Neutrino fast flavor instability in three dimensions

Sherwood Richers^{✉*}

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

Donald Willcox and Nicole Ford

Computational Research Division, Lawrence Berkeley National Lab, California 94720, USA

(Received 17 September 2021; accepted 26 October 2021; published 18 November 2021)

Neutrino flavor instabilities have the potential to shuffle neutrinos between electron, mu, and tau flavor states, possibly modifying the core-collapse supernova mechanism and the heavy elements ejected from neutron star mergers. Analytic methods indicate the presence of so-called fast flavor transformation instabilities, and numerical simulations can be used to probe the nonlinear evolution of the neutrinos. Simulations of the fast flavor instability to date have been performed assuming imposed symmetries. We perform simulations of the fast flavor instability that include all three spatial dimensions and all relevant momentum dimensions in order to probe the validity of these approximations. If the fastest growing mode has a wave number along a direction of imposed symmetry, then the instability can be suppressed. The late-time equilibrium distribution of flavor, however, seems to be little affected by the number of spatial dimensions. This is a promising hint that the results of lower-dimensionality simulations to date have predictions that are robust against their the number of spatial dimensions, though simulations of a wider variety of neutrino distributions need to be carried out to support this claim more generally.

DOI: [10.1103/PhysRevD.104.103023](https://doi.org/10.1103/PhysRevD.104.103023)

I. INTRODUCTION

Neutrinos are produced in immense numbers in core-collapse supernovae and neutron star mergers, but the neutrino's elusive nature and behavior currently limits our understanding of these explosive astrophysical phenomena. Core-collapse supernovae, produced when a massive star collapses after exhausting its nuclear fuel, rely on neutrinos to carry energy from the collapsed core to the infalling material below the shock front, and do so on a short enough timescale that they are able to propel the shock through the star (see, e.g., [1,2] for reviews). When two neutron stars or a neutron star and a black hole merge, the neutrinos emitted from the resulting hot accretion disk can enhance outflows that form heavy elements (see [3] for a recent review). In both cases, the matter ejected pollutes the surrounding environment with metals that later form more metal-rich stars, planets, and life. Furthermore, although mergers are the prime candidate source of some of the heaviest elements in the Universe, neutrinos play a deciding role in determining whether the ejecta is

sufficiently neutron rich to efficiently form these elements (e.g., [4,5]).

Since we cannot directly see the interiors of supernovae and mergers and have only seen neutrinos from a single supernova [6–8], we resort to simulation to evaluate whether fundamental physics as it is understood today is capable of reproducing what we see in nature. Understanding supernova and merger dynamics without allowing neutrino flavor change is still an extremely active and productive field, as some of the arguments that neutrino flavor oscillations are not deeply important for the dynamics of supernovae have been rather convincing (e.g., [9,10]). Promisingly, there is increasingly strong evidence that such simulations may in fact be capable of yielding explosions with energies trending toward those observed in nature (e.g., [11–14]). Neutron star merger simulations are also becoming increasingly sophisticated tools for predicting the amount and composition of ejecta from compact object mergers, their electromagnetic and gravitational wave signals, and the type of compact central object left behind (e.g., [15–21]). These simulations have matured to include effects of general relativity, inviscid hydrodynamics, magnetic fields (or viscosity mimicking the action of magnetic fields), a dense nuclear equation of state, nuclear reactions, and transport of energy by neutrinos.

However, understanding supernovae and mergers may not be a simple matter of general relativistic neutrino radiation magnetohydrodynamics; we must also include

*srichers@berkeley.edu

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

the quantum nature of neutrinos in this already long list of relevant physics. There are three known species (flavors) of neutrinos, and a given neutrino is generally in a quantum superposition of all three states. Terrestrial experiments and observations of solar neutrinos show that this quantum state changes as the neutrino propagates according to the Schrödinger equation (e.g., [22], among many other detections since). Contributing to the Hamiltonian that drives this flavor change are the neutrino masses, the interactions of neutrinos with the background matter, and the interaction of neutrinos with other neutrinos (see, e.g., [23]). The latter of these is special in that it leads to a fascinating and poorly understood nonlinear evolution of the flavor whenever neutrino-neutrino interactions are significant. Since different neutrino flavors have significantly different rates of interaction with matter, flavor-changing neutrinos can complicate our already limited understanding of the dynamics of supernovae and neutron star mergers (e.g., [20,24–27]).

Recent developments have suggested the exciting possibility that neutrinos are unstable to flavor change deep in the engine of a supernova or neutron star merger [28,29]. In fact, these instabilities are likely ubiquitous in both systems [30–32]. The so-called fast flavor instability simply requires that, at a given location, there be an overabundance of neutrinos in some directions and an overabundance of antineutrinos moving in other directions [33–35]. Such conditions imply the “crossing” of the energy-integrated lepton differential number density through zero, since antineutrinos carry negative lepton number. This differential number density is often referred to as electron lepton number (ELN), though the argument is equally valid for any flavor. Searches for these conditions in data from numerical simulations of both events that neglect flavor transforming physics have suggested that there are unstable conditions outside the supernova shock, inside the supernova shock, inside the turbulence supernova core, and around the accretion disk from a neutron star merger [25–27,30,31,36–44].

A first-principles global simulation of neutrino quantum kinetics in a supernova or a merger event remains intractable. However, it is still valuable to see how large an effect neutrino flavor transformation could have based on some computationally feasible prescription so as to help bound the realm of possibility. Wu *et al.* [25] postprocess simulation snapshots and tracer particle trajectories from a simulation of a neutron star-black hole (NS-BH) merger disk, modifying the neutrino interaction rates in locations with unstable neutrino distributions and assuming complete flavor mixing. The result was that the flavor conversion resulted in a more neutron-rich neutrino-driven wind and the formation of more heavy elements, though a similar analysis of a NS-NS merger disk showed the opposite [26]. Li and Siegel [20] use the assumption of complete flavor mixing for unstable neutrino distributions and dynamically couple it to the hydrodynamics in simulations of a neutron

star merger disk (a scenario similar to [25]). This, too, resulted in a more neutron-rich outflow than when neutrinos are assumed not to change flavor. Xiong *et al.* [27] imposed multiple flavor transformation prescriptions on parametric models of the neutrino-driven wind following a core-collapse supernova. The flavor transformation resulted in more mass lost and a more proton-rich ejecta. Although a more robust parametric way to describe neutrino flavor transformation is needed, these results already show that neutrino flavor transformation can have a significant impact on galactic chemical evolution, supernova neutrino signals, and the properties of compact object mergers inferred from observation.

To understand the effects of the fast flavor instability, recent efforts have focused on directly simulating neutrino quantum kinetics. To make the problem computationally tractable, these simulations are performed using some combination of a limited physical domain, imposed symmetries, fewer neutrino flavors, and finite resolution. Dasgupta and Mirizzi [45] and Capozzi *et al.* [46] assume a two-beam model, where neutrinos can only move radially inward or radially outward. Abbar *et al.* [47] and Abbar and Volpe [48] use a “line model” as a toy geometry that permits one free spatial and one free angular dimension. Padilla-Gay *et al.* [49] use a line-model-like symmetry but also allow for temporal variations. References [50–57] allow asymmetry in only one momentum dimension, assuming homogeneity in the spatial dimensions. In [58] they extend this to two momentum dimensions. Mirizzi *et al.* [59] use a small number of neutrino directions, assume spatial translational symmetry in one direction and periodic symmetry in the other, and allow neutrinos to move in both dimensions parallel to an emitting neutrino plane. Bhattacharyya and Dasgupta [60] and Wu *et al.* [61] assume translational symmetry in two spatial directions and azimuthal symmetry in neutrino momentum around the “free” spatial direction, and later extend the method to consider two spatial dimensions and two momentum dimensions [62]. Richers *et al.* [63] assume translational symmetry in two dimensions and periodic boundaries in the third, and allow both neutrino direction dimensions.

In this work, we attempt to take some of the guesswork out of interpreting whether simulation results are valid when symmetries are imposed. We simulate all dimensions (three spatial, two direction dimensions) relevant for the fast flavor instability where the neutrino potential dominates other potentials, imposing only periodic boundary conditions in all directions. In Sec. II we review the particle-in-cell method for neutrino quantum kinetics and describe enhancements made to the code since we published the one-dimensional results in [63]. We describe the three-dimensional nature of the instability growth, instability saturation, and eventual quasiequilibrium in Sec. III. Finally, we conclude with Sec. IV and describe the major results from this work that can be applied to global

simulations of neutron star mergers and supernovae. We discuss the most important aspects of numerical convergence in the main text, but provide additional convergence details in the Appendix.

II. METHODS

We simulate neutrino flavor transformation using the particle-in-cell code Emu [63,64]. In this method the neutrino field is represented by a large number of individual computational particles distributed throughout the domain. Each computational particle carries several quantities: the number of neutrinos N and antineutrinos $\bar{N}$ that the particle represents, the 3×3 density matrices for neutrinos (ρ) and antineutrinos ($\bar{\rho}$), the four-position x^α of the particle, and the four-momentum p^α of each neutrino or antineutrino in the particle. The particles move at the speed of light and the density matrices evolve following the Schrödinger equation as a collisionless approximation to the quantum kinetic equations:

$$\begin{aligned} \frac{dx^\alpha}{dt} &= c \frac{p^\alpha}{p^0}, \\ \frac{dN}{dt} &= \frac{d\bar{N}}{dt} = \frac{dp^\alpha}{dt} = 0, \\ \frac{d\rho}{dt} &= \frac{-i}{\hbar} [\mathcal{H}, \rho], \\ \frac{d\bar{\rho}}{dt} &= \frac{-i}{\hbar} [\bar{\mathcal{H}}, \bar{\rho}]. \end{aligned} \quad (1)$$

The potentials $\mathcal{H}$ and $\bar{\mathcal{H}}$ contain contributions from the neutrino masses, interactions with the surrounding matter, and interactions with other neutrinos as detailed in [63], although here we set the matter term to 0. All neutrinos in these simulations have an energy of 10 MeV. We somewhat arbitrarily assign the neutrino parameters as $m_1 = m_2 = 0$, $m_3 = 0.049$ eV, $\theta_{12} = 10^{-6}$, $\theta_{13} = 48.3^\circ$, $\theta_{23} = 8.61^\circ$, $\alpha_1 = 0$, $\alpha_2 = 0$, and $\delta_{\text{CP}} = 222^\circ$. These parameters have a small effect on the initial transient fluctuations as described in Sec. III A 1, but they are completely overpowered by the fast flavor instability. We have verified that the conclusions in this paper are independent of these choices as long as the vacuum potential remains small, and also note that the vacuum potential is too weak to significantly change the net lepton density (see the Appendix). The neutrino contribution to the potential for a neutrino moving in direction Ω is

$$\mathcal{H}_{\text{neutrino},ab} = \sqrt{2}G_F[(n_{ab} - \bar{n}_{ab}^*) - (\mathbf{f}_{ab} - \bar{\mathbf{f}}_{ab}^*) \cdot \Omega], \quad (2)$$

where a and b are flavor indices. Each particle contributes to the number density n and number flux $\mathbf{f}$ in nearby grid cells according to a third-order shape function. n and $\mathbf{f}$ are then interpolated at third order from the background grid. All of these steps are embedded in a fourth-order Runge Kutta time integration scheme.

Although the original version of Emu was capable of three dimensional simulations, we change the time-integration procedure to be optimized for a larger number of particles per cell. Previously, particles were transferred between message passing interface (MPI) blocks at each substep of the Runge-Kutta (RK) integration. We now instead add an additional layer of ghost zones to the outside of the domain (which do still get updated at every RK substep). This increases the number of grid cells that must be communicated between MPI ranks but decreases the frequency with which particles need to be communicated by allowing the particles to avoid transferring to a new block until the end of the full time step.

For our spatial grid we use Cartesian coordinates and uniform grid spacing. The particles are initialized in the center of each cell and are distributed approximately uniformly in direction such that each particle represents the same amount of solid angle as described in [63]. We place 64 particles in the $\hat{x}$ - $\hat{y}$ plane in our production simulations, which results in a total of 1506 particles per cell.

A. Initial conditions

We focus on three different physical conditions to elucidate the role of symmetries in the outcome of the fast flavor instability. In all cases, the simulations begin with only electron neutrinos and antineutrinos, and no heavy lepton neutrino content. The Fiducial case is the same as that in [63], is described in the top grouping of Table I, and is depicted in the top panel of Fig. 1. There are an equal number of electron neutrinos and antineutrinos, and the density of each depends linearly on the cosine of the angle from the z axis:

$$\begin{aligned} \frac{dn_{\nu_e}}{d\Omega} &= \frac{n_{\nu_e}}{4\pi} (1 + \Omega \cdot \hat{z}), \\ \frac{dn_{\bar{\nu}_e}}{d\Omega} &= \frac{n_{\bar{\nu}_e}}{4\pi} (1 - \Omega \cdot \hat{z}), \end{aligned} \quad (3)$$

where $d\Omega$ is the solid angle differential and Ω is again the direction vector. The 90Degree simulations have an identical initial electron neutrino distribution, but the electron antineutrinos are instead distributed proportional to the cosine of the angle from the $\hat{x}$ direction:

$$\begin{aligned} \frac{dn_{\nu_e}}{d\Omega} &= \frac{n_{\nu_e}}{4\pi} (1 + \Omega \cdot \hat{z}), \\ \frac{dn_{\bar{\nu}_e}}{d\Omega} &= \frac{n_{\bar{\nu}_e}}{4\pi} (1 + \Omega \cdot \hat{x}). \end{aligned} \quad (4)$$

This creates a distribution that is symmetric in the number of neutrinos and antineutrinos, but there is no axis of symmetry in momentum space. Finally, the TwoThirds simulations have an initially isotropic distribution of electron neutrinos with the same integrated number density as the other two cases. The electron antineutrinos have the

TABLE I. List of simulation parameters. $\mathcal{N}_x$, $\mathcal{N}_y$, and $\mathcal{N}_z$ are the number of grid cells in each direction. All boundary conditions are periodic, so a grid size of 1 implies imposed homogeneity in the corresponding direction. Similarly, L_x , L_y , and L_z denote the physical extent of the domain in each direction, although the domain size in the direction of unit grid size carries no actual meaning. The 1D simulations have a larger domain size to be in line with the convergence criteria established in [63], and the multidimensional simulations are based on the convergence criteria established in the Appendix. All simulations were done with an angular resolution of 64 particles in the $\hat{x}$ - $\hat{y}$ plane, corresponding to 1506 particles per cell. n_{ν_e} , $n_{\bar{\nu}_e}$, and n_{ν_χ} denote the initial number density of electron neutrinos, electron antineutrinos, and heavy lepton neutrinos, respectively. These are all three-flavor calculations, and both mu and tau neutrino and antineutrino densities are set to 0. The initial flux factor and flux direction of each neutrino species is listed in the final column. Graphical representations of the Fiducial, 90Degree, and TwoThirds initial conditions are shown in Fig. 1.

Name	$\mathcal{N}_x \times \mathcal{N}_y \times \mathcal{N}_z$	$L_x \times L_y \times L_z$ (cm)	$n_{\{\nu_e, \bar{\nu}_e, \nu_\mu\}}$ (10^{32} cm^{-3})	$(\mathbf{f}/n)_{\{\nu_e, \bar{\nu}_e, \nu_\mu\}}$
Fiducial_3D	$128 \times 128 \times 128$	$8 \times 8 \times 8$		
Fiducial_2D	128×128	8×8	$\{4.89, 4.89, 0\}$	$\{1/3\hat{z}, -1/3\hat{z}, -\}$
Fiducial_1D	1024	64		
90Degree_3D	$128 \times 128 \times 128$	$8 \times 8 \times 8$		
90Degree_2D	128×128	8×8	$\{4.89, 4.89, 0\}$	$\{1/3\hat{z}, 1/3\hat{x}, -\}$
90Degree_1D	1024	64		
TwoThirds_3D	$128 \times 128 \times 128$	$32 \times 32 \times 32$		
TwoThirds_2D	128×128	32×32	$\{4.89, 3.26, 0\}$	$\{0, -1/3\hat{z}, -\}$
TwoThirds_1D	1024	256		

same linear dependence as in the Fiducial simulations but with an overall density scaled by $2/3$:

$$\begin{aligned} \frac{dn_{\nu_e}}{d\Omega} &= \frac{n_{\nu_e}}{4\pi}, \\ \frac{dn_{\bar{\nu}_e}}{d\Omega} &= \frac{n_{\bar{\nu}_e}}{4\pi} (1 - \mathbf{\Omega} \cdot \hat{z}). \end{aligned} \quad (5)$$

This makes for an asymmetry between the numbers of neutrinos and antineutrinos and makes the depth and size of the ELN crossing significantly smaller than in the other cases. Although none of these calculations are directly representative of realistic astrophysical conditions, they elucidate which features are likely accurately captured by simulations with artificially imposed symmetries.

In practice, this means that the weight of a computational particle (i.e., the number of physical neutrinos it represents) is proportional to the angular distribution described above, and we initialize the quantum density matrix for each particle such that $\rho_{ee} = \bar{\rho}_{ee} = 1$ and all other components are zero. We then place a uniform random number in the range $(-10^{-6}, 10^{-6})$ in the real and imaginary components of each off-diagonal matrix element and normalize the diagonal elements to ensure a unit flavor vector length. This choice of random initial conditions prevents us from evaluating exact convergence, especially once the system saturates and becomes chaotic. Changing the number of particles necessarily changes the initial conditions in a random way. However, in the same manner, this does also prevent the results of the simulation from depending on the precise functional form of the initial conditions, as all modes are present at some amplitude in the randomized initial conditions. It was also pointed out by [61] that the

final equilibrium distributions depend increasingly little on the precise form of the initial perturbations, a point that we reinforce in this work.

The spatial resolution of 1024 cells in the 1D simulations follows the resolution requirements established in [63]. However, we found that the results of multidimensional simulations are both less sensitive to spatial resolution and more sensitive to angular resolution (see the Appendix). For this reason, all of our production simulations for this work have four times the number of particles per cell, but the multidimensional simulations are only 128 cells on a side. We find similar results down to a domain size of 8 cm on a side for the multidimensional Fiducial and 90Degree simulations. The TwoThirds simulation set has a smaller antineutrino density and crossing strength, and therefore the dynamics are considerably slower and have longer characteristic length scales. Because of this, we extend the domain to 32 cm on a side for the multidimensional TwoThirds simulations and 256 cm for the 1D TwoThirds simulation.

III. RESULTS

We present a series of simulations described in Sec. II A designed to elucidate the differences between simulations with and without imposed spatial symmetries. All three distributions are unstable to the fast flavor instability, so their evolution is initially characterized by a linear growth phase, where unstable modes grow exponentially. Following this, the amplitude of the growing modes approach their maximal values, the evolution equations stop being well approximated by their linearized form, and power moves from the fastest growing mode to many other modes. Eventually, the distribution settles into a

FIG. 1. Graphical representation of the three initial conditions considered in this work. The radial extent of the curves represents the differential number density of that species of neutrino in the corresponding direction. The real initial conditions are distributed in all three directions and we only show the part in the $\hat{x}$ - $\hat{z}$ plane; the full distribution can be visualized by rotating each curve around its axis of symmetry.

quasiequilibrium with increasingly small fluctuations around average values. We will go through each of these three phases in detail and show that while artificially imposed spatial and directional symmetries can suppress instability by limiting which modes can be expressed, the final abundances of each flavor are the same between 1D, 2D, and 3D simulations for each of our three initial neutrino distributions. There are also some more subtle differences

in the evolution of the Fourier and angular power spectra, but interpreting these details requires significant care in the context of numerical limitations.

A. Overall appearance

We found that the complex phase of $n_{e\mu}$ is particularly useful in visualizing the complex matrix-valued solution, especially during the growth phase, as it clearly shows wavefronts of the flavor-transforming modes. In Fig. 2 we render surfaces of constant complex phase of this quantity at three points in time (each column corresponds to a snapshot time) and for 1D (top row), 2D (middle row), and 3D (bottom row) simulations. The Fiducial_1D simulation (top row) only has a finite computational extent in the $\hat{z}$ direction, but it is visualized by extruding the data in the $\hat{x}$ and $\hat{y}$ directions. Similarly, the Fiducial_2D simulation (middle row) is simulated with finite computational extent in the $\hat{y}$ and $\hat{z}$ directions, but it is visualized by extruding in the $\hat{x}$ direction.

1. Linear growth

During the growth phase of the Fiducial_1D simulation (top left panel), we see evenly spaced flat sheets of constant phase, indicating that the phase is varying in the $\hat{z}$ direction with a wavelength corresponding to the fastest growing mode (this was demonstrated in greater detail in [63]). Due to the symmetry between the $\pm\hat{z}$ directions of the Fiducial initial conditions, the real part of the frequency of the fastest growing mode is zero and this is a standing wave; even as the mode amplitude grows, the surfaces of constant phase do not move until the instability saturates.

The same structure exists in the multidimensional simulations, though the planes are distorted and interconnected. The center-left panel shows the equivalent data from the Fiducial_2D simulation. In the bottom left panel is the data from the Fiducial_3D simulation. The wavelength of the fastest growing mode in the $\hat{z}$ direction in the multidimensional simulations matches that expected from the 1D analysis in [63] even though the one-dimensional dispersion analysis does not account for the existence of variations in the x and y directions. Animations of these visualizations show that the kinks and holes in the phase contours of the multidimensional simulations are also essentially unchanging until the instability saturates. The exact form of these features is nonunique and randomly determined by the perturbations to the initial conditions. In summary, shortly after the start of the simulation, the distribution settles into a particular multidimensional eigenmode that grows with the same characteristic growth rate and wavelength as the purely planar solution in the 1D simulation.

There are also small transient effects due to our choice of neutrino mixing parameters. The data plotted in the left column of Fig. 2 is entirely unpolluted by neutrino mass effects, since we choose $m_1 = m_2 = 0$. However, $n_{e\tau}$ (not

FIG. 2. Volume rendering of contours of constant phase of $n_{e\mu}$ for the Fiducial_1D (top row), Fiducial_2D (center row), and Fiducial_3D (bottom row) simulations. Phases of $-2\pi/3$, 0 , and $2\pi/3$ are shown in blue, white, and red, respectively. The left column shows the results at $t = 0.29$ ns during the linear growth phase of the fast flavor instability, the center column shows the results at $t = 0.77$ ns after the instability saturates, and the right column shows the results at $t = 2.2$ ns as the distribution is building power on small scales. The phase of $n_{e\mu}$ demonstrates wavefronts of the fastest growing unstable mode. The Fiducial_1D data is copied into the x and y dimensions and the Fiducial_2D data is copied into the x direction for visualization purposes. Although there is significant multidimensional structure, the 3D results are qualitatively and quantitatively similar to the 1D and 2D results.

shown) quickly overcomes its random initial phases and establishes a constant phase throughout the domain instead of the planar structure seen in $n_{e\mu}$. This is a result of the $1 \leftrightarrow 3$ vacuum mixing that grows linearly regardless of the amplitude of the initial perturbations. The fastest growing mode, however, still grows on top of the vacuum oscillations. Within 0.2 ns the fastest growing mode overtakes the vacuum oscillations and creates a phase pattern just like

in $n_{e\mu}$. On the other hand, $n_{\mu\tau}$ (also not shown) has a phase distribution that is negative the phase $n_{e\mu}$ for the first 0.2 ns. This reflects the fact that some of $n_{e\mu}$ induced by the fast flavor instability is subsequently being pushed into $n_{\mu\tau}$ by the vacuum potential. After $t = 0.2$ ns, the phase of $n_{\mu\tau}$ also transforms to an altogether different distribution, instead varying on length scales comparable to 8 cm size of the domain.

The 90Degree_3D simulation (not shown) proceeds quite similarly, except that the wave vector of the fastest growing neutrino mode points in the direction of $(\hat{z} - \hat{x})/\sqrt{2}$ (and in the opposite direction for antineutrinos). The phase pattern is again stationary, but with a longer wavelength due to the smaller magnitude of the self-interaction potential resultant from neutrino and antineutrino distributions with a smaller angle between them. In the 1D and 2D simulations, the $\hat{x}$ direction is assumed to be homogeneous, so the fastest growing is instead one with a wave vector in the $\hat{z}$ direction. The TwoThirds simulation instead has a fastest growing mode with nonzero real frequency. The phase pattern (also with a wave vector parallel to $\hat{z}$) of $n_{e\mu}$ and $n_{e\tau}$ drift along $\hat{z}$ with time. The effects of the vacuum potential are identical in these simulations as in the Fiducial simulation described above.

2. Saturation

Flavor transformation is only able to rotate each neutrino's quantum density matrix. Since our particles start in a pure electron flavor state, the nonzero components of the eight-dimensional SU(3) polarization vector (i.e., coefficients of the Gell-Mann matrices) are $P_3 = 1/2$ and $P_8 = 1/2\sqrt{3}$. The magnitude of this polarization vector is then $|P| = 1/\sqrt{3}$, so the maximum possible value for a flavor off-diagonal component of the particle's density matrix is also $1/\sqrt{3}$.

When the particle quantum states approach this limit in the Fiducial simulations at $t \approx 0.3$ ns, the evolution equations begin to manifest their nonlinearity and the fastest growing modes cease to approximate the solution. The bottom-center panel of Fig. 2 shows the state of the Fiducial_3D simulation at $t = 0.77$ ns, well after the saturation of the instability has dismantled the fastest growing mode. The main features have a size scale larger than the wavelength of the fastest growing mode, but with no preferred direction. We will describe this more quantitatively in Sec. III C. Quasiplanar structures occasionally spontaneously form in the phase of $n_{e\mu}$ with a random orientation, but are again destroyed within a few tenths of a nanosecond. The Fiducial_2D (middle center panel) simulation qualitatively (and we will later see, also quantitatively) reproduces this analogously in two dimensions. Although the Fiducial_1D (top center panel) simulation is still restricted to one spatial dimension, this chaotic, nonlinear behavior is also consistent with multidimensional simulations, though it is more obvious in the curves plotted in [63] than in the volume renderings in Fig. 2.

After $t \approx 1$ ns, the distribution develops significantly more structure at smaller scales (e.g., bottom right panel of Fig. 2). Although we are able to demonstrate convergence of the shape of this distribution (see the Appendix), the time at which this high-frequency structure arises depends on the angular resolution. We will try to argue in

Secs. III C and III D that this is a result of the interplay between modes of small spatial and angular scales, and that adding more particles delays the onset of these high-frequency fluctuations by reducing the initial amplitude of the angular modes. However, a full understanding of the origin of this feature requires further investigation. The same process occurs in the 1D (top right) and 2D (center right) simulations.

The same comments can be made about the 90_Degree simulations, except that the smaller potentials lead to a later time of saturation at $t \approx 0.5$ ns and the onset of high-frequency structures at $t \approx 1.5$ ns. The TwoThirds simulations, in turn, saturate at $t \approx 1.5$ ns and do not develop a strong high-frequency component in $n_{e\mu}$ over the 5 ns of the simulation. The fastest growing mode in the TwoThirds simulations is however not completely destroyed and continues to dominate the solution of $n_{e\mu}$ and $n_{e\tau}$, though with some random fluctuations on top of it. This is a reflection of the fact that the amount of possible flavor transformation is much smaller, so the postsaturation distribution remains much more similar to the presaturation distribution than in the Fiducial or 90Degree simulations.

B. Average flavor evolution

Perhaps the most important metric of flavor transformation is the expectation value of the amount of each flavor of neutrino. We represent this by averaging the density matrices of particles as

$$\langle \rho \rangle = \frac{\sum_p N \rho}{\sum_p N}, \quad (6)$$

where the sum is over computational particles. As described in Sec. II, each particle has a unit-trace density matrix ρ representing the flavor state of each neutrino and a scalar N indicating the number of physical neutrinos that particle represents (and similarly for antineutrino quantities). In these simulations, all neutrinos and antineutrinos begin in electron flavor states, so $\langle \rho_{ee} \rangle$ is equivalent to the survival probability.

The evolution of $\langle \rho_{ee} \rangle$ plotted in the top panel of Fig. 3 further shows the dimension independence of the solution for the Fiducial simulations. The values start on the left side of the plot at 1 because all neutrinos begin in electron neutrino states (modulo small perturbations). The amplitude of these perturbations grows exponentially until $t \approx 0.3$ ns, where the amplitudes have grown to order unity, characterized by a steep drop in $\langle \rho_{ee} \rangle$. This happens at a very similar time for 1D, 2D, and 3D simulations, as the fastest growing mode can manifest in all three with the same growth rate. Progressing farther to the right on the plot, the averaged survival probabilities of the 1D, 2D, and 3D simulations all approach the equilibrium value of $1/3$ (horizontal green line indicating complete flavor mixing) at the same rate, albeit with random fluctuations. The nonlinearity of the equations after the saturation of

FIG. 3. Time evolution of fraction of electron flavor neutrinos for the Fiducial (top), 90Degree (center), and TwoThirds (bottom) simulations. The 1D, 2D, and 3D simulation data are plotted in gray, black, and blue, respectively. The green line marks the equilibrium fraction of electron neutrinos ($1/3$ for Fiducial and 90Degree, 0.7 for TwoThirds). Equilibrium abundances (at late times) of each neutrino species seems to be independent of simulation dimensionality. The instability growth rates show weak dimensionality dependence in the 90Degree simulations, indicated by a slight offset in the time of the drop. The offsets in the drop for the TwoThirds simulation reflects initial mode amplitude differences in 1D, 2D, and 3D.

the instability cause the precise solution to be chaotic, so this randomness is expected. The simulations of lower dimensionality have larger fluctuations, though this is simply a result of the fact that there are fewer grid cells to average over in the domain.

We note that for the 2D and 3D Fiducial simulations, the time to saturation is remarkably close to the light crossing

time of the domain ($8 \text{ cm}/c = 0.27 \text{ ns}$), though this turns out to be coincidental. We investigate the dependence of the domain size in 1D simulations in [63] and in 2D simulations in the Appendix, and find no effect. Indeed, even the 1D simulations presented here live on a domain size of 64 cm . In addition, the 90Degree simulations (center panel) live on the same grid structure but show a different time to saturation. The domain size of the multidimensional simulations is effectively determined by requiring that the fastest growing modes be able to fit in the domain. The ratio of the growth rate and the wave number for the fastest growing mode in the Fiducial simulations is $\text{Im}(\omega)/k = 2.3 \times 10^{10} \text{ cm s}^{-1}$. This is indeed close to the speed of light and likely to cause such coincidences when the domain size is comparable to the fastest growing mode's wavelength.

The slight spread in the time of the drop for the Fiducial simulations is due to differences in the initial amplitude of the fastest growing mode and not due to differences in the instability growth rate (all three clock in at $6.3 \times 10^{10} \text{ s}^{-1}$). Higher-dimensional simulations have more grid cells, which implies a larger number of possible modes, but the contribution to each of those modes from the random perturbations is smaller. The TwoThirds simulation (bottom panel of Fig. 3) is much more sensitive to this effect, as the saturation of the instability in the 3D simulation is about 0.2 ns later than in the 1D simulations, even though the growth rate in all three TwoThirds simulations is measured at $1.3 \times 10^{10} \text{ s}^{-1}$. Following the saturation of the instability in the TwoThirds simulations, about 30% of the neutrinos and 45% of the antineutrinos have converted to a heavy lepton flavor. We will get to the 90Degree simulations shortly.

The inability of the TwoThirds simulations to undergo complete flavor transformation is understandable intuitively. Since the exponential growth from the fast flavor instability is entirely a result of the self-interaction terms (assuming of course that other contributions to the Hamiltonian are small), we can rely on the symmetries in that portion of the Hamiltonian. Specifically, the self-interaction term has no preferred flavor, owing to the fact that it arises only from neutral current interactions. The antineutrino Hamiltonian is also trivially related to the neutrino Hamiltonian ($\bar{H}_{\text{neutrino}} = -H_{\text{neutrino}}^*$). As a result, the system must conserve the total flavor charge $q_a = n_a - \bar{n}_a$ for each flavor a and the evolution of the antineutrinos mirrors that of neutrinos. In the TwoThirds simulations, there is an excess of electron neutrinos over antineutrinos in the $+\hat{z}$ direction. If neutrinos and antineutrinos moving in this direction begin transforming flavor to a heavy lepton flavor, the total number of electron neutrinos would decrease faster than the number of electron antineutrinos, which on its own would result in a change in q_a . This transformation must be complemented by flavor transformation on the other side of the crossing, where

electron antineutrinos are more abundant than electron neutrinos, in a way that prevents q_a from changing. This is another way of stating that the fast flavor instability is fundamentally a multidirection phenomena. However, it is not yet clear to us how to predict the final abundances without carrying out a simulation.

The 1D and 2D 90Degree simulations do, however, differ slightly in their measured growth rate ($3.6 \times 10^{10} \text{ s}^{-1}$) from the 3D simulation ($4.3 \times 10^{10} \text{ s}^{-1}$) because the former are unable to support the true fastest growing mode. The electron antineutrino distribution is pointed in the $\hat{x}$ direction, but the 90Degree_2D simulation only allows inhomogeneity in the $\hat{y}$ and $\hat{z}$ directions and the 90Degree_1D simulation allows inhomogeneity only along $\hat{z}$. That is, the antineutrino distribution is pointing out of the plane of the 2D computational domain, and the fastest growing mode has a wave vector in a direction between the two distributions. As part of our numerical tests (see the Appendix), we did try simulations with a grid instead in the $\hat{x} - \hat{z}$ plane, which resulted in a growth rate that matched the 3D simulation. Despite the differences in growth rate, the 1D, 2D, and 3D simulations all approach the same equilibrium value of $\langle \rho_{ee} \rangle = 1/3$.

The Fiducial and 90Degree simulations of all dimensionalities approach an equilibrium characterized by complete flavor mixing. This is reflected in the tendency of $\langle \rho_{ee} \rangle$ toward $1/3$ in Fig. 3. We also show the azimuthally integrated equilibrium distributions for the Fiducial simulations in the top panel of Fig. 4. The gray vertical line shows the polar angle where the crossing in the original distribution was, and the blue curve, black marks, and gray marks, respectively show the 3D, 2D, and 1D averages at that polar angle at $t \approx 5 \text{ ns}$. At this point, there still are some fluctuations, but it is clear that all polar angles show even flavor mixing. The asymmetry of the 90Degree simulations would require a Mollweide projection color plot to show the spatially averaged survival probability for each direction, since there is no reason to expect azimuthal symmetry around any particular axis. We do not show the resulting plot because it is exceedingly boring—the survival probabilities also all lie near $1/3$ for all dimensionalities.

The TwoThirds simulation (bottom panel) is more interesting. The location of the crossing in the original distribution is closer to the $-\hat{z}$ direction (also apparent in Fig. 1). The result is a distribution that has a much smaller region where there is an excess of antineutrinos over neutrinos, and in that region the difference between the differential neutrino and antineutrino densities are smaller. We see that in simulations of all three dimensionalities, the region “inside” the crossing completely mixes, and the region outside the crossing compensates in a way that preserves the total neutrino-antineutrino asymmetry. This independently corroborates similar results found in [61]. However, an explanation of the precise functional form of

FIG. 4. Probability that a neutrino (top) or antineutrino (bottom) starting in the electron flavor ends up in the electron (left), muon (center), or tauon (right) flavor state as a function of polar angle. Blue lines, black crosses, and gray circles are from 3D, 2D, and 1D simulations, respectively. The vertical gray line shows the angle of the crossing; to the left of this line there are more antineutrinos than neutrinos. The green horizontal line is at $\langle \rho_{ee} \rangle = 1/3$, marking complete flavor mixing. Flavor content equilibrates between the flavors in the antineutrino-dominated region, but no flavor transformation occurs far from the crossing.

the compensating flavor transformation to the right of the gray line still eludes us. The 1D results seem to have significant scatter from the 2D and 3D results, but this appears to be because the 1D simulations have more difficulty relaxing. If the simulation is run for a longer period of time (not shown to preserve the common time snapshot), then the points do fluctuate around the multi-dimensional results.

C. Power spectrum

We now turn our attention to the evolution of the Fourier power spectrum of the neutrino distribution. In particular, we assume the convention that

$$\tilde{n}_{ab}(\mathbf{k}) = \sum_{l=0}^{n_x-1} \sum_{m=0}^{n_y-1} \sum_{n=0}^{n_z-1} e^{-i\mathbf{k} \cdot \mathbf{x}_{lmn}} n_{ab}(\mathbf{x}_{lmn}), \quad (7)$$

where we use l , m , and n to denote the index of the x , y , and z position of the grid cell and a and b to denote the flavor

index. The above is evaluated as a discrete fast Fourier transform using SciPy [65], and we use the resulting grid of wave numbers k from 0 to π/L_z with spacing $\Delta k = \pi/\Delta z$ (where $\Delta z = L_z/n_z$ is the spatial grid size). We integrate power in shells of magnitude of $\mathbf{k}$ according to

$$|\tilde{n}_{ab}|^2(k) = \sum_{k_{\min} < |\mathbf{k}| \leq k_{\max}} |\tilde{n}_{ab}(\mathbf{k})|^2. \quad (8)$$

$k_{\min} = k - \Delta k/2$ and $k_{\max} = k + \Delta k/2$ define the inner and outer radius of the spherical shell (in $\mathbf{k}$ space) over which power is integrated for each value of k .

Returning to the Fiducial_3D simulation, the spectra of n_{ee} and $n_{e\mu}$ are shown in the left column of Fig. 5. The color denotes the time at which the spectrum was calculated; all of the dark blue curves are from $t \in [0, 0.25)$ ns and hence all represent points in time in the linear growth phase (which ends at $t \approx 0.3$ ns). The results are essentially identical to those in the 1D simulations of [63] during this phase. The fastest growing mode, characterized by a wavelength of $\lambda = 2.2$ cm, is visible as an exponentially growing bump in $|\tilde{n}_{e\mu}|^2$. $|\tilde{n}_{ee}|^2$ also grows sympathetically, since an increase in $|n_{e\mu}|$ requires a decrease in n_{ee} (the quantum state “rotates” away from the pure electron flavor state). Meanwhile, the noise floor ($k \gtrsim 1.5 \text{ cm}^{-1}$) increases linearly (indicated by horizontal lines that become increasingly dense) at a scale that is not visible on the bottom panel, and slows down toward $|\tilde{n}_{ee}| \approx 10^{44} \text{ cm}^{-6}$. This reflects the behavior of the spectrum of $n_{\mu\tau}$ (not shown),

which grows because the vacuum part of the Hamiltonian rotates perturbations into $n_{e\mu}$. The subsequent exponential growth of the flat part of the spectrum (indicated by evenly spaced light blue horizontal lines) is likely a numerical artifact that tracks the exponential growth of the physical instability peak. The increasingly large amplitude of the physical instability excites artificial vibrations throughout the domain, but at amplitudes more than ten orders of magnitude smaller than the peak itself. Higher-resolution simulations have a lower noise floor, and the floor stops growing exactly when the amplitude of the peak ceases to be able to grow.

When the instability saturates (light blue curves), power rapidly spreads to larger wave numbers and immediately establishes an exponential distribution that intersects the noise floor at $k \approx 25 \text{ cm}^{-1}$ (for the resolution described in Table I). This (still light blue) tail remains static for around a tenth of a nanosecond before the high- k part of the spectrum suddenly kicks out to intersect the noise floor at $k \approx 33 \text{ cm}^{-1}$ (rightmost light blue curves). As time progresses (orange curves), this feature travels up the slope.

The multidimensional simulation spectra differ significantly from the 1D simulation spectra after saturation. First, the much larger number of cells results in a significantly smoother spectrum due to the summation in Eq. (8). Second, the multidimensional simulations are both much more sensitive to the fidelity of the simulation (especially the angular resolution; see the Appendix), and converge more quickly with resolution. Because of this, the

FIG. 5. Time evolution of the Fourier power spectra as defined in Eq. (8) for the Fiducial_3D (left column), 90Degree_3D (center column), and TwoThirds_3D (right column) simulations. The color indicates the time of the snapshot for which the power spectrum was evaluated. Early times are characterized by a rising bump with a peak at the fastest growing wave number, and late times show an exponential distribution centered at $k = 0$. The slow spread of the exponential tail is a numerical artifact (see the Appendix).

long-term spread of the exponential tail (orange and beyond in the left column of Fig. 5) seems to be a numerical artifact. As the angular resolution increases this spread slows, and including more spatial dimensions makes the difference between simulations of different resolutions more severe.

We will argue in Sec. III D that this feature is a result of high- k modes interacting with high- l angular modes, though we still lack a complete description. Even so, at this point we can provide some numerical evidence suggesting the credibility of the tail kick. The size, shape, and speed, and timing of this feature is consistent between the 2D and 3D simulations of the same resolution. This is indicated by the overlap of the solid black and blue curves in Fig. 6, which shows both spectra at $t \approx 1$ ns. The size, shape, and speed are consistent between simulations of different resolutions as indicated by the series of black curves in Fig. 6, which show the spectra from the Fiducial_2D simulation, along with simulations at higher angular resolution (Fiducial_2D_128d and Fiducial_2D_256d). Note that the slope of the spectrum on either side of the bump is the same. The sole difference is the location of the bump, an indication that the feature was launched sooner in the lower-resolution simulations. Looking at the 1D data (gray), this feature is not at all

FIG. 6. Snapshot of the power spectrum of $n_{e\mu}$ at $t = 1$ ns for the 1D, 2D, and 3D Fiducial simulations under variations in the number of particles per cell; 32d means 32 equatorial directions or 378 particles per cell, 128d means 6022 particles per cell, and 256 means 24088 particles per cell. Solid curves reproduce a single snapshot of the data shown in Fig. 5. The 2D and 3D results at similar resolutions have very similar peak behavior, and the resolution dependence is similar for 2D and 3D simulations (though 1D simulations behave differently). The four black curves show that the slope of the exponential tail on either side of the upward-propagating feature is robust, but the time at which that feature is launched is increasingly late with increasing resolution.

apparent, though the (significantly noisier due to fewer cells to average over) spectrum appears to already match the slope in the multidimensional simulations from after the tail kick is launched. The lower resolution 2D and 3D simulations (dashed black and blue, respectively) differ greatly for $k \gtrsim 5 \text{ cm}^{-1}$, even though they were quite similar for the standard resolution. This is because, since even at this early time, the artificial rapid spread of the exponential tail is significantly faster in the 2D simulation than in the 3D simulation (see the Appendix).

A similar process occurs in the 90Degree_3D simulations (middle column of Fig. 5). The fastest growing mode has a slightly longer wavelength (smaller k). This is a result of the smaller average angle between neutrinos and anti-neutrinos compared to the Fiducial simulation, which results in a smaller average strength of the self-interaction Hamiltonian. The instability saturates at $t \approx 0.5$ ns (dark orange curves), and a similar tail kick travels up the exponential tail.

The TwoThirds simulation (right column of Fig. 5) shows two separate modes growing, and in this case the azimuthally asymmetric mode dominates the symmetric mode. In the top right panel, the amplitude of the $k \approx 0.4 \text{ cm}^{-1}$ mode is initially imperceptible for $t \leq 0.5$ ns (dark blue and light blue curves), but by $t = 0.8$ ns (light orange), it has overtaken the amplitude of the mode at $k \approx 1.5 \text{ cm}^{-1}$. That is to say that the mode causing the low- k feature is initially smaller amplitude than the high- k feature but has a faster growth rate. However, the bottom right panel shows that the mode amplitude remains very small in $\tilde{n}_{e\mu}$.

This is naturally explained by a rapidly growing, azimuthally odd mode with a long wavelength. If the value of $\langle \rho_{e\mu} \rangle(\hat{x}) = -\langle \rho_{e\mu} \rangle(-\hat{x})$, when $n_{e\mu}$ is calculated these two contributions cancel out. However, in both cases $\langle \rho_{ee} \rangle$ must decrease as the flavor vector rotates away from the electron flavor axis, so the mode invisible in $\tilde{n}_{e\mu}$ is still represented in $\tilde{n}_{ee}$. If we compute the cell-integrated transverse neutrino flux $f_{e\mu,x} = \sum N \rho_{e\mu} p_x / |p|$ (where the sum is over all particles in a grid cell), the factor $p_x / |p|$ itself is odd in direction, so $f_{e\mu,x}$ is not sensitive to isotropic modes and only shows direction-odd modes. Indeed, plots of the phase of this quantity (not shown) exhibit a wave with a wavelength of 16 cm and wave vector along $\hat{z}$, consistent with the low- k peak.

The behavior of the exponential tail in the TwoThirds simulation is similar to that of the Fiducial and 90Degree calculations in that it also temporarily freezes ($t \approx 1.8$ ns, salmon curves) before the high- k part of the tail kicks out and sends a bump traveling up the slope. Finally, the main peak in the bottom right plot is slowly moving toward $k = 0$ as in the other simulations. The movement of the peak is slower in the TwoThirds simulation because of the overall longer average timescale from the weaker self-interaction potential.

D. Angular structure

None of the simulations in this work restrict any angular dimensions; every simulation in Table I has 1506 particles in each cell with momenta distributed evenly over the full 4π steradians of solid angle. Even so, restrictions in spatial dimensions can have some impact on the angular structure of the results, and restrictions in angular dimensions can prevent modes from growing even if all three spatial dimensions are included.

The imposed symmetries significantly affect the structure of the fastest growing mode in the 90Degree simulations, since the distribution itself does not have any axis of symmetry. We show in Fig. 7 a mollweide projection of the spatially averaged value of ρ_{ee} for each particle direction in the three 90Degree simulations at $t = 0.56$ ns, shortly before the instability saturates. The precise colors should not

FIG. 7. Mollweide projection of the electron neutrino abundance in the 90Degree simulation at $t = 0.56$ ns, or shortly before the flavor instability saturates. The 1D (top) and 2D (center) show similar mode structure, though the mode in the 3D simulation (bottom) much more closely tracks the initial ELN crossing (gray curve).

be compared directly, as they can differ between panels due to slightly different simulation output times. Instead, the valuable information is in the shape of the regions that exhibit strong flavor transformation. In the 3D simulation (bottom panel), the fastest growing mode exhibits neutrinos in the directions near the original ELN crossing (gray curve) transforming flavor most quickly. In the 1D calculation (top panel), the band of maximal flavor transformation does tend slightly toward the ELN crossing, but appears much more tightly bound to the equatorial plane as a result of the influence of the imposed translational symmetry. The 2D calculation (center panel) is quite similar to the 1D calculation, since the imposed symmetry results in the majority of the antineutrino distribution being pointed in a direction ($\hat{x}$) that is assumed to be homogeneous. Indeed, repeating the 90Degree_2D simulation with the imposed symmetry instead in the $\hat{y}$ direction (not shown) yields a growing mode with an angular structure matching that of the 3D simulation. This corroborates the argument in Sec. III B that because the fastest growing mode cannot exist in the 90Degree_1D and 90Degree_2D simulations, a different, more slowly growing mode compatible with the imposed symmetry takes the mantle during the growth of the instability.

We can describe the angular structure of the neutrino radiation field in terms of a spherical harmonic power spectrum. A continuous angular distribution function can be decomposed into spherical harmonics as

$$f_{ab}(\cos \theta, \phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^l f_{ab,lm} Y_{lm}(\cos \theta, \phi). \quad (9)$$

In this section, f_{ab} refers to the angular distribution function flavor matrix, which is related to the number density matrix via $\int d\Omega f_{ab} = n_{ab}$, rather than the number flux vector. We can approximately evaluate the coefficients $f_{ab,lm}$ as [66]

$$f_{ab,lm} = \int d\Omega f_{ab}(\cos \theta, \phi) Y_{lm}^*(\cos \theta, \phi), \\ \approx \frac{1}{V} \sum_p N_p \rho_{p,ab} Y_{lm}^*(\cos \theta_p, \phi_p). \quad (10)$$

We have assumed that all particles occupy the same solid angle in this direction, which they do approximately by construction. We then evaluate a one-dimensional angular power spectrum as

$$|f_{ab,l}|^2 = \sum_{m=-l}^l |f_{ab,lm}|^2. \quad (11)$$

Looking first at the power spectrum of $f_{e\mu}$ in the 90Degree_3D simulation (bottom panel of Fig. 8), we see that the initial random perturbations have a power

FIG. 8. Evolution of the angular power spectrum of $f_{e\mu}$ as defined in Eq. (11) for the 1D (top panel) and 3D (bottom panel) 90Degree simulations. Solid lines show data from the standard simulations listed in Table I with 1506 computational particles per cell, while dashed lines show data from otherwise identical simulations with 378 particles per cell. The standard simulations contain significant information up to $l = 32$, while the low-resolution simulations only contain significant information up to $l = 16$. The time of the snapshot used to make the curve is indicated by color in the same way as in Fig. 5. The “tail kick” in Fig. 5 corresponds to the saturation of low- l angular power (dark orange curves) and not to power cascading to smaller scales than the angular resolution can support (light orange curves).

spectrum characterized by less power on smaller angular scales (lowest dark blue curve), as low- l modes effectively average over a large number of incoherently random values. The fastest growing mode begins to emerge in the next two dark blue curves corresponding to $t < 0.25$ ns, exhibited by an angular spectrum peaked at $l = 0$ and that extends out to $l = 10$. The mode then grows in amplitude with the same spectral shape for the next 0.25 ns (light blue curves), bringing up an artificial the high- l plateau with it, similar to the exponential of the Fourier plateau discussed in Sec. III C. It is once again important to note that the level of the artificial plateau is several orders of magnitude below the peak. After $t = 0.5$ ns (dark orange curves), the angular power spectrum stops growing and power cascades away from $l = 0$ to smaller angular scales.

The time during which the angular spectrum is extending to higher l (dark orange curves) corresponds to the time of the “tail-kick” feature in Fig. 5 (see Sec. III C). It is not until $t \approx 0.75$ ns (light orange curves) that this cascade hits $l = 32$, the highest- l mode representable by our 1506 particles per cell (64 particles in the $\hat{x}$ - $\hat{y}$ plane in momentum space). By this point, the tail-kick feature is already halfway up the Fourier power spectrum slope in Fig. 5. Hence, we believe the tail-kick feature is a physical result of modes of high l interacting with modes of high k , rather than a numerical artifact of our limited angular resolution.

We commented in Sec. III C that the tail kick occurs later for simulations with higher angular resolution (more particles per cell). In the top panel of Fig. 8 we also see dashed lines that depict the spectra from otherwise equivalent simulations performed with fewer particles per cell. Since these low-resolution simulations have only 32 particles in the $\hat{x}$ - $\hat{y}$ plane in momentum space, the highest- l mode that is represented in the data is $l = 16$. The initial noise level (bottom dashed blue curve) is a factor of 4 higher because each mode now represents an average over a factor of 4 fewer particles. This factor of 4 is maintained throughout the growth phase (dark blue and light blue, note that there are fewer dashed curves because the data output frequency was a factor of 2 lower), so it reaches its maximal amplitude sooner than the standard resolution simulation.

For $t \gtrsim 0.75$ (light orange and later solid curves), power continues flowing from low l to high l modes. By the end of the simulation (solid black curve), the angular power spectrum again takes on the shape representative of incoherent random noise, much like the random initial conditions (lowest dark blue curve), but with a much larger amplitude. Once again, the power in the lower-resolution simulation (dashed black curve) is a factor of 4 larger because each mode averages over a factor of 4 fewer

incoherent particles. This spread to small angular scales is consistent with that observed by Johns *et al.* [50], though our method is not susceptible to the same kind of spurious oscillations associated with a boundary condition imposed at small angular scales. Since all of these results are distributed from a set of computational particles, there is no explicit angular boundary condition enforced. Rather, it just arises naturally (though no less artificially) as a limited number of particles are less able to provide the angular detail that a larger number of particles could.

Looking briefly at the top panel of Fig. 8, we see that the same line of reasoning applies to the 90Degree_1D simulation. The exception is the form of the angular power spectrum at the end of the simulation (solid black curve). The shape is mostly flat, unlike the monotonically rising power spectrum indicative of incoherent random noise (like the bottom most dark blue curve). This is a result of the fact that the restricted spatial dimensions suppresses angular anisotropy, as was reported in [63].

The Fiducial and TwoThirds simulations have initial conditions that are initially symmetric around the $\hat{z}$ direction, but the evolution of the angular spectrum (not shown) is still very similar to the 90Degree simulation. In [63], we reported that in the Fiducial_1D simulation, this azimuthal symmetry is very well preserved through the entirety of the simulation, even though the initial perturbations are random and not axisymmetric. When we relax the assumption of homogeneity in $\hat{x}$ and $\hat{y}$, the Fiducial_2D and Fiducial_3D simulations are similarly characterized by a fastest growing mode that is azimuthally symmetric around the $\hat{z}$ axis. However, the high degree of azimuthal symmetry begins to break down even before the instability saturates, and fluctuations begin to exist on increasingly smaller angular scales. The fastest growing mode in the TwoThirds simulations is azimuthally asymmetric (see Sec. III C), but fluctuations once again move from large to small angular scales upon saturation of the instability. In the case of the TwoThirds simulations, it is clear that if axial symmetry were imposed, only the slower mode would be able to grow.

IV. CONCLUSIONS

We perform the first simulations of the fast flavor instability in three spatial dimensions and two momentum dimensions (i.e., all dimensions relevant when the vacuum Hamiltonian is negligible). We simulate each of three idealized initial conditions in one, two, and three spatial dimensions in order to probe for effects of artificially imposed symmetries. In all cases, the simulations are characterized by a linear growth phase where unstable modes grow exponentially, followed by saturation of the instability, where power cascades to smaller spatial (Fig. 2) and angular (Fig. 8) scales. For our choice of initial conditions, the predictions of the abundance of each neutrino flavor after the fast flavor instability saturates are independent of the dimensionality of the simulation

(Figs. 3 and 4). This is a very welcome result that supports the validity of the many previous simulations performed with imposed symmetries, though our set of three initial conditions is far from exhaustive.

However, imposed symmetries are not without impact. The growth rate of the instability is mildly dependent on simulation dimensionality when the fastest growing mode has a wave vector with some component along a direction assumed to be homogeneous (Fig. 5). When spatial symmetries are imposed, the angular structure of the fastest growing mode can also be significantly different (Fig. 7). In addition, the fastest growing mode in our “TwoThirds” set of simulations is axially asymmetric. Imposing axial symmetry would preclude this mode from existing, and one would only see the more slowly growing axially symmetric mode (right column of Fig. 5).

Similar to the 1D calculations in [63], Fourier transforms of the neutrino distribution show exponentially growing perturbations in agreement with the dispersion relation. When the instability saturates, power quickly jumps to higher wave numbers, establishing an exponential tail (Fig. 5). Our multidimensional simulations show a slower growth of this tail and stronger dependence on the angular resolution. From these data, we expect that physical distributions would be characterized by a power distribution at late time centered at $k = 0$ and with a static exponential tail. The much larger number of grid cells also make the power spectrum much clearer, such that we can identify a “tail-kick” feature in the spectrum, where shortly after saturation of the instability, the high- k part of the spectrum suddenly grows, sending a bump up the slope of the power spectrum toward small wave numbers. We believe this is related to the growth of modes at small angular scales interacting physically with modes of small spatial scales, though more work is required to confirm and explain this on a fundamental level.

The particle-in-cell method keeps some features of a many-body quantum system in that the flavor vector length of each particle cannot change or average out, but only rotate. However, the largest physical shortcoming of this work is the lack of many-body entanglement. There have been significant advances on this front recently (e.g., [67–71]), and if entanglement proves to be important in general in the limit of many particles, these results will likely need to be reconsidered.

In the future it will be important to simulate a larger variety of initial conditions to map out the steady-state flavor distribution induced by each. A confidence that the abundances of each neutrino flavor is independent of any numerical choices will make sub-grid modeling a much more attractive and feasible prospect. In addition, it will be interesting to simulate cases where the matter and vacuum potentials are relevant in order to determine if multidimensional effects are relevant more broadly. Finally, as with any local simulation, it is not clear how well simulations

performed with a particular choice of initial perturbations and boundary conditions represent reality. Although we do ensure that our flavor abundances are independent of our domain size, simulations following an single local perturbation (as opposed to our random perturbations) lead to qualitatively different behavior if not allowed to interact with other perturbations [55,61]. While we believe the randomized perturbations are more easily interpreted as natural, more work is required to understand how to map simulation results to natural intuition.

ACKNOWLEDGMENTS

We are very grateful to Gail McLaughlin, Samuel Flynn, Evan Grohs, and James Kneller for many useful discussions. This material is based upon work supported by the National Science Foundation under Grant No. 2001760. 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-2 system, which is supported by NSF Grant No. 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 [72], NumPy [73], MATPLOTLIB [74], SciPy [65], and YT [75].

APPENDIX: CONVERGENCE STUDY

The symmetries of the self-interaction Hamiltonian guarantee that, if the vacuum and matter parts of the Hamiltonian are neglected, the difference between the number of neutrinos and antineutrinos remains a constant. This is not explicitly enforced in our code, and in fact we do include a vacuum contribution to the Hamiltonian, but it is so weak to be of little consequence. Figure 9 shows the violation of this conservation of electron lepton number, computed as $(\delta n_{ee} - \delta \bar{n}_{ee})/\text{Tr}(n + \bar{n})$ and $\delta n = n(t) - n(0)$. In all simulations, the integrated electron lepton number is preserved to approximately one part in 10^6 .

The most obvious relevant prediction from these simulations is the expectation value of the number of electron neutrinos present. In Fig. 10 we show the evolution of the abundance of electron neutrinos for 2D (top) and 3D (bottom) simulations of all three initial conditions, similar to Fig. 3. We performed a much larger set of numerical studies for 2D simulations because of their relatively lower computational cost. The growth rates of the instability and

FIG. 9. Violation of the conservation of integrated electron flavor content, computed as $(\delta n_{ee} - \delta \bar{n}_{ee})/\text{Tr}(n + \bar{n})$ and $\delta n = n(t) - n(0)$. The symmetry of the neutrino self-interaction part of the Hamiltonian requires this to be zero, but this is not enforced by the code.

the equilibrium abundances are quite robust to changes in the number of particles per cell (comparing to simulations labeled “32d,” “128d,” and “256d” after the number of directions in the $\hat{x}$ - $\hat{y}$ plane in momentum space), the size of the domain (comparing to simulations labeled “16 cm” and “64 cm” indicating the cube size), and the number of grid cells in each direction (comparing to simulations labeled “nx256” and “nx1024” indicating the number of grid cells in each nonhomogeneous direction). The results seem to be most strongly dependent on the number of particles per cell. For the 2D 90Degree simulations, the growth rate of the instability can be enhanced by making the computational plane in the $\hat{x}$ - $\hat{z}$ direction to match the principal directions of the neutrino and antineutrino distributions (labeled by “inplane”). This permits the growth of the true fastest growing mode, which has a wave vector that is between those two directions.

The evolution of the power spectrum is a telling indicator of numerical effects, so we investigate it in some detail. Figure 11 shows the evolution of the power spectrum of $n_{e\mu}$ for the standard-resolution Fiducial_3D simulation (bottom panel), along with an otherwise identical simulation where we set the number of equatorial directions to 32 (top panel, corresponding to 378 particles per cell and the “32d” dashed blue curve in the bottom left panel of Fig. 10). The top panel has a higher line density simply as a result of more frequent data output. The most obvious difference is that with more particles per cell, the exponential tail spreads more slowly. In addition, looking at the lowest orange curve in each panel, it is clear that the “tail-kick” feature described in Sec. III C occurs at a later time in the higher-resolution simulations (e.g., the bump in the lowest orange

FIG. 10. Evolution of the domain-averaged electron neutrino abundance considering various numerical choices. Solid curves reproduce data shown in Fig. 3. The top row shows results from 2D simulations while the bottom shows those from 3D simulations. Each column shows results from each of our three initial conditions. Labels with a “d” refer to equatorial directions to be compared to the standard resolution of 64: 32d means 32 equatorial directions or 378 particles per cell, 128 means 6022 particles per cell, and 256d means 24088 particles per cell. “nx” labels the number of grid cells on each extended dimension of the simulation, keeping the domain size constant, where “nx128” is the standard resolution for the multidimensional simulations. “cm” labels the domain size of each dimension of the cube in centimeters. The horizontal green line marks complete flavor mixing. “inplane” for the 2D 90Degree simulations means the computational domain extends over the $\hat{x} - \hat{z}$ plane instead of the $\hat{y} - \hat{z}$ plane, slightly increasing the instability growth rate. Overall, our numerical choices have little effect on the instability growth rates and the late-time electron neutrino abundances.

curve is farther up the slope in the top panel than in the bottom panel). We argue in Sec. III D that this timing difference is because more particles per cell means that the low- l part of the angular power spectrum has a smaller amplitude. It then takes slightly longer for the angular power spectrum to grow to the point that it interacts significantly with the high- k part of the Fourier spectrum. In both cases, the peak of the distribution drifts toward $k = 0$ at the same rate for the first 1 ns, but beyond that the Fiducial_3D simulation develops a peak at a single k not evident in the Fiducial_3D_32d simulation. We do not believe that feature is physical. The details of the spectrum this close to $k = 0$ show limitations from the domain size—a larger domain would afford finer resolution in k .

The same can be said about the 2D simulations shown in Fig. 12. From top to bottom we show simulations with 32, 64, 128, and 256 equatorial directions (378, 1506, 6022, and 24088 particles per cell, respectively). As before, higher angular resolution causes the exponential tail to drift more slowly. In addition, higher angular resolution

also causes more choppy features in the spectrum near $k = 0$ at the scale of the k resolution. The limiting numerical factor for the $0 \leq k \leq 10$ part of the spectrum for the Fiducial_2D_256d simulation at $t = 5$ ns is likely the domain size.

We will focus on the drift of the exponential tail as a probe of numerical artifacts. Interestingly, for the same angular resolution the 3D results in Fig. 11 perform better than the 2D results in Fig. 12 in that the exponential tail drifts more slowly. The spread of the exponential tail after $t = 1$ ns is significantly smaller for the Fiducial_3D simulation than for the Fiducial_2D simulation. In addition, the difference between the Fiducial_3D_32d and Fiducial_3D exponential tails is much greater than the difference between the Fiducial_2D_32d and Fiducial_2D tails. This pattern extends down to 1D simulations—there is a small difference between simulations of different angular resolution, which led us to expect that the exponential drift was not a numerical artifact in [63]. The numerical origin of this drift is not clear.

FIG. 11. Time evolution of the power spectrum of $n_{e\mu}$ for the Fiducial_3D simulation (bottom panel) employing 64 equatorial directions (i.e., 1506 particles per cell) and then an otherwise identical simulation employing only 32 equatorial directions (top panel, 378 particles per cell). The spread of the exponential tail toward higher values of k is more strongly affected by this angular resolution than any other numerical choice.

Finally, we look at the evolution of the peak of the power spectrum with time in Fig. 13. Initially, the randomized initial conditions result in a peak at high wave number, but almost immediately shoots down to the value of $k \approx 3 \text{ cm}^{-1}$ corresponding to the fastest growing wavelength of $\lambda = 2.2 \text{ cm}$. After the instability saturates at $t \approx 0.3 \text{ ns}$, the peak begins to drift toward $k = 0$ over the next few tenths of a nanosecond, and should remain there for the rest of the simulation. Looking first at the standard resolutions (solid curves), all three dimensionalities agree well until $t \approx 2 \text{ ns}$, at which point all three dimensionalities show peaks drifting away from $k = 0$. The higher-resolution 2D results (red and green dashed curves) tend to stay much closer to $k = 0$ at late times, while the lower-resolution 2D and 3D results (dashed black and blue, respectively) drift upward as early as $t \approx 1 \text{ ns}$.

FIG. 12. Time evolution of the power spectrum of $n_{e\mu}$ for the Fiducial_2D simulation (second panel) employing 64 equatorial directions (i.e., 1506 particles per cell). The other panels show the equivalent results from simulations employing only 32, 128, and 256 equatorial directions (378, 6022, and 24088 particles per cell, respectively). The spread of the exponential tail toward higher values of k is more strongly affected by this angular resolution than any other numerical choice.

FIG. 13. Evolution of the value of k at the peak of the power spectrum of $n_{e\mu}$ for 1D, 2D, and 3D simulations of varying angular resolution. The low-resolution (32d) simulations stay tightly packed until about 1 ns, while the standard resolution results are very similar until about 2 ns. The highest-resolution 2D simulation predicts a peak very close to $k = 0$ throughout the postsaturation phase.

-
- [1] H.-T. Janka, Explosion mechanisms of core-collapse supernovae, *Annu. Rev. Nucl. Part. Sci.* **62**, 407 (2012).
 - [2] D. Branch and J. Craig Wheeler, *Supernova Explosions*, Astronomy and Astrophysics Library (Springer Berlin Heidelberg, Berlin, Heidelberg, 2017).
 - [3] D. Radice, S. Bernuzzi, and A. Perego, The dynamics of binary neutron star mergers and GW170817, *Annu. Rev. Nucl. Part. Sci.* **70**, 95 (2020).
 - [4] J. Lippuner, R. Fernandez, L. F. Roberts, F. Foucart, D. Kasen, B. D. Metzger, and C. D. Ott, Signatures of hyper-massive neutron star lifetimes on r-process nucleosynthesis in the disc ejecta from neutron star mergers, *Mon. Not. R. Astron. Soc.* **472**, 904 (2017).
 - [5] I. Kullmann, S. Goriely, O. Just, R. Ardevol-Pulpillo, A. Bauswein, and H.-T. Janka, Dynamical ejecta of neutron star mergers with nucleonic weak processes I: Nucleosynthesis, [arXiv:2109.02509](https://arxiv.org/abs/2109.02509).
 - [6] E. N. Alexeyev, L. N. Alexeyeva, I. V. Krivosheina, and V. I. Volchenko, Detection of the neutrino signal from SN 1987A in the LMC using the INR Baksan underground scintillation telescope, *Phys. Lett. B* **205**, 209 (1988).
 - [7] R. M. Bionta *et al.*, Observation of a Neutrino Burst in Coincidence with Supernova 1987A in the Large Magellanic Cloud, *Phys. Rev. Lett.* **58**, 1494 (1987).
 - [8] K. Hirata *et al.*, Observation of a Neutrino Burst from the Supernova SN1987A, *Phys. Rev. Lett.* **58**, 1490 (1987).
 - [9] H. Duan and A. Friedland, Self-Induced Suppression of Collective Neutrino Oscillations in a Supernova, *Phys. Rev. Lett.* **106**, 091101 (2011).
 - [10] B. Dasgupta, E. P. O'Connor, and C. D. Ott, Role of collective neutrino flavor oscillations in core-collapse supernova shock revival, *Phys. Rev. D* **85**, 065008 (2012).
 - [11] E. O'Connor, R. Bollig, A. Burrows, S. Couch, T. Fischer, H.-T. Janka, K. Kotake, E. J. Lentz, M. Liebendrfel, O. E. Bronson Messer, A. Mezzacappa, T. Takiwaki, and D. Vartanyan, Global comparison of core-collapse supernova simulations in spherical symmetry, *J. Phys. G* **45**, 104001 (2018).
 - [12] 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, *Mon. Not. R. Astron. Soc.* **481**, 4786 (2018).
 - [13] A. Burrows, D. Radice, D. Vartanyan, H. Nagakura, M. Aaron Skinner, and J. C. Dolence, The overarching framework of core-collapse supernova explosions as revealed by 3D foranx simulations, *Mon. Not. R. Astron. Soc.* **491**, 2715 (2020).
 - [14] M. A. Sandoval, W. Raphael Hix, O. E. Bronson Messer, E. J. Lentz, and J. Austin Harris, Three dimensional core-collapse supernova simulations with 160 isotopic species evolved to shock breakout, [arXiv:2106.01389](https://arxiv.org/abs/2106.01389).
 - [15] P. Mösta, D. Radice, R. Haas, E. Schnetter, and S. Bernuzzi, A magnetar engine for short GRBs and kilonovae, *Astrophys. J.* **901**, L37 (2020).
 - [16] V. Nedora, F. Schianchi, S. Bernuzzi, D. Radice, B. Daszuta, A. Endrizzi, A. Perego, A. Prakash, and F. Zappa, Mapping dynamical ejecta and disk masses from numerical relativity simulations of neutron star mergers, [arXiv:2011.11110](https://arxiv.org/abs/2011.11110).
 - [17] F. Foucart, P. Moesta, T. Ramirez, A. J. Wright, S. Darbha, and D. Kasen, Estimating outflow masses and velocities in

- merger simulations: Impact of r-process heating and neutrino cooling, [arXiv:2109.00565](#).
- [18] Z. Zhu and L. Rezzolla, Fully general-relativistic simulations of isolated and binary strange quark stars, *Phys. Rev. D* **104**, 083004 (2021).
 - [19] B. D. Metzger and R. Fernández, From neutrino- to photon-cooled in three years: Can fallback accretion explain the x-ray excess in GW170817?, *Astrophys. J. Lett.* **916**, L3 (2021).
 - [20] X. Li and D. M. Siegel, Neutrino Fast Flavor Conversions in Neutron-Star Postmerger Accretion Disks, *Phys. Rev. Lett.* **126**, 251101 (2021).
 - [21] O. Just, S. Goriely, H.-T. Janka, S. Nagataki, and A. Bauswein, Neutrino absorption and other physics dependencies in neutrino-cooled black-hole accretion disks, [arXiv:2102.08387](#).
 - [22] Q. R. Ahmad *et al.*, Measurement of the Rate of $\nu_e + d \rightarrow p + p + e^-$ Interactions Produced by B 8 Solar Neutrinos at the Sudbury Neutrino Observatory, *Phys. Rev. Lett.* **87**, 071301 (2001).
 - [23] H. Duan, G. M. Fuller, and Y.-Z. Qian, Collective neutrino oscillations, *Annu. Rev. Nucl. Part. Sci.* **60**, 569 (2010).
 - [24] A. Malkus, G. C. McLaughlin, and R. Surman, Symmetric and standard matter neutrino resonances above merging compact objects, *Phys. Rev. D* **93**, 045021 (2016).
 - [25] M.-R. Wu, I. Tamborra, O. Just, and H.-T. Janka, Imprints of neutrino-pair flavor conversions on nucleosynthesis in ejecta from neutron-star merger remnants, *Phys. Rev. D* **96**, 123015 (2017).
 - [26] M. George, M.-R. Wu, I. Tamborra, R. Ardevol-Pulpillo, and H.-T. Janka, Fast neutrino flavor conversion, ejecta properties, and nucleosynthesis in newly-formed hypermassive remnants of neutron-star mergers, *Phys. Rev. D* **102**, 103015 (2020).
 - [27] Z. Xiong, A. Sieverding, M. Sen, and Y.-Z. Qian, Potential impact of fast flavor oscillations on neutrino-driven winds and their nucleosynthesis, *Astrophys. J.* **900**, 144 (2020).
 - [28] R. F. Sawyer, Speed-up of neutrino transformations in a supernova environment, *Phys. Rev. D* **72**, 045003 (2005).
 - [29] R. F. Sawyer, Neutrino Cloud Instabilities Just above the Neutrino Sphere of a Supernova, *Phys. Rev. Lett.* **116**, 081101 (2016).
 - [30] M.-R. Wu and I. Tamborra, Fast neutrino conversions: Ubiquitous in compact binary merger remnants, *Phys. Rev. D* **95**, 103007 (2017).
 - [31] T. Morinaga, H. Nagakura, C. Kato, and S. Yamada, Fast neutrino-flavor conversion in the preshock region of core-collapse supernovae, *Phys. Rev. Research* **2**, 012046(R) (2020).
 - [32] H. Nagakura, L. Johns, A. Burrows, and G. M. Fuller, Where, when and why: Occurrence of fast-pairwise collective neutrino oscillation in three-dimensional core-collapse supernova models, *Phys. Rev. D* **104**, 083025 (2021).
 - [33] B. Dasgupta, A. Mirizzi, and M. Sen, Fast neutrino flavor conversions near the supernova core with realistic flavor-dependent angular distributions, *J. Cosmol. Astropart. Phys.* **02** (2017) 019.
 - [34] S. Abbar and H. Duan, Fast neutrino flavor conversion: Roles of dense matter and spectrum crossing, *Phys. Rev. D* **98**, 043014 (2018).
 - [35] T. Morinaga, Fast neutrino flavor instability and neutrino flavor lepton number crossings, [arXiv:2103.15267](#).
 - [36] I. Tamborra, L. Hdeppohl, G. G. Raffelt, and H.-T. Janka, Flavor-dependent neutrino angular distribution in core-collapse supernovae, *Astrophys. J.* **839**, 132 (2017).
 - [37] S. Abbar, H. Duan, K. Sumiyoshi, T. Takiwaki, and M. C. Volpe, On the occurrence of fast neutrino flavor conversions in multidimensional supernova models, *Phys. Rev. D* **100**, 043004 (2019).
 - [38] M. D. Azari, S. Yamada, T. Morinaga, W. Iwakami, H. Okawa, H. Nagakura, and K. Sumiyoshi, Linear analysis of fast-pairwise collective neutrino oscillations in core-collapse supernovae based on the results of Boltzmann simulations, *Phys. Rev. D* **99**, 103011 (2019).
 - [39] M. Delfan Azari, S. Yamada, T. Morinaga, H. Nagakura, S. Furusawa, A. Harada, H. Okawa, W. Iwakami, and K. Sumiyoshi, Fast collective neutrino oscillations inside the neutrino sphere in core-collapse supernovae, *Phys. Rev. D* **101**, 023018 (2020).
 - [40] 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, *Phys. Rev. D* **101**, 043016 (2020).
 - [41] H. Nagakura, T. Morinaga, C. Kato, and S. Yamada, Fast-pairwise collective neutrino oscillations associated with asymmetric neutrino emissions in core-collapse supernovae, *Astrophys. J.* **886**, 139 (2019).
 - [42] R. Glas, H. T. 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, *Phys. Rev. D* **101**, 063001 (2020).
 - [43] F. Capozzi, S. Abbar, R. Bollig, and H.-T. Janka, Fast neutrino flavor conversions in one-dimensional core-collapse supernova models with and without muon creation, *Phys. Rev. D* **103**, 063013 (2021).
 - [44] S. Abbar, F. Capozzi, R. Glas, H.-T. Janka, and I. Tamborra, On the characteristics of fast neutrino flavor instabilities in three-dimensional core-collapse supernova models, *Phys. Rev. D* **103**, 063033 (2021).
 - [45] B. Dasgupta and A. Mirizzi, Temporal instability enables neutrino flavor conversions deep inside supernovae, *Phys. Rev. D* **92**, 125030 (2015).
 - [46] F. Capozzi, B. Dasgupta, A. Mirizzi, M. Sen, and G. Sigl, Collisional Triggering of Fast Flavor Conversions of Supernova Neutrinos, *Phys. Rev. Lett.* **122**, 091101 (2019).
 - [47] S. Abbar, H. Duan, and S. Shalgar, Flavor instabilities in the multiangle neutrino line model, *Phys. Rev. D* **92**, 065019 (2015).
 - [48] S. Abbar and M. C. Volpe, On fast neutrino flavor conversion modes in the nonlinear regime, *Phys. Lett. B* **790**, 545 (2019).
 - [49] I. Padilla-Gay, S. Shalgar, and I. Tamborra, Multi-dimensional solution of fast neutrino conversions in binary neutron star merger remnants, *J. Cosmol. Astropart. Phys.* **01** (2021) 017.
 - [50] L. Johns, H. Nagakura, G. M. Fuller, and A. Burrows, Fast oscillations, collisionless relaxation, and spurious evolution

- of supernova neutrino flavor, *Phys. Rev. D* **102**, 103017 (2020).
- [51] I. Padilla-Gay and S. Shalgar, Fast flavor conversion of neutrinos in presence of matter bulk velocity, [arXiv:2108.00012](#).
- [52] S. Shalgar and I. Tamborra, A change of direction in pairwise neutrino conversion physics: The effect of collisions, *Phys. Rev. D* **103**, 063002 (2021).
- [53] S. Shalgar and I. Tamborra, Three flavor revolution in fast pairwise neutrino conversion, *Phys. Rev. D* **104**, 023011 (2021).
- [54] Z. Xiong and Y.-Z. Qian, Stationary solutions for fast flavor oscillations of a homogeneous dense neutrino gas, *Phys. Lett. B* **820**, 136550 (2021).
- [55] J.D. Martin, C. Yi, and H. Duan, Dynamic fast flavor oscillation waves in dense neutrino gases, *Phys. Lett. B* **800**, 135088 (2020).
- [56] J.D. Martin, J. Carlson, V. Cirigliano, and H. Duan, Fast flavor oscillations in dense neutrino media with collisions, *Phys. Rev. D* **103**, 063001 (2021).
- [57] H. Duan, J.D. Martin, and S. Omanakuttan, Flavor isospin waves in one-dimensional axisymmetric neutrino gases, [arXiv:2110.02286](#).
- [58] S. Shalgar and I. Tamborra, Symmetry breaking induced by pairwise conversion of neutrinos in compact sources, [arXiv:2106.15622](#).
- [59] A. Mirizzi, G. Mangano, and N. Saviano, Self-induced flavor instabilities of a dense neutrino stream in a two-dimensional model, *Phys. Rev. D* **92**, 021702 (2015).
- [60] S. Bhattacharyya and B. Dasgupta, Late-time behavior of fast neutrino oscillations, *Phys. Rev. D* **102**, 063018 (2020).
- [61] M.-R. Wu, M. George, C.-Y. Lin, and Z. Xiong, Collective fast neutrino flavor conversions in an 1D box: (I) Initial condition and long-term evolution, [arXiv:2108.09886](#) [Phys. Rev. D (to be published)].
- [62] S. Bhattacharyya and B. Dasgupta, Fast Flavor Depolarization of Supernova Neutrinos, *Phys. Rev. Lett.* **126**, 061302 (2021).
- [63] S. Richers, D.E. Willcox, N.M. Ford, and A. Myers, Particle-in-cell simulation of the neutrino fast flavor instability, *Phys. Rev. D* **103**, 083013 (2021).
- [64] D. Willcox, S. Richers, and N. Ford, Emu v1.3, Zenodo, <https://doi.org/10.5281/zenodo.5514018> (2021).
- [65] P. Virtanen *et al.*, SciPy 1.0: Fundamental algorithms for scientific computing in Python, *Nat. Methods* **17**, 261 (2020).
- [66] B. Fornberg and J. M. Martel, On spherical harmonics based numerical quadrature over the surface of a sphere, *Adv. Comput. Math.* **40**, 1169 (2014).
- [67] M. J. Cervia, A. V. Patwardhan, A. B. Balantekin, S. N. Coppersmith, and C. W. Johnson, Entanglement and collective flavor oscillations in a dense neutrino gas, *Phys. Rev. D* **100**, 083001 (2019).
- [68] E. Rrapaj, Exact solution of multiangle quantum many-body collective neutrino-flavor oscillations, *Phys. Rev. C* **101**, 065805 (2020).
- [69] A. Roggero, Dynamical phase transitions in models of collective neutrino oscillations, [arXiv:2103.11497](#).
- [70] A. Roggero, Entanglement and many-body effects in collective neutrino oscillations, [arXiv:2102.10188](#) [Phys. Rev. D (to be published)].
- [71] B. Hall, A. Roggero, A. Baroni, and J. Carlson, Simulation of collective neutrino oscillations on a quantum computer, *Phys. Rev. D* **104**, 063009 (2021).
- [72] A. Meurer *et al.*, SymPy: Symbolic computing in Python, *PeerJ Comput. Sci.* **3**, e103 (2017).
- [73] S. van der Walt, S.C. Colbert, and G. Varoquaux, The NumPy array: A structure for efficient numerical computation, *Comput. Sci. Eng.* **13**, 22 (2011).
- [74] J.D. Hunter, MATPLOTLIB: A 2D graphics environment, *Comput. Sci. Eng.* **9**, 90 (2007).
- [75] M.J. Turk, B.D. Smith, J.S. Oishi, S. Skory, S.W. Skillman, T. Abel, and M.L. Norman, YT: A multi-code analysis toolkit for astrophysical simulation data, *Astrophys. J. Suppl. Ser.* **192**, 9 (2011).