for Nuclear Astrophysics Center for the Evolution of the Elements). S.M.C. is supported by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics, Early Career Research Program under Award Number DE-SC0015904. J.F. is supported by the Network for Neutrinos, Nuclear Astrophysics, and Symmetries (N3AS), through the National Science Foundation Physics Frontier Center grant No. PHY-2020275 and the Heising-Simons Foundation grant No. 2017-228. S.R. is supported by a National Science Foundation Astronomy and Astrophysics Postdoctoral Fellowship under award AST-2001760. F.F. is supported by the Department of Energy, Office of Science, Office of Nuclear Physics, under contract number DE-AC02-05CH11231, by NASA through grant 80NSSC22K0719, and by the National Science Foundation through grant AST-2107932. The calculations presented in this paper were undertaken, in part, on the *Payne* machine at North Carolina State University, which is supported, in part, by the

Research Corporation for Science Advancement. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231 using NERSC award ERCAP0021631. This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research and Office of Nuclear Physics, Scientific Discovery through Advanced Computing program under Award Number DE-SC0017955. The authors would like to acknowledge the use of the following software: Matplotlib (Hunter 2007), Numpy (van der Walt et al. 2011), and SciPy (Jones et al. 2001).

Appendix A 1D Test Problems

To test the algorithms for simultaneous flavor transformation and advection we implemented in FLASH, we have designed and conducted three tests in 1D to quantify the numerical accuracy and precision of the code. This appendix contains the results of those tests, which we term Vacuum; MSW; and Bipolar Oscillations. We compare to analytic solutions for the vacuum and MSW tests and to an analytic prediction for the period of the bipolar test.

A.1. Vacuum Oscillations

In this test problem, we consider only the vacuum potential H_V in the QKEs of Equations (24) and (25), but with a transformation into spherical coordinates. Neutrinos and antineutrinos are emitted isotropically from a point source at the origin. The domain has an outer radius of 4 km. We use a spherical geometry with an inner radius of $r_0 = 100$ m to avoid dividing by zero errors at the origin, and assume the flux factors at r_0 are all unity as the neutrinos stream outward in the radial direction. The number of grid points is 640, all evenly spaced in radius. Initially, all cells have the same density of the four particle species scaled by $1/r^2$. The flux factors are set to 1.0. The energy densities of the off-diagonal species are set to zero. We have set the outer boundary condition to an outflow, and the inner boundary to that of an oven emitting ν_e , $\bar{\nu}_e$, ν_x , and $\bar{\nu}_x$. At time $t = 0$, we switch on the emission of neutrinos and antineutrinos at a constant rate from a source located at r_0 moving in the positive outward direction. The neutrinos are 99% the electron flavor and 1% x flavor; the antineutrinos are 90% the electron antineutrino flavor and 10% the x -antineutrino flavor. For $r < c t$, the solution of the moment transport equations is that

$$\frac{E_{ee}^{(th)}(r)}{\text{Tr}[E^{(th)}(r)]} = [1 - P_T^{(V)}(r)] \frac{E_{ee}(r_0)}{\text{Tr}[E(r_0)]} + P_T^{(V)}(r) \frac{E_{xx}(r_0)}{\text{Tr}[E(r_0)]} \quad (\text{A1})$$

$$\frac{E_{xx}^{(th)}(r)}{\text{Tr}[E^{(th)}(r)]} = [1 - P_T^{(V)}(r)] \frac{E_{xx}(r_0)}{\text{Tr}[E(r_0)]} + P_T^{(V)}(r) \frac{E_{ee}(r_0)}{\text{Tr}[E(r_0)]} \quad (\text{A2})$$

where the function $P_T^{(V)}$ is the flavor transition probability

$$P_T^{(V)}(r) = \sin^2(2\theta) \sin^2 \left[\frac{\delta m^2}{4p} (r - r_0) \right]. \quad (\text{A3})$$

The parameter θ is the mixing angle, δm^2 is the squared neutrino mass difference, and p is the neutrino energy. Analogous expressions to Equations (A1) and (A2) for the

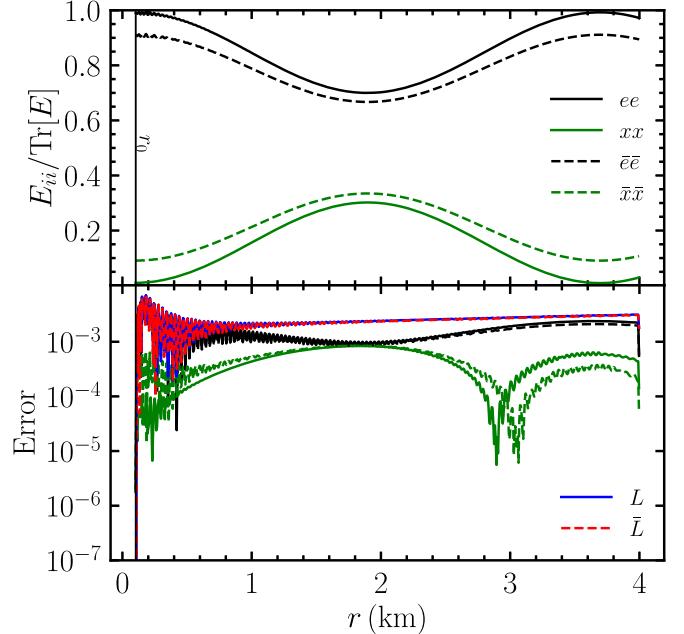


Figure 11. Results from the vacuum test of FLASH plotted against radial coordinate. [Top] Diagonal density matrix components of the energy-density moment (E_{ii}) normalized by the trace of E . Barred components denote the antineutrino counterparts, i.e., $E_{\bar{e}\bar{e}} \equiv \bar{E}_{ee}$ and are scaled by the trace of \bar{E} . Bottom: errors in various components. Green and black lines correspond to difference between the calculated E_{ii} and the theoretical predictions for Equations (A1) and (A2) scaled by the trace of E . Antineutrino counterparts follow from similar equations. Solid blue (dashed red) correspond to the difference between the calculated length of the polarization vector and the initial length of the polarization vector for neutrinos (antineutrinos) scaled by the initial polarization vector length.

theoretical antineutrino energy-density moments exist with an identical transition probability $\bar{P}_T^{(V)}(r) = P_T^{(V)}(r)$.

Figure 11 shows the results of this test at a time $t = 3.3 \times 10^{-5}$ s. We use a mixing angle $\theta = 0.28818$, a mass-squared difference $\delta m^2 = 6.9 \times 10^{-4}$ eV², and an energy $p = 1$ MeV. The top panel of the figure shows the flavor content evolving with the familiar oscillatory pattern. We use the shorthand notation of $E_{\bar{e}\bar{e}} \equiv \bar{E}_{ee}$ for the electron antineutrino energy-density moment, and similarly for the $\bar{x}\bar{x}$ component. In the lower panel, we plot the relative error of the numerical solution compared to the analytic results in Equations (A1), (A2), and the antineutrino analogs, namely

$$\text{Error} = \frac{|Q - Q^{(\text{th})}|}{Q^{(\text{th})}}, \quad (\text{A4})$$

for the relevant quantity Q . In addition to the errors in the energy-density moments, we give the relative error in the Bloch polarization vector for neutrinos, L , and antineutrinos, \bar{L} . The dominant source of the error arises from the finite grid spacing in radius. We have verified that decreasing/increasing the number of grid points increases/reduces the errors. The size of the error scales inversely with the number of grid points. Notice that the largest errors occur at small radii for E_{ee} and \bar{E}_{ee} (and consequently also for L and \bar{L}) at a level of a few parts in 10^3 . This error is due to numerical oscillations at small radii, which are a result of the inner boundary condition in spherical coordinates and not related to flavor transformation. Overall,

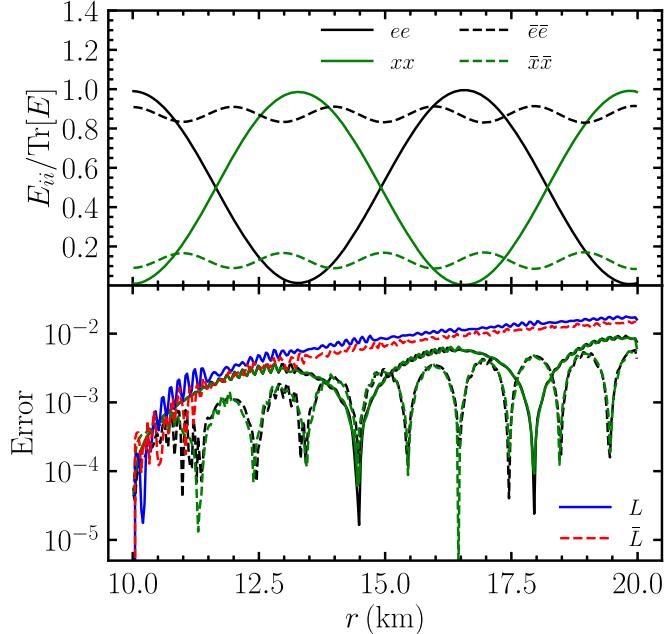


Figure 12. Results from the MSW test of FLASH plotted against radial coordinate. Notation for both panels is the same as Figure 11. Error for the E_{ii} components corresponds to the difference between the calculated E_{ii} and the theoretical predictions using the transition probability in Equation (A5) for neutrinos, and a similar expression for antineutrinos.

the FLASH calculation is able to successfully reproduce the analytic results to a precision better than $\sim 1\%$.

A.2. MSW Oscillations

Our second 1D test case is similar to the vacuum test case described in the previous section. In addition to the vacuum potential, we now add a constant matter density to observe MSW oscillations (Bethe 1986; Haxton 1986). Our domain ranges from the inner boundary at $r_0 = 10$ km to the outer boundary at $r = 20$ km. We use 320 cells evenly spaced in the radial coordinate. The evolution equations for this test can actually be transformed such that the analytic solution to the energy-density moments is the same as Equations (A1) and (A2) with a different MSW transition probability

$$P_T^{(M)}(r) = \sin^2(2\theta_{\text{eff}}) \sin^2 \left[\frac{\delta m_{\text{eff}}^2}{4p} (r - r_0) \right], \quad (\text{A5})$$

$$\sin^2(2\theta_{\text{eff}}) = \frac{\sin^2(2\theta)}{\sin^2(2\theta) + C^2}, \quad (\text{A6})$$

$$\delta m_{\text{eff}}^2 = \delta m^2 \sqrt{\sin^2(2\theta) + C^2}, \quad (\text{A7})$$

$$C = \cos(2\theta) - \frac{2\sqrt{2} p G_F Y_e \rho}{\delta m^2 m_u}, \quad (\text{A8})$$

where ρ is the matter density, Y_e the electron fraction, and m_u the atomic mass unit. For antineutrinos $\tilde{P}_T^{(M)}$ is calculated the same as above but with the placement of a minus sign on the electron fraction, $Y_e \rightarrow -Y_e$. We can then compare this solution to the FLASH calculation, which simply adds H_M to H_V in Equations (24) and (25).

The results of this test case at $t = 1.67 \times 10^{-4}$ s are shown in Figure 12, where the panels have the same plotting conventions as those in Figure 11. We use a matter density of

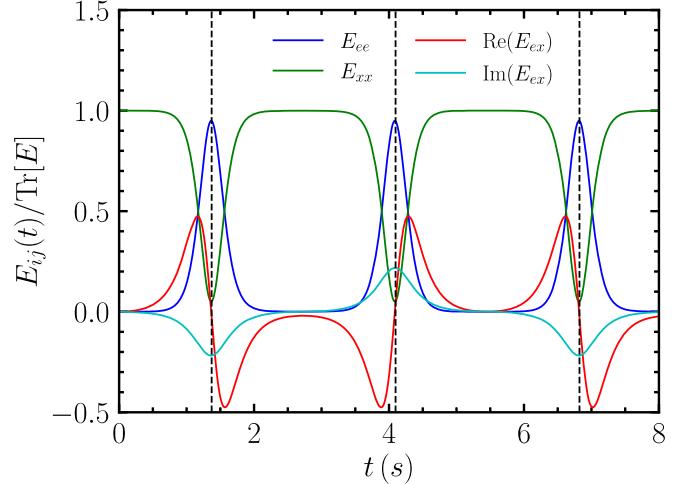


Figure 13. Results from the bipolar test of FLASH plotted against time. Only density matrix components for neutrinos are shown. We include vertical dashed black lines to indicate the period of the oscillations based on the analytical result of Hannestad et al. (2006).

$\rho = 8 \times 10^3 \text{ g cm}^{-3}$ and an electron fraction $Y_e = 0.5$. The familiar oscillatory evolution of the neutrinos and antineutrinos as a function of distance is apparent. The difference in period between the neutrinos and antineutrinos is due to the replacement of Y_e by $-Y_e$ in Equation (A8). The errors for this MSW test case are of a similar order of magnitude as the vacuum case.

A.3. Bipolar Oscillations

Our final test case is a study of bipolar neutrino oscillations. This system isolates and tests the numerical convergence in our flavor-mixing subroutines by obviating the advection processes. When using negative energy eigenvalues for the antineutrino density matrices, systems that are represented by two nearly oppositely directed flavor polarization vectors are termed bipolar systems (Duan et al. 2010). As they evolve, the neutrino polarization vector nearly flips in direction, while the antineutrino polarization vector exhibits the same behavior but π radians out of phase. In our treatment, we use positive energy eigenvalues for the antineutrinos in Equation (8), and so we begin with aligned flavor polarization vectors to study the bipolar behavior. As a result, we would expect bipolar oscillations in phase between the neutrinos and antineutrinos. We study a simple case where the system is homogeneous and monochromatic. Only the H_V and H_v terms are included in the Hamiltonian for Equations (24) and (26). The flux moments are all zero. We pick $\delta m^2/p$ such that the period of vacuum oscillations is 2π s, and take a small vacuum mixing angle $\theta = 0.01$. The energy density is picked such that the strength of self-interactions is 10 rad s^{-1} .

Figure 13 shows the results of our bipolar test for neutrinos; the antineutrinos evolve in the same manner. The various color curves label the individual components of the energy-density matrix. Initially the neutrinos are completely in x -flavor eigenstates. They subsequently flip to a state which is $\sim 95\%$ electron flavor, before reverting back to the initial state. This pattern repeats indefinitely as the system is always unstable. Vertical dashed black lines give predictions for the locations in time of the period of oscillations in E_{xx} , centered on the first peak of the E_{ee} curve. We use a closed-form expression

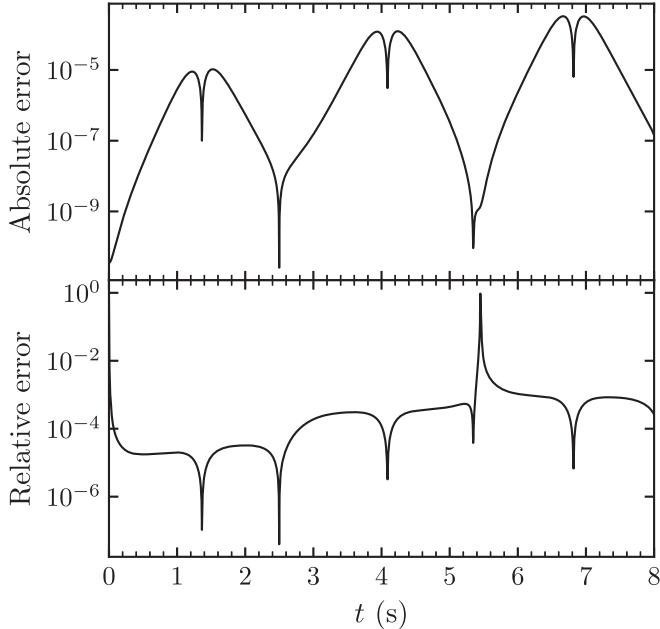


Figure 14. Relative and absolute errors for $E_{ee}(t)$ between the FLASH calculation and a direct numerical resolution of Equation (12) in Hannestad et al. (2006).

involving elliptic integrals to predict the period of oscillation in the unstable bipolar system. Note that the period for the flavor off-diagonal components, $\text{Re}(E_{ex})$ and $\text{Im}(E_{ex})$ have double the period as the diagonal components. There exists a slight decrease in the period of the FLASH results, which becomes more apparent as time continues. At the first full period near $t \sim 4.1$ s, the relative difference in expected versus calculated time is 1.7×10^{-3} . At the second full period, near $t \sim 6.8$ s, the drift is 2.1×10^{-3} . Both of these drift values are within the size of the timestep.

We compare this calculation using the FLASH architecture to a more direct and straightforward resolution of the equation of motion,¹⁰ based on the study of Hannestad et al. (2006). Using the `solve_ivp` function from the Scipy library with high-precision parameters ($\text{rtol} = \text{atol} = 10^{-8}$), we solve Equation (12) from Hannestad et al. (2006) for the tilt angle φ . We can solve for the trace-normalized value of E_{ee} by taking the tilt angle and solving for the Bloch polarization vector [see Equations (6) and (8) of Hannestad et al. (2006)].¹¹ Figure 14 gives the absolute (top panel) and relative (bottom panel) errors between FLASH and the Scipy function `solve_ivp` for the quantity $E_{ee}/\text{Tr}[E]$. The relative error is the same expression as Equation (A4), but we substitute the results from `solve_ivp` for $Q^{(\text{th})}$. The absolute error is the numerator of the rhs of Equation (A4). We see an overall growth in the peak absolute error with increasing cycle number, consistently with the drift

¹⁰ For the interested reader: Xiong et al. (2023d) gives an analytical formulation of the curves in Figure 13.

¹¹ Since we consider a normal hierarchy scenario with an initial system of x (anti)neutrinos instead of electronic ones, there are some minor changes to the equations in Hannestad et al. (2006). Borrowing their notations:

$$(8) \rightarrow Q = \left[4 + \left(\frac{\omega}{\mu} \right)^2 - 4 \frac{\omega}{\mu} \cos 2\theta_0 \right]^{1/2},$$

$$\varphi(0) \simeq \pi + (\omega/\mu Q) 2\theta_0,$$

with here $\mu/\omega = 10$ and $\theta_0 = 0.01$.

in the period of oscillations. Both the absolute and relative errors are the result of numerical convergence in our RKCK solver. We have verified that these errors scale linearly with the tolerance criterion of the RKCK algorithm.

Appendix B Moments in an Orthonormal Tetrad

In order to perform simulations of realistic neutrino distributions in a small domain using codes that assume a Minkowski metric, we need to transform the radiation field quantities output by SpEC-Hydro into an orthonormal tetrad comoving with the background fluid. The radiation field is given in terms of the lab-frame energy density $E^{(v_a)}$, energy flux $F_i^{(v_a)}$, the average neutrino energy $\epsilon^{(v_a)}$, and the fluid transport velocity \tilde{v}^i at each point in space and for each neutrino flavor a . The spacetime metric is given in standard 3 + 1 quantities: the lapse α , shift β^i , and three-metric γ_{ij} . This is a common procedure within general-relativistic truncated moment simulation codes (Shibata et al. 2011; Foucart et al. 2016b; Kuroda et al. 2016), but we make the procedure explicit here for completeness. Note that, in the main text, we revert to the more conventional symbols E and F to indicate tetrad quantities, but they refer to the lab-frame quantities in this appendix for consistency with Shibata et al. (2011).

The radiation stress-energy tensor can be defined in either the lab or comoving frames as (Shibata et al. 2011)

$$\begin{aligned} T^{\mu\nu} &= En^\mu n^\nu + F^\mu n^\nu + n^\mu F^\nu + P^{\mu\nu} [\text{lab}] \\ &= Ju^\mu u^\nu + H^\mu u^\nu + u^\mu H^\nu + S^{\mu\nu} [\text{comoving}]. \end{aligned} \quad (\text{B1})$$

To get the moments in the comoving frame, we evaluate the full stress-energy tensor and project onto suitably chosen tetrad basis vectors. This is made more complicated by the fact that the pressure tensor needs to be expressed as a linear combination of the optically thick and thin limits according to the two-moment closure.

The index on the shift vector can be lowered with the three-metric ($\beta_i = \beta^j \gamma_{ij}$). The vector that is normal to the hypersurface of constant time t can be constructed as $n_\mu = \alpha(-1, 0, 0, 0)$. The transport velocity is defined in terms of the four-velocity as $\tilde{v}^i = u^i/u^t$, where $W = \alpha u^t$ is the generalized Lorentz factor. From these, it is straightforward to construct the fluid four-velocity u^α and the three-velocity $v^i = u^i/W + \beta^i/\alpha$.

The optically thin pressure tensor is straightforwardly defined as $P_{\text{thin}}^{ij} = (E/F^k F_k) F^i F^j$, where contractions are naturally evaluated using the lab-frame three-metric γ_{ij} . Obtaining P_{thick}^{ij} is more involved, but can be determined analytically as follows. It is straightforward to write the optically thick pressure tensor in the lab frame as $S_{\text{thick}}^{ij} = J h^{ij}/3$, where $h^{ij} = g^{ij} + u^i u^j$ is the comoving-frame three metric. By contracting Equation (B1) with various combinations of u^α and n^α , one can work out that the comoving-frame radiation energy density is (i.e., starting with Equation (87) in Foucart et al. (2016b))

$$\begin{aligned} J &= T^{\mu\nu} u_\mu u_\nu \\ &= W^2 [E - 2v^i F_i + JW^2(v_i v^i)^2 + 2WH_i v^i(v_i v^i) + S^{ij} v_i v_j]. \end{aligned} \quad (\text{B2})$$

Plugging in S^{ij} from above, this in turn yields the comoving-frame energy density in the optically thick limit:

$$J_{\text{thick}} = \frac{3}{2W^2 + 1} [(2W^2 - 1)E - 2W^2 v^i F_i]. \quad (\text{B3})$$

Similarly,

$$H_{\text{thick}}^i = -T^{\mu\nu}u_\mu(g_\nu^i + u_\nu u^i) \\ = \frac{F^i}{W} + \frac{v^i W}{2W^2+1}[(4W^2+1)F_i v^i - 4W^2 E]. \quad (\text{B4})$$

Finally, these can be used to calculate the pressure tensor in the optically thick limit as

$$P_{\text{thick}}^{ij} = T^{\mu\nu}(g_\mu^i + u_\mu u^i)(g_\nu^j + u_\nu u^j) \\ = \frac{J_{\text{thick}}}{3}(4W^2 v^i v^j + \gamma^{ij}) + W(H_{\text{thick}}^i v^j + v^i H_{\text{thick}}^j). \quad (\text{B5})$$

By plugging E , F^i , and P^{ij} [from Equation (36)] into Equation (B1), we now have the complete stress-energy tensor $T^{\mu\nu}$.

For a comoving orthonormal tetrad, the time-like basis vector is the four-velocity ($\hat{t}_\mu = u_\mu$). We then choose a spacelike trial vector $(0, 1, 0, 0)$ (and similarly for the y and z trial vectors) and apply Gram-Schmidt orthonormalization to get the spacelike basis vectors $x_\mu^{(i)}$. The moments in the orthonormal tetrad are then

$$J_{\text{tet}} = T^{\mu\nu}\hat{t}_\mu\hat{t}_\nu, \quad (\text{B6})$$

$$H_{\text{tet}}^i = T^{\mu\nu}\hat{t}_\mu\hat{x}_\nu^{(i)}. \quad (\text{B7})$$

From this point forward, we do not have to think about spacetime curvature, as at each location, spacetime is locally flat. Specifically, the moments can now be used more intuitively. For example, the flux factor is simply $f = \sqrt{\sum_i (H_{\text{tet}}^i)^2} / J_{\text{tet}}$.

Appendix C Differences in the Initial Conditions between FLASH and EMU

We must specify the initial conditions for all six of our test and NSM simulations shown in Section 5 across the entire domain. Tables 1 and 2 give the electron and x -flavor number

Table 5
Initial Values of Domain-averaged Off-diagonal Magnitude of N and Saturation Times for All Six FLASH and EMU (2f) FFC Simulations

Name	$\langle N_{ex} / N_{ee} \rangle _{t=0}$	$t_{\text{sat}} (10^{-9} \text{ s})$
Fiducial		
FLASH	7.65×10^{-7}	0.257
EMU (2f)	9.85×10^{-9}	0.365
90Degree		
FLASH	7.65×10^{-7}	0.346
EMU (2f)	9.85×10^{-9}	0.511
TwoThirds		
FLASH	7.65×10^{-7}	0.859
EMU (2f)	8.53×10^{-9}	1.73
NSM 1		
FLASH	7.65×10^{-7}	0.202
EMU (2f)	1.19×10^{-8}	0.348
NSM 2		
FLASH	7.65×10^{-7}	0.380
EMU (2f)	6.19×10^{-9}	1.86
NSM 3		
FLASH	7.65×10^{-7}	0.184
EMU (2f)	1.42×10^{-8}	0.491

densities and flux factors for the three tests and NSM points, respectively. For the off-diagonal coherence terms, we use random numbers per the procedure outlined in Section 4.3. As a result of the randomization process, the complex valued ϱ_{ex} terms add coherently, which acts to reduce the domain-averaged values of the initial perturbation in $\langle |N_{ex}| \rangle$ for the EMU simulations. The first column of Table 5 gives $\langle |N_{ex}| / N_{ee} \rangle$ at the start of all six simulations for FLASH and EMU (2f).

Starting from smaller initial values, the EMU simulations will undergo different time evolution as compared to the FLASH simulations. We give the evolution of $\langle N_{ee} \rangle$ and $\langle N_{ex} \rangle$ for the

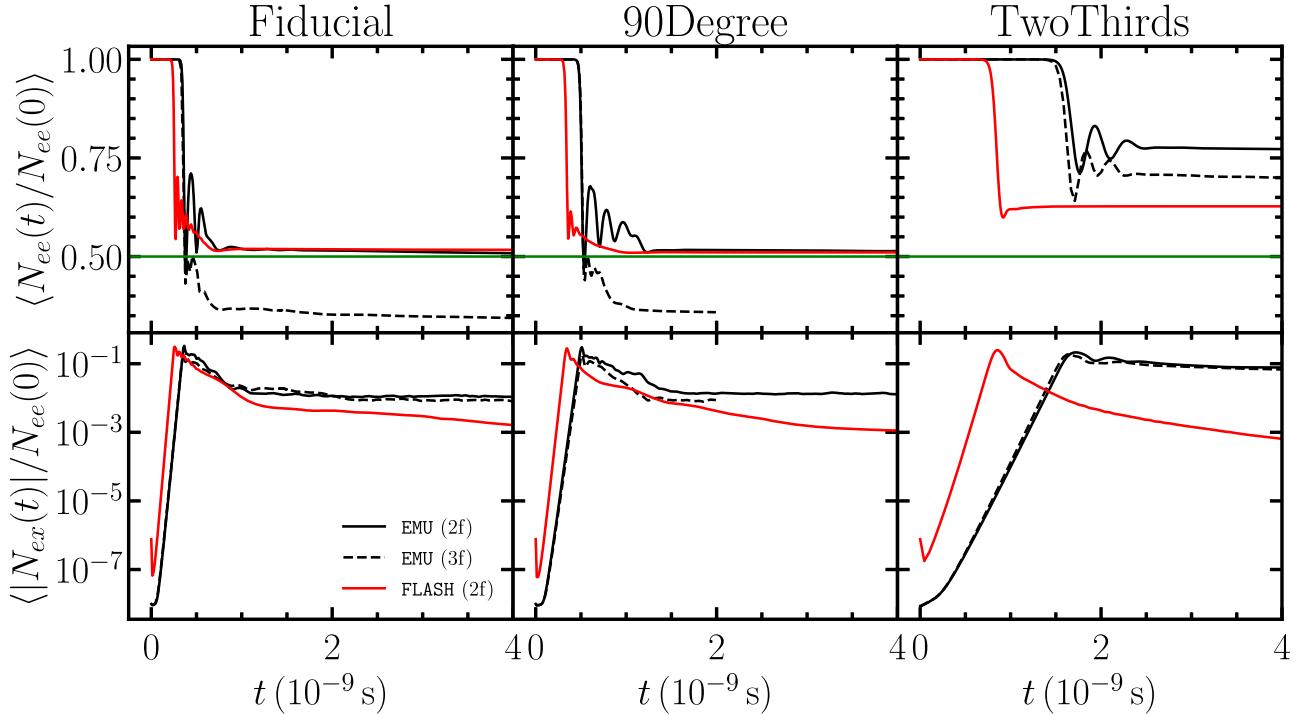


Figure 15. Density matrix elements versus time for the three tests in Section 5.1. Figure content is the same as Figure 2 except all curves begin at simulation time $t = 0$.

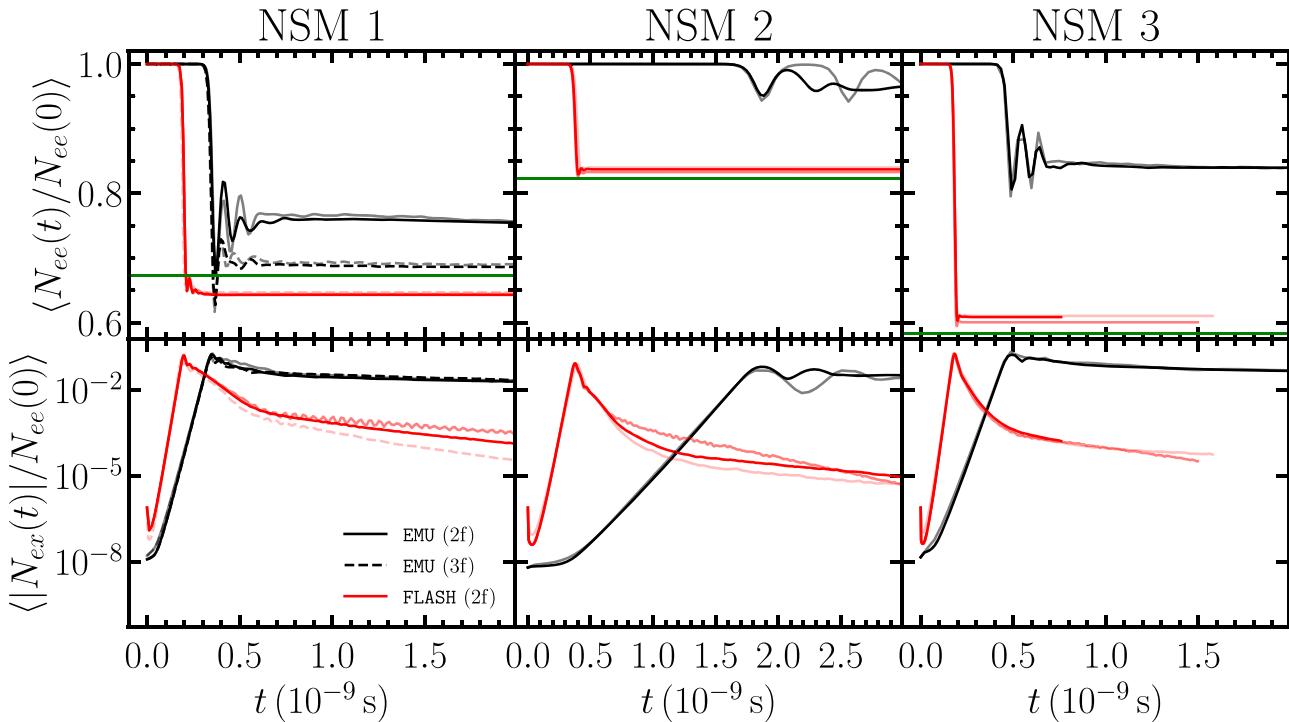


Figure 16. Density matrix elements versus time for the three NSM points in Section 5.2. Figure content is the same as Figure 7 except all curves begin at simulation time $t = 0$.

three tests of Section 5.1 (Figure 15) and the three NSM points of Section 5.2 (Figure 16). The content in Figures 15 and 16 is the same as Figures 2 and 7, respectively, except for the definition of the horizontal axis, where we use simulation time for the figures in this appendix. The smaller initial perturbations of N_{ex} coupled with the smaller growth rates for EMU lead to later saturation points. The second column of Table 5 gives the saturation times for all six simulations.

ORCID iDs

Evan Grohs <https://orcid.org/0000-0002-9409-3468>
 Sherwood Richers <https://orcid.org/0000-0001-5031-6829>
 Sean M. Couch <https://orcid.org/0000-0002-5080-5996>
 Francois Foucart <https://orcid.org/0000-0003-4617-4738>
 Julien Froustey <https://orcid.org/0000-0002-6466-8232>
 James P. Kneller <https://orcid.org/0000-0002-3502-3830>
 Gail C. McLaughlin <https://orcid.org/0000-0001-6811-6657>

References

Abbar, S. 2023, *PhRvD*, **107**, 103006
 Abbar, S., & Duan, H. 2018, *PhRvD*, **98**, 043014
 Abbar, S., Duan, H., Sumiyoshi, K., Takiwaki, T., & Volpe, M. C. 2019, *PhRvD*, **100**, 043004
 Ardevol-Pulillo, R., Janka, H. T., Just, O., & Bauswein, A. 2019, *MNRAS*, **485**, 4754
 Balantekin, A. B., Cervia, M. J., Patwardhan, A. V., Rrapaj, E., & Siwach, P. 2023, *EPJA*, **59**, 186
 Balantekin, A. B., & Pehlivan, Y. 2007, *J. Phys. G*, **34**, 47
 Baxter, A. L., et al. 2022, *ApJ*, **925**, 107
 Bellm, E. C., Kulkarni, S. R., Graham, M. J., et al. 2018, *PASP*, **131**, 018002
 Bethe, H. A. 1986, *PhRvL*, **56**, 1305
 Bhattacharyya, S., & Dasgupta, B. 2020, *PhRvD*, **102**, 063018
 Bhattacharyya, S., & Dasgupta, B. 2022, *PhRvD*, **106**, 103039
 Blaschke, D. N., & Cirigliano, V. 2016, *PhRvD*, **94**, 033009
 Bollig, R., Yadav, N., Kresse, D., et al. 2021, *ApJ*, **915**, 28
 Bruenn, S. W., Blondin, J. M., Hix, W. R., et al. 2020, *ApJS*, **248**, 11
 Burrows, A., Radice, D., & Vartanyan, D. 2019, *MNRAS*, **485**, 3153
 Cernohorsky, J., & Bludman, S. A. 1994, *ApJ*, **433**, 250
 Cervia, M. J., Siwach, P., Patwardhan, A. V., et al. 2022, *PhRvD*, **105**, 123025
 Cirigliano, V., Fuller, G. M., & Vlasenko, A. 2015, *PhLB*, **747**, 27
 Cowperthwaite, P. S., Berger, E., Villar, V. A., et al. 2017, *ApJL*, **848**, L17
 Curtis, S., Ebinger, K., Fröhlich, C., et al. 2019, *ApJ*, **870**, 2
 Curtis, S., Miller, J. M., Fröhlich, C., et al. 2023, *ApJL*, **945**, L13
 Cusinato, M., Guercilena, F. M., Perego, A., et al. 2022, *EPJA*, **58**, 99
 Dasgupta, B. 2022, *PhRvL*, **128**, 081102
 Dasgupta, B., Dighe, A., Raffelt, G. G., & Smirnov, A. Y. 2009, *PhRvL*, **103**, 051105
 Dasgupta, B., Mirizzi, A., & Sen, M. 2017, *JCAP*, **02**, 019
 Dasgupta, B., Mirizzi, A., & Sen, M. 2018, *PhRvD*, **98**, 103001
 Duan, H., Friedland, A., McLaughlin, G., & Surman, R. 2011, *J. Phys. G*, **38**, 035201
 Duan, H., Fuller, G. M., & Qian, Y.-Z. 2006, *PhRvD*, **74**, 123004
 Duan, H., Fuller, G. M., & Qian, Y.-Z. 2010, *ARNPS*, **60**, 569
 Dubey, A., Antypas, K., Ganapathy, M. K., et al. 2009, *ParC*, **35**, 512
 Ehring, J., Abbar, S., Janka, H.-T., Raffelt, G., & Tamborra, I. 2023a, *PhRvD*, **107**, 103034
 Ehring, J., Abbar, S., Janka, H.-T., Raffelt, G., & Tamborra, I. 2023b, *PhRvL*, **131**, 061401
 Fernández, R., Richers, S., Mulyk, N., & Fahlman, S. 2022, *PhRvD*, **106**, 103003
 Foucart, F., Duez, M. D., Haas, R., et al. 2023, *PhRvD*, **107**, 103055
 Foucart, F., Duez, M. D., Hebert, F., et al. 2021, *ApJ*, **920**, 82
 Foucart, F., O'Connor, E., Roberts, L., et al. 2016a, *PhRvD*, **D94**, 123016
 Foucart, F., Haas, R., Duez, M. D., et al. 2016b, *PhRvD*, **93**, 044019
 Freiburghaus, C., Rosswog, S., & Thielemann, F.-K. 1999, *ApJ*, **525**, L121
 Frensel, M., Wu, M.-R., Volpe, C., & Perego, A. 2017, *PhRvD*, **95**, 023011
 Froustey, J., Pitrou, C., & Volpe, M. C. 2020, *JCAP*, **12**, 015
 Froustey, J., Richers, S., Grohs, E., et al. 2023, arXiv:2311.11968
 Fryxell, B., Olson, K., Ricker, P., et al. 2000, *ApJS*, **131**, 273
 Fujibayashi, S., Kiuchi, K., Wanajo, S., et al. 2023, *ApJ*, **942**, 39
 Fujimoto, S.-i., & Nagakura, H. 2023, *MNRAS*, **519**, 2623
 George, M., Lin, C.-Y., Wu, M.-R., Liu, T. G., & Xiong, Z. 2023, *CoPhC*, **283**, 108588
 George, M., Wu, M.-R., Tamborra, I., Ardevol-Pulillo, R., & Janka, H.-T. 2020, *PhRvD*, **102**, 103015
 GIZZI, D., O'Connor, E., Rosswog, S., et al. 2019, *MNRAS*, **490**, 4211

Grohs, E., Fuller, G. M., Kishimoto, C. T., Paris, M. W., & Vlasenko, A. 2016, *PhRvD*, **93**, 083522

Grohs, E., Richers, S., Couch, S. M., et al. 2023, *PhRvB*, **846**, 138210

Hannestad, S., Raffelt, G. G., Sigl, G., & Wong, Y. Y. Y. 2006, *PhRvD*, **74**, 105010

Harten, A., Lax, P. D., & Leer, B. 1983, *SIAMR*, **25**, 35

Haxton, W. C. 1986, *PhRvL*, **57**, 1271

Hayashi, K., Fujibayashi, S., Kiuchi, K., et al. 2022, *PhRvD*, **106**, 023008

Holmbeck, E. M., Hansen, T. T., Beers, T. C., et al. 2020, *ApJS*, **249**, 30

Hunter, J. D. 2007, *CSE*, **9**, 90

Iwakami, W., Okawa, H., Nagakura, H., et al. 2020, *ApJ*, **903**, 82

Izaguirre, I., Raffelt, G., & Tamborra, I. 2017, *PhRvL*, **118**, 021101

Johns, L. 2023, *PhRvL*, **130**, 191001

Johns, L., & Nagakura, H. 2021, *PhRvD*, **103**, 123012

Johns, L., Nagakura, H., Fuller, G. M., & Burrows, A. 2020a, *PhRvD*, **101**, 043009

Johns, L., Nagakura, H., Fuller, G. M., & Burrows, A. 2020b, *PhRvD*, **102**, 103017

Johns, L., & Xiong, Z. 2022, *PhRvD*, **106**, 103029

Jones, E., Oliphant, T., & Peterson, P., 2001 SciPy: Open Source Scientific Tools for Python, <http://www.scipy.org/>

Just, O., Abbar, S., Wu, M.-R., et al. 2022, *PhRvD*, **105**, 083024

Just, O., Obergaulinger, M., & Janka, H.-T. 2015, *MNRAS*, **453**, 3386

Kalogera, V., Sathyaprakash, B. S., Bailes, M., et al. 2021, arXiv:2111.06990

Kuroda, T., Takiwaki, T., & Kotake, K. 2016, *ApJS*, **222**, 20

Lacroix, D., Balantekin, A. B., Cervia, M. J., Patwardhan, A. V., & Siwach, P. 2022, *PhRvD*, **106**, 123006

Li, X., & Siegel, D. M. 2021, *PhRvL*, **126**, 251101

Lien, A., & Fields, B. D. 2009, *JCAP*, **2009**, 047

Lippuner, J., & Roberts, L. F. 2015, *ApJ*, **815**, 82

LSST Science Collaboration, Abell, P. A., Allison, J., et al. 2009, arXiv:0912.0201

Malkus, A., Friedland, A., & McLaughlin, G. C. 2014, arXiv:1403.5797

Malkus, A., Kneller, J. P., McLaughlin, G. C., & Surman, R. 2012, *PhRvD*, **86**, 085015

Martin, J. D., Neill, D., Roggero, A., Duan, H., & Carlson, J. 2023a, *PhRvD*, **108**, 123010

Martin, J. D., Roggero, A., Duan, H., & Carlson, J. 2023b, arXiv:2301.07049

Martin, J. D., Yi, C., & Duan, H. 2020, *PhLB*, **800**, 135088

McLaughlin, G. C., Fuller, G. M., & Wilson, J. R. 1996, *ApJ*, **472**, 440

Meurer, A., Smith, C. P., Paprocki, M., et al. 2017, *PeerJ Comp. Sci.*, **3**, e103

Miller, J. M., Ryan, B. R., & Dolence, J. C. 2019a, *ApJS*, **241**, 30

Miller, J. M., Ryan, B. R., Dolence, J. C., et al. 2019b, *PhRvD*, **100**, 023008

Miller, J. M., Sprouse, T. M., Fryer, C. L., et al. 2020, *ApJ*, **902**, 66

Minerbo, G. N. 1978, *JQSRT*, **20**, 541

Morinaga, T. 2022, *PhRvD*, **105**, L101301

Mukhopadhyay, P. 2022, PhD thesis, Stanford Univ.

Myers, M., Cooper, T., Warren, M., et al. 2022, *PhRvD*, **105**, 123036

Nagakura, H. 2022, *PhRvD*, **106**, 063011

Nagakura, H. 2023, *PhRvD*, **108**, 103014

Nagakura, H. 2023, *PhRvL*, **130**, 211401

Nagakura, H., Iwakami, W., Furusawa, S., et al. 2018, *ApJ*, **854**, 136

Nagakura, H., & Johns, L. 2021, *PhRvD*, **104**, 063014

Nagakura, H., Johns, L., Burrows, A., & Fuller, G. M. 2021, *PhRvD*, **104**, 083025

Nagakura, H., Sumiyoshi, K., & Yamada, S. 2014, *ApJS*, **214**, 16

Nagakura, H., & Zaizen, M. 2022, *PhRvL*, **129**, 261101

Nagakura, H., & Zaizen, M. 2023, *PhRvD*, **108**, 123003

Neilsen, D., Liebling, S. L., Anderson, M., et al. 2014, *PhRvD*, **89**, 104029

O'Connor, E. P., & Couch, S. M. 2018, *ApJ*, **854**, 63

Padilla-Gay, I., Shalgar, S., & Tamborra, I. 2021, *JCAP*, **01**, 017

Patwardhan, A. V., Cervia, M. J., & Balantekin, A. B. 2021, *PhRvD*, **104**, 123035

Pehlivan, Y., Balantekin, A. B., Kajino, T., & Yoshida, T. 2011, *PhRvD*, **84**, 065008

Perego, A., Cabezón, R., & Käppeli, R. 2016, *ApJS*, **223**, 22

Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 1992, *Numerical Recipes in FORTRAN: The Art of Scientific Computing* (2nd edn; New York: Cambridge Univ. Press)

Radice, D., Bernuzzi, S., Perego, A., & Haas, R. 2022, *MNRAS*, **512**, 1499

Reichert, M., Obergaulinger, M., Eichler, M., Aloy, M. Á., & Arcones, A. 2021, *MNRAS*, **501**, 5733

Richers, S. 2020, *PhRvD*, **102**, 083017

Richers, S. 2022, *PhRvD*, **106**, 083005

Richers, S., Willcox, D., & Ford, N. 2021a, *PhRvD*, **104**, 103023

Richers, S., Willcox, D. E., Ford, N. M., & Myers, A. 2021b, *PhRvD*, **103**, 083013

Richers, S. A., McLaughlin, G. C., Kneller, J. P., & Vlasenko, A. 2019, *PhRvD*, **99**, 123014

Roggero, A., Rrapaj, E., & Xiong, Z. 2022, *PhRvD*, **106**, 043022

Rosswog, S., & Liebendörfer, M. 2003, *MNRAS*, **342**, 673

Rrapaj, E. 2020, *PhRvC*, **101**, 065805

Ruffert, M., & Janka, H.-T. 1999, *A&A*, **344**, 573

Sawyer, R. F. 2005, *PhRvD*, **72**, 045003

Shibata, M., Kiuchi, K., Sekiguchi, Y.-I., & Suwa, Y. 2011, *PTPh*, **125**, 1255

Sigl, G., & Raffelt, G. 1993, *NuPhB*, **406**, 423

Siwach, P., Suliga, A. M., & Balantekin, A. B. 2023, *PhRvD*, **107**, 023019

Skinner, M. A., Dolence, J. C., Burrows, A., Radice, D., & Vartanyan, D. 2019, *ApJS*, **241**, 7

Stapleford, C. J., Fröhlich, C., & Kneller, J. P. 2020, *PhRvD*, **102**, 081301

Surman, R., McLaughlin, G. C., & Hix, W. R. 2006, *ApJ*, **643**, 1057

Tartaglia, L., Sand, D., Valenti, S., et al. 2018, *ApJ*, **853**, 62

Tian, J. Y., Patwardhan, A. V., & Fuller, G. M. 2017a, *PhRvD*, **95**, 063004

Tian, J. Y., Patwardhan, A. V., & Fuller, G. M. 2017b, *PhRvD*, **96**, 043001

van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, *CSE*, **13**, 22

Vlasenko, A., Fuller, G. M., & Cirigliano, V. 2014, *PhRvD*, **89**, 105004

Vlasenko, A., & McLaughlin, G. C. 2018, *PhRvD*, **97**, 083011

Volpe, C., Vääränen, D., & Espinoza, C. 2013, *PhRvD*, **87**, 113010

Volpe, M. C. 2023, arXiv:2301.11814

Wanajo, S., Sekiguchi, Y., Nishimura, N., et al. 2014, *ApJL*, **789**, L39

Willcox, D. E., & Richers, S. 2021, AMReX-Astro/Emu v1.3, Zenodo, doi: 10.5281/zenodo.5514018

Wu, M.-R., Duan, H., & Qian, Y.-Z. 2016, *PhLB*, **752**, 89

Wu, M.-R., & Tamborra, I. 2017, *PhRvD*, **95**, 103007

Xiong, Z., Johns, L., Wu, M.-R., & Duan, H. 2023a, *PhRvD*, **108**, 083002

Xiong, Z., Wu, M.-R., Abbar, S., et al. 2023b, *PhRvD*, **108**, 063003

Xiong, Z., Wu, M.-R., Martínez-Pinedo, G., Fischer, T., et al. 2023c, *PhRvD*, **107**, 083016

Xiong, Z., Wu, M.-R., & Qian, Y.-Z. 2023d, *PhRvD*, **108**, 043007

Zaizen, M., & Morinaga, T. 2021, *PhRvD*, **104**, 083035

Zaizen, M., & Nagakura, H. 2023, *PhRvD*, **107**, 103022

Zhang, Y., & Burrows, A. 2013, *PhRvD*, **88**, 105009