Particle-in-cell Simulation of the Neutrino Fast Flavor Instability

Sherwood Richers*
Department of Physics, University of California Berkeley, CA 94720 and
Nuclear Science Division, Lawrence Berkeley National Lab, CA 94720

Don E. Willcox, Nicole M. Ford, and Andrew Myers
Computational Research Division, Lawrence Berkeley National Lab, CA 94720

Neutrinos drive core-collapse supernovae, launch outflows from neutron star merger accretion disks, and set the ratio of protons to neutrons in ejecta from both systems that generate heavy elements in the universe. Neutrinos of different flavors interact with matter differently, and much recent work has suggested that fast flavor instabilities are likely ubiquitous in both systems, but the final flavor content after the instability saturates has not been well understood. In this work we present particle-in-cell calculations which follow the evolution of all flavors of neutrinos and antineutrinos through saturation and kinematic decoherence. We conduct one-dimensional three-flavor simulations of neutrino quantum kinetics to demonstrate the outcome of this instability in a few example cases. We demonstrate the growth of both axially symmetric and asymmetric modes whose wavelength and growth rate match predictions from linear stability analysis. Finally, we vary the number density, flux magnitude, and flux direction of the neutrinos and antineutrinos and demonstrate that these factors modify both the growth rate and post-saturation neutrino flavor abundances. Weak electron lepton number (ELN) crossings in these simulations produce both slow growth of the instability and little difference between the flavor abundances in the initial and final states. In all of these calculations the same number of neutrinos and antineutrinos change flavor, making the least abundant between them the limiting factor for post-saturation flavor change. Many more simulations and multi-dimensional simulations are needed to fully probe the parameter space of the initial conditions.

I. INTRODUCTION

Supernovae represent the explosive birth and neutron star mergers the cataclysmic destruction of neutron stars (see [1–7] for recent reviews). Neutrinos are the dominant couriers of energy and lepton number in both cases and will be observable from the next galactic event. They carry the energy that explodes the star in a core-collapse supernova, drive outflows, and modify the ratio of neutrons to protons available for nucleosynthesis, thereby determining the abundances of elements that enrich the nearby universe. These heavy elements mix with ambient hydrogen that forms later generations of stars and planets.

Although state of the art models in many cases yield qualitatively correct explosion energies, ejecta masses, neutrino signals, electromagnetic transients, and gravitational waves (e.g., [8–12]), there remain many holes in the details. Simulations are growing increasingly sophisticated, but still suffer from low resolution (e.g., [13, 14]) and uncertain initial conditions (e.g., [15–17]). Simulations also require as input an equation of state for matter beyond nuclear densities (e.g., [15, 18–20]) and the nuclear reaction rates of heavy and very unstable elements (e.g., [21]), both of which are poorly constrained. The final dominant model uncertainty is the treatment of neutrinos. The weak interactions between neutrinos and matter and other neutrinos cause them to be largely out of equilibrium, necessitating an expensive kinetic treatment. There is a natural trade-off between inexpensive approximate methods that allow for a larger number of simulations or a focus on specific aspects of the problem and high-accuracy methods that result in different explosion dynamics, ejecta properties, and neutrino signals [9, 12, 22–27].

In addition, no global simulation dynamically treats the full effects of neutrino flavor transformation. Individual neutrinos exist in a quantum superposition of flavor states, so neutrino masses and potentials due to interactions with matter and other neutrinos can drive rapid and large transformations between quantum states. The neutrino-neutrino interaction term makes the evolution equations nonlinear, resulting in an exponentially difficult many body problem [28]. Because of this complexity the majority of the work in the field has been done under the mean field approximation (e.g., [29, 30]), though the history and phenomenology is quite rich even with this approximation.

Vacuum and matter-induced oscillations were applied to core-collapse supernovae shortly after they were proposed as a solution to the solar neutrino problem [31–33] and are still used to map supernova simulation output spectra to estimate signals detectable at earth (e.g., [34]). So-called collective oscillations occur when the potential neutrinos feel from interactions with other neutrinos is comparable to the neutrino mass energy scale and has a long history in the core-collapse supernova context (see [35, 36] for recent reviews). State of the art bulb-model simulations indicate that this transformation

* srichers@berkeley.edu

mode is likely unimportant for the explosion mechanism (e.g., [37]), but is still impactful for ejecta nucleosynthesis (e.g., [38–40]). The matter-neutrino resonance can occur under certain conditions when the neutrino potential is comparable to the matter potential and may impact neutrino signals and nucleosynthesis from both supernovae [41] and neutron star mergers [42]. In addition, a small fraction of neutrinos emitted from an ongoing core-collapse supernova explosion will scatter off of heavy nuclei in the collapsing stellar envelope far outside of the shock front, forming a diffuse halo of scattered neutrinos, some of which are moving inward. These diffuse neutrinos can drive significant flavor changes in the much larger number of outgoing neutrinos [43–45].

Although proposed more than a decade ago [46], it was only appreciated in recent years that a new flavor instability could drive neutrino flavor transformation orders of magnitude more quickly than the previous transformation mechanisms [47–53]. While the magnitude of the effect of this fast flavor instability (FFI) on the abundances of each neutrino flavor is still uncertain, straightforward arguments suggest that the FFI should be a ubiquitous feature of both CCSNe and NSMs [45, 54]. Although the associated small length and timescales make simulations of the full quantum kinetic equations (QKE) presently impossible in global multidimensional simulations, neutrino transport without flavor transformation is much more tractable. Many authors have used the neutrino distributions in these multidimensional neutrino transport simulations to diagnose the presence or lack of the FFI in CCSNe [40, 45, 55–63] and NSMs [38, 54, 64]. Even if flavor transformation does not affect the dynamics, it has been suggested that it can significantly modify the elements formed in the ejecta, hampering the production of heavy elements in NSMs [38, 64] and enhancing the production of light-p nuclei in neutrino-driven winds from CCSNe [40].

Motivated by the potential impact of the FFI and the hope that a general understanding of the final state of an unstable distribution can eventually be applied to global simulations, several authors have begun working on direct simulations of the instability. The majority of the simulations so far assume homogeneity ("one-zone" models) and are discretized in only angle (or angular moments) and time. These simulations have demonstrated a consistency between linear stability analysis and direct evolution of the nonlinear equations [65–68] and have shed some insight into the late-time angular turbulence and kinematic decoherence [69–71]. Other simulations of the FFI have included inhomogeneity [72, 73], a simplified treatment of collisional processes [74], or both [50].

The capabilities of existing simulation methods are currently limited by various imposed symmetries, a reduced number of neutrino species, and/or a requirement for a very large number of grid cells or basis elements. To unify and expand on these models, a more general framework for simulating kinetics is needed. Neutrino transport methods such as moment methods [75–82], discrete ordinates [83, 84], and Monte Carlo [85–88] could in principle all be extended to treat coherent flavor effects [89, 90], but the strong dependence of the evolution of each neutrino on integrals of nearby neutrino distributions is challenging to implement efficiently.

In this paper, we utilize technology from the plasma physics community and describe a particle-in-cell (PIC) method that formally solves the mean field quantum kinetic equations in an efficient, scalable manner. In Section II we outline a particle-in-cell implementation of the neutrino quantum kinetic equations. In Section III we demonstrate the exponential growth, saturation, and kinematic decoherence of a toy neutrino distribution on a one-dimensional mesh. Finally, in Section IV we vary the neutrino distributions to begin the parameter study needed to build a sub-grid model of the FFI. We provide some concluding remarks in Section V.

We have developed the new code `Emu` [91] to implement this PIC method for solving the QKEs. `Emu` is fully open-source and is available at https://github.com/AMReX-Astro/Emu. All parameter files and select data from this study are publicly available [92] and further data is available upon request.

II. EMU: PIC NEUTRINO FLAVOR KINETICS

`Emu` solves the mean-field quantum kinetic equations without collisions (Section II A) by evolving the position and quantum state of a collection of computational particles moving through a background grid (Section II B). During each simulation timestep the particles aggregate their quantum states to construct a distribution within each grid cell (Section II B 1). The neutrino and background matter distribution are next interpolated to each particle's position in order to construct time derivatives of the position and quantum state (Section II B 2). All particles are then integrated forward in time using a high-order integrator and a performance portable domain decomposed parallelization scheme using the `AMReX` framework (Section II C). This is repeated until the simulation is evolved for the desired amount of time. We discuss each of these steps in more detail below.

A. Quantum Kinetic Equations

The quantum kinetic equations that describe the transport of relativistic quantum particles read [29, 30]

$$\frac{\partial f_{ab}}{\partial t} + c\mathbf{\Omega}\cdot\nabla f_{ab} = \mathcal{C}_{ab} - \frac{i}{\hbar}\left[\mathcal{H}, f\right]_{ab} \ . \tag{1}$$

$f_{ab}(\mathbf{x}, t, \mathbf{p})$ is a $N_F \times N_F$ Hermitian matrix, where N_F is the number of neutrino flavors. Throughout this paper we use the convention that a, b, c, d are flavor indices, and are $\in \{e, \mu, \tau\}$ for quantities in the flavor basis or $\in \{1, 2, 3\}$ for quantities in the mass basis. The diagonals of f_{ab} represent the occupation probability for neutrinos

of flavor a located at position $\mathbf{x}$ and time t, moving with momentum $|\mathbf{p}|$ in direction $\mathbf{\Omega} = \mathbf{p}/|\mathbf{p}|$ at approximately the speed of light c. The neutrino energy is determined by $\epsilon^2 = \mathbf{p}^2c^2 + m^2c^4 \approx \mathbf{p}^2c^2$ with energy $h\nu$. One can write a similar equation for the antineutrino distribution $\bar{f}_{ab}$ that involves a distinct collision integral $\bar{\mathcal{C}}_{ab}$ and Hamiltonian $\bar{\mathcal{H}}_{ab}$. Throughout this work we neglect the collision terms $\mathcal{C}_{ab}$ and $\bar{\mathcal{C}}_{ab}$, which are weak enough to not significantly affect the distributions on the timescales simulated here. Implicit in this form of the equations is the assumption that there is no spin coherence, there are no right-handed neutrinos or left-handed antineutrinos, spacetime is flat, and the neutrino momentum is a constant of motion.

The Hamiltonian term $\mathcal{H}_{ab}$ is also a $N_F \times N_F$ Hermitian matrix that encodes potential energy in the form of mass and interactions with other particles. It is usually broken down into a sum of the vacuum potential (due to the neutrino mass), the matter potential (due to interactions with non-neutrino particles), and the self-interaction potential (due to interactions with other neutrinos). In the convention we use these can be written as

$$\begin{aligned}
\mathcal{H}_{\text{vaccum},ab} &= U_{ac} H^{(m)}_{\text{vaccum},cd} U^{\dagger}_{db} \\
\mathcal{H}_{\text{matter},ab} &= \sqrt{2} G_F (\hbar c)^3 \left[(n_a - \bar{n}^*_a) - \mathbf{\Omega} \cdot (\mathbf{f}_a - \bar{\mathbf{f}}^*_a) \right] \delta_{ab} \\
\mathcal{H}_{\text{neutrino},ab} &= \sqrt{2} G_F (\hbar c)^3 \left[(n_{ab} - \bar{n}^*_{ab}) - \mathbf{\Omega} \cdot (\mathbf{f}_{ab} - \bar{\mathbf{f}}^*_{ab}) \right]
\end{aligned} \tag{2}$$

where $H^{(m)}_{\text{vacuum},ab} = \sqrt{\mathbf{p}^2c^2 + m_a^2c^4}\delta_{ab} \approx |\mathbf{p}|c + m_a^2c^4\delta_{ab}/2|\mathbf{p}|$. The first term ($|\mathbf{p}|c$) is the same for all flavors and therefore cancels out in the commutator in Equation 1, so it can be ignored here. The magnitude of the momentum vector is $|\mathbf{p}| \approx h\nu$, so the vacuum Hamiltonian is usually written as $\mathcal{H}_{\text{vacuum},ab} \approx m_a^2c^4/2h\nu$. The number density and flux are angular integrals over the neutrino distribution:

$$\begin{aligned}
n_{ab} &= \frac{1}{c^3} \int d\mathbf{\Omega} \int d\left(\frac{\nu^3}{3}\right) f_{ab} \ , \\
\mathbf{f}_{ab} &= \frac{1}{c^3} \int d\mathbf{\Omega} \int d\left(\frac{\nu^3}{3}\right) f_{ab} \mathbf{\Omega} \ .
\end{aligned} \tag{3}$$

Note that we denote a scalar lepton number density with n_a and a Hermitian matrix neutrino number density with n_{ab}. $\mathbf{f}_{ab}$ is the (vector of Hermitian matrices) neutrino number flux density, not to be confused with the distribution function f_{ab} or the lepton flux $\mathbf{f}_a$. $\bar{n}_a$ and $\bar{\mathbf{f}}_a$ are real, but we leave the complex conjugation on them in Equation 2 for comparison to the neutrino potential. Our simulations are performed in the fluid comoving frame, making $\mathbf{f}_a = \bar{\mathbf{f}}_a = 0$. For antineutrinos, $\bar{\mathcal{H}}_{\text{vacuum},ab} = \mathcal{H}^*_{\text{vacuum},ab}$, $\bar{\mathcal{H}}_{\text{matter},ab} = -\mathcal{H}^*_{\text{matter},ab}$, and $\bar{\mathcal{H}}_{\text{neutrino},ab} = -\mathcal{H}^*_{\text{neutrino},ab}$. Note also that there are other conventions in the literature that affect the form of the antineutrino Hamiltonians depending on whether one writes the evolution equation for $\bar{f}_{ab}$ or $\bar{f}^*_{ab}$.

B. PIC Algorithm

The left hand side of Equation 1 is a comoving derivative moving along the neutrino trajectory. Therefore rather than evaluating the derivative as fluxes between adjacent grid cells as is done in a finite volume scheme, we can simulate the distribution as a set of particles such that the advection term is accounted for by moving the particle's location. A computational particle carries two scalars N and $\bar{N}$ that represent the number of physical neutrinos and antineutrinos contained in the computational particle, along with two $N_F \times N_F$ Hermitian unit-trace quantum density matrices ρ and $\bar{\rho}$ that represent the quantum state of each neutrino and antineutrino. In practice, we only store and evolve the real and imaginary components of each element in the upper right half and the diagonal of these Hermitian matrices in the flavor basis. Wherever the lower left corners of Hermitian matrices are required, we express them in terms of the components of the upper right half in our implementation.

For a given particle labeled by index p, the equations of motion can then be expressed as

$$\begin{aligned}
\frac{\partial \rho_{ab,p}}{\partial t} &= -\frac{i}{\hbar} \left[\mathcal{H}_p, \rho_p \right]_{ab} \\
\frac{\partial \bar{\rho}_{ab,p}}{\partial t} &= -\frac{i}{\hbar} \left[\bar{\mathcal{H}}_p, \bar{\rho}_p \right]_{ab} \\
\frac{\partial \mathbf{x}_p}{\partial t} &= c\mathbf{\Omega}_p \\
\frac{\partial \nu_p}{\partial t} &= 0 \\
\frac{\partial \mathbf{\Omega}_p}{\partial t} &= 0 \\
\frac{\partial N_p}{\partial t} &= \frac{\partial \bar{N}_p}{\partial t} = 0
\end{aligned} \tag{4}$$

Formally, these equations are only true in the limit of a large neutrino energy compared to the neutrino mass and potentials, but the same is true of Equation 1. In both forms, more terms would need to be included to faithfully treat low-energy neutrinos (e.g. [29]).

Because the neutrino self-interaction couples the flavor evolution of neutrinos and antineutrinos to their local distributions, we integrate Equation 4 for all particles together in step. To compute the right hand side of Equation 4 for each particle, we need to evaluate the Hamiltonian terms in Equation 2. The matter and self-interaction potentials respectively depend on the local lepton and neutrino density and flux. Because evaluating the local neutrino density and flux requires determining the local neutrino distributions represented by the particles, we evaluate the Hamiltonian in two steps: deposition and interpolation.

1. *Evaluating the Hamiltonian: Depositing Particle Information to the Grid*

The background neutrino density stored in each grid cell must reflect the particle distributions themselves. A straightforward way would be to simply add up all of the particles within a cell, which would assume that each particle is shaped as a delta function in position and only interacts with particles in the same grid cell. However, we found that this causes the particle evolution to be unstable to artificially growing perturbations with wavelengths on the length scale of the grid resolution. We eliminated this numerical instability by instead modelling each particle with an extended shape as is usually done in PIC methods in order to smooth particle-mesh deposition and interpolation operations with higher order stencils.

In plasma PIC simulations, computational particles can be designed to represent a distribution of charge centered at the particle position but with an extent comparable to the size of a grid cell that defines the particle's interpolating "shape" function. This yields higher order spatial convergence and suppression of spurious numerical instabilities. We apply the same method in `Emu` for particle flavor matrices instead of particle charge. The particle shape routines in `Emu` are copied directly from the open-source plasma physics code `WarpX`[1] [93] and include options from a delta function (zeroth order interpolant) to a cubic distribution (third order interpolant). We generally use the quadratic shape, since the cubic shape requires more ghost zones (see Section II C) and higher-order interpolators do not necessarily correspond to higher-order convergence in PIC codes [94, 95]. We integrate the total neutrino density and neutrino number flux at grid cell j as

$$\begin{aligned} n_{ab,j} &= \frac{1}{\Delta V} \sum_{p=1}^{N_{\text{particles}}} N_p w_{pj} \rho_{ab,p} \\ \mathbf{f}_{ab,j} &= \frac{1}{\Delta V} \sum_{p=1}^{N_{\text{particles}}} N_p w_{pj} \rho_{ab,p} \mathbf{\Omega}_p \end{aligned} \tag{5}$$

where w_{pj} is the fraction of particle p that contributes to grid cell j according to the particle's shape function. $N_{\text{particles}}$ is the total number of computational particles in the simulation, N_p is the number of physical neutrinos that particle p represents, and ΔV is the volume of the grid cell.

2. *Evaluating the Hamiltonian: Interpolating from the Grid to Calculate Neutrino Potentials*

Now that the lepton and neutrino number density and number flux are known at the center of each grid cell, we interpolate them from the grid to each particle's position $\mathbf{x}_p$ using the particle shape function as in Equation 6:

$$\begin{aligned} n_{a,p} &= \sum_{j=1}^{N_{\text{zones}}} w_{pj} n_{a,j} \ , \\ n_{ab,p} &= \sum_{j=1}^{N_{\text{zones}}} w_{pj} n_{ab,j} \ , \\ \mathbf{f}_{ab,p} &= \sum_{j=1}^{N_{\text{zones}}} w_{pj} \mathbf{f}_{ab,j} \ . \end{aligned} \tag{6}$$

This yields the lepton density ($n_{a,p}$), the neutrino densities ($n_{ab,p}$ and $\bar{n}_{ab,p}$) and fluxes ($\mathbf{f}_{ab,p}$ and $\bar{\mathbf{f}}_{ab,p}$). We then obtain the neutrino and antineutrino Hamiltonians $\mathcal{H}_{ab,p}$ and $\bar{\mathcal{H}}_{ab,p}$ for each particle by applying the particle state and these interpolated quantities to Equation 2. We finally use Equation 4 to calculate the particle's density matrix time derivatives $\partial\rho_{ab,p}/\partial t$ and $\partial\bar{\rho}_{ab,p}/\partial t$.

C. PIC Implementation

We implement this algorithm in modern C++ using the exascale computing framework AMReX [97], which provides high level abstractions for domain decomposition and both distributed and shared-memory parallelization. We thus represent the domain as a union of non-overlapping rectangular blocks, each containing particles and Cartesian grid data. These blocks are distributed across MPI ranks, which allocate memory for their local grid and particle data. Because higher order particle shape functions require particles to deposit and interpolate from neighboring cells, when a particle is near a block's boundary it needs to access data stored on other blocks. Another way to say this is that the particle-mesh operations have a spatial stencil width, where the stencil width refers to the number of grid zones in each dimension a particular operation requires. Piecewise constant shape functions have a stencil width of 1, linear shape functions have a stencil width of 2, quadratic shape functions have a stencil width of 3, etc. (centered on the grid zone containing the particle). Rather than communicate data between blocks every time an individual particle requires it, we instead place N_G layers of ghost zones around the boundary of each block to serve as temporary storage. Ghost zones attached to a given block line up with valid zones within the domain of other blocks, so MPI calls can be used to transfer data to and from all ghost zones and their corresponding valid zones all at once. We choose the number of ghost cells consistent with the stencil width of the particle shape function desired (i.e. $N_G = 1$ for quadratic shape functions with a stencil size of 3, since the particle-mesh operations require data from the particle's current grid cell and one cell to the left and right in each coordinate direction.).

We illustrate how we implement the deposition, communication, and interpolation steps discussed in Sec-

[1] https://github.com/ECP-WarpX/WarpX

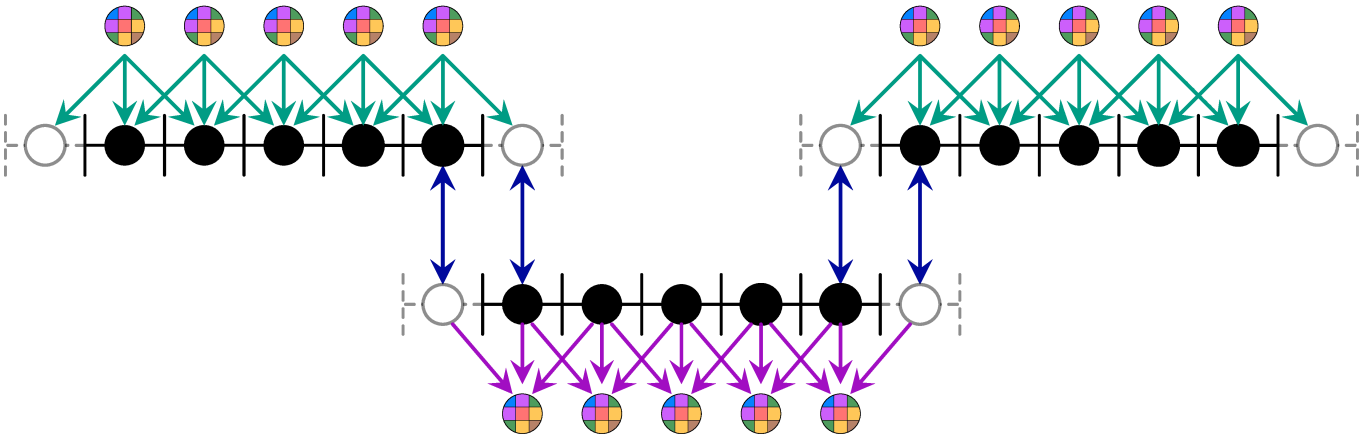

FIG. 1. Illustration of the `Emu` algorithm for evaluating the time derivatives of all particles' density matrices. Particles are represented by the colored circles denoting a 3×3 matrix. Each block of grid cells on the domain is represented by a contiguous array of cells composed of interior cells (black circles) and ghost cells (white circles). In the first step of the right hand side evaluation, we deposit particle states onto the underlying grids using second order particle-in-cell shape functions (green arrows, omitted for the center block to avoid cluttering the illustration). This yields neutrino density and fluxes on the grid. We next sum ghost cells into the valid cells on neighboring blocks and then fill ghost cells from the valid cells (blue arrows). Finally, we interpolate neutrino density and fluxes from the grid to particles (purple arrows, omitted for the left and right blocks to avoid cluttering the illustration), construct particle Hamiltonians, and evaluate neutrino density matrix time derivatives. Figure available under CC BY 4.0 license [96].

tion II B in Figure 1. We carry out the deposition in Equation 5 by first summing (green arrows) particle contributions to neutrino density and flux into local grids, including both interior grid cells (black circles) and neighboring ghost cells (white circles). This happens on all blocks, but we omit green arrows from the center block to avoid cluttering the figure. We next communicate (blue arrows) ghost cell values to the MPI ranks containing the corresponding valid cells and sum the ghost cell contributions into the valid cells. We then copy the grid data from the valid cells to all corresponding ghost cells. We apply periodic boundary conditions to ghost cells at the domain boundaries, mapping them to valid cells on the other side of the domain. Finally, we interpolate (purple arrows) the neutrino density and flux from each grid to particle locations and compute the density matrix right hand sides in Equation 4, as described in Section II B 2. This happens on all blocks, but we omit purple arrows from the left and right blocks to avoid cluttering the figure.

We use a fourth-order explicit Runge-Kutta method to integrate Equation 4 for all particles. We choose a timestep based on two criteria as:

$$\begin{aligned} \Delta t_{\text{advection}} &= C \frac{\min(\Delta x, \Delta y, \Delta z)}{c} \\ \Delta t_{\text{matter},\nu} &= C_F \frac{\hbar}{\sqrt{2} G_F \max(n_{aa}, \bar{n}_{aa}, n_a)} \\ \Delta t &= \min(\Delta t_{\text{advection}}, \Delta t_{\text{matter},\nu}) \ , \end{aligned} \tag{7}$$

where the maximum is taken over grid cells and matrix components. The two constants C and C_F controlling the timestep correspond to the advection and commutator terms in Equation 1, respectively. C and C_F are thus analogous to hydrodynamic CFL factors, and for the simulations in this paper we use $C = C_F = 0.5$. We find that at least a fourth-order Runge-Kutta method is necessary to adequately control the size of the error for the flavor transformation of individual particles. However, `Emu`'s overall order of space-time convergence is formally limited by the second order PIC method as mentioned in Section II B 1. We assess `Emu`'s convergence properties for the simulations of Section III in Appendix B.

We integrate particle positions concurrently with their flavor states, so we do not split the advection and flavor-changing operators in Equation 1. Thus, when constructing the state at which to evaluate a Runge-Kutta stage right hand side, we first calculate particle positions along with their density matrices at that time. As a result, a particle may move from one block on the domain to another, which may reside on a different MPI rank. Such a particle will therefore need access to grid data on a different block before we can evaluate Equation 4 for this Runge-Kutta stage. Thus, immediately after calculating the new particle positions and density matrices corresponding to a Runge-Kutta stage, if any particle position now resides within a different block we redistribute its data for all Runge-Kutta stages to its new block and MPI rank. We then resume calculating the right hand side defined by Equation 4 for this Runge-Kutta stage using the deposition and interpolation operations discussed in Section II B 1 and Section II B 2. We repeat this redistribution procedure at all Runge-Kutta stages and at the end of the time step. This ensures that we are able to evaluate all Runge-Kutta stage right hand sides for the particles even if previous stage right hand side terms were

originally evaluated on different MPI ranks.

Local particle arithmetic operations (deposition, interpolation, evolution) constitute the vast majority of the computational cost. We implement these kernels using AMReX's performance-portable strategy for targeting either CPU threading (with OpenMP) or GPUs (with CUDA). The quantities evolved in Equation 4 are arranged in a struct for each particle and each block stores an array of these particle structs. For CPUs, we distribute blocks with particles to threads using OpenMP. Each thread loops over particles within the block, though we find in practice that increasing the number of MPI ranks can give better performance on CPUs than OpenMP threading. For GPUs, we assign one MPI rank to each GPU and allocate local grid and particle memory using CUDA Unified Memory. Each MPI rank loops over its local blocks and asynchronously launches a CUDA kernel parallelizing the loop over particles within the block. We then synchronize GPU kernels only after looping over all local blocks to ensure high kernel occupancy.

III. RESULTS

In this work we restrict ourselves to simulations performed in one spatial dimension and two angular dimensions. That is, we arbitrarily assume all neutrinos have an energy of 50 MeV, though this only influences the vacuum Hamiltonian, which has negligible effect on the timescale of these simulations. In this section we examine in detail only a single choice of initial conditions. This simplified setup will serve as a basis for comparison with other simulations in this and future work. For our fiducial simulation we choose initial conditions that exhibit a strong electron lepton number (ELN) crossing (i.e., a direction along which there are equal numbers of electron neutrinos and antineutrinos), characterized by electron neutrino and electron antineutrino distributions with a flux factor of 1/3 in opposite directions. The distribution is thus quite unstable to the FFI. We randomly perturb the initial conditions to seed all modes, as we will describe in Section III A. The amplitudes of the unstable modes grow exponentially precisely in the way predicted by linear stability analysis as shown in Section III B. The instability then saturates, diffusing neutrino flavor fluctuations away from the fastest growing wavelength into modes of all wavelengths as described in Section III C. Select data and all parameter files from this study are publicly available [92]. Additional data is available upon request.

A. Simulation Setup

In our fiducial simulation we use a simulation domain 64 cm in length with a grid resolution of $1 \times 1 \times 1024$. The unit size in the x and y directions makes the simulation behave in a one-dimensional manner. That is, the neutrino distributions are homogeneous in the x and y directions, but are not restricted to isotropy or axial symmetry. We place many particles at the center of each grid cell with directions distributed approximately isotropically using the "Polar Coordinate Subdivision" method of Ref. [98], except that we specify the number of points along the equator rather than the total number of points. Requesting 16 directions in the $\hat{x}-\hat{y}$ plane results in 92 particles per grid cell. In contrast, PIC codes often distribute particles initially randomly. We elect to distribute particles regularly so we do not have to consider Monte Carlo numerical noise when interpreting the results.

We arbitrarily choose an initial neutrino distribution with a number density of $n = 4.89 \times 10^{32}\,\mathrm{cm}^{-3}$ of each electron neutrinos and electron antineutrinos, and zero density for heavy lepton neutrinos. The fastest growing mode of a two-beam model with these neutrino densities has a wavelength of 1 cm, allowing a more straightforward comparison to the two-beam test in Appendix A 4. To create the desired anisotropic neutrino distribution using an isotropic distribution of computational particles, we give each particle weights (i.e. the number of physical neutrinos or antineutrinos that a computational particle represents) of

$$\begin{aligned} N_p &= \frac{1}{N_{\mathrm{ppz}}}(n_{\mathrm{input}} + 3\mathbf{f}_{\mathrm{input}} \cdot \mathbf{\Omega})\Delta V \ , \\ \bar{N}_p &= \frac{1}{N_{\mathrm{ppz}}}(\bar{n}_{\mathrm{input}} + 3\bar{\mathbf{f}}_{\mathrm{input}} \cdot \mathbf{\Omega})\Delta V \ , \end{aligned} \tag{8}$$

where n_{input} ($\bar{n}_{\mathrm{input}}$) is the input number density of (anti-)neutrinos and $\mathbf{f}_{\mathrm{input}}$ ($\bar{\mathbf{f}}_{\mathrm{input}}$) is the input number flux of (anti-)neutrinos and N_{ppz} is the number of computational particles per grid zone. We use $\mathbf{f}_{\mathrm{input}} = n_{\mathrm{input}}\hat{z}/3$ and $\bar{\mathbf{f}}_{\mathrm{input}} = -\bar{n}_{\mathrm{input}}\hat{z}/3$ for this fiducial simulation. This defines $\hat{z}$ as our axial symmetry axis and results in a healthy ELN crossing in the $\hat{x}-\hat{y}$ plane. The initial state of each particle is set to be in a nearly pure electron flavor state with a random perturbation in the off-diagonal elements. That is,

$$\rho_{ab,p} = \begin{bmatrix} 1-\epsilon_\mu-\epsilon_\tau & \alpha(U+Ui) & \alpha(U+Ui) \\ \rho^*_{e\mu} & \epsilon_\mu & 0 \\ \rho^*_{e\tau} & 0 & \epsilon_\tau \end{bmatrix} . \tag{9}$$

$U \in [-1,1]$ is a uniform random number generated individually each time it appears and $\alpha = 10^{-6}$ is the strength of the random perturbation. ϵ_μ and ϵ_τ are determined after the random numbers are generated in order to ensure unit trace and a flavor vector length of one. Since the initial distributions have almost no ν_τ or ν_μ content, we set the $\rho_{\mu\tau}$ off-diagonal components to zero.

The left column of Figures 2-4 show these initial conditions at $t = 0$. The top left panel of Figure 2 shows the initial number density of each neutrino species plotted over 16 cm of the 64 cm domain divided by the trace of the number density matrix. All neutrinos (solid curves)

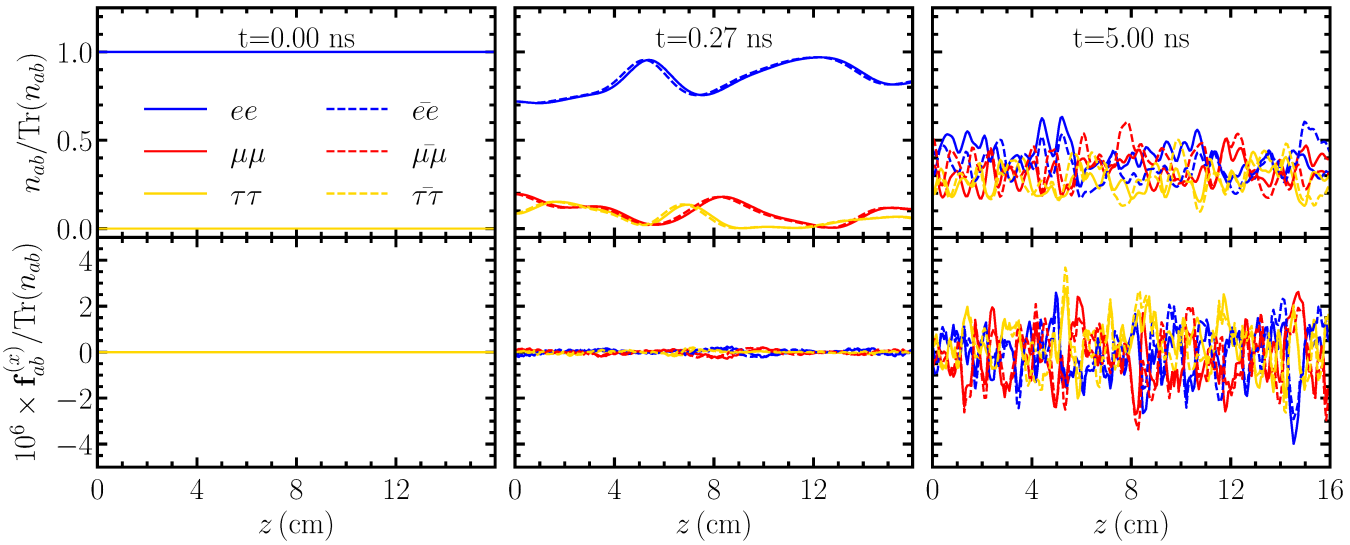

FIG. 2. Snapshots of flavor-diagonal components of direction integrated quantities in the fiducial 1D simulation. The three columns show snapshots at the initial time, near the end of the linear growth phase, and the post-saturation phase, respectively. The top row shows neutrino number density and the bottom row shows the flux in the $\hat{x}$ direction. Blue, red, and gold solid (dashed) curves show electron, muon, and tauon flavor (anti-)neutrinos, respectively. The abundances of all flavors are roughly equal in the post-saturation state.

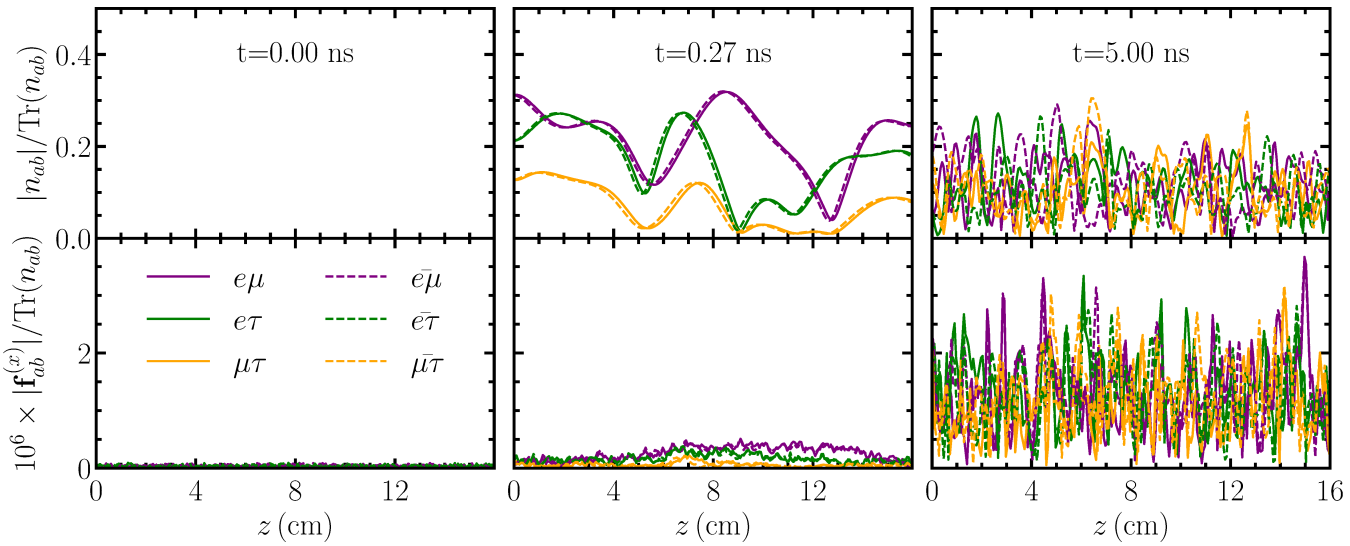

FIG. 3. Snapshots of magnitude of flavor off-diagonal components of direction integrated quantities in the fiducial 1D simulation. The three columns show snapshots at the initial time, near the end of the linear growth phase, and the post-saturation phase, respectively. The top row shows neutrino number density and the bottom row shows the flux in the $\hat{x}$ direction. Purple, green, and orange solid (dashed) curves show $e\mu$, $e\tau$, and $\mu\tau$ components for (anti-)neutrinos, respectively. Modes with wavelength near that of the fastest growing mode dominate the linear evolution, but break up into a broad spectrum of wavelengths in the post-saturation phase.

and antineutrinos (dashed curves) begin in the electron flavor state (blue), and hence the electron neutrinos and antineutrinos constitute essentially all of the number density. All neutrino flux is in the $\pm\hat{z}$ direction, so the x component of the flux shown in the bottom panel at $t = 0$ is zero everywhere. The left column of Figures 3 and 4 show the magnitude and complex phase, respectively, of the $e\mu$ (purple), $e\tau$ (green), and $\mu\tau$ (orange) components of the neutrino density. The perturbations are too small to be visible in the top left panel of Figure 3. Although the perturbations of the density matrix of each particle individually are $\mathcal{O}(10^{-6})$, they average together to make the perturbation of the x component of the flux factor (bottom panel) much smaller than 10^{-6}. The left panels

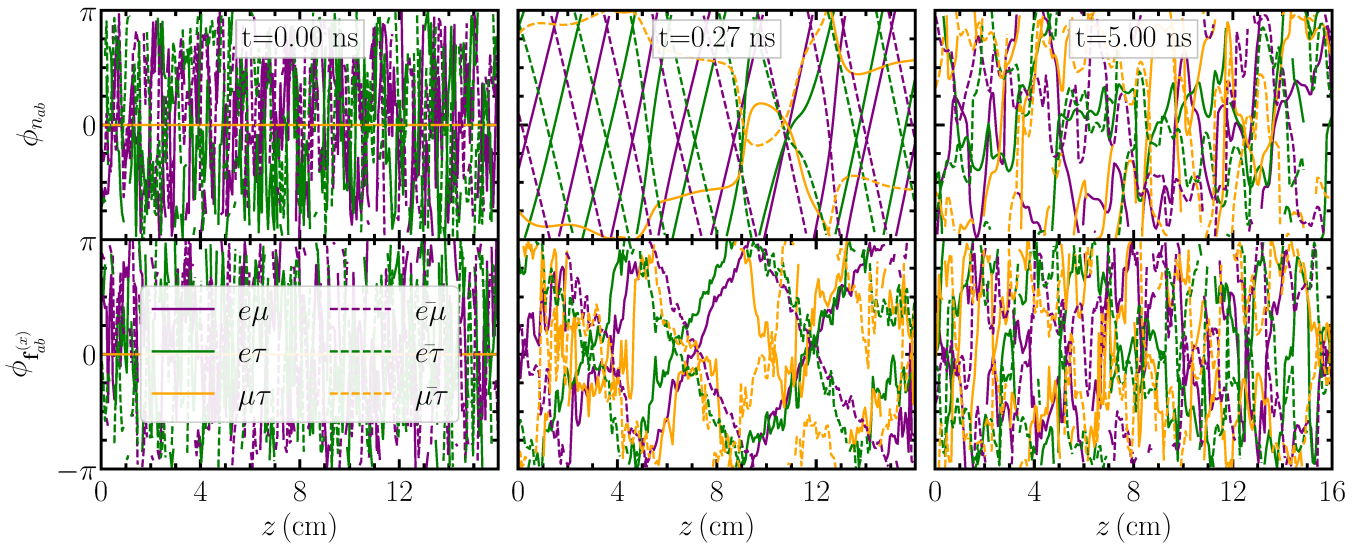

FIG. 4. Snapshots of complex phase of flavor off-diagonal components of direction integrated quantities in the fiducial 1D simulation. The three columns show snapshots at the initial time, near the end of the linear growth phase, and the post-saturation phase, respectively. The top row shows neutrino number density and the bottom row shows the flux in the $\hat{x}$ direction. Purple, green, and orange solid (dashed) curves show $e\mu$, $e\tau$, and $\mu\tau$ components for (anti-)neutrinos, respectively. The wavelength (distance between points of equal phase) of the fastest growing axially symmetric mode is visible in the top panel of the center column, while the wavelength of the fastest growing axially asymmetric mode is visible in the bottom panel of the center column.

of Figure 4 simply show a random distribution of phases of the $e\mu$ and $e\tau$ components of both the number density and flux flavor matrices. The complex phase of the $\mu\tau$ component is shown as zero because it has exactly zero magnitude.

B. Linear Growth Phase

Before discussing the simulation results, we briefly go through the standard linear stability analysis to point out what modes should be unstable and the rate at which they are expected to grow. Following [49], we calculate the imaginary component of the frequency for modes of the form $N_{ex} = \alpha e^{-i(\omega t - kz)}$ described by a frequency ω, wavenumber k, and initial perturbation amplitude α. x here represents either μ or τ. This is a two-flavor analysis, so the results are applicable to $e-\mu$ and $e-\tau$ flavor transformation but does not describe $\mu-\tau$ transformation [99]. The cosine dependence of the distribution on the polar angle makes computing the integrals in the polarization tensor analytic, though a numerical root find still needs to be performed to extract the growth rate of each mode. We plot the results in Figure 5. Since the results scale with the neutrino density (assuming the flux factors and directions are held constant), we plot both the growth rate and wavenumber scaled by $\sqrt{2}G_F n_{\nu_e}$ as a metric of the self-interaction potential. There are two branches, one of which has a larger growth rate than the other at all wavenumbers. The faster growing mode

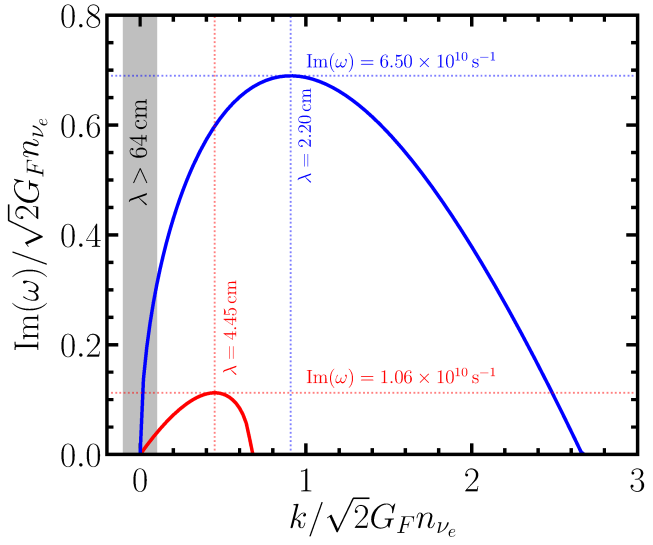

FIG. 5. Dispersion relation of modes with frequency ω and wavenumber k. The blue curve corresponds to modes that do not break axial symmetry, while the red curve corresponds to modes that do. Both are normalized by a characteristic self-interaction potential so the diagram is independent of neutrino density as long as the self-interaction potential is much larger than the vacuum potential. Unstable modes exist for wavelengths $\lambda \gtrsim 0.75\,\mathrm{cm}$ ($k/\sqrt{2}G_F n_{\nu_e} \lesssim 2.66$), though the growth rate for the homogeneous mode ($k = 0$) is also 0. The gray shadow shows the region of wavenumbers where the corresponding mode cannot exist on our domain because of our limited domain size of 64 cm.

(blue) has a peak growth rate of $\mathrm{Im}(\omega) = 6.50 \times 10^{10}\,\mathrm{s}^{-1}$ and a wavelength of 2.20 cm. This is still somewhat slower and longer wavelength than the fastest growing mode in a two-beam model with equivalent densities (Appendix A 4). The slower growing mode has a maximum growth rate of $\mathrm{Im}(\omega) = 1.06 \times 10^{10}\,\mathrm{s}^{-1}$ at a wavelength of 4.45 cm. We will later show that the blue curve represents the axial symmetry preserving mode and the red curve represents the axial symmetry breaking mode.

Note that some of the unstable modes have wavelengths longer than can fit on our domain (shaded gray region), but the growth rate of these modes is significantly smaller than the peak growth rate. In addition, the finite number of grid points limits the shortest possible wavelength to 0.125 cm, or a maximum wavenumber $k/\sqrt{2}G_F n_{\nu_e} \approx 16$, so that all modes shown in Figure 5 outside the gray shaded region are within the grid's Nyquist limit. More importantly, our finite resolution limits the number of grid cells supporting each wavelength, and we find that to resolve the instability we need many grid cells per wavelength of the fastest growing mode. We explore the resolution requirements to obtain a converged instability growth rate in Appendix B and verify that the quoted results of our simulation are robust against changes to the domain size and grid resolution.

The simulation immediately exhibits rapid growth of the flavor off-diagonal components of the neutrino density matrix. As a metric of the net growth of all modes, we calculate the magnitude of the off-diagonal components of the neutrino density and the x component of the flux averaged over the entire domain as

$$\begin{aligned} \langle |n_{ab}| \rangle &= \frac{1}{N_{\mathrm{zones}}} \sum_{\mathrm{zones}} |n_{ab}| \,, \\ \langle |\mathbf{f}^{(x)}_{ab}| \rangle &= \frac{1}{N_{\mathrm{zones}}} \sum_{\mathrm{zones}} |\mathbf{f}^{(x)}_{ab}| \,. \end{aligned} \tag{10}$$

$\langle |n_{ab}| \rangle$ probes the presence of modes that preserve axial symmetry, while $\langle |\mathbf{f}^{(x)}_{ab}| \rangle$ probes modes that break axial symmetry because $\langle |\mathbf{f}^{(x)}_{ab}| \rangle = 0$ would result from axially symmetric modes alone. The purple and green solid curves in Figure 6 show growth of the $e\mu$ and $e\tau$ components of the neutrino density. For the first ~ 0.1 ns they both exhibit rapid, but not precisely exponential growth. In this phase, there are several competing modes growing at different rates, all of which are initially large enough to contribute to $\langle |n_{e\mu}| \rangle$ and $\langle |n_{e\tau}| \rangle$. After 0.1 ns the average amplitudes grow exponentially until they saturate at $\sim$ 0.27 ns. After saturation, the average flavor off-diagonal densities hover around the saturation level of a few tenths of the total neutrino density.

The top panel of the center column of Figure 2 shows the flavor diagonal components of the neutrino density at the end of the linear phase. Electron neutrinos are still dominant, though there is now a significant fraction of ν_μ and ν_τ, showing that the instability indeed leads to significant flavor transformation. The amplitude of the x-component of the flux has grown as well, though it

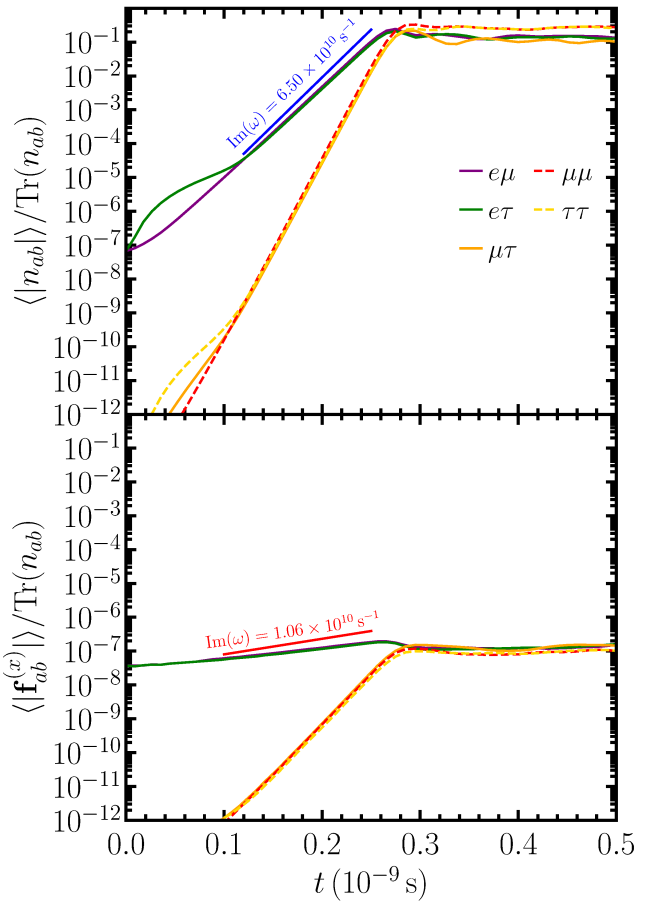

FIG. 6. Time evolution of domain-averaged absolute-value of the neutrino number density (top panel) and number flux in the $\hat{x}$ direction (bottom panel). The labeled line segments show the slopes corresponding to the growth rates of the fastest growing axially symmetric (blue line segment) and axially asymmetric modes (red line segment) shown in Figure 5. The $e\mu$ (purple) and $e\tau$ (green) components of the number density grow at approximately the axially symmetric maximum growth rate, and the number fluxes grow at approximately the axially asymmetric growth rate. The $\mu\tau$ (orange) components grow parasitically at the same rate as the $\mu\mu$ (red dashed) and $\tau\tau$ (gold dashed) components.

remains at only $\sim 10^{-7}$ of the total neutrino density. Although during the linear growth phase the anti-neutrino curves (dashed) match the neutrino curves (solid), as the simulation approaches saturation the nonlinear interaction between modes drives chaotic behavior and the neutrino and antineutrino distributions begin to show slight differences. The center column of Figure 3 correspondingly shows the magnitude of the flavor off-diagonal components of the neutrino density and flux. By this point in time, the off-diagonal elements have become comparable in magnitude to the diagonal components, again indicating that saturation is approaching.

The center column of Figure 4 most clearly shows that the wavelength of the growing modes matches those pre-

dicted in the dispersion relation. The top panel shows the complex phase ϕ of the flavor off-diagonal components of the neutrino density. The complex phase of $n_{e\mu}$ and $n_{e\tau}$ increase with z and have a wavelength (distance between points of equal phase) consistent with the 2.20 cm wavelength of the fastest growing mode on the blue branch in Figure 5. The antineutrinos (dashed) are similar, except that they have a negative wavenumber, resulting in phase decreasing with increasing z. The bottom panel shows the complex phase of the flavor off-diagonal x components of the neutrino flux. Unlike the number density, this quantity is sensitive to axial symmetry breaking and is not sensitive to modes that preserve axial symmetry. The wavelength is consistent with the 4.45 cm prediction from the peak of the red curve in Figure 5. Thus, we expect that the blue curve in Figure 5 represents the symmetry-preserving mode and the red curve represents the symmetry-breaking mode.

The mode identification is further corroborated by the growth rates of both modes. The growth rate of the flavor off-diagonal number density (solid purple and green curves) and x-flux (dashed purple and green curves) in Figure 6 are consistent with the growth rates predicted by the peaks of the blue and red curves, respectively, shown in Figure 5. The simulations do, however, show a slightly slower growth rate than predicted by the peak of the dispersion relation. This may be partially due to the fact that the domain-averaged quantities reflect many modes, all but the fastest of which grow slower than the fastest growing mode. However, the simulated growth rate also does depend somewhat on numerical considerations. We perform a convergence study in Appendix B to show this dependence.

The $\mu-\tau$ mixing is more difficult to interpret. The $n_{\mu\tau}$ component (solid yellow curve in Figure 6) grows with twice the growth rate of $n_{e\mu}$ and $n_{e\tau}$. We expect that this simply follows the growth of the $n_{\mu\mu}$ (brown) and $n_{\tau\tau}$ (salmon) elements, each of which also rise with this doubled growth rate. This can be understood intuitively in terms of a rotating flavor vector $\vec{\rho}$ in SU(3) flavor space. The eight components of this vector are

$$\rho_i = \frac{1}{2}\lambda^i_{ab}\rho_{ab} \ , \tag{11}$$

where λ^i_{ab} are Gell-Mann matrices and i is the vector index. Flavor transformation rotates this vector, preserving $\mathrm{Tr}(\rho_{ab}) = 1$ with constant $|\vec{\rho}|$. In this case, $|\vec{\rho}| = 1/\sqrt{3}$ because we start in a nearly pure electron flavor state in Equation 9. Combining these two conditions yields a relationship

$$|\rho_{e\mu}|^2+|\rho_{e\tau}|^2+|\rho_{\mu\tau}|^2 = (\rho_{\tau\tau}+\rho_{\mu\mu})-(\rho^2_{\mu\mu}+\rho^2_{\tau\tau}+\rho_{\mu\mu}\rho_{\tau\tau}) \ . \tag{12}$$

If, as in the linear growth phase of the simulation, the neutrinos are only slightly perturbed from the electron flavor state, the second term on the right hand side is much smaller than the first. Additionally, $\rho_{\mu\tau}$ is much smaller than $\rho_{e\mu}$ and $\rho_{e\tau}$ throughout the linear growth phase. Combining these approximations,

$$\eta \equiv \rho_{xx} - \rho^2_{ex} \approx 0 \ , \tag{13}$$

where $\rho^2_{ex} = \rho^2_{e\mu} + \rho^2_{e\tau}$ and $\rho_{xx} = \rho_{\mu\mu} + \rho_{\tau\tau}$. Thus, if $\rho_{ex} \propto \exp(\mathrm{Im}(\omega)t)$ (where x refers to either μ or τ) then $\rho_{xx} \propto \exp(2 \cdot \mathrm{Im}(\omega)t)$. Since the growth of $\rho_{\mu\tau}$ requires non-zero difference between $\rho_{\mu\mu}$ and $\rho_{\tau\tau}$, the growth of the initial random differences between $\rho_{\mu\mu}$ and $\rho_{\tau\tau}$ give $\rho_{\mu\tau}$ the room to grow at the same rate as $\rho_{\mu\mu}$ and $\rho_{\tau\tau}$. Note that this shows that the growth of $n_{\mu\tau}$ can be up to twice that of $n_{e\mu}$ and $n_{e\tau}$ but does not show that it must be so. It also does not explain the rapid growth of all fluxes except $\mathbf{f}_{e\mu}$ and $\mathbf{f}_{e\tau}$. More work is required for a complete description of the growth of all modes and flavors (though see [99] for work in this direction).

The differing growth rates of modes of different wavelengths can be seen in the evolution of the Fourier transforms of neutrino distribution shown in Figure 7. Each curve in the plot is the Fourier transform of a grid quantity time-averaged over a 0.1 ns interval to reduce noise in the plot. The color indicates the end of the time interval for which the transform was evaluated. Beginning with the middle panel, the bottom-most curve shows the Fourier transform of $n_{e\mu}$ from $0-0.1$ ns and the next curve (salmon) shows the same for the interval $0.1-0.2$ ns. We use salmon for the second curve to highlight the time interval squarely within the linear growth phase. The wavenumber of the peak, the minimum unstable wavenumber, and the maximum unstable wavenumber are at the locations predicted by the blue curve in Figure 5, as both of these intervals lie within the linear growth phase of the instability. The third curve shows the time interval $0.2-0.3$ ns, during which the instability saturates. The peak location is consistent with the wavenumber of the fastest growing modes, but additional features develop away from the peak that we will discuss in Section III C. The top panel of the figure shows the Fourier transform of n_{ee} on the same time intervals. Since this component is purely real, the Fourier transform is symmetric in k, but still reflects the characteristic wavelengths in the unstable modes. The bottom panel shows the Fourier transform of $\mathbf{f}^{(x)}_{e\mu}$. Once again, the location and width of the peak in the salmon curve matches the expectations from the symmetry-breaking (red) branch of Figure 5.

The symmetry breaking and preserving modes are also apparent in the angular structure of the neutrino flavor. The top panel of Figure 8 shows a Mollweide projection of the directions of all neutrinos located within a single grid cell for a simulation with high angular resolution. Each point on the plot represents the direction of an individual computational particle. In this high-resolution simulation we request 128 directions in the $\hat{x}-\hat{y}$ plane, visible as 128 points along the equator in the plot, for a total of 6022 particles per cell. Our fiducial simulation uses only 16 points along the equator, but we use the high angular resolution data for better visualization. We show in Appendix B that the growth rates and satura-

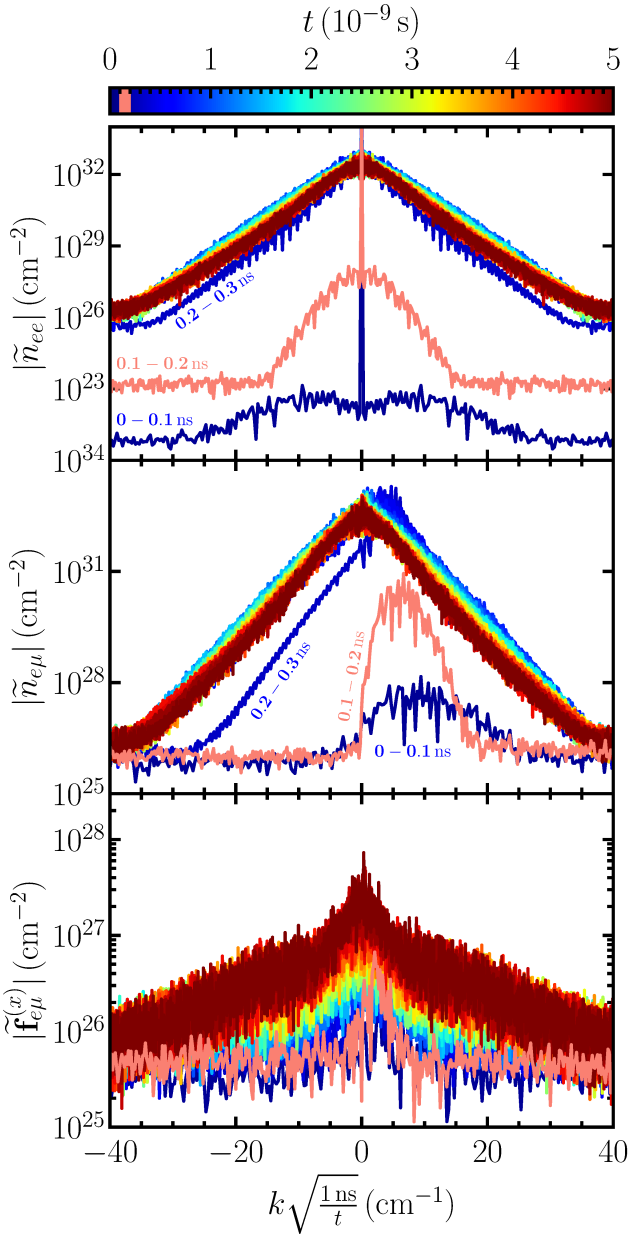

FIG. 7. Fourier transforms of n_{ee} (top panel), $n_{e\mu}$ (middle panel), and $\mathbf{f}^{(x)}_{e\mu}$ (bottom panel), each time-averaged over 0.1 ns intervals and colored by the final time of the corresponding interval. The second curve in each plot is salmon colored to highlight the time interval squarely within the linear growth phase ($0.1-0.2\,\mathrm{ns}$). The $\tilde{n}_{ee}$, $\tilde{n}_{e\mu}$, and $\tilde{\mathbf{f}}^{(x)}_{e\mu}$ presaturation peaks align well with the wavenumber/wavelength locations predicted in Figure 5, and spread out at late times to form a chaotic mixture of many modes.

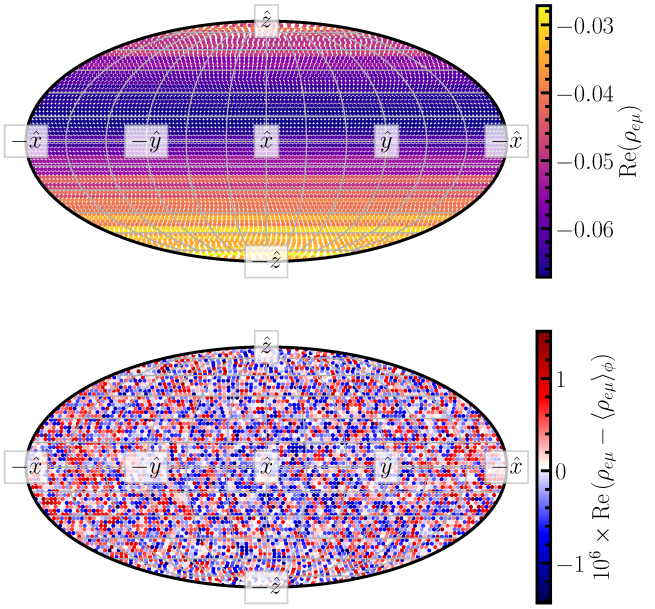

FIG. 8. Mollweide projection of the neutrino particle directions in a single grid cell near the end of the linear growth phase at $t = 0.27\,\mathrm{ns}$ in the "128 Equatorial Direction" simulation in Appendix B. *Top panel:* the color of each point encodes the value of $\mathrm{Re}(\rho_{e\mu})$ for each particle and shows the large amplitude of the fastest growing axially symmetric mode. *Bottom panel:* the color encodes the deviation of $\mathrm{Re}(\rho_{e\mu})$ from the axial average to emphasize the axial symmetry breaking mode. The axially asymmetric mode has an amplitude comparable to the random perturbations of the initial conditions, seen as an overabundance of blue in the $\hat{x}$ direction and an overabundance of red in the $-\hat{x}$ direction.

tion states are the same in both simulations. The colors of the points show the real part of $\rho_{e\mu}$ for each packet at $t = 0.27\,\mathrm{ns}$. Although the FFI has greatly amplified these numbers from their initially random values of $\mathcal{O}(10^{-6})$, the flavor structure is very axially symmetric. If we subtract the axial average (bottom panel), more structure becomes apparent. What remains at this particular location is a slight excess of positive values in the $-\hat{x}$ direction and a slight excess of negative values in the $\hat{x}$ direction, corresponding to a small but significant flux of off-diagonal flavor $\mathbf{f}^{(x)}_{e\mu} = -1.62 \times 10^{24}\hat{x}\,\mathrm{cm}^{-3}$. The axial asymmetry is very weak because the growth rate of the axial symmetry breaking mode is so much slower. By the time the fastest growing axially symmetric mode saturates, the axially asymmetric mode has only grown by a factor of $\exp(\mathrm{Im}(\omega_{\max})\Delta t) \approx 20$.

Overall, the simulations confirm the predictions of the linear growth phase from linear stability analysis. Modes of the expected wavenumbers all grow at the expected rates. The modes with the fastest growth rates appear to be the most important, as they saturate and lead to nonlinear evolution long before the slower modes can build significant amplitude.

C. Saturated Phase

At $t \approx 0.27\,\mathrm{ns}$ the exponential growth of both the axially symmetric and asymmetric modes shown in Figure 6 cease as the instability saturates. After saturation, the

number density of all three neutrino flavors tends toward equipartition at late times. The right column of Figure 2 shows a snapshot of the flavor diagonal elements of n_{ab} (top) and $\mathbf{f}^{(x)}_{ab}$ (bottom) at $t = 5\,\mathrm{ns}$. In both cases, the flavor content is widely variable in z, but when averaged over the domain, the neutrino density and flux are very close to flavor equipartition. The right column of Figure 3 shows that the flavor off-diagonal components of the density and flux remain large compared to the flavor diagonal components. Finally, the right column of Figure 4 shows that the phase is effectively randomized in the saturation phase, indicating that the coherent modes that grew in the linear phase are all but disrupted.

Figures 2-4 also show variation on much shorter wavelengths (higher wave numbers) than the fastest growing modes that dominated the linear growth phase. This evolution to short wavelengths can be seen in Figure 7, which shows Fourier transforms of n_{ee}, $n_{e\mu}$ and $\mathbf{f}^{(x)}_{e\mu}$ from the $N_z = 2048$ simulation in Appendix B. We display this higher-resolution calculation because it allows us to track the evolution out to larger wavenumbers than our fiducial simulation, while yielding the same instability growth rates and final averaged neutrino flavor abundances as in the fiducial simulation. We described the amplification of the unstable modes shown in the earliest two curves (the lowest purple curve and the salmon curve) in Section III B. By 0.3 ns (third curve from the bottom of each panel, though it is not clearly visible in the lowest panel) the amplitudes of the fastest growing modes have reached their maximal values and power in all three quantities begins to spread to larger and smaller wavenumbers. Though the noise in the flux spectra (bottom panel) makes them difficult to interpret, each curve of $|\widetilde{n}_{ee}|$ and $|\widetilde{n}_{e\mu}|$ varies linearly with $|k|\sqrt{1\,\mathrm{ns}/t}$ in the logarithmic plot, indicating that

$$|\widetilde{n}_{ee}| \sim |\widetilde{n}_{e\mu}| \sim \exp\left[-|k|\left(\frac{t}{10^{-9}\,\mathrm{s}}\right)^{-1/2}\right] . \tag{14}$$

Some understanding of the scaling with time and wavenumber of the post-saturation spectrum can be obtained by Fourier transforming the QKEs in Equation 1. Again ignoring the collision term and only considering the neutrino potential, this becomes

$$\frac{\partial \widetilde{f}_k}{\partial t} + ic\widetilde{f}_k \mathbf{\Omega}\cdot\mathbf{k} = \frac{-i\sqrt{2}G_F}{\hbar c^3\sqrt{2\pi}} \int d\mathbf{\Omega}' \int d\left(\frac{\nu'^3}{3}\right)(1-\mathbf{\Omega}\cdot\mathbf{\Omega}') \int dx\, e^{-ikx}\left[f'-\bar{f}', f\right] , \tag{15}$$

where we use the convention that $\widetilde{f}_k = \int_{-\infty}^{\infty} dx\, e^{-ikx} f(x)/\sqrt{2\pi}$ and denote $f' = f(\mathbf{x}, t, \mathbf{\Omega}', \nu')$. The primed quantities are integrated over; a neutrino with direction $\mathbf{\Omega}$ and frequency ν is moving through a distribution of neutrinos, each of which has a particular value of $\mathbf{\Omega}'$ and ν'. The second term on the left is the advection term that only changes the phase of each Fourier component as waves propagate, and hence it does not affect the magnitude of the component. Representing f' in terms of its Fourier components $f' = \int_{-\infty}^{\infty} dk'\, e^{ik'x}\widetilde{f}'_{k'}/\sqrt{2\pi}$, writing for simplicity $\widetilde{A}'_{k'} = \widetilde{f}'_{k'} - \widetilde{\bar{f}}'_{k'}$ and rearranging the integrals, the term on the bottom line of Equation 15 becomes

$$\int_0^{\infty} dk' \left(\left[\widetilde{A}'_{k'}, \widetilde{f}_{k-k'}\right] + \left[\widetilde{A}'_{-k'}, \widetilde{f}_{k+k'}\right]\right) . \tag{16}$$

Next we note that after saturation the distribution is approximately symmetric in k, making $\widetilde{A}'_{k'} \approx \widetilde{A}'_{-k'}$. In addition, $\widetilde{f}_k$ is strongly peaked near $k = 0$. Taking inspiration from [70], this allows us to approximate $\widetilde{f}_{k-k'} + \widetilde{f}_{k+k'} \approx 2\widetilde{f}_k + k'^2\partial_k^2\widetilde{f}_k$. The integral can now be expressed as

$$\int_0^{\infty} dk' \left[\widetilde{A}'_{k'}, \left(2\widetilde{f}_k + k'^2\partial_k^2\widetilde{f}_k\right)\right] . \tag{17}$$

Similar to the behavior of the angular power spectrum in [70], the behavior of the second term is invariant to transformations where $t \to \alpha^2 t$ and $k \to \alpha k$, matching what we see at late times. Note that we do not explain why the distribution migrates from the fastest growing wavenumber to be centered at $k = 0$, and we use the fact that the distribution is so sharply peaked in deriving Expression 17. We leave a more detailed investigation of the saturation and decoherence of spectral power to future work.

By $t = 5\,\mathrm{ns}$ the base of the wings in Figure 7 have reached a wavenumber of $|k| \approx 100\,\mathrm{cm}^{-1}$, which is the maximum representable wavenumber on our grid of 2048 spatial points. Power begins to artificially build up at these higher wavenumbers, resulting in the slight upswing in the final (red) curves at $k\sqrt{1\,\mathrm{ns}/t} \approx \pm 40$. In the lower resolution fiducial simulation this upswing occurs after $t \approx 2\,\mathrm{ns}$ when the wings reach that grid's maximum $|k|$ value of $50\,\mathrm{cm}^{-1}$. This suggests that in nature a true steady state solution will only be achieved by waiting long enough for the power to diffuse from the unstable wavelength to the inter-particle distance (representing the maximum physical wavenumber) multiple times, such that the distribution in wavenumber can equilibrate. However, the process of kinematic decoherence likely behaves differently in nature's three dimensions than in our simulation's single spatial dimension, so the decoherence rates demonstrated in this simulation are not necessarily realistic.

Finally, we show the angular structure of the post-saturation distribution in Figure 9 corresponding to $t = 4.95\,\mathrm{ns}$. Similarly to Figure 8, the top panel of Figure 9 shows $\mathrm{Re}(\rho_{e\mu})$ for each particle in a single grid cell, where particle directions are plotted using a Mollweide projection. Amazingly, the saturated distribution remains highly axially symmetric, with deviations from the axial average (shown in the bottom panel) only slightly larger than they were at the end of the linear growth phase.

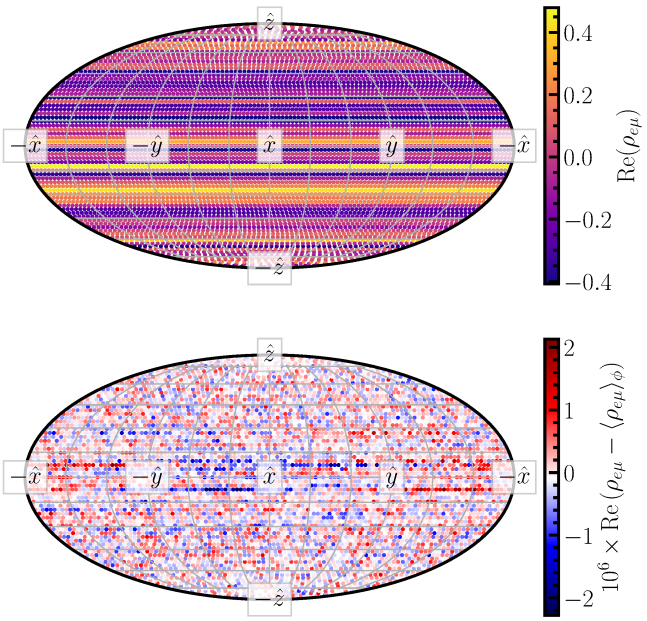

FIG. 9. Mollweide projection of the neutrino particle directions in a single grid cell during the saturated phase at $t = 4.95\,\mathrm{ns}$ in the "128 Equatorial Direction" simulation in Appendix B. *Top panel:* the color of each point encodes the value of $\mathrm{Re}(\rho_{e\mu})$ for that particle and shows that even in the saturated state the distribution remains largely axially symmetric. *Bottom panel:* the color encodes the deviation of $\mathrm{Re}(\rho_{e\mu})$ from the axial average to emphasize that the amplitudes of axial symmetry breaking modes remain small.

This can be understood from the dashed curves in Figure 6 that show the growth of axially asymmetric modes. Although these modes are initially unstable, the axially symmetric modes are much faster and saturate while the asymmetric modes still have small amplitude. This saturation seems not only to stop the growth of the symmetric modes, but also the asymmetric modes. However, this picture also likely changes in three spatial dimensions, so three dimensional simulations will be required for a more complete understanding of axial symmetry breaking. In addition, because the initial conditions in this simulation are themselves axially symmetric, the presence of asymmetries in the initial conditions would likely make this process rich in phenomena not expressed in our fiducial simulation.

In addition to the similarity to the angular power spectrum in Ref. [70], these results also bear a striking resemblance to the evolution of the power spectrum in Ref. [100]. In the latter case, the simulations were performed in the regime of slow mode flavor transformations (i.e., the matter potential and the neutrino self-interaction potential are both relevant), but the nonlinear interaction between modes still causes power to cascade to higher modes. This was not seen in Ref. [72], which shows the power remaining at low wavenumbers for the duration of the simulation. We suspect this to be an effect of the choice of neutrino distribution, the structure of the initial perturbation, and the duration of the simulation. For example, Martin et al. 72 use a well-controlled localized initial perturbation and allow evolution only as long as the perturbation does not travel more than the size of their periodic domain, while we use random perturbations throughout the domain and simply verify invariance of the results for different choices of the periodic domain size (Appendix B). While we expect simulations using these boundary conditions to provide valuable insight into statistical properties of the instability growth and saturation, a more complete understanding of nature's realization of neutrino quantum kinetics will require a globally resolved quantum kinetics simulation of the full supernova or neutron star merger. Given our analysis above, we expect that in any situation where the initial perturbation grows to become nonlinear, power will cascade to all angular and spatial modes, though the rate at which this happens is likely to depend on the neutrino distribution itself. It is possible that the simulations of Ref. [72] would see the same behavior on longer timescales. Given the approximate azimuthal symmetry of the fiducial simulation, we suspect that all three codes would produce statistically consistent results for the same simulation parameters, though this remains to be checked. However, initial conditions and perturbations that break the azimuthal symmetry in Section IV will require a simulation method that simulates neutrinos in the relevant number of dimensions.

IV. PARAMETER STUDY

Our final goal in this work is to provide a first look into the dependence of the FFI on the initial matter and neutrino distributions. The initial conditions described in Section III A rely on nine input parameters: the number density of neutrinos n_{input} and antineutrinos $\bar{n}_{\mathrm{input}}$, the number flux vector of neutrinos $\mathbf{f}_{\mathrm{input}}$ and antineutrinos $\bar{\mathbf{f}}_{\mathrm{input}}$ (three components each), and the background electron density n_e. Although we have shown that `Emu` simulates the FFI with high fidelity, it is impractical to densely sample the full parameter space, even under this limited subset of possible initial conditions. Instead, we scratch the surface here by performing a parameter sweep in a single input variable at a time, holding the other input parameters the same as in the fiducial simulation.

The results of these simulations are summarized in Figure 10, where we plot the time evolution of the domain average abundance of electron neutrinos. Each curve is a distinct simulation with the input parameter varied from the fiducial simulation indicated by the color. In the top panel, we show the results of seven simulations identical to the fiducial simulation except that we have included a nonzero background electron density n_e. The fiducial simulation has $n_e = 0$ and is plotted as the darkest curve, though it is occulted by the other curves. As described in Section III, the instability saturates after $\sim 0.3\,\mathrm{ns}$, re-

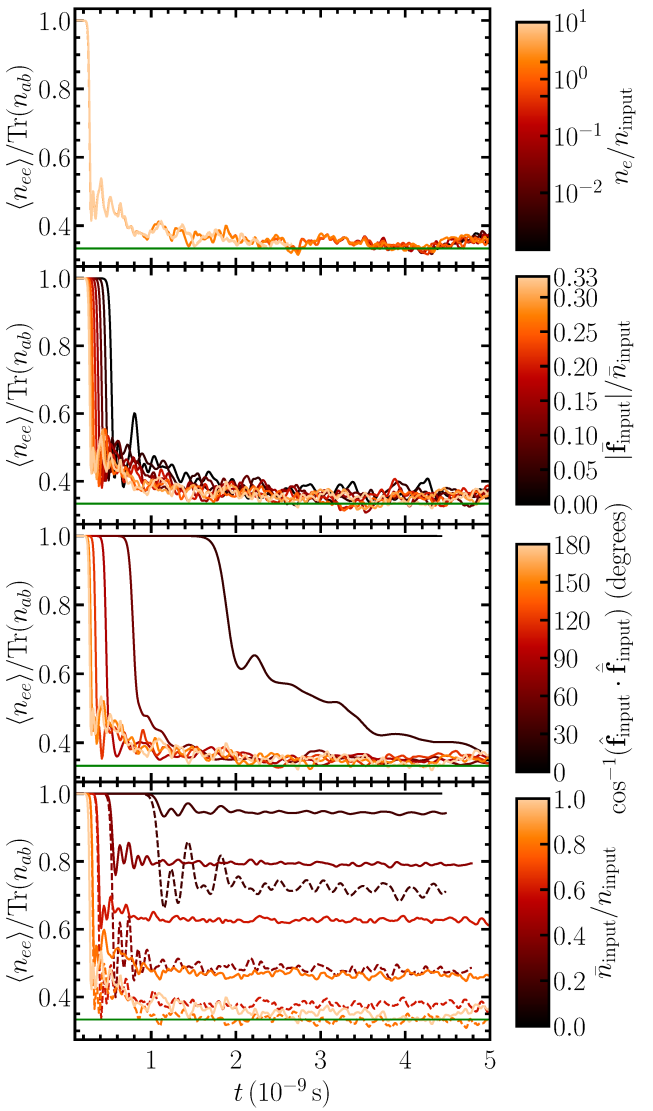

FIG. 10. Time evolution of the electron diagonal component of the domain-averaged neutrino density (solid lines) and antineutrino density (dashed lines) in our parameter study. From top to bottom, the simulations shown varied the background electron density (first panel), the antineutrino flux factor (second panel), the antineutrino flux direction (third panel), and the antineutrino density (fourth panel). The color indicates the value of the input parameter that is varied from the fiducial simulation and the horizontal green line is at a value of 1/3. The electron density affects neither the growth rate nor the saturation state of the instability. All other factors affect the growth rate. That only the antineutrino density affects the final saturation state is an artifact of the neutrino-antineutrino symmetry of the fiducial simulation.

sulting in a significant decrease in the number of electron neutrinos. At late times, the number density of electron neutrinos (solid curve) and antineutrinos (dashed curve) approaches one third of the number density (green line) of all neutrinos or antineutrinos, respectively, indicating flavor equilibration. As expected from the stability analysis, the electrons modify the real component of the mode frequencies but do not affect the growth rate of the unstable modes. Thus, the simulations varying n_e proceed very similarly, diverging at later times due to the chaotic nature of the saturation phase.

In the second panel from the top we show a series of eight simulations where we hold the number densities and flux directions of neutrinos and antineutrinos constant, but vary the magnitude of the antineutrino flux from 0 (isotropic, black curve) to $\bar{n}_{\mathrm{input}}/3$ (fiducial simulation, light orange curve). In the third panel from the top we show results from a series of seven simulations where we hold the neutrino densities and flux magnitudes constant, but change the direction of the antineutrino flux from aligned with the neutrino flux (black curve) to opposed to the neutrino flux (fiducial simulation, light orange curve). In both cases, the parameter variation reduces the growth rate of the instability, increasing the time when the instability saturates. However, these parameters do not appear to change the final abundances of neutrino flavors. All simulations approach flavor equilibration except for the case where the antineutrino and neutrino fluxes align and no flavor transformation occurs. Also note that when we vary either the magnitude or direction of antineutrino flux, the neutrinos and antineutrinos follow the same trajectory and are plotted one on top of the other.

Finally, in the bottom panel we show the results of six simulations where we vary the number density of antineutrinos, keeping the direction and flux factor ($|\bar{\mathbf{f}}_{\mathrm{input}}|/\bar{n}_{\mathrm{input}} = 1/3$) constant. The fiducial simulation has $\bar{n}_{\mathrm{input}}/n_{\mathrm{input}} = 1$ and is plotted as the lightest orange curve. As we decrease the number of electron antineutrinos in the initial conditions (moving toward darker colors), the growth rate of the instability slows, indicated by a longer time before the instability saturates. In addition, the final abundances of both neutrinos and antineutrinos increase relative to the fiducial simulation, but neutrinos increase at a more rapid rate than antineutrinos. When there are no antineutrinos present (uppermost black line) the distribution shows no flavor transformation.

However, the fact that only the ratio of antineutrino to neutrino density changes the final flavor abundances is an artifact of our choice of initial conditions for our fiducial simulation. We performed another set of simulations with initial parameters $\bar{n}_{\mathrm{input}}/n_{\mathrm{input}} = 2/3$, $\mathbf{f}_{\mathrm{input}} = 0$, and $\hat{\bar{\mathbf{f}}}_{\mathrm{input}} = -\hat{z}$ and varied $|\bar{\mathbf{f}}_{\mathrm{input}}|$ from 0 to $\bar{n}_{\mathrm{input}}/3$. Using these parameters, there is an ELN crossing only for $3|\bar{\mathbf{f}}_{\mathrm{input}}|/\bar{n}_{\mathrm{input}} \geq 1/2$. We run each simulation for 25 ns to ensure that all unstable simulations had sufficient time to reach an approximately stationary final state. In Figure 11 we plot time averages of the electron neutrino (solid) and antineutrino (dashed) abundances in blue for the post-saturation phase. As expected, when there is no ELN crossing ($3|\bar{\mathbf{f}}_{\mathrm{input}}|/\bar{n}_{\mathrm{input}} < 1/2$) the distribution exhibits no flavor transformation. Beyond this point, increasing the flux factor increases the size of the ELN crossing, increases the growth rate of the instability, and

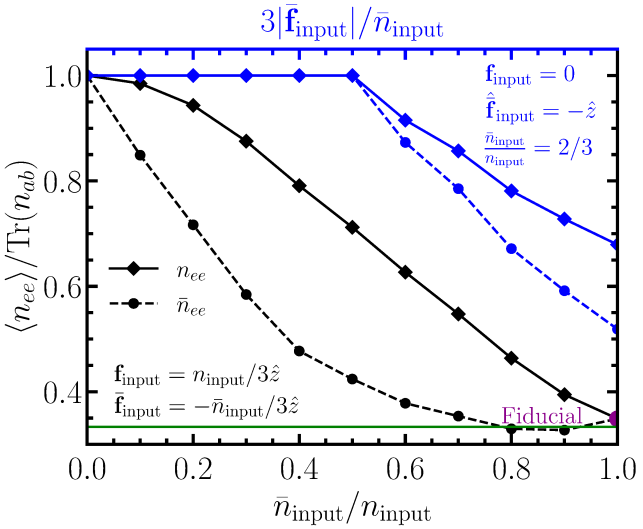

FIG. 11. Time averages of the electron neutrino (solid) and antineutrino (dashed) abundances in the post-saturation phase for parameter sweeps of $|\bar{\mathbf{f}}_{\rm input}|$ and $\bar{n}_{\rm input}/n_{\rm input}$, run for 25 ns. In blue, we show the results of a set of simulations with fixed $\bar{n}_{\rm input}/n_{\rm input}$, $\mathbf{f}_{\rm input}$, and $\hat{\mathbf{f}}_{\rm input}$ and vary $|\bar{\mathbf{f}}_{\rm input}|$ from 0 to $\bar{n}_{\rm input}/3$. The distribution is unstable only when there is an ELN crossing ($3|\bar{\mathbf{f}}_{\rm input}|/\bar{n}_{\rm input} > 1/2$). In black, we show the results of simulations where we set $\mathbf{f}_{\rm input}$ and $\bar{\mathbf{f}}_{\rm input}$ and vary $\bar{n}_{\rm input}/n_{\rm input}$ from $0-1$ (the same data as in the bottom panel of Figure 10). The fiducial simulation is plotted as a purple point, and the horizontal green line marks complete flavor equilibration ($\langle n_{ee}\rangle = {\rm Tr}(n_{ab})/3$). In both parameter sweeps, the antineutrinos consistently experience a greater degree of relative flavor transformation.

increases the amount of eventual flavor transformation. For none of our simulations in this parameter sweep do we see complete flavor equilibration ($\langle n_{ee}\rangle = {\rm Tr}(n_{ab})/3$), though this may occur for larger antineutrino flux factors. For comparison, we also plot the same quantities for the $\bar{n}_{\rm input}/n_{\rm input}$ parameter sweep (i.e., calculated from the results shown in the bottom panel of Figure 10) in black and label the fiducial simulation in purple. In both parameter sweeps the antineutrinos consistently experience a greater degree of relative flavor transformation. This is simply a reflection of the fact that in all of our simulations, the same number of neutrinos and antineutrinos oscillate. That is, we find that in our simulations

$$n_{ee}(0) - \langle n_{ee}\rangle = \bar{n}_{ee}(0) - \langle \bar{n}_{ee}\rangle \ . \tag{18}$$

As elsewhere in the paper, all variations of n in the equation are number densities with units cm^{-3}, and in our simulations $n_{\rm input}$ and $n_{ee}(0)$ are equivalent. Since there are fewer antineutrinos than neutrinos, a larger fraction must change flavor to match the total number of oscillating neutrinos. This is not entirely surprising. Since the neutrino self-interaction potential dominates the total potential, neutrinos and antineutrinos experience the same potential (up to a sign) and the flavor vectors of an individual particle rotate the same amount for neutrinos and antineutrinos, though in opposite directions. Another way to say this is that the total lepton number is approximately preserved when the equations of motion are approximately matter-antimatter symmetric. This is in line with previous predictions derived assuming a negligible contribution from the vacuum potential (e.g. Equation 5 in Martin et al. 72), but in situations where the vacuum potential is relevant, the matter-antimatter symmetry of the evolution equations is broken and this relationship no longer holds.

Equation 18 does not predict the amount of neutrinos that oscillate, but if this result holds in general for fast flavor transformations, it does say that the net electron flavor content (i.e. number of electron neutrinos minus the number of antineutrinos) does not change under the fast flavor instability even in saturation. In addition, it provides a bound on the amount of neutrinos that change flavor. In the limit of zero heavy lepton neutrinos, $n_{ee}(0) - \langle n_{ee}\rangle = \min(n_{ee}(0), \bar{n}_{ee}(0))$, since it is not possible for more electron neutrinos to change flavor than there exist electron neutrinos. We can extend this to three flavors, noting that the amplitude of possible flavor transformation is given by $\sqrt{3}l$, where l is the length of the SU(3) flavor isospin vector (analogous to Equation 11 for n_{ab} instead of ρ_{ab}). That is,

$$|n_{ee}(0) - \langle n_{ee}\rangle| \leq \sqrt{3}\min(l, \bar{l}) \tag{19}$$

and similarly for antineutrinos. Explicitly, when the flavor off-diagonal elements are zero, these can be written

$$l^2 = \frac{1}{6}\left[(n_{ee} - n_{\mu\mu})^2 + (n_{ee} - n_{\tau\tau})^2 + (n_{\mu\mu} - n_{\tau\tau})^2\right] \tag{20}$$

and similarly for antineutrinos. If the antineutrino density flavor vector length is smaller than that of the neutrinos, the flavor rotation of the antineutrinos is not bound, but the flavor rotation of neutrinos must match. Complete flavor equilibration of both neutrinos and antineutrinos is possible only if both flavor vector lengths are equal. Once again, these bounds do not actually predict the final abundances; the abundances of electron neutrinos in Figure 11 show an actual amount of flavor change that is significantly smaller than the bound. For instance, the bound for the fiducial simulation only says that the $\langle n_{ee}\rangle \geq 0$.

In summary, we show that the details of the neutrino distribution affect both the growth rate of the FFI and the final abundances of each neutrino flavor. Although we do not have a general way to map initial conditions to final neutrino abundances, we show that a small ELN crossing can produce minimal flavor change. Clearly, the amount of flavor transformation present in the final saturated state must be mapped out by a more comprehensive parameter sweep. We save this for future work.

V. CONCLUSIONS

We present the new open-source particle-in-cell (PIC) neutrino quantum kinetics code `Emu`. The PIC algorithm evolves the quantum states of individual computational particles by accumulating number density and flux on a background grid and then interpolating the net number density and flux from the background grid to evaluate the potential at each individual neutrino's location. The code is based on the AMReX framework and is performance portable to CPU and GPU hardware. We perform a series of test problems and a resolution study in the appendix.

We use `Emu` to perform simulations of the neutrino fast flavor instability (FFI) in one spatial dimension and two momentum (direction) dimensions. We simulate and analyze in detail a distribution of initially electron neutrinos and antineutrinos with number fluxes in opposite directions and a flux factor of 1/3 (Figures 2-4). We demonstrate the simultaneous evolution of the axial symmetry-preserving and symmetry-breaking modes with growth rates (Figure 6) and wavelengths (Figure 4) that match predictions from linear stability theory (Figure 5). The symmetry-preserving mode has the faster growth rate, resulting in a high degree of axial symmetry during instability growth (Figure 8) that persists into the post-saturation phase (Figure 9). After the instability saturates, spatial fluctuations in neutrino flavor cascade from the length scale corresponding to the wavelength of the fastest growing axial symmetry mode to high wavenumbers (Figure 7) following a simple analytic relationship (Equation 14). The final abundances of neutrinos and antineutrinos in this simulation approach flavor equipartition at late times.

We then vary the initial distribution of antineutrinos from the fiducial simulation in a series of simulations. We show that that the antineutrino number density, flux magnitude, and flux direction all modify the growth rate of the instability (Figure 10). In addition, we show that varying antineutrino flux factor and number density result in different final abundances of each neutrino flavor after the instability saturates (Figure 11). We furthermore show that a small ELN crossing leads to slower instability and little flavor change between the initial and post-saturation states, and that the growth rate and flavor change increase with larger ELN crossings (Figure 11). Although we do not have a way of mapping the initial distributions to final flavor abundances, we do demonstrate that in all of our simulations the same number of neutrinos and antineutrinos change flavor. That is, the final abundances of neutrino flavors imply the final abundances of antineutrino flavors (Equation 18). This relationship results in a bound on the amount of neutrino and antineutrino flavor change (Equation 19), though the simulation results never come close to this bound.

Our initial parameter study based on our highly symmetric fiducial simulation (Figure 10) demonstrated how a sparse exploration of the parameter space can lead to incomplete conclusions. Exploring more variations of the initial conditions (blue curve in Figure 11) revealed that the antineutrino flux factor affects the final abundances of neutrino flavor, even though the parameter sweep in Figure 10 showed all antineutrino flux factors resulting in flavor equilibration. We expect that it will be possible to build a mapping from initial distributions to final neutrino abundances, but this will require many more simulations. One major shortcoming of this work is that the simulations in this work were performed in a single spatial dimension. Although a dense covering of the parameter space using three-dimensional simulations is not currently feasible, it will be important to anchor the 1D simulations with select 3D ones. This, too, we leave to future work.

We also expect that the one-dimensional nature of the simulations in this work will significantly affect the process of kinematic decoherence and cascade of power to small scales due to the artificially reduced number of degrees of freedom. In addition, all of the neutrinos in these simulations were of a single energy. Although we include neutrino mass terms, the neutrino potential in these simulations is much larger than the vacuum potential, so all of these simulations are well outside the realm of slow collective oscillations. Many of these conclusions likely do not apply when the vacuum potential becomes significant. This is yet another avenue for future study.

With these prospects in mind, we note that the `Emu` implementation of modern and established numerical methods from the plasma physics community enables efficient numerical predictions with even modest computational resources. For scale, each of the simulations in the parameter survey in Figure 10 requires about 30 minutes on an NVIDIA P100 GPU. The development of this code was greatly accelerated by open access to the open-source plasma physics code `Warp-X`. This goes to show the importance of open-source software beyond the code's intended purpose. We are hopeful that the development of numerical techniques and open-source software will accelerate the field's convergence to a phenomenological theory of fast flavor transformations.

VI. ACKNOWLEDGEMENTS

We are very grateful to Gail McLaughlin, Samuel Flynn, Evan Grohs, Revathi Jambunathan, Luke Johns, James Kneller, Hiroki Nagakura, and MacKenzie Warren for many useful discussions. This material is based upon work supported by the National Science Foundation under Award No. 2001760. SR is supported by the N3AS Fellowship under National Science Foundation grant PHY-1630782 and Heising-Simons Foundation grant 2017-228. This research was supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy Office of Science and the National Nuclear Security Administration. This work was supported in part by the U.S. Department of

Energy, Office of Science, Office of Workforce Development for Teachers and Scientists (WDTS) under the Science Undergraduate Laboratory Internship (SULI) program. This work used the Bridges system, which is supported by NSF award number ACI-1445606, at the Pittsburgh Supercomputing Center (PSC). This research used the Cori supercomputer 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.

In this work we make use of `SymPy` [101], `NumPy` [102], `MatPlotLib` [103], `SciPy` [104], `yt` [105], and `Mathematica` [106].

[1] Brian D. Metzger, "Kilonovae," Living Reviews in Relativity **20**, 3 (2017).

[2] Nikhil Sarin and Paul D. Lasky, "The evolution of binary neutron star post-merger remnants: a review," arXiv:2012.08172 (2020).

[3] Jennifer Barnes, "The Physics of Kilonovae," Frontiers in Physics **8**, 355 (2020).

[4] David Radice, Sebastiano Bernuzzi, and Albino Perego, "The Dynamics of Binary Neutron Star Mergers and GW170817," Annual Review of Nuclear and Particle Science **70**, 95–119 (2020).

[5] Ondřej Pejcha, "The Explosion Mechanism of Core-Collapse Supernovae and Its Observational Signatures," Reviews in Frontiers of Modern Astrophysics (2020), 10.1007/978-3-030-38509-5_7.

[6] Ernazar Abdikamalov, Giulia Pagliaroli, and David Radice, "Gravitational Waves from Core-Collapse Supernovae," arXiv:2010.04356 (2020).

[7] B. Müller, "Hydrodynamics of core-collapse supernovae and their progenitors," Living Reviews in Computational Astrophysics **6**, 3 (2020).

[8] Adam Burrows, David Radice, David Vartanyan, Hiroki Nagakura, M Aaron Skinner, and Joshua C Dolence, "The overarching framework of core-collapse supernova explosions as revealed by 3D fornax simulations," Monthly Notices of the Royal Astronomical Society **491**, 2715–2735 (2020).

[9] Evan O'Connor, Robert Bollig, Adam Burrows, Sean Couch, Tobias Fischer, Hans-Thomas Janka, Kei Kotake, Eric J Lentz, Matthias Liebendörfer, O E Bronson Messer, Anthony Mezzacappa, Tomoya Takiwaki, and David Vartanyan, "Global comparison of core-collapse supernova simulations in spherical symmetry," Journal of Physics G: Nuclear and Particle Physics **45**, 104001 (2018).

[10] A. Mezzacappa, P. Marronetti, R. E. Landfield, E. J. Lentz, K. N. Yakunin, S. W. Bruenn, W. R. Hix, O.E.Bronson Messer, E. Endeve, J. M. Blondin, and J. A. Harris, "Gravitational-wave signal of a core-collapse supernova explosion of a 15 Msun star," Physical Review D **102**, 023027 (2020).

[11] Bernhard Müller and Vishnu Varma, "A 3D simulation of a neutrino-driven supernova explosion aided by convection and magnetic fields," Monthly Notices of the Royal Astronomical Society: Letters **498**, L109–L113 (2020).

[12] Robert Glas, "Three-dimensional Core-collapse Supernova Simulations with Multidimensional Neutrino Transport Compared to the Ray-by-ray-plus Approximation," The Astrophysical Journal **873**, 19 (2019).

[13] David Radice, Christian D. Ott, Ernazar Abdikamalov, Sean M. Couch, Roland Haas, and Erik Schnetter, "Neutrino-driven convection in core-collapse supernovae: high-resolution simulations," The Astrophysical Journal **820**, 76 (2016).

[14] Hiroki Nagakura, Adam Burrows, David Radice, and David Vartanyan, "Towards an understanding of the resolution dependence of Core-Collapse Supernova simulations," Monthly Notices of the Royal Astronomical Society **490**, 4622–4637 (2019).

[15] André da Silva Schneider, Evan O'Connor, Elvira Granqvist, Aurore Betranhandy, and Sean M. Couch, "Equation of State and Progenitor Dependence of Stellar-mass Black Hole Formation," The Astrophysical Journal **894**, 4 (2020).

[16] H Andresen, E Müller, H-Th Janka, A Summa, K Gill, and M Zanolin, "Gravitational waves from 3D core-collapse supernova models: The impact of moderate progenitor rotation," Monthly Notices of the Royal Astronomical Society **486**, 2238–2253 (2019).

[17] Bernhard Müller, Tobias Melson, Alexander Heger, and Hans-Thomas Janka, "Supernova simulations from a 3D progenitor model – Impact of perturbations and evolution of explosion properties," Monthly Notices of the Royal Astronomical Society **472**, 491–513 (2017).

[18] David Radice, Albino Perego, Francesco Zappa, and Sebastiano Bernuzzi, "GW170817: Joint Constraint on the Neutron Star Equation of State from Multimessenger Observations," The Astrophysical Journal **852**, L29 (2018).

[19] Ben Margalit and Brian D. Metzger, "The Multi-messenger Matrix: The Future of Neutron Star Merger Constraints on the Nuclear Equation of State," The Astrophysical Journal **880**, L15 (2019).

[20] Sherwood Richers, Christian D. Ott, Ernazar Abdikamalov, Evan O'Connor, and Chris Sullivan, "Equation of state effects on gravitational waves from rotating core collapse," Physical Review D **95**, 063019 (2017).

[21] M. R. Mumpower, R. Surman, D.-L. Fang, M. Beard, P. Moller, T. Kawano, and A. Aprahamian, "Impact of individual nuclear masses on r -process abundances," Physical Review C **92**, 035807 (2015).

[22] Sherwood Richers, Hiroki Nagakura, Christian D. Ott, Joshua Dolence, Kohsuke Sumiyoshi, and Shoichi Yamada, "A Detailed Comparison of Multidimensional Boltzmann Neutrino Transport Methods in Core-collapse Supernovae," The Astrophysical Journal **847**, 133 (2017).

[23] M. Liebendorfer, M. Rampp, H.-Th. Janka, and A. Mezzacappa, "Supernova Simulations with Boltzmann Neutrino Transport: A Comparison of Methods," The Astrophysical Journal **620**, 840–860 (2005).

[24] O Just, R Bollig, H-Th Janka, M Obergaulinger, R Glas, and S Nagataki, "Core-collapse supernova simulations

in one and two dimensions: comparison of codes and approximations," Monthly Notices of the Royal Astronomical Society **481**, 4786–4814 (2018).
[25] Sherwood Richers, Daniel Kasen, Evan O'Connor, Rodrigo Fernández, and Christian D. Ott, "Monte Carlo neutrino transport through remnant disks from neutron star mergers," The Astrophysical Journal **813**, 38 (2015).
[26] F. Foucart, M. D. Duez, L. E. Kidder, R. Nguyen, H. P. Pfeiffer, and M. A. Scheel, "Evaluating radiation transport errors in merger simulations using a Monte Carlo algorithm," Physical Review D **98**, 063007 (2018).
[27] Haakon Andresen, Robert Glas, and H.-Thomas Janka, "Gravitational-wave Signals From Three-dimensional Supernova Simulations With Different Neutrino-Transport Methods," arXiv:2011.10499 (2020).
[28] Ermal Rrapaj, "Exact solution of multiangle quantum many-body collective neutrino-flavor oscillations," Physical Review C **101**, 065805 (2020).
[29] Alexey Vlasenko, George M. Fuller, and Vincenzo Cirigliano, "Neutrino quantum kinetics," Physical Review D **89**, 105004 (2014).
[30] Cristina Volpe, "Neutrino quantum kinetic equations," International Journal of Modern Physics E **24**, 1541009 (2015).
[31] S. P. Mikheyev and A. Yu. Smirnov, "Resonant neutrino oscillations in matter," Progress in Particle and Nuclear Physics **23**, 41–136 (1989).
[32] L. Wolfenstein, "Neutrino oscillations and stellar collapse," Physical Review D **20**, 2634–2635 (1979).
[33] L. Wolfenstein, "Neutrino oscillations in matter," Physical Review D **17**, 2369–2374 (1978).
[34] Shunsaku Horiuchi and James P Kneller, "What can be learned from a future supernova neutrino detection?" Journal of Physics G: Nuclear and Particle Physics **45**, 043002 (2018).
[35] Huaiyu Duan, George M. Fuller, and Yong-Zhong Qian, "Collective Neutrino Oscillations," Annual Review of Nuclear and Particle Science **60**, 569–594 (2010).
[36] G. Bellini, L. Ludhova, G. Ranucci, and F. L. Villante, "Neutrino Oscillations," Advances in High Energy Physics **2014**, 1–28 (2014).
[37] Huaiyu Duan and Alexander Friedland, "Self-Induced Suppression of Collective Neutrino Oscillations in a Supernova," Physical Review Letters **106**, 091101 (2011).
[38] Manu George, Meng-Ru Wu, Irene Tamborra, Ricard Ardevol-Pulpillo, and Hans-Thomas Janka, "Fast neutrino flavor conversion, ejecta properties, and nucleosynthesis in newly-formed hypermassive remnants of neutron-star mergers," Physical Review D **102**, 103015 (2020).
[39] E. Grohs, George M. Fuller, and Manibrata Sen, "Consequences of neutrino self interactions for weak decoupling and big bang nucleosynthesis," Journal of Cosmology and Astroparticle Physics **2020**, 001–001 (2020).
[40] Zewei Xiong, Andre Sieverding, Manibrata Sen, and Yong-Zhong Qian, "Potential Impact of Fast Flavor Oscillations on Neutrino-driven Winds and Their Nucleosynthesis," The Astrophysical Journal **900**, 144 (2020).
[41] Alexey Vlasenko and G. C. McLaughlin, "Matter-neutrino resonance in a multiangle neutrino bulb model," Physical Review D **97**, 083011 (2018).
[42] A. Malkus, G. C. McLaughlin, and R. Surman, "Symmetric and standard matter neutrino resonances above merging compact objects," Physical Review D **93**, 045021 (2016).
[43] John F. Cherry, J. Carlson, Alexander Friedland, George M. Fuller, and Alexey Vlasenko, "Halo modification of a supernova neutronization neutrino burst," Physical Review D **87**, 085037 (2013).
[44] John F. Cherry, George M. Fuller, Shunsaku Horiuchi, Kei Kotake, Tomoya Takiwaki, and Tobias Fischer, "Time of flight and supernova progenitor effects on the neutrino halo," Physical Review D **102**, 023022 (2020).
[45] Taiki Morinaga, Hiroki Nagakura, Chinami Kato, and Shoichi Yamada, "Fast neutrino-flavor conversion in the preshock region of core-collapse supernovae," Physical Review Research **2**, 012046(R) (2020).
[46] R. F. Sawyer, "Speed-up of neutrino transformations in a supernova environment," Physical Review D **72**, 045003 (2005).
[47] R. F. Sawyer, "Neutrino Cloud Instabilities Just above the Neutrino Sphere of a Supernova," Physical Review Letters **116**, 081101 (2016).
[48] S. Chakraborty, R. S. Hansen, I. Izaguirre, and G.G. Raffelt, "Self-induced neutrino flavor conversion without flavor mixing," Journal of Cosmology and Astroparticle Physics **2016**, 042–042 (2016).
[49] Ignacio Izaguirre, Georg Raffelt, and Irene Tamborra, "Fast Pairwise Conversion of Supernova Neutrinos: A Dispersion Relation Approach," Physical Review Letters **118**, 021101 (2017).
[50] Francesco Capozzi, Basudeb Dasgupta, Alessandro Mirizzi, Manibrata Sen, and Günter Sigl, "Collisional Triggering of Fast Flavor Conversions of Supernova Neutrinos," Physical Review Letters **122**, 091101 (2019).
[51] Basudeb Dasgupta and Manibrata Sen, "Fast neutrino flavor conversion as oscillations in a quartic potential," Physical Review D **97**, 023017 (2018).
[52] Francesco Capozzi, Georg Raffelt, and Tobias Stirner, "Fast neutrino flavor conversion: collective motion vs. decoherence," Journal of Cosmology and Astroparticle Physics **2019**, 002–002 (2019).
[53] Changhao Yi, Lei Ma, Joshua D Martin, and Huaiyu Duan, "Dispersion relation of the fast neutrino oscillation wave," PHYS. REV. D , 11 (2019).
[54] Meng-Ru Wu and Irene Tamborra, "Fast neutrino conversions: Ubiquitous in compact binary merger remnants," Physical Review D **95**, 103007 (2017).
[55] Irene Tamborra, Lorenz Hüdepohl, Georg G. Raffelt, and Hans-Thomas Janka, "Flavor-dependent Neutrino Angular Distribution in Core-collapse Supernovae," The Astrophysical Journal **839**, 132 (2017).
[56] Sajad Abbar, Huaiyu Duan, Kohsuke Sumiyoshi, Tomoya Takiwaki, and Maria Cristina Volpe, "On the occurrence of fast neutrino flavor conversions in multidimensional supernova models," Physical Review D **100**, 043004 (2019).
[57] M. D. Azari, Shoichi Yamada, Taiki Morinaga, Wakana Iwakami, Hirotada Okawa, Hiroki Nagakura, and Kohsuke Sumiyoshi, "Linear analysis of fast-pairwise collective neutrino oscillations in core-collapse supernovae based on the results of Boltzmann simulations," Physical Review D **99**, 103011 (2019).
[58] M. Delfan Azari, S. Yamada, T. Morinaga, H. Nagakura, S. Furusawa, A. Harada, H. Okawa, W. Iwakami, and K Sumiyoshi, "Fast collective neu-

trino oscillations inside the neutrino sphere in core-collapse supernovae," Physical Review D **101**, 023018 (2020).
[59] S. Abbar, H. Duan, K. Sumiyoshi, T. Takiwaki, and M. C. Volpe, "Fast neutrino flavor conversion modes in multidimensional core-collapse supernova models: The role of the asymmetric neutrino distributions," Physical Review D **101**, 043016 (2020).
[60] Hiroki Nagakura, Taiki Morinaga, Chinami Kato, and Shoichi Yamada, "Fast-pairwise Collective Neutrino Oscillations Associated with Asymmetric Neutrino Emissions in Core-collapse Supernovae," The Astrophysical Journal **886**, 139 (2019).
[61] Robert Glas, H. Thomas Janka, F. Capozzi, M. Sen, B. Dasgupta, A. Mirizzi, and G. Sigl, "Fast neutrino flavor instability in the neutron-star convection layer of three-dimensional supernova models," Physical Review D **101**, 063001 (2020).
[62] Francesco Capozzi, Sajad Abbar, Robert Bollig, and H.-Thomas Janka, "Fast neutrino flavor conversions in one-dimensional core-collapse supernova models with and without muon creation," arXiv:2012.08525 (2020).
[63] Sajad Abbar, Francesco Capozzi, Robert Glas, H.-Thomas Janka, and Irene Tamborra, "On the characteristics of fast neutrino flavor instabilities in three-dimensional core-collapse supernova models," arXiv:2012.06594 (2020).
[64] Meng-Ru Wu, Irene Tamborra, Oliver Just, and Hans-Thomas Janka, "Imprints of neutrino-pair flavor conversions on nucleosynthesis in ejecta from neutron-star merger remnants," Physical Review D **96**, 123015 (2017).
[65] Basudeb Dasgupta, Alessandro Mirizzi, and Manibrata Sen, "Fast neutrino flavor conversions near the supernova core with realistic flavor-dependent angular distributions," Journal of Cosmology and Astroparticle Physics **2017**, 019–019 (2017).
[66] Sajad Abbar and Maria Cristina Volpe, "On fast neutrino flavor conversion modes in the nonlinear regime," Physics Letters B **790**, 545–550 (2019).
[67] Shashank Shalgar and Irene Tamborra, "Dispelling a myth on dense neutrino media: fast pairwise conversions depend on energy," arXiv:2007.07926 (2020).
[68] Lucas Johns, Hiroki Nagakura, George M. Fuller, and Adam Burrows, "Fast oscillations, collisionless relaxation, and spurious evolution of supernova neutrino flavor," Physical Review D **102**, 103017 (2020).
[69] Lucas Johns, Hiroki Nagakura, George M. Fuller, and Adam Burrows, "Neutrino oscillations in supernovae: Angular moments and fast instabilities," Physical Review D **101**, 043009 (2020).
[70] Soumya Bhattacharyya and Basudeb Dasgupta, "Fast Flavor Depolarization of Supernova Neutrinos," arXiv:2009.03337 (2020).
[71] Soumya Bhattacharyya and Basudeb Dasgupta, "Late-time behavior of fast neutrino oscillations," Physical Review D **102**, 063018 (2020).
[72] Joshua D. Martin, Changhao Yi, and Huaiyu Duan, "Dynamic fast flavor oscillation waves in dense neutrino gases," Physics Letters B **800**, 135088 (2020).
[73] Ian Padilla-Gay, Shashank Shalgar, and Irene Tamborra, "Multi-Dimensional Solution of Fast Neutrino Conversions in Binary Neutron Star Merger Remnants," arXiv:2009.01843 (2020).
[74] Shashank Shalgar and Irene Tamborra, "A change of direction in pairwise neutrino conversion physics: The effect of collisions," arXiv:2011.00004 (2020).
[75] Ran Chu, Eirik Endeve, Cory D. Hauck, Anthony Mezzacappa, and Bronson Messer, "thornado-transport: IMEX schemes for two-moment neutrino transport respecting Fermi-Dirac statistics," Journal of Physics: Conference Series **1225**, 012013 (2019).
[76] Luke F. Roberts, Christian D. Ott, Roland Haas, Evan P. O'Connor, Peter Diener, and Erik Schnetter, "General-relativistic three-dimensional multi-group neutrino radiation-hydrodynamics simulations of core-collapse supernovae," The Astrophysical Journal **831**, 98 (2016).
[77] Evan O'Connor, "An open-source neutrino radiation hydrodyanamics code for core-collapse supernovae," The Astrophysical Journal Supplement Series **219**, 24 (2015).
[78] Evan P. O'Connor and Sean M. Couch, "Two-dimensional Core-collapse Supernova Explosions Aided by General Relativity with Multidimensional Neutrino Transport," The Astrophysical Journal **854**, 63 (2018).
[79] N.M.H. Vaytet, E. Audit, B. Dubroca, and F. Delahaye, "A numerical model for multigroup radiation hydrodynamics," Journal of Quantitative Spectroscopy and Radiative Transfer **112**, 1323–1335 (2011).
[80] Francois Foucart, Evan O'Connor, Luke Roberts, Lawrence E. Kidder, Harald P. Pfeiffer, and Mark A. Scheel, "Impact of an improved neutrino energy estimate on outflows in neutron star merger simulations," Physical Review D **94**, 123016 (2016).
[81] Yuichiro Sekiguchi, Kenta Kiuchi, Koutarou Kyutoku, and Masaru Shibata, "Dynamical mass ejection from binary neutron star mergers: Radiation-hydrodynamics study in general relativity," Physical Review D **91**, 064059 (2015).
[82] O. Just, M. Obergaulinger, and H.-T. Janka, "A new multidimensional, energy-dependent two-moment transport code for neutrino-hydrodynamics," Monthly Notices of the Royal Astronomical Society **453**, 3387–3414 (2015).
[83] Hiroki Nagakura, Kohsuke Sumiyoshi, and Shoichi Yamada, "Three-dimensional Boltzmann-hydro Code for Core-collapse in Massive Stars. III. A New Method for Momentum Feedback from Neutrino to Matter," The Astrophysical Journal **878**, 160 (2019).
[84] Matthias Liebendorfer, O. E. Bronson Messer, Anthony Mezzacappa, Stephen W. Bruenn, Christian Y. Cardall, and F.-K. Thielemann, "A Finite Difference Representation of Neutrino Radiation Hydrodynamics in Spherically Symmetric General Relativistic Spacetime," The Astrophysical Journal Supplement Series **150**, 263–316 (2004).
[85] Sherwood Richers, "Rank-3 moment closures in general relativistic neutrino transport," Physical Review D **102**, 083017 (2020).
[86] Francois Foucart, Matthew D. Duez, Francois Hebert, Lawrence E. Kidder, Harald P. Pfeiffer, and Mark A. Scheel, "Monte-Carlo Neutrino Transport in Neutron Star Merger Simulations," The Astrophysical Journal **902**, L27 (2020).
[87] Jonah M. Miller, Benjamin R. Ryan, Joshua C. Dolence, Adam Burrows, Christopher J. Fontes, Christopher L. Fryer, Oleg Korobkin, Jonas Lippuner, Matthew R.

Mumpower, and Ryan T. Wollaeger, "Full transport model of GW170817-like disk produces a blue kilonova," Physical Review D **100**, 023008 (2019).
[88] Chinami Kato, Hiroki Nagakura, Yusuke Hori, and Shoichi Yamada, "Neutrino transport with Monte Carlo method: I. Towards fully consistent implementation of nucleon recoils in core-collapse supernova simulations," The Astrophysical Journal **897**, 43 (2020).
[89] Yunfan Zhang and Adam Burrows, "Transport equations for oscillating neutrinos," Physical Review D **88**, 105009 (2013).
[90] Sherwood A. Richers, Gail C. McLaughlin, James P. Kneller, and Alexey Vlasenko, "Neutrino quantum kinetics in compact objects," Physical Review D **99**, 123014 (2019).
[91] Donald Willcox and Sherwood Richers, "AMReX-Astro/Emu: Emu v1.1.1 (Version v1.1.1)," Zenodo (2021), http://doi.org/10.5281/zenodo.4609363.
[92] Sherwood Richers, Donald Willcox, Nicole Ford, and Andrew Myers, "Particle-in-cell Simulation of the Neutrino Fast Flavor Instability [Data set]," Zenodo (2021), 10.5281/zenodo.4423095.
[93] J.-L. Vay, A. Almgren, J. Bell, L. Ge, D.P. Grote, M. Hogan, O. Kononenko, R. Lehe, A. Myers, C. Ng, J. Park, R. Ryne, O. Shapoval, M. Thévenet, and W. Zhang, "Warp-X: A new exascale computing platform for beam–plasma simulations," Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **909**, 476–479 (2018).
[94] B. Wang, G. H. Miller, and P. Colella, "A Particle-in-cell Method with Adaptive Phase-space Remapping for Kinetic Plasmas," SIAM Journal on Scientific Computing **33**, 3509–3537 (2011).
[95] Essex Edwards and Robert Bridson, "A high-order accurate particle-in-cell method," International Journal for Numerical Methods in Engineering **90**, 1073–1088 (2012).
[96] "Emu Illustrations," (2021), publisher: figshare.
[97] Weiqun Zhang, Ann Almgren, Vince Beckner, John Bell, Johannes Blaschke, Cy Chan, Marcus Day, Brian Friesen, Kevin Gott, Daniel Graves, Max Katz, Andrew Myers, Tan Nguyen, Andrew Nonaka, Michele Rosso, Samuel Williams, and Michael Zingale, "AMReX: a framework for block-structured adaptive mesh refinement," Journal of Open Source Software **4**, 1370 (2019).
[98] Ali Katanforoush and Mehrdad Shahshahani, "Distributing Points on the Sphere, I," Experimental Mathematics **12**, 199–209 (2003).
[99] Madhurima Chakraborty and Sovan Chakraborty, "Three flavor neutrino conversions in supernovae: slow & fast instabilities," Journal of Cosmology and Astroparticle Physics **2020**, 005–005 (2020).
[100] Joshua D. Martin, Sajad Abbar, and Huaiyu Duan, "Nonlinear flavor development of a two-dimensional neutrino gas," Physical Review D **100**, 023016 (2019).
[101] Aaron Meurer, Christopher P. Smith, Mateusz Paprocki, Ondřej Čertík, Sergey B. Kirpichev, Matthew Rocklin, AMiT Kumar, Sergiu Ivanov, Jason K. Moore, Sartaj Singh, Thilina Rathnayake, Sean Vig, Brian E. Granger, Richard P. Muller, Francesco Bonazzi, Harsh Gupta, Shivam Vats, Fredrik Johansson, Fabian Pedregosa, Matthew J. Curry, Andy R. Terrel, Štěpán Roučka, Ashutosh Saboo, Isuru Fernando, Sumith Kulal, Robert Cimrman, and Anthony Scopatz, "SymPy: symbolic computing in Python," PeerJ Computer Science **3**, e103 (2017), publisher: PeerJ Inc.
[102] S. van der Walt, S. C. Colbert, and G. Varoquaux, "The NumPy Array: A Structure for Efficient Numerical Computation," Computing in Science Engineering **13**, 22–30 (2011), conference Name: Computing in Science Engineering.
[103] J. D. Hunter, "Matplotlib: A 2D Graphics Environment," Computing in Science Engineering **9**, 90–95 (2007), conference Name: Computing in Science Engineering.
[104] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C. J. Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, and Paul van Mulbregt, "SciPy 1.0: fundamental algorithms for scientific computing in Python," Nature Methods **17**, 261–272 (2020), number: 3 Publisher: Nature Publishing Group.
[105] Matthew J. Turk, Britton D. Smith, Jeffrey S. Oishi, Stephen Skory, Samuel W. Skillman, Tom Abel, and Michael L. Norman, "yt: a multi-code analysis toolkit for astrophysical simulation data," The Astrophysical Journal Supplement Series **192**, 9 (2011).
[106] Wolfram Research Inc., "Mathematica," (2020).
[107] Steen Hannestad, Georg G Raffelt, Gunter Sigl, and Yvonne Y Y Wong, "Self-induced conversion in dense neutrino gases: Pendulum in flavor space," Physical Review D **74**, 105010 (2006).

Appendix A: Test Problems

1. Mikheyev-Smirnov-Wolfenstein (MSW) Oscillations

To test our implementation of the vacuum and matter potentials we follow [90] and perform a simple two-flavor one-zone test of neutrino oscillations in a regime where the neutrino potential is negligible and the matter density is tuned to put the neutrino oscillations in resonance. Specifically, we set up a 1 cm cube in a $1 \times 1 \times 1$ domain filled with matter at a density of $1.35 \times 10^9\,\mathrm{g\,cm^{-3}}$ and an artificial electron fraction of $Y_e = 1$. The neutrino mass eigenstates have masses of $m_1 = 8.60 \times 10^{-3}\,\mathrm{eV}$ and $m_2 = 0\,\mathrm{eV}$, and the mixing angle is set to $\theta_{12} = 33.82$ degrees. We only initialize two particles, each of which represents a single neutrino and a single antineutrino at an energy of $h\nu = (m_2^2 - m_1^2)c^4 \sin(2\theta_{12})/(8\pi\hbar c) = 0.138\,\mathrm{eV}$. This combination of background matter and neutrino properties make the antineutrinos undergo complete flavor oscillations with a period of $ct = 1\,\mathrm{cm}$. The flavor-diagonal values of the density matrix of each neutrino and antineutrino are plotted as colored curves in Fig-

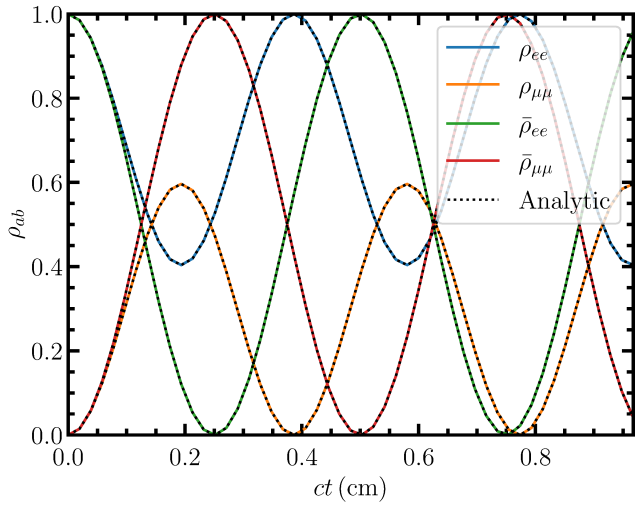

FIG. 12. A two-flavor (ν_e and ν_μ) one-zone MSW test. We set the background matter, neutrino masses, and flavor mixing angles such that antineutrinos undergo complete flavor oscillations with a period of $ct = 1\,\text{cm}$. The density matrix of each (anti-)neutrino is shown as a function of time, with the electron flavor in blue (green) and the muon flavor in orange (red). The analytic results are overplotted as dotted black curves and show good agreement. The antineutrinos undergo a complete transformation, but since the neutrinos are not in resonance they do not completely swap flavors.

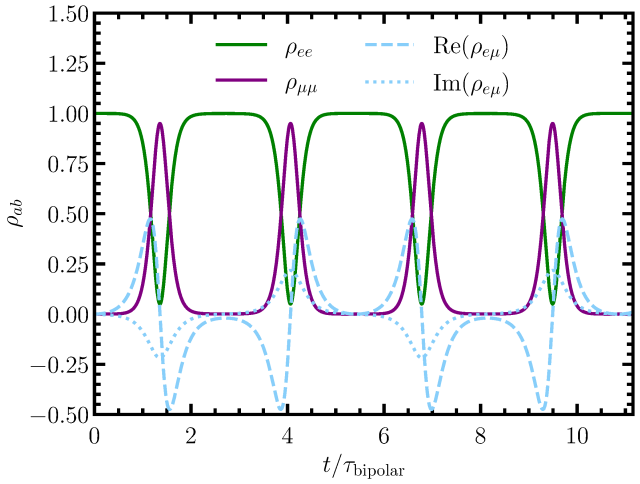

FIG. 13. An isotropic one-zone bipolar oscillation test using two counter-propagating beams. We set the mixing angle θ_{12} and background matter density such that $\mathcal{H}_{\text{neutrino}}/\mathcal{H}_{\text{vacuum}} = 10$ with characteristic oscillation time $\tau_{\text{bipolar}} = 8.96 \times 10^{-4}\,\text{s}$. The solid curves show the evolution of the diagonal components of the neutrino density matrix (electron flavor in green, muon flavor in purple). The dashed and dotted blue curves show the evolution of the real and imaginary parts of the off-diagonal element, respectively. Our results match those of [107] and [90].

ure 12 and the analytic results (e.g. [36]) are overplotted as dotted black curves. The antineutrino values (green and red) undergo complete transformation, but since the neutrinos (blue and orange) are not in resonance they do not completely swap. The maximum difference between the analytic and simulated quantities is 3.7×10^{-4} when using a timestep of $6.5 \times 10^{-13}\,\text{s}$.

2. Bipolar Oscillations

To test our implementation of the isotropic components of the self-interaction potentials we follow [90] and perform an isotropic bipolar oscillation test. Once again, we use only one particle moving in the $+z$ direction and one moving in the $-z$ direction, since this exactly replicates the behavior of an isotropic distribution in the context of this test. We choose particle weights to ensure the neutrino and antineutrino number densities are $n = \bar{n} = 10(m_2 - m_1)^2c^4/(2\sqrt{2}G_F E)$ with energy $E = 50$ MeV. We use the same neutrino masses as in Section A 1, but we instead choose a mixing angle of $\theta_{12} = 0.01 = 0.573\,\text{degrees}$ and a background mass density of zero. This corresponds to the $\mathcal{H}_{\text{neutrino}}/\mathcal{H}_{\text{vacuum}} = 10$ case shown in Ref. [107] with a characteristic oscillation time of $\tau = 8.96 \times 10^{-4}\,\text{s}$ (see their Equation 21). The solid curves in Figure 13 show the evolution of the diagonal components of the neutrino density matrix, while the dashed and dotted curves show the evolution of the real and imaginary parts of the off-diagonal element, respectively. We do not perform a detailed error calculation since the analytic result in Ref. [107] is itself approximate, but the curves in Figure 13 match those in Ref. [107] and [90]. Note that unlike [90] we use an inverted mass hierarchy, so we start in an electron flavor eigenstate instead of a muon flavor eigenstate.

3. Homogeneous Two Beam Fast Flavor Instability

To test the anisotropic terms in the self-interaction potential, we perform a two-flavor two-beam fast flavor conversion test. The growth rate of the fast flavor instability in this simplified setup is easily calculated analytically as described by [48]. We use the same neutrino masses as in Section A 1, but choose a mixing angle of $\theta_{12} = 10^{-6}$ degrees. We use a single grid cell containing one particle moving in the $+z$ direction representing only neutrinos, one moving in the $-z$ direction representing only antineutrinos, and a vacuum matter background. The number of neutrinos each particle represents is set such that the total number density of each of neutrinos and antineutrinos is $n = (m_2^2 - m_1^2)c^4/(4\sqrt{2}G_F h\nu) = 2.92 \times 10^{24}\,\text{cm}^{-3}$ (where G_F here has units of $\text{erg}\,\text{cm}^3$). This yields the fastest growing instability given the neutrino parameters according to Equation 2.10 in Ref. [48].

The results are shown in Figure 14. The blue curve on top shows the evolution of the electron diagonal compo-

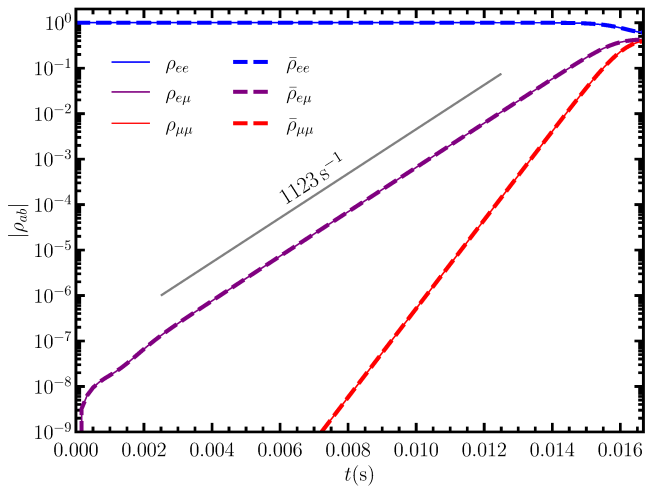

FIG. 14. Homogeneous ($k = 0$) two-flavor two-beam fast flavor instability prepared in an initially pure electron flavor state. The instability grows exponentially, as electron (blue) neutrinos (solid) and antineutrinos (dashed) rotate toward the pure muon state (red) and a quantum superposition of electron and muon flavor states (purple). The growth rate of the flavor off-diagonal component matches the prediction from linear stability theory well (gray line segment).

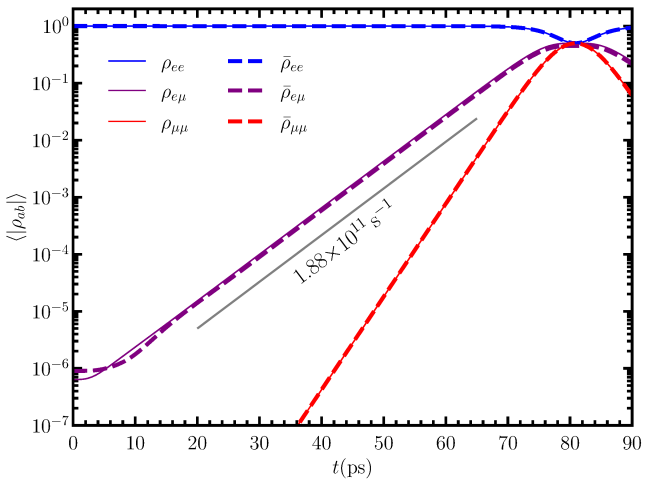

FIG. 15. Inhomogeneous ($k \neq 0$) two-flavor two-beam fast flavor instability prepared in an initially nearly pure electron flavor state. The instability grows exponentially, as electron (blue) neutrinos (solid) and antineutrinos (dashed) rotate toward the pure muon state (red) and a quantum superposition of electron and muon flavor states (purple). The growth rate of the flavor off-diagonal component matches the prediction from linear stability theory well (gray line segment).

nent of the density matrix for neutrinos (solid) and antineutrinos (dashed). As time progresses, the neutrinos mix into the the muon flavor state (red) and a quantum superposition of muon and electron flavor states (purple). Linear stability theory predicts an exponential increase in the off-diagonal component with a growth rate of $\mathrm{Im}(\omega) = (m_2^2 - m_1^2)c^4/(2\hbar h\nu) = 1123\,\mathrm{s}^{-1}$. The fitted growth rates between $t = 8.4\,\mathrm{ms}$ and $11.7\,\mathrm{ms}$ match the theoretical rate with a relative error of at most 3.2×10^{-5}.

4. Inhomogeneous Fast Flavor Instability

Finally, we perform a simulation of the inhomogeneous two-beam two-flavor fast flavor instability to test all of the previous components of the code together with the spatial advection of neutrinos. We choose the same neutrino masses and mixing angle as in Section A 3. In order to get the wavelength of the fastest growing mode to match the size of our box (i.e. $k = 2\pi/(1\,\mathrm{cm})$), we set each of the neutrino and antineutrino number densities to

$$n = \frac{(m_2^2 - m_1^2)c^4/(2h\nu) + \hbar ck}{2\sqrt{2}G_F} = 4.89\times10^{32}\,\mathrm{cm}^{-3} \tag{A1}$$

such that the fastest growing mode has a wavelength of 1 cm for a neutrino energy of 50 MeV (see [48] Equation 2.7). To isolate the fastest growing mode, we also perturb the initial conditions using $\mathrm{Re}(f_{e\mu}) = \mathrm{Re}(\bar{f}_{e\mu}) = 10^{-6}\sin(kz_0)$, where z_0 is the initial position of the particle. The results are shown in Figure 15. As for the homogenous case in Section A 3, we see neutrinos (solid) and antineutrinos (dashed) mix from the nearly pure initial electron flavor state (blue) into the muon flavor state (red) and a quantum superposition of muon and electron flavor states (purple). We see the growth of the off-diagonal component is exponential, with a growth rate matching the prediction of linear stability theory of $\mathrm{Im}(\omega) = (m_2^2 - m_1^2)c^4/(2\hbar h\nu) + ck = 1.88\times10^{11}\,\mathrm{s}^{-1}$ between $t = 33\,\mathrm{ps}$ and $53\,\mathrm{ps}$ with a relative error of at most 4.3×10^{-3}.

Appendix B: Convergence

The calculations we perform here are on a very small length and time scale (typically 64 cm domain size) compared to a supernova or neutron star merger. Thus, the quantity we are primarily interested in is the net amount of electron neutrinos that remain after the instability saturates, averaged over the entire domain and over time. We compute the spatial-temporal average of the neutrino and antineutrino density matrices $\langle N_{ab}\rangle$ and $\langle \bar{N}_{ab}\rangle$ as

$$\langle N_{ab}\rangle = \frac{\sum_{\mathrm{i}=1}^{N_{\mathrm{zones}}}\sum_{\mathrm{j}=\mathrm{j}_{\mathrm{start}}}^{N_{\mathrm{snapshots}}} n_{ab,ij}}{N_{\mathrm{zones}}(N_{\mathrm{snapshots}} - j_{\mathrm{start}} + 1)} \tag{B1}$$

where i labels a grid zone and j labels the simulation snapshot. j_{start} is the first snapshot after $t = 3\,\mathrm{ns}$. This allows us to average only over times unaffected by the initial transitory phase of the instability growing and saturating. After saturation, the system is chaotic and thus impossible to demonstrate exact convergence. However,

since the quantity significant for core-collapse and merger simulations is precisely $\langle N_{ab} \rangle$, this is the quantity we ensure is converged.

There are three parameters dictating the fidelity of these simulations. The first is the domain size of the simulation. Larger domains can host modes with longer wavelengths, which even if they are not the fastest growing unstable modes, they may still be unstable or interact with shorter-wavelength modes after saturation. The second is the number of grid cells. The background grid discretizes the neutrino potentials, so finer grid cells are required to allow modes with shorter wavelengths to be represented in the background potential. Finally, there is the number of computational particles per cell. More particles means that a larger number of unique particle directions are represented on the grid and modes with finer angular structure can be expressed by the simulation. We use only one particle per direction per cell, since we have found that placing multiple particles at different locations in each cell but moving in the same direction does not result in a solution significantly different from the one calculated using only one particle per direction per cell. In the simulations discussed below, we initialize the neutrino distribution as described in Section III A to assess what simulation parameters lead to a numerically converged result.

The top left panel of Figure 16 shows the space- and time-averaged fraction of ν_e in the post-saturation state under variations in the number of grid zones for a fixed domain size of 64 cm. Using eight or fewer grid zones causes the instability to never develop, leaving the final fraction of electron neutrinos equal to the initial conditions. As the number of grid zones increases beyond that, the saturated state appears to move increasingly toward an evenly mixed flavor distribution until 32 grid zones. Then the separation between the flavors increases slightly until 1024 grid zones, beyond which the results are mostly constant. Note that the numbers shown here are the values averaged over a simulation with only a duration of 5×10^{-9} s, so this gap between electron and heavy lepton neutrinos is at least in part a result of not allowing the simulations to extend until momentum-space decoherence is complete. In the bottom left panel, we plot the fitted growth rate of the axially symmetric (blue) and asymmetric (red) modes along with the corresponding theoretical results from a linear dispersion analysis (gray horizontal lines). The growth rate of the symmetric mode appears to asymptote to the theoretical value. The growth rate of the asymmetric mode approaches the theoretical value until $n_z = 256$ and then appears to slowly decrease. We expect that this is due in part to contamination from the underlying random perturbations, since the amplitude of the asymmetric mode is never large, and in part due to the limited resolution. The asymmetric mode does not significantly contribute to the post-saturation state, so we leave a more finely resolved simulation of the asymmetric mode to future work.

A similar progression occurs for changes in the domain size while keeping the size of each grid cell constant (middle panels in Figure 16). If the domain has an extent of only 1cm (4 grid cells) only a weak instability occurs, since the grid cannot contain the fastest growing mode. The amount of transformation over the course of the simulation varies up until a domain size of 64 cm, after which it remains roughly constant. The growth rate of the symmetric mode is nearly correct at a 2 cm domain size, since it can nearly fit the fastest growing mode with a wavelength of 2.2 cm. There are once again small variations until a domain size of 64 cm, after which the growth rate remains relatively constant. The growth rate of the asymmetric mode once again proves to be difficult, peaking at a domain size of 16 cm and then slowly decreasing with increasing domain size. We expect this to be the result of the same reasons as for the behavior in the n_z convergence test above.

Finally, we test the robustness of the results while varying the number of particles per cell, which also allows the particles to probe a larger number of directions. The leftmost point in the right panels of Figure 16 represents a two-beam calculation. In each grid cell one particle is moving in the $+z$ direction and one particle is moving in the $-z$ direction. For the rest of the points, the horizontal axis describes the number of particle directions along the equatorial $(x - y)$ plane. Particles are given directions according to Section III A. 4, 8, 16, 32, 64, and 128 equatorial directions correspond to 6, 24, 92, 378, 1056, and 6022 particles per cell, respectively. The amount of post-saturation flavor transformation (top panel) does not appear to be sensitive to the angular resolution of the simulation beyond 8 equatorial directions. The symmetric mode growth rate appears to vary little after 16 equatorial directions. Once again, the asymmetric mode growth rate is more difficult to resolve, but approaches the theoretical line with increasing angular resolution. It is interesting that such sparse sampling of the angular distribution can yield such accurate simulations of the FFI. This is in contrast to simulations of collective oscillations that require thousands of angular bins [37].

We also investigate our integration method's order of space-time convergence. We use the data set shown in the leftmost column of Figure 16 where we kept the extent of the domain constant while varying the number of grid zones, thus varying the grid resolution Δz. We measure convergence by measuring the instability growth rate during the linear phase in the time interval $[0.15, 0.25]$ ns, firmly centered in the linear regime as shown in Figure 6. We take as a reference point for computing the convergence error the highest resolution simulation shown in Figure 16 corresponding to $n_z = 2048$ grid cells in the domain. We denote the grid resolution of our $n_z = 2048$ simulation as Δz_{2048} and each other grid resolution where we measured convergence error as Δz. We then denote the convergence error of the growth rates in the lower-resolution simulations as $\delta\omega = |\mathrm{Im}(\omega - \omega_{2048})|$ and compute this error for the same neutrino densities n_{ab} and fluxes $\mathbf{f}_{ab}^{(x)}$ shown growing in Figure 6. In the top panel of

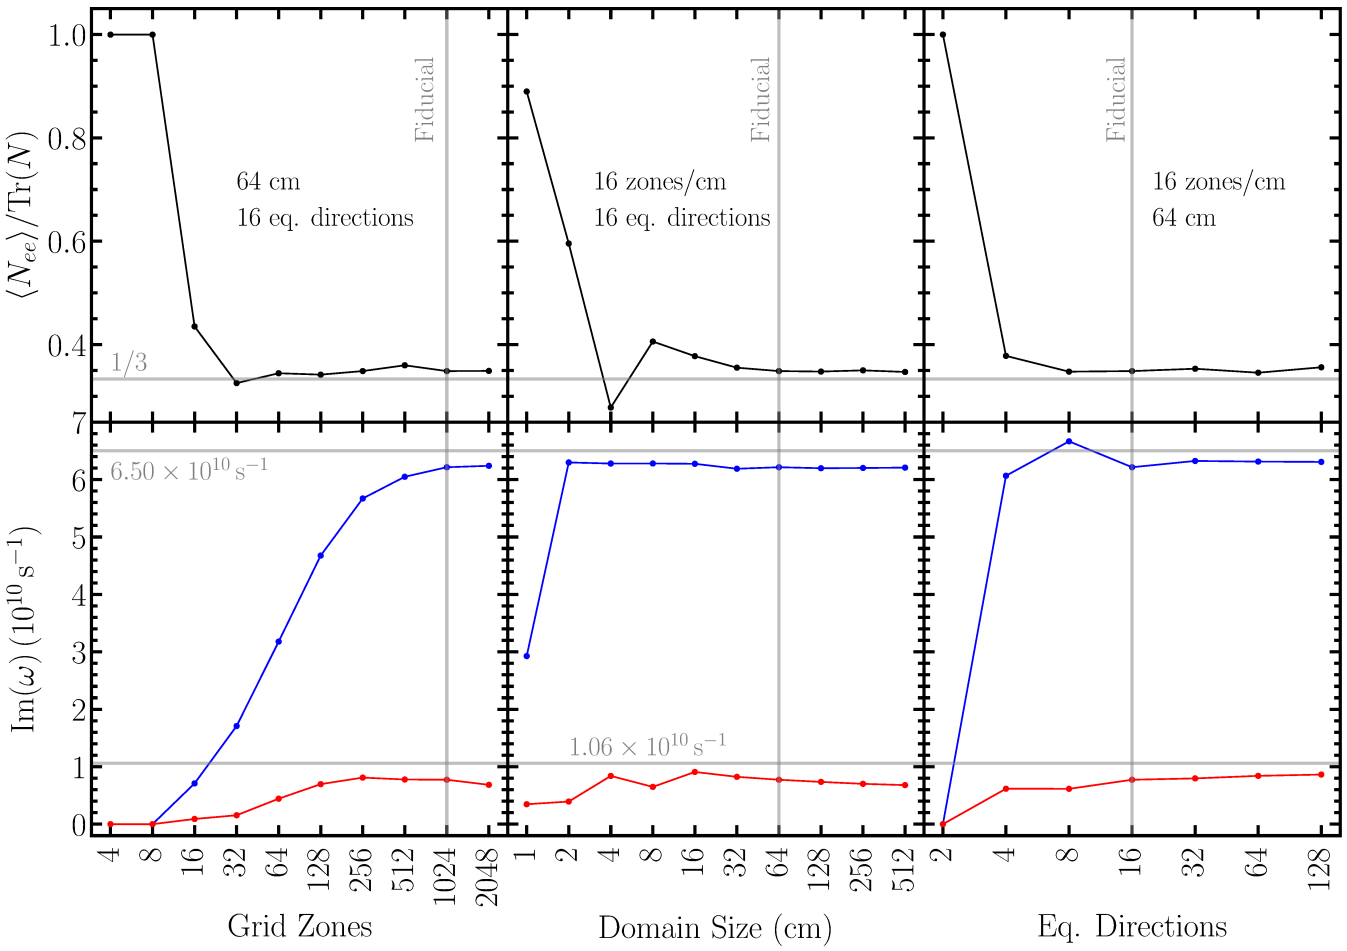

FIG. 16. Convergence of the post-saturation electron neutrino flavor transformation (top) and fitted growth rate of the axially symmetric (bottom, blue) and asymmetric (bottom, red) modes in the fiducial simulation setup. The left column shows a varying number of grid cells covering a fixed domain of 64 cm, the center column shows a varying domain size with a fixed number of grid cells per centimeter, and the right column shows a varying number of particles per cell. This number of directions is expressed in terms of the number of particle directions on the equator (in the $x-y$ plane). The total number of particles per cell is approximately C^2/π, where C is the number of equatorial directions. 16 cells/cm, a domain size of 64 cm (1024 total cells), and 16 equatorial directions (92 particles per cell) are sufficient for convergence of the averaged post-saturation electron flavor content.

Figure 17, we shift all logarithmic growth rate errors of each component of n_{ab} by the minimum logarithmic error for that component across all resolutions. This permits us to clearly compare our achieved convergence order to representative lines (dashed) for first (red), second (magenta), third (cyan), and fourth (blue) order convergence rates. We see that towards high resolutions, density components $n_{e\mu}$ (purple), $n_{\mu\tau}$ (yellow), and $n_{\mu\mu}$ (brown) all converge to second order or better. Exceptions are $n_{e\tau}$ (green) and $n_{\tau\tau}$ (salmon), which converge between first and second order. We note that although we are using a fourth order Runge-Kutta method as our time integrator, the PIC deposition and interpolation shape functions we use reduce the achievable space-time convergence to second order, as we discuss in Section II B 1.

In the bottom panel of Figure 17 we next show the unscaled convergence error $\delta\omega$ for the axially asymmetric flux $\mathbf{f}_{ab}^{(x)}$ with the same colors representing each component as in the top panel. Here we see that at higher resolutions, $\mathbf{f}_{\mu\tau}^{(x)}$, $\mathbf{f}_{\mu\mu}^{(x)}$, and $\mathbf{f}_{\tau\tau}^{(x)}$ converge to second order, although $\mathbf{f}_{e\mu}^{(x)}$ and $\mathbf{f}_{e\tau}^{(x)}$ do not. We note nevertheless that at the fiducial resolution of $n_z = 1024$ (the leftmost set of points in the lower panel of Figure 17), the size of the error $\delta\omega$ corresponds to $\approx 10\%$ error compared to the axially asymmetric growth rate controlling $\mathbf{f}_{e\mu}^{(x)}$ and $\mathbf{f}_{e\tau}^{(x)}$. The relative error for other components is lower, as other components grow at either the symmetric growth rate or twice that rate. Convergence studies that explore the joint grid resolution and angular resolution space may shed additional light on what determines the convergence requirements for $\mathbf{f}_{e\mu}^{(x)}$ and $\mathbf{f}_{e\tau}^{(x)}$, and we leave this for future work.

These results justify our choice of fiducial simulation

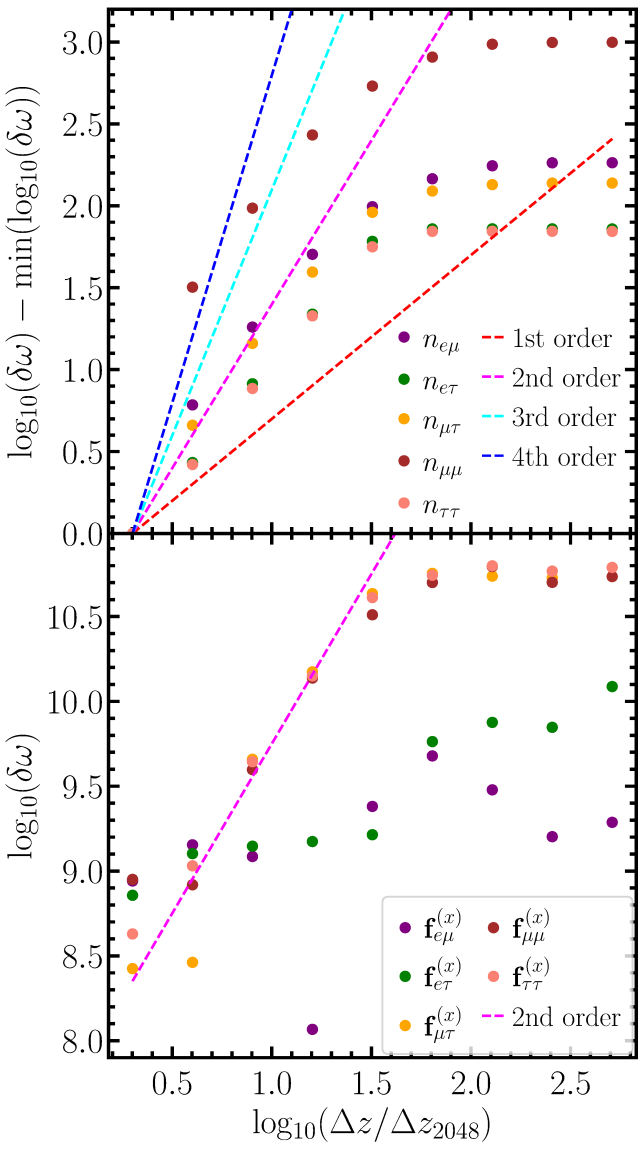

FIG. 17. Convergence of the error in the growth rate of the neutrino number density (top panel) and $\hat{x}$ component of the flux (bottom panel) with increasing grid cell size Δz relative to the cell size in the $n_z = 2048$ simulation ($\Delta z_{2048} = 0.03125\,\mathrm{cm}$). Here $\delta\omega = |\mathrm{Im}(\omega - \omega_{2048})|$, where ω_{2048} is the fitted growth rate in the $n_z = 2048$ simulation. The color of each point denotes the component of the flavor matrix. For reference, scaling relations for convergence of order 1, 2, 3, and 4 are shown as dashed lines. `Emu` achieves better than first order convergence for the growth rate of all distribution quantities except $\mathbf{f}_{e\mu}$ and $\mathbf{f}_{e\tau}$.

parameters. However, it should still be noted that these simulations are too short to probe the long-term behavior of the post-saturation nonlinear evolution.