Neutrino thermalization via randomization on a quantum processor

Oriel Kiss,[1, 2, 3, *] Ivano Tavernelli,[4] Francesco Tacchino,[4] Denis Lacroix,[5] and Alessandro Roggero[3, 6]

[1] *Department of Nuclear and Particle Physics, University of Geneva, Geneva 1211, Switzerland*
[2] *European Organization for Nuclear Research (CERN), Geneva 1211, Switzerland*
[3] *Physics Department, University of Trento, Via Sommarive 14, I-38123 Trento, Italy*
[4] *IBM Quantum, IBM Research Europe – Zurich, 8803 Rueschlikon, Switzerland*
[5] *Université Paris-Saclay, CNRS/IN2P3, IJCLab, 91405 Orsay, France*
[6] *INFN-TIFPA Trento Institute of Fundamental Physics and Applications, Trento, Italy*

(Dated: October 30, 2025)

The dynamical evolution of neutrino flavor in supernovae can be modeled by an all-to-all spin Hamiltonian with random couplings. Simulating such two-local Hamiltonian dynamics remains a major challenge, as methods with controllable accuracy require circuit depths that increase at least linearly with system size, exceeding the capabilities of current quantum devices. The eigenstate thermalization hypothesis predicts that these systems should thermalize, a behavior confirmed in small-scale classical simulations. In this work, we investigate flavor thermalization in much larger systems using random quantum circuits as an empirical tool to emulate the non-local dynamics, and demonstrate that the thermal behavior can be reproduced using a depth independent of the system size. By simulating dynamics of over one hundred qubits, we find that the thermalization time grows approximately as the square root of the system size, consistent with predictions from semi-classical methods. Beyond this specific result, our study illustrates that near-term quantum devices are useful tools to test and validate empirical classical methods. It also highlights a new application of random circuits in physics, providing insight into complex many-body dynamics that are classically intractable.

I. INTRODUCTION

In astrophysical scenarios marked by high temperature and density, such as core collapse supernovae and binary neutron star mergers, the emission of neutrinos in substantial fluxes plays a pivotal role in the dynamical evolution of the system and the nucleosythesis that takes place in these environments [1–4]. Neutrinos interact through weak currents with local nucleons, influencing the proton-to-neutron ratio. A comprehensive understanding of neutrino flavor content is essential for detailed investigations into the evolution of these systems [5–9]. In dense regions, neutrinos can experience coherent forward scattering, a process highly sensitive to their quantum mechanical flavor states. Flavor-dependent relative phases may arise during scattering on charged leptons, leading to the well-known MSW effect, but they can also be generated through neutral current interactions with other neutrinos in the system [1, 6, 10]. In scenarios with a sufficiently high neutrino number density, this flavor exchange effect is expected to dominate the neutrino gas's flavor evolution, giving rise to novel coherent phenomena like flavor spectrum splits and swaps [11–13]. The phenomenology of collective neutrino oscillations has been traditionally explored by means of mean-field methods (see [13–15] for reviews) but in recent years a renewed interest in the role played by quantum correlations have emerged (see [16–18] for recent reviews).

In this paper, building on Ref. [19], we focus on investigating the flavor evolution of a neutrino gas under the condition where the potential arising from coherent $\nu-\nu$ forward scattering prevails as the dominant effect. This scenario is commonly referred to in the literature as the "fast" flavor oscillation limit [20–25]. In this limit, the redistribution of flavor among neutrinos is anticipated to transpire over length scales of approximately one meter. Our exploration is centered on this specific regime, aiming to comprehensively analyze the associated dynamics. We study this regime using randomized quantum simulations strategies based on product formula, previously tackled using standard Trotter Suzuki decomposition by Refs [26, 27].

We carry out our computational investigations by leveraging, besides classical benchmarks, the opportunities offered by state-of-the-art digital quantum simulators [28, 29]. In recent years, this technology has advanced significantly, enabling the first so-called "quantum utility" experiments [30, 31], in which noisy quantum processors augmented by error mitigation techniques successfully executed quantum circuits beyond the reach of exact classical simulation. These pioneering demonstrations marked the initial steps toward employing quantum computers as genuine research tools in fundamental physics, first to validate theoretical predictions and eventually to explore regimes previously inaccessible to classical methods. Several utility-scale experiments have since been reported in high-energy physics [32, 33], with a primary focus on simulating scattering dynamics [34–39] and in general on lattice gauge theories [40–42]. In this work, we introduce quantum simulation techniques into the domain of neutrino physics, opening a new research direction and presenting the first results at the 100-qubit scale in this context.

* oriel.kiss@unige.ch

The remainder of the paper is organized as follows. Section II introduces the physical model and the observables of interest. Exact small-scale simulations using converged Trotterization are presented in Section III. Section IV describes our randomization strategy, while large-scale experiments are discussed in Section V A for classical methods and in Section V B for implementations on IBM's superconducting quantum hardware.

II. MODEL AND THERMALIZATION

The $\nu-\nu$ coherent forward scattering potential takes the form of an all-to-all Heisenberg-like interaction [43, 44], and such Hamiltonians are generically expected to be nonintegrable except in special cases [19, 45–48]. Furthermore, non-integrable Hamiltonians are expected to "thermalize" in the sense that expectation values of few-body operators tend to equilibrate to a value that can be predicted from an appropriate thermodynamic partition function (see [49] for a review). For the rest of this work, we assume that the vacuum oscillation potential is negligibly small relative to the other two potentials. We also assume that the matter profile is uniform in the regions of the environment under consideration, allowing us to consider a corotating frame to remove its overall effect on each neutrino. After these modifications, only the $\nu-\nu$ coherent forward scattering Hamiltonian remains. We assume axial symmetry for the velocity coupling term, enforcing that the velocity components orthogonal to the momentum symmetry axis (denoted by z) average to zero. This leads to the following Hamiltonian for N neutrinos [43, 44]

$$\begin{aligned}\hat{H}_{\nu\nu} &= \frac{\mu}{2N}\sum_{i<j}(1-\mathbf{v}_i\cdot\mathbf{v}_j)\hat{\vec{\sigma}}_i\cdot\hat{\vec{\sigma}}_j \\ &\equiv H_0+\frac{\mu}{N}\sum_{i<j}(1-\mathbf{v}_i\cdot\mathbf{v}_j)P_{i,j}.\end{aligned} \tag{1}$$

Here, $\mathbf{v}_k$ is a three-dimensional real vector containing the velocity of the k-th neutrino, $\hat{\vec{\sigma}}$ is a vector of Pauli operators, μ is the scale of the $\nu-\nu$ interaction, $P_{i,j}$ is a swap operation which can be expressed as

$$P_{i,j}=\frac{1}{2}(\mathbb{1}+\vec{\sigma}_i\cdot\vec{\sigma}_j), \tag{2}$$

and H_0 an irrelevant constant. Given a typical core-collapse supernova flux at a radius of 50 km, a simple estimate gives $\mu\approx 1\,\mathrm{cm}^{-1}$. As it is the only dimensional parameter in the problem, we hereafter measure all distances and times in units of μ and set $\mu=1$. We consider a forward peak distribution with $v_z\sim 1-\frac{2}{\sqrt{\pi}\sigma}\Theta(x)e^{-x^2/2\sigma^2}$ with $\sigma=0.5$. The remaining coordinates are then fixed as

$$\begin{aligned}v_x &= \sqrt{1-v_z^2}\cos\theta \\ \text{and } v_y &= \sqrt{1-v_z^2}\sin\theta \\ \text{where } \theta &\sim U(0,2\pi).\end{aligned} \tag{3}$$

This choices for modeling the angular distribution of neutrinos were previously used also in Ref. [19] where the connection between the chaotic nature of the Hamiltonian and flavor thermalization was first investigated for small systems with $N\leq 16$. The rationale behind the use of a random distribution for the velocities is twofold: on one hand it captures the expectation for the geometry of a core collapse supernova with mostly outgoing neutrinos and, at the same time, allows to avoid possible spurious symmetries that can be introduced in the calculation when using uniform grids of angles [19, 50].

In this work we are interested in the typical dynamics generated by the neutrino interactions and, given the presence of random couplings, we will consider an ensemble of Hamiltonians $\{H^{(k)}\}_{k=1,\ldots,N_H}$ defined explicitly as follows

$$\hat{H}^{(k)}_{\nu\nu}=\frac{\mu}{N}\sum_{i<j}(1-\mathbf{v}_i^{(k)}\cdot\mathbf{v}_j^{(k)})P_{i,j}\equiv\sum_{i<j}\mu^{(k)}_{i,j}P_{i,j}\ , \tag{4}$$

with the associated unitary time evolution denoted as

$$\mathcal{U}^{(k)}(t)=\exp\left(-it\hat{H}^{(k)}_{\nu\nu}\right)\ . \tag{5}$$

Previous work [19] estimated the thermalization time τ_{th}, or at least a lower bound, using single-qubit observables such as the flavor expectation value for a product state $|\psi_0\rangle$:

$$Z_j^{(k)}(t)=\langle\psi_0|\mathcal{U}^{(k)}(t)^\dagger Z_j\mathcal{U}^{(k)}(t)|\psi_0\rangle. \tag{6}$$

A characteristic time was defined as the first inversion point [51]. Additional estimates were obtained from the time when the single-qubit entropy reached 95% of its maximum, or from global measures such as the Loschmidt echo,

$$L^{(k)}(t)=|\langle\psi_0|\mathcal{U}^{(k)}(t)|\psi_0\rangle|^2, \tag{7}$$

with τ_{th} identified from the first minimum of $L^{(k)}(t)$.

In contrast, the randomization approach used here does not track individual neutrinos, but instead allows us to evaluate observables invariant under qubit relabeling. We therefore introduce the flavor variance,

$$V^{(k)}(t)=\frac{1}{N}\sum_j Z_j^{(k)}(t)^2\ . \tag{8}$$

Since every neutrino thermalizes to the same asymptotic state, the variance vanishes at long times. As a proxy for the typical time scale, we define τ_{th} as the first time $V^{(k)}(t)$ drops below $\Gamma=0.1$.

Fig. 1 shows the characteristic time of the single-qubit observable, the entropy, as well as the variance and Loschmidt echo, where the evolution is nearly exact at using a converged Trotter decomposition with time step $\delta t=0.5$, acting on the domain-wall configuration, $|\phi_0\rangle=|0\rangle^{\otimes N/2}\otimes|1\rangle^{\otimes N/2}$. The errorbars represent a 10%–90% confidence interval over different Hamiltonian realizations.

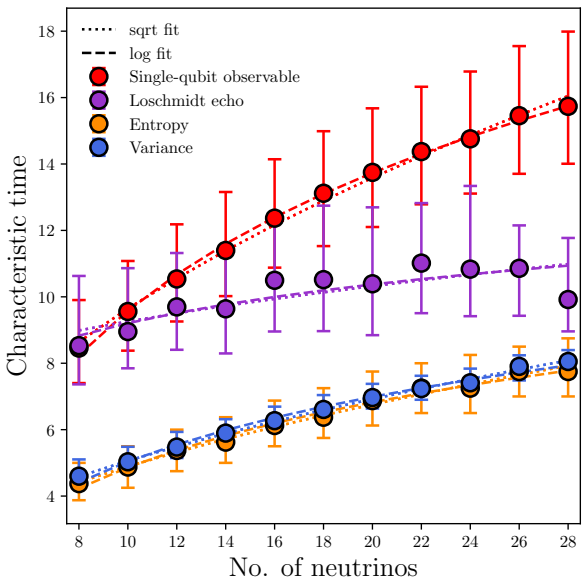

Figure 1: **Trotterized dynamics.** Scaling of the characteristic time as function of the system size for the single qubit observables (red), Loschmidt echo (purple), single qubit entropy (orange) and variance (blue), alongside square root and logarithm fits (dotted and dashed lines). The calculation are performed with a converged second-order Trotter product formula ($\delta t = 0.5$).

Two key observations emerge. First, the variance and entropy exhibit consistent behavior, which is significant since our random quantum channel only allows access to permutation-invariant observables. Therefore, the variance is a reasonable observable to look for thermalization, since it is consistent with the predictions using the entropy. Second, in the absence of simulation of larger systems, the current data is insufficient to distinguish between a square-root and a logarithmic scaling; both remain statistically compatible. Consequently, simulations at larger system sizes are required to resolve this ambiguity. Below we discuss several method to treat approximately the problem on a NISQ devices, such as Trotterization and our randomized channel. We do not delve into asymptotical optimal methods such as Quantum Singular Value Transformation [52] as the overhead is too high for current platforms.

III. TROTTERIZED DYNAMICS

As was shown in previous works [27, 53] a Trotter decomposition of the evolution operator can be implemented on a qubit chain by using a SWAP network where every pair interaction is followed by a SWAP. With this strategy a full Trotter step requires $N/2$ layers of nearest neighbor interactions. In order to express this construction in a compact way, let us first define a single layer unitary which applies nearest neighbor Heisenberg interactions on the $N/2$ even pairs followed by the $N/2-1$ odd pairs

$$W(t,\mathbf{h},\mathbf{p}) = \left[\prod_{j\ even} e^{-ith_j P_{j,j+1}} P_{j,j+1}^{p_j}\right] \times \left[\prod_{j\ odd} e^{-ith_j P_{j,j+1}} P_{j,j+1}^{p_j}\right], \tag{9}$$

where the vector $\mathbf{p}$ has $N-1$ components that can be either zero or one: if $p_j = 1$ then we perform a SWAP operation between qubits j and $j+1$ while we implement the interaction between them. Instead, the $N-1$ component vector $\mathbf{h}$ is used to retrieve the correct coupling constants for interactions. Using the fact that $P_{i,j}$ is idempotent, a more convenient expression for $W(t,\mathbf{h},\mathbf{p})$ can be found to be

$$W(t,\mathbf{h},\mathbf{p}) = \left[\prod_{j\ even} e^{-i\left(th_j + p_j\frac{\pi}{2}\right)P_{j,j+1}}\right] \times \left[\prod_{j\ odd} e^{-i\left(th_j + p_j\frac{\pi}{2}\right)P_{j,j+1}}\right], \tag{10}$$

which now makes it apparent that the addition of possible SWAP operations does not need to incur in additional gates being implemented [53]. Using this layer unitary as building block, and for an even number of particles N, a single first order Trotter step for the Hamiltonian in Eq. (4) can then be written as

$$\mathcal{U}_T^{(k)}(t) = \prod_{l=1}^{N/2} W(t,\mathbf{h}^{(k,l)},\mathbf{1}), \tag{11}$$

where $\mathbf{1}$ is a $N-1$ component vector containing 1 in every entry and the coupling constant vectors are given by

$$h_j^{(k,l)} = \frac{\mu}{N}\left(1 - \mathbf{v}_{i(j,l)}^{(k)} \cdot \mathbf{v}_{ip(j,l)}^{(k)}\right) = \mu_{i(j,l),ip(j,l)}^{(k)}. \tag{12}$$

The label $i(j,l)$ denotes the label of the neutrino associated to qubit j at the $l-th$ layer and similarly $ip(j,l)$ denotes the label of the neutrino associated to qubit $j+1$ at the same layer. The mapping between $i(j,l)$ and the qubit j depends on the permutations that have been used in previous layers, taking into account that for odd j we have applied already the even qubit layer in W. For instance, when $N=4$ and we perform a SWAP at every step, i.e. $\vec{p}^{(l)} = \vec{1}$, we have

$$\begin{array}{llll} i(0,1)=0 & ip(0,1)=1 & i(2,1)=2 & ip(2,1)=3 \\ & i(1,1)=0 & ip(1,1)=3 & \\ i(0,2)=1 & ip(0,2)=3 & i(2,2)=0 & ip(2,2)=2 \\ & i(1,2)=1 & ip(1,2)=2. & \end{array} \tag{13}$$

The Trotter unitary can then be used to approximate the correct time evolution operator leading to

$$\left\| \mathcal{U}^{(k)}(t) - \mathcal{U}_T^{(k)}(t/M)^M \right\| \leq \frac{t^2}{M} C^{(k)} \ , \tag{14}$$

where $C^{(k)}$ is a sum of norms of commutators [53], and M is the number of so-called Trotter steps.

IV. RANDOMIZED DYNAMICS

Simulating the flavor dynamics of neutrinos using product formulas like those described in the previous section is challenging on current and near-term quantum hardware due to the required depth growing at least linearly with system size [53]. In recent works, interacting neutrino systems have been shown to display intricate many-body correlations (see e.g. [24, 50, 54–58]) from a phenomenological point of view however the most important physical observables are single-neutrino properties (these dominate reactions in the supernova and are those measured on Earth). Our goal in this work is to design an approximate method that is able to track the evolution of single neutrino (qubit) observables in an efficient manner. We shall think of our method as a physically motivated ansatz, which is empirically able to capture permutation-invariant observables using only a circuit depth of $\mathcal{O}(t)$ —independent of the system size—at least up to intermediate time [59]. Although rigorous guarantees are beyond the scope of this work, establishing them remains an important avenue for future research.

From the discussion on the Trotterized dynamics in the previous section it is apparent that the full evolution circuit is very similar to the evolution of a, time-dependent, random Heisenberg chain. Correlations among the, otherwise random, rotation angles are critical to allow to uniquely associate a specific neutrino to an individual qubit at any given point in the various layers. The coupling constants $\mu_{i,j}^{(k)}$ in the Hamiltonian are random and uncorrelated with the neutrino labels (i, j). Consequently, when focusing on a single qubit, the evolution of its density matrix depends solely on the number of partial SWAP operations it has undergone since the initial state at $t = 0$. However, because these partial SWAPs only involve neutrinos with misaligned spins, the resulting qubit-level spin dynamics will generally depend on the initial assignment of qubits to neutrinos, even when the total flavor (i.e., the spin magnetization) is conserved. For instance, if we consider two states

$$\begin{aligned} |\phi_0\rangle &= |0\rangle^{\otimes N/2} \otimes |1\rangle^{\otimes N/2} \\ |\phi_1\rangle &= \left(|0\rangle \otimes |1\rangle\right)^{\otimes N/2} \ , \end{aligned} \tag{15}$$

in the first layer of nearest-neighbor Heisenberg interaction encoded in the unitary W from Eq. (10), only the central qubits will evolve for the state $|\phi_0\rangle$ while all of them will be affected for the state $|\phi_1\rangle$. The full Trotter step $\mathcal{U}_T^{(k)}(t)$ from Eq. (11) is invariant for relabeling of qubits to neutrinos (at least for short times t in general) so the final evolution is the same for both states.

Our strategy to compress the time-evolution unitary, while keeping the average evolution of one-body observables correct, relies on randomizing the coefficients in the individual layers where W acts and looking at the resulting dynamics for a varying number of layers which is in general now independent on the system size N. The resulting unitary takes then the following form

$$\mathcal{U}_R^{(k)}(\tau, L) = \prod_{l=1}^{L} W(\tau, \widetilde{\mathbf{h}}^{(k,l)}, \mathbf{1}) \ , \tag{16}$$

where the components of the coupling constant vectors $\widetilde{\mathbf{h}}^{(k,l)}$ are now chosen uniformly at random from the entire $N(N-1)/2$ components of the $\mu_{i,j}^{(k)}$ coupling matrix.

For a given choice of step size, the time parameter τ is then connected to the actual total time evolution but these two are not identical. Our rationale for decoupling the number of layers in the evolution unitary $\mathcal{U}_R^{(k)}(\tau, L)$ from the system size N is that, on average, each qubit undergoes $2L$ random pairwise interactions across the L alternating even–odd layers. Within the Trotterized dynamics, one application of $\mathcal{U}_T^{(k)}(t)$ involves $N/2$ sequential layers, each implementing pairwise rotations with angles of order $\mathcal{O}(\mu t/N)$. For systems described by initial states with a constant fraction of orthogonal flavor state as a function of N, the cumulative effect of these layers produces an overall flavor rotation of order $\mathcal{O}(\mu t)$. Since the randomized protocol implemented using $\mathcal{U}_R^{(k)}(\tau, L)$ is designed to approximate the average long-time dynamics, in the limit of large L we expect to be able to reproduce the overall rotation angle $\mathcal{O}(\mu t)$ using L layers of rotation of magnitude $\mathcal{O}(\mu\tau/N) = \mathcal{O}(\mu t/L)$ each. In other words, we expect a physical time evolution for time t will correspond to a value of the time-parameter τ given by

$$\tau = \alpha t \quad \text{with} \quad \alpha = \mathcal{O}\left(\frac{N}{L}\right) \ . \tag{17}$$

This argument however neglects the effect coming from the fact that, as the partial swaps proceed, the qubits progressively rotate away from their original polarization reducing the effect of the partial swaps. Indeed for two neutrinos with the same flavor polarization, the partial swap acts simply as the identity. We therefore expect that the total accumulated rotation angle will grow slower than linearly in the depth and we expect that $\alpha < N/L$ in practice. The actual dependence of α from N and L also strongly depends on the initial state used in the simulation as discussed previously when we considered the initial states in Eq. (15). In our simulations we account for the dependence on the initial state by performing, for each individual run, a random permutation to the qubits in the state $|\phi_0\rangle$ obtaining a state $|\phi_s\rangle$. Physical observable are then recovered by looking at the

distribution of outcomes generated by sampling both the Hamiltonian $\hat{H}^{(k)}_{\nu\nu}$ as well as the initial state.

For a given initial state and Hamiltonian we then evaluate the flavor variance as

$$V_R^{(k,s)}(\tau,L)=\frac{1}{N}\sum_j\left(\langle\phi_s|\mathcal{U}_R^{(k)}(\tau,L)^\dagger Z_j\mathcal{U}_R^{(k)}(\tau,L)|\phi_s\rangle\right)^2 . \tag{18}$$

The matching between the discrete values $V_R^{(k,s)}(\tau,L)$ to the corresponding variance at a physical time t is obtained through the proportionality constant α. In order to find this factor we impose that for short times the randomized expectation value

$$V_R^{(k,s)}(\tau,L)=1+g_R^{(k,s,L)}\frac{\tau^2}{2}+\mathcal{O}(\tau^3) \tag{19}$$

matches the ideal variance

$$V^{(k)}(t)=1+g^{(k)}\frac{t^2}{2}+\mathcal{O}(t^3)\,, \tag{20}$$

where the linear term is zero by construction for our initial states while

$$\begin{aligned} g^{(k)}&\equiv\frac{d^2}{dt^2}V^{(k)}(t)\Big|_{t=0}\\ &=\frac{1}{2N}\sum_{ij}\left(\mu_{i,j}^{(k)}\right)^2(1-\langle Z_i\rangle\langle Z_j\rangle)\,, \end{aligned} \tag{21}$$

where the expectation values are evaluated on the initial state [60]. For more details on the derivation see Appendix A and Refs. [19, 46]. Using this strategy, the rescaling factor between t and τ is then

$$\alpha^{(k,s,L)}=\sqrt{\frac{g^{(k)}}{g_R^{(k,s,L)}}}\,. \tag{22}$$

We show in Fig. 2 that the rescaling factor is in fact constant as we increase the number of layers.

We now explore two separate protocols to perform this short time matching which are equivalent in the asymptotics but behave differently in the intermediate regime of interest. While our protocols share similarities in spirit with the qDrift protocol [61–63], the key distinction lies in the channel distance we are interested in minimizing. The goal of the qDrift protocol is to implement a random unitary channel $\Lambda_{qDrift}(t)[\rho]=\mathcal{V}_R(t)\rho\mathcal{V}_R^\dagger(t)$ whose averages matches the exact unitary dynamics at fixed couplings for short times

$$\left\|\mathbb{E}\left[\Lambda_{\rm qDrift}(t)[\rho]\right]-\mathcal{U}^{(k)}(t)\rho\mathcal{U}^{(k)}(t)^\dagger\right\|=\mathcal{O}(t^2)\,. \tag{23}$$

In our case instead we are first only interested in the distance between the following comvex superposition of reduced density matrices

$$\rho_{\rm rnd}[\rho]=\frac{1}{N!}\sum_{\sigma\in S_N}\bigotimes_{j=1}^N{\rm Tr}_{N/\sigma(j)}[\rho]\,, \tag{24}$$

where $\sigma(j)$ is an arbitrary permutation of the qubit labels taken for the symmetric group S_N and ${\rm Tr}_{N/\sigma(j)}[\rho]$ is the single qubit reduced density matrix for the $\sigma(j)$ neutrino. This is then a much weaker condition than reproducing the full density matrix itself. The new distance reads

$$\left\|\rho_{\rm rnd}\left[\mathcal{U}^{(k)}(t)\rho\mathcal{U}^{(k)}(t)^\dagger\right]-\mathbb{E}_l\left[\rho_{\rm rnd}\left[\mathcal{U}_R^{(k)}(\tau)\rho^l\mathcal{U}_R^{(k)}(\tau)^\dagger\right]\right]\right\|\,, \tag{25}$$

and the average is over initial states obtained by qubit permutations, which give rise to the same final density matrix under exact evolution. As explained above the value of τ is found by matching the short time behavior and the goal is for the randomized procedure to provide a small value for this distance, at least one averaged over the random set of Hamiltonians $\hat{H}^{(k)}_{\nu\nu}$, for times comparable to the thermalization time.

We will show below numerical evidence suggesting that our choices for implementing the randomized channel indeed are able to keep the distance in Eq. (25) small. An interesting and important extension of the present work would be to provide more rigorous bounds on it.

A. Protocol A

The first possibility is to pick a fixed value of L for all times t and consider

$$V_A^{(k,s,L)}(t)=V_R^{(k,s)}\left(\frac{t}{L}\alpha^{(k,s,L)},L\right)\,. \tag{26}$$

This approach resembles a Trotter decomposition with a fixed depth where different evolution times are obtained by changing the time-step t/L used for each individual layer. In the limit $L\to\infty$ the randomized evolution approximates progressively well the exact evolution under the average Hamiltonian. To see this we can look for the effective generator $H_{\rm eff}$ by expanding at small times since we can write

$$\mathcal{U}_R^{(k)}\left(\frac{t}{L},L\right)=\prod_{m=1}^M\mathcal{U}_R^{(k)}\left(\frac{t}{L},\frac{L}{M}\right)\,, \tag{27}$$

and for $L\to\infty$ at fixed M we have

$$\begin{aligned} \mathcal{U}_R^{(k)}\left(\frac{t}{L},\frac{L}{M}\right)&=\prod_{l=1}^{2L/(NM)}\prod_{n=1}^{N/2}W\left(\frac{\tau}{L},\tilde{\mathbf{h}}^{(k,Nl/2+n)},\mathbf{1}\right)\\ &\approx\mathbb{1}-i\frac{t}{L}\sum_{l=1}^{2L/(NM)}\sum_{i<j}^N\mu_{i,j}^{(k,l)}P_{i,j} \end{aligned} \tag{28}$$

so that we have as effective Hamiltonian,

$$H_{\rm eff}=\sum_{i<j}^N\left(\sum_{l=1}^{2L/(NM)}\mu_{i,j}^{(k,l)}\right)P_{i,j}\to\frac{2L}{NM}\mathbb{E}[\hat{H}^{(k)}_{\nu\nu}]\,. \tag{29}$$

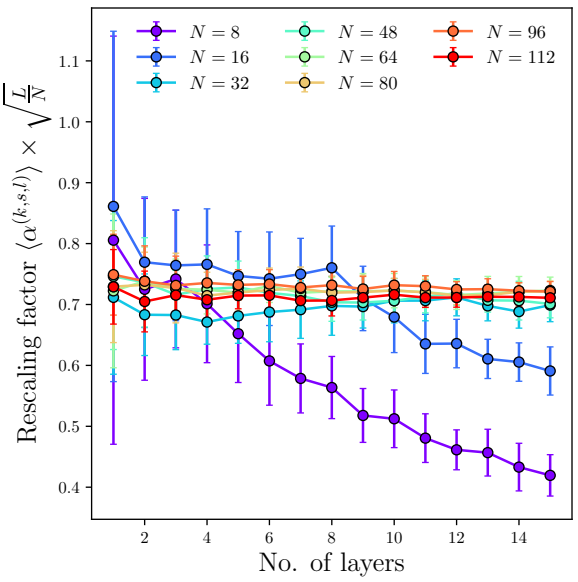

Figure 2: **Rescaling factor.** The rescaling factor normalized with $\sqrt{N/l}$ shown as a function of the number of layers for different system sizes. The expectation value and standard deviation is computed across Hamiltonian realization and initial state (k, s) at fixed depth l.

This is then exactly the average Hamiltonian and we have

$$\mathcal{U}_R^{(k)}\left(\frac{t}{L}, L\right) \to \exp\left(-i\frac{2t}{N}\mathbb{E}[\hat{H}_{\nu\nu}^{(k)}]\right), \tag{30}$$

for a sufficiently large number of layers. This provides insight into the convergence of our random channel. Yet, because the dynamics under the average Hamiltonian can be completely characterized analytically and is in general different from the correct dynamical evolution, we must modify the protocol to capture the more interesting typical evolution of a random Hamiltonian.

B. Protocol B

This time, we instead fix a time-step δt and vary the number of layers L to reach different final evolution times. In this way, the physical flavor variance as a function of time is estimated as

$$V_B^{(k,s,\tau)}(t = l\delta t) = V_R^{(k,s)}\left(\delta t \alpha_B^{(k,s,l)}, l\right), \tag{31}$$

for discrete times of step δt. The limit $L \to \infty$ now corresponds to $t \to \infty$ in steps of constant size.

C. Comparison between the two protocols

We aim to compute the variance of the neutrino flavor for a typical Hamiltonian $\hat{H} \sim H$. To this end, we repeatedly sample a Hamiltonian $\hat{H}$ together with an initial state (a randomly shuffled wall configuration) and evaluate the median and 25%-75% percentiles of the resulting dynamics. Consequently, a randomly chosen Hamiltonian from this family has a 50% probability of lying within the shaded region. The simulation of the time evolution is performed stochastically, as described below.

We first analyze protocol A, shown in Fig. 3. We observe that the results converge, as the number of layers increases (colors in the figure), toward the dynamics governed by the time-averaged Hamiltonian $\bar{H} = \mathbb{E}[\hat{H}_{\nu\nu}^{(k)}]$ (black dashed line). Moreover, the protocol remains accurate even beyond the validity of the second-order Taylor expansion (grey dashed line). However, protocol A reproduces the dynamics of a typical Hamiltonian only at short times.

We next consider protocol B, shown in Fig. 4. Here too, the dynamics converge in the limit of a large number of layers to those of the time-averaged Hamiltonian (black dashed line) and remain accurate beyond the breakdown of the second-order Taylor expansion (grey dashed line). In contrast to protocol A, convergence occurs more slowly, leading to a region around $L = t_{\text{end}}$ where the protocol closely follows the exact dynamics up to intermediate times. Thus, protocol B provides a useful approximation for the dynamics of a typical Hamiltonian within this intermediate-time regime. The behavior of the evolution of the average Hamiltonian at $N = 10$ is attributed to finite size effect.

V. RESULTS

We are interested in understanding the scaling of the thermalization time in systems of forward-scattering neutrinos. While a rigorous lower bound of $\log N$ can be established from the Lieb–Robinson velocity [64], a meaningful upper bound remains unknown. To probe thermalization, we define a practical proxy: the time required for the system to reach a threshold value $\Gamma = 0.1$ (see Eq. (8)). This criterion is analogous to the use of single-qubit entropy as a thermalization indicator in Ref. [19].

In what follows, we systematically study this proxy for thermalization as a function of system size. Our analysis is carried out using the randomized protocol B with $L = t$, implemented across multiple computational backends: tensor-network simulations [65], Pauli propagation methods [66], and quantum processing units (QPUs) provided by IBM. This multi-platform approach enables us to benchmark the protocol's performance, compare classical and quantum resources, and gain insight into the scaling of thermalization in regimes beyond purely analytic control. We will also compare our results with the aforementioned phase-space approach, which differs in that it does not simulate the randomized protocol directly, and is therefore an independent method.

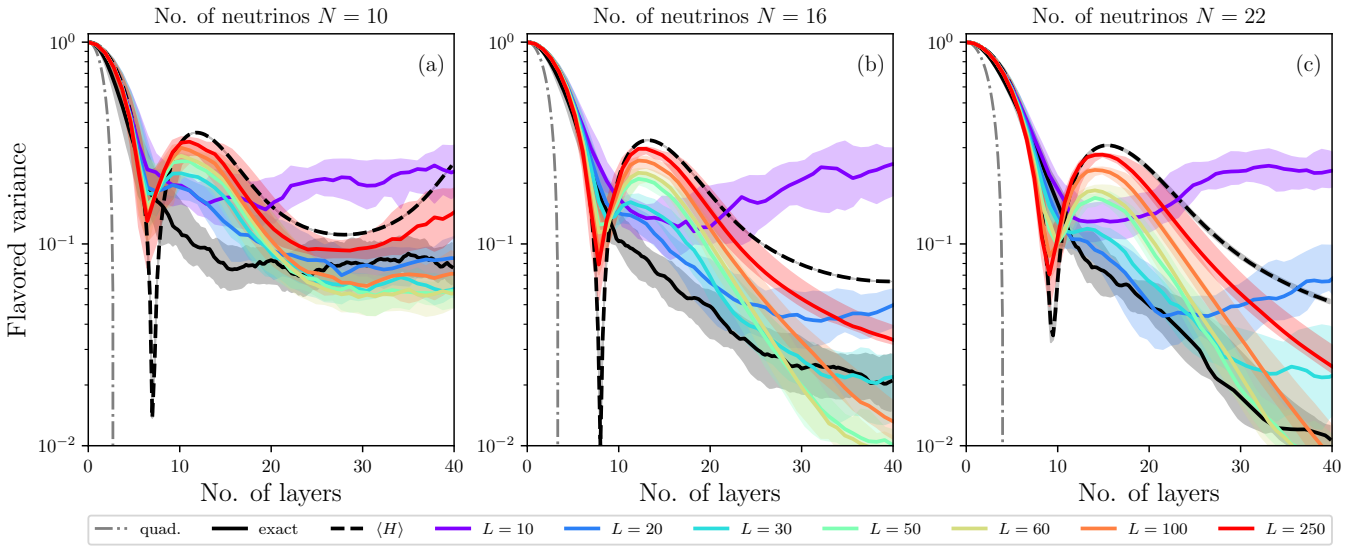

Figure 3: **Protocol A.** The variance of a single-qubit observable, denoted $V(t)$, is shown as a function of time under protocol A for different number of (fix) layers, time step $\tau = t/L$ and system sizes: (a) $N = 10$, (b) $N = 16$, (c) $N = 22$. The black solid line represents the exact dynamics computed using a converged product formula, while the colored line corresponds to the evolution with the randomized protocol. The shaded band around the latter indicates the interquartile range across disorder realizations.

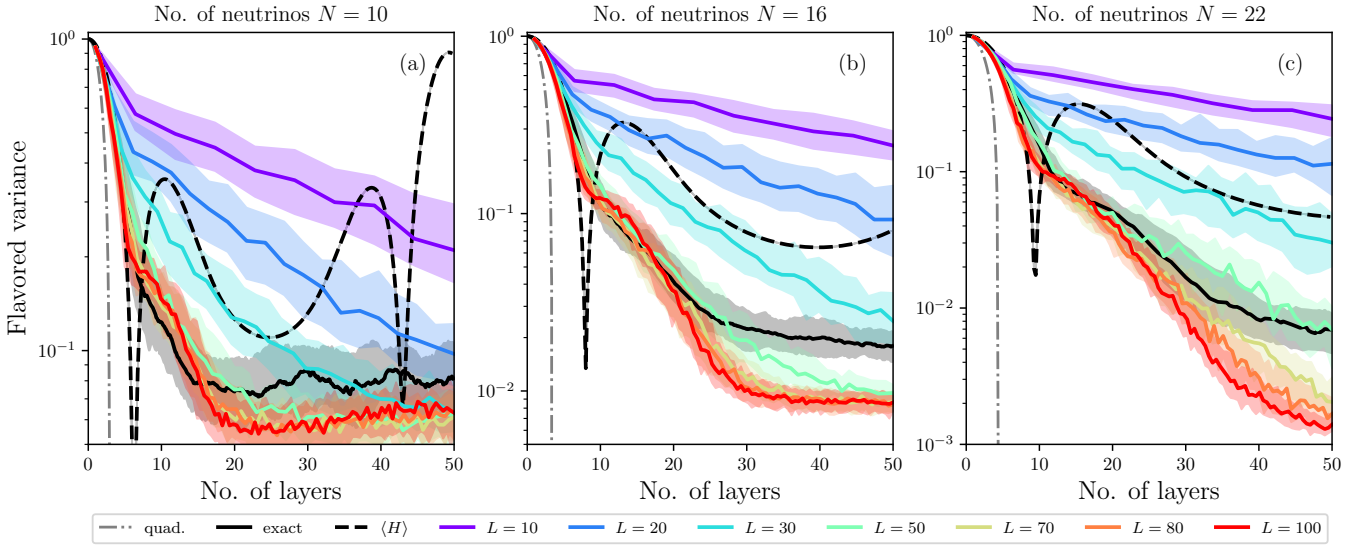

Figure 4: **Protocol B.** The variance of a single-qubit observable, denoted $V(t)$, is shown as a function of time under protocol B for different number of layers L, time step $\tau = t_{\text{end}}/L$ and system sizes: (a) $N = 10$, (b) $N = 16$, (c) $N = 22$. The black solid line represents the exact dynamics computed using a converged product formula, while the colored line corresponds to the evolution under the random channel. The shaded band around the latter indicates the interquartile range across disorder realizations.

A. Classical simulation

When attempting to use a quantum computer, it is prudent to first assess whether the same task can be efficiently and reliably performed on classical hardware. In the following, we will review state of the art protocols based on tensor network, phase space approximation and pauli propagation.

a. Tensor networks methods provide a powerful framework for simulating many-body quantum systems, particularly those with limited entanglement. Among these, matrix product states (MPS) [65] are especially ef-

fective for one-dimensional systems due to their efficient representation of quantum states with area-law entanglement. In essence, MPS encode a quantum state as a sequence of tensors connected linearly, with the bond dimension χ controlling the amount of entanglement the representation can capture.

In the following, we present numerical results obtained using MPS and demonstrate that this method is not suitable for our task, as the entanglement entropy grows linearly with time. We simulate our random channel using the *quimb* library [67], varying the bond dimension $\chi \in 2^{\{3,\dots,11\}}$, and show in Fig. 5 that convergence requires saturating the bond dimension to $\chi = 2^{N/2-1}$. The inset instead shows the relative error between the MPS simulation and the exact one, where for the larger system casewe take the calculation with the largest bond dimension as a proxy.

To mitigate this limitation, one can extrapolate observables to the limit $1/\chi \to 0$ in an attempt to approximate the exact result. To this end, we fit the inverse bond dimension $\vec{x}_t = 1/\chi$ to the variance $y_t = V(t;\chi)$ using the exponential form

$$f(\chi) = a \exp\left(-b\chi^{-c}\right), \quad \text{with } a, b, c > 0. \tag{32}$$

To improve robustness, we repeat the fit over random initial conditions and report the mean and standard deviation as a black dotted line in Fig. 5. Notably, the error bars increase as entanglement builds up over time, indicating that the extrapolation becomes unstable. Consequently, we do not expect this approach to yield accurate dynamics at larger system sizes or longer evolution times. In Fig. 6, we compare the MPS simulation against the other simulation protocols.

b. The Phase-Space Approximation (PSA) was originally introduced in nuclear physics [68–70] to treat many-body fermionic problems beyond the independent particle approximation, also called the mean-field approximation. Instead of considering a single mean-field evolution, consider a statistical ensemble of independent mean-field evolutions that differ from their initial conditions. In spirit, the PSA approach can be seen as a mapping from a quantum problem to a classical problem similar to what can be done with general quantum mechanical problems [71], except that here the mapping is made at the many-body level. The main hypothesis of the PSA is that (i) at initial time, a quantum system can be mapped to the initial ensemble by imposing that at least the first and second moments of one-body observables of the quantum systems match the statistical average over the ensemble. Eventually, higher moments can also be imposed [72]; (ii) Each initial condition evolves in time using the mean-field, which is considered here as the most "classical" equation of motion for the set of interacting particles; (iii) Since the trajectories are supposed to be independent, interferences between trajectories that might affect the long-time evolution are neglected.

Despite these assumption, the PSA turns out to significantly improves over the mean-field approach by providing in general rather accurate evolution of both one- and two-body observables evolutions for a wide range of coupling strengths. It was successfully applied to a wide range of physical systems ranging from permutation invariant systems presenting a quantum phase transition [69, 73, 74], small superfluid systems, or electronic systems on a lattice [70].

As shown recently [75], this technique offers a competitive approach to interacting-neutrino dynamics under both time-independent and time-dependent couplings. Most recent developments of the PSA include the possibility to treat an arbitrary number of neutrino beams [72], assuming that each neutrinos have only two flavors, or the more realistic case where neutrinos can have three flavors [76]. One key aspect of the PSA is that it inherits the simplicity of the mean-field approach, and the number of coupled equations to solve along each trajectory scales linearly with the number of neutrinos N (as $3N$ and $8N$ for the 2- and 3-flavor case). Since the trajectories are independent, it can also be easily parallelized on a CPU farm. In Refs. [72] and [76], the approach was shown to apply to simulations of several hundred interacting neutrinos.

Here we use the PSA approach employing directly the technique developed in Ref. [72] for multi-beams, with the specific limit of one neutrino per beam. All technical details are already given in [72] and here we only give the results obtained using the random Hamiltonian proposed in this article. Focusing first on the Hamiltonian (9), for N neutrinos having two possible flavor states, the PSA evolution along each path reduces to a set of $3N$ coupled equation of motion given by:

$$\partial_t \vec{\Sigma}_i^{\lambda}(t) = \sum_{j\neq i}^{N} \mu_{ij} \vec{\Sigma}_j^{\lambda}(t) \wedge \vec{\Sigma}_i^{\lambda}(t), \tag{33}$$

where $\vec{\Sigma}_i^{\lambda}(t)$ is associated to the neutrino i and has three real components $\left[X_i^{\lambda}(t), Y_i^{\lambda}(t), Z_i^{\lambda}(t)\right]$, where $\lambda = 1, \cdots, N_{\rm evt}$ labels the events. Denoting generically by $\mathcal{P}_i$ one of the Pauli matrices associated to the neutrino "i", within the phase-space approach any quantum expectation value of a Pauli string is replaced by a statistical average over the events, i.e.,

$$\langle \mathcal{P}_{i_1} \cdots \mathcal{P}_{i_k} \rangle_t \xleftrightarrow{PSA} \overline{\mathcal{P}_{i_1}^{\lambda}(t) \cdots \mathcal{P}_{i_k}^{\lambda}(t)}.$$

$\mathcal{P}_i$ in the left side denotes one of the Pauli operators, while $\mathcal{P}_i^{\lambda}(t)$ denotes one of the $\vec{\Sigma}_i^{\lambda}(t)$ components. The average on the right-hand side of a generic observable A, denoted by $\overline{A^{\lambda}}$ stands for the statistical average over events, e.g. $\overline{A^{\lambda}} = (N_{\rm evt})^{-1} \sum_{\lambda} A^{\lambda}$. The PSA trajectories given by (33) are solved assuming randomized initial conditions, such that the statistical average over events at initial time, reproduces at least the mean one-body observables and their fluctuations. The sampling method we used assumes bi-valued probabilities as prescribed in Ref. [72]. With this method, average evolution of one-

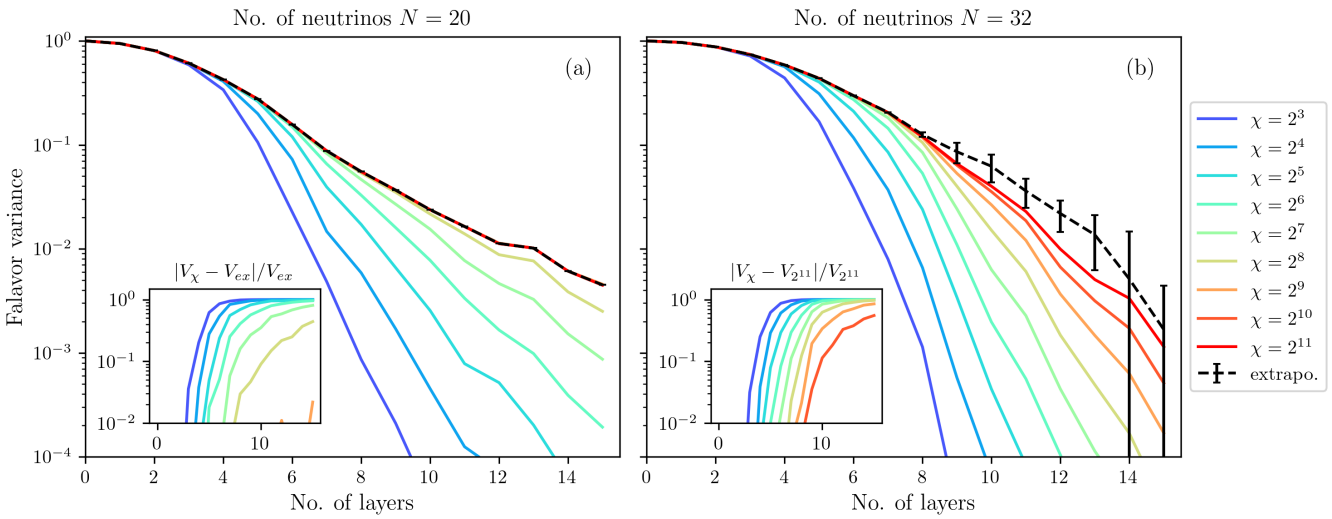

Figure 5: **MPS simulations.** The variance $V(t)$ is shown as a function of the number of layers using a fixed time step $\delta t = 1$ for different bond dimension and system sizes $N = 20$ (a) and $N = 32$ (b). The insets shows the relative absolute error with the exact evolution, where for the larger system we use the largest bond dimension as a proxy.

neutrino or two-neutrinos observables, as well as associated one- and two-neutrinos entanglement entropies have been evaluated in general with good accuracy compared to available exact results (see [72, 75]). In the present work, we use the PSA method as a reference for classical simulation of large scale systems. Even if PSA does not have guarantees, it serves as an empirical independent benchmark.

c. Pauli propagation (PP) [66] is another powerful classical simulation technique, which operates in the Heisenberg picture by propagating observables backward through the quantum circuit. The method approximates the observable by expressing it as a sum over Pauli strings and truncating the expansion by discarding terms whose Pauli weight exceeds a maximum threshold W, or whose coefficients fall below a fixed threshold ε.

We perform simulations for system size $N = 20$, circuit depth $L = 15$, and parameter choices $W \in \{4, 6\}$ and $\varepsilon \in \{10^{-4}, 10^{-6}\}$, shown in the left inset of Fig. 6. The resulting absolute error remains large in the region of interest, which is insufficient in this context. While the use of this method out of the box does not yield satisfactory results, a more systematic investigation would be required to definitively assess the applicability of PP in this setting.

We compare all of our numerical simulation strategies in Fig. 6. In (a) we display the exact (black) MPS (blue), Pauli Propagation (green), PSA (violet) and statevector (red) for $N = 20$ while in (b) we show results on the QPU without (orange) and with (red) error mitigation compared to MPS (blue) for $N = 32$. We observe that MPS is only accurate when the bond dimension saturates the size of the Hilbert space ($\chi = 2^{N/2}$), whereas PP captures only short-time dynamics and would require higher-order terms to achieve useful accuracy in this setting. On the other hand, PSA is able to reproduce the exact evolution up to long simulation times, and is thus expected to be a strong benckmark when scaling up to larger systems.

B. Results on quantum hardware

Error mitigation is an essential technique for harnessing the computational capabilities of noisy quantum devices. Our protocol builds on a noise renormalization approach, which has demonstrated strong performance in previous studies [77–79]. The core idea is to model the noise as a depolarizing channel, estimate its parameters directly on hardware, and then invert its effect in post-processing.

We recall that a depolarizing noise channel $\mathcal{N}_p[\cdot]$, parameterized by p, affects the expectation value of a Pauli observable σ under a quantum state ρ as follows:

$$\mathrm{Tr}\left(\sigma \mathcal{N}_p[\rho]\right) = (1-p)\mathrm{Tr}\left(\sigma\rho\right). \tag{34}$$

When the depolarizing parameter p is known, noisy expectation values can be corrected by dividing them by the factor $(1-p)$. In our implementation, we estimate p using the all-zero initial state $|0\rangle^{\otimes N} \equiv |\bar{0}\rangle$, which remains invariant under the Hamiltonian evolution

$$\langle\bar{0}|\mathcal{U}(t)^\dagger Z_j \mathcal{U}(t)|\bar{0}\rangle = 1. \tag{35}$$

We can then use this initial state to estimate the survival probability $(1-p)$ and employ it to renormalize the noisy estimates. Selecting an initial state that is invariant under the dynamics is, in general, a more robust

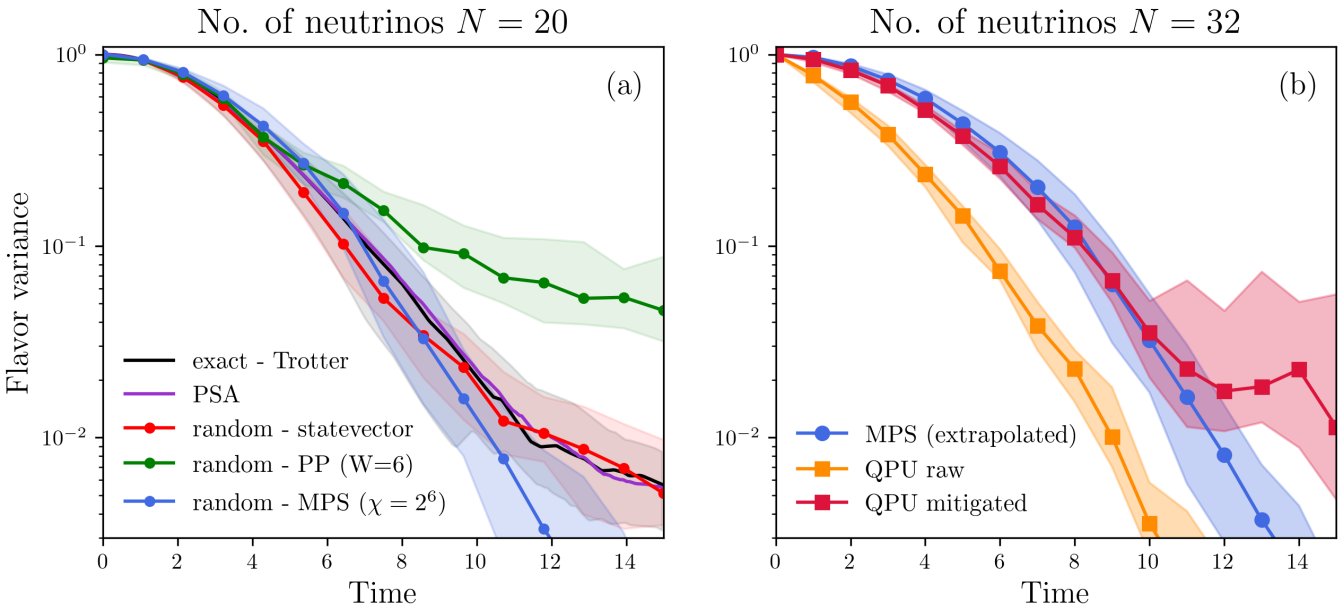

Figure 6: **Comparing different backend.** The variance $V(t)$ is shown as a function of the number of layers using a fixed time step $\delta t = 1$ for different simulation methods. (a) we compare the exact (black), MPS (blue), Pauli Propagation (green), PSA (violet) and statevector (red) for $N = 20$ while in (b) we show results on the QPU without (orange) and with (red) error mitigation compared to MPS (blue) for $N = 32$.

strategy than attempting to invert the evolution using $\mathcal{U}(t/2)\mathcal{U}(-t/2)$. Indeed, such an inversion does not reproduce the same set of Pauli paths as the forward evolution in the target computation. As a result, the noise acts differently on the system, which constitutes a significant limitation of this approach. It is also worth noting that this method is most effective under the assumption of depolarizing noise, which may not fully capture the behavior of real quantum devices. Nevertheless, this limitation can be alleviated through Pauli twirling [80]. Measurement (readout) errors are addressed via device calibration [81] and by twirling the measurement channel [82].

To validate our error mitigation protocol, we evaluate it on a smaller system with $N = 32$, as shown Fig. 6 (b). This choice ensures that the MPS calculation is effectively exact, as, at this size, we can work at the maximal bond dimension. We find that our results successfully reproduce the MPS benchmark within the error bars, except for the last few points where accumulated noise becomes dominant.

Finally, we compute the characteristic time of the variance for $N \in \{32, 48, 64, 80, 96, 112\}$ on the 156 qubits ibm devices ibm_fez and ibm_aachen, using the experimental setting described in Tab. I. The characteristic time is shown as a function of the system size in Fig. 7 for the different methods (exact in black, MPS in blue, QPU in red, and PSA in orange), alongside square-root (dotted) and logarithm (dashed) fits. We observe that the PSA and QPU results are compatible with each other within errorbars and exhibit a square-root scaling. For small system sizes ($N < 32$), where exact classical simulations are feasible, all methods (QPU, PSA, and MPS) show excellent agreement. For larger N, the MPS results start to deviate and systematically underestimate Γ, highlighting the difficulty of accurately simulating quantum dynamics. For large system sizes, both the QPU and PSA results continue to increase, exhibiting a scaling behavior compatible with $\sqrt{N}$. The PSA yields slightly higher values of Γ than the QPU. It remains unclear whether this discrepancy originates from the approximate nature of the PSA or from residual errors in the QPU results despite mitigation. As indicated in Fig. 6 (b), unmitigated evolutions (orange curve) tend to underestimate Γ, suggesting that imperfect error mitigation could account for the observed difference. Notably, as seen in the inset of Fig. 7, the QPU data points (red) systematically lie below the exact results (black) even for small N.

We can nevertheless conclude from these results that the scaling of the thermalization time obtained using both PSA and our randomized methods is consistent, as the quadratic fit is clearly the best in both cases.

VI. CONCLUSIONS

We have introduced a randomized quantum approach capable of efficiently capturing global properties in typical random systems. Although we do not provide an analytical proof, our numerical results for small systems show that the method indeed reproduces on average the

Table I: Experimental settings.

No. of configurations	$10 \times L$
No. of twirls (total)	75
No. of shots	500
Quantum devices used	`ibm_fez`, `ibm_aachen`
Maximum number of layers	$L = 15$
Time step	$\tau = 1$
CNOT depth	$3 \times 2 \times L \leq 90$
Error mitigation overhead	$\times 2$
QPU time	22-24 [s] per config.

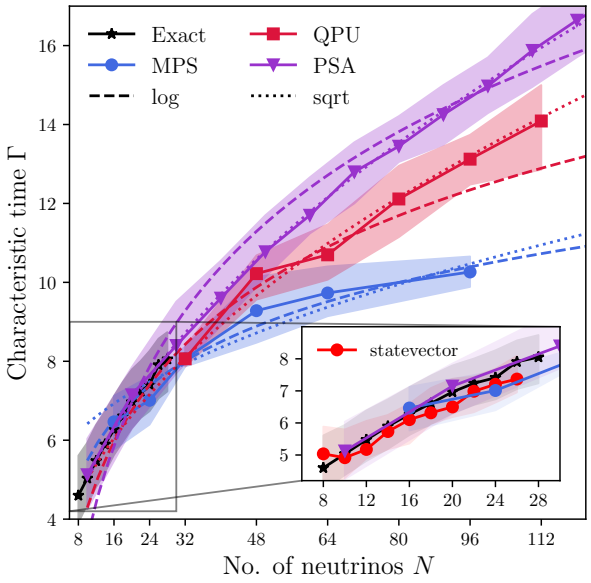

Figure 7: **Scaling of the thermalization time.** The characteristic time Γ is shown as a function of the system size N for the different methods (exact in black, MPS in blue, QPU in red, and PSA in violet), alongside square-root (dotted) and logarithmic (dashed) fits.

exact flavor dynamics for single qubit observables while requiring only a $\mathcal{O}(t)$ maximal circuit depth. Moreover, we demonstrate that this class of problems is challenging for tensor-network methods such as MPS, due to the linear growth of entanglement with time.

Applying our approach to systems of neutrinos interacting only through forward-scattering we find that the thermalization time scales as the square root of the system size. This behavior is consistent with the the semi-classical PSA approach and thus providing strong supporting evidence for the validity of the semi-classical treatment of this problem. This result illustrates how NISQ devices can serve as empirical validators for classical approximation methods. Thanks to their better scaling with system size, these semi-classical methods can then be used to both explore larger systems or to tackle more challenging simulation problems. In the context of flavor evolution of neutrinos this includes both the extension to the full three flavor case [83–87] and the inclusion of non-forward scattering between neutrinos [88].

More broadly, our work proposes a novel application of random quantum circuits (RQC) in physics. While RQC are known to be difficult to simulate classically [89], they have so far found limited use in physical modeling. Here, we demonstrate how RQCs can provide physical insight by connecting them to the dynamics of random Hamiltonians.

Finally, our results fit within quantum utility experiments: using NISQ devices as experimental platforms to test, benchmark, and refine empirical classical methods. Although our current implementation is limited to specific observables and lacks rigorous guarantees, it opens the door to extending such quantum-assisted benchmarking to more complex quantities, such as two-point correlation functions [90].

ACKNOWLEDGMENTS

We thank J. Carlson, N. Mueller and D. Neill for discussions about the topics covered in this work. We acknowledge the use of IBM Quantum services for this work. The views expressed are those of the authors and do not necessarily reflect the official policy or position of IBM or the IBM Quantum team. OK is supported by CERN through the CERN Quantum Technology Initiative, as well by a Doc. Mobility fellowship from the University of Geneva. This project has received financial support from the CNRS through the AIQI-IN2P3 project. This work is part of HQI initiative (www.hqi.fr) and is supported by France 2030 under the French National Research Agency award number "ANR-22-PNQC-0002". This research was also partially funded by the Swiss National Science Foundation through the project RESQUE (grant number 225229).

[1] J. Pantaleone, Neutrino oscillations at high densities, Physics Letters B **287**, 128–132 (1992).
[2] H.-T. Janka, K. Langanke, A. Marek, G. Martínez-Pinedo, and B. Müller, Theory of core-collapse supernovae, Physics Reports **442**, 38 (2007).
[3] S. Woosley and T. Janka, The physics of core-collapse supernovae, Nature Physics **1**, 147 (2005).
[4] R. D. Hoffman, S. E. Woosley, and Y. Qian, Nucleosynthesis in neutrino-driven winds. ii. implications for heavy element synthesis, The Astrophysical Journal **482**, 951 (1997).
[5] G. Sigl and G. Raffelt, General kinetic description of relativistic mixed neutrinos, Nuclear Physics B **406**, 423 (1993).

[6] S. Pastor and G. Raffelt, Flavor oscillations in the supernova hot bubble region: Nonlinear effects of neutrino background, Phys. Rev. Lett. **89**, 191101 (2002).

[7] S. Pastor, G. Raffelt, and D. V. Semikoz, Physics of synchronized neutrino oscillations caused by self-interactions, Phys. Rev. D **65**, 053011 (2002).

[8] N. F. Bell, A. A. Rawlinson, and R. Sawyer, Speed-up through entanglement—many-body effects in neutrino processes, Physics Letters B **573**, 86–93 (2003).

[9] A. B. Balantekin and Y. Pehlivan, Neutrino–neutrino interactions and flavour mixing in dense matter, Journal of Physics G: Nuclear and Particle Physics **34**, 47–65 (2006).

[10] G. G. Raffelt and A. Y. Smirnov, Adiabaticity and spectral splits in collective neutrino transformations, Physical Review D **76**, 10.1103/physrevd.76.125008 (2007).

[11] H. Duan, G. M. Fuller, and Y.-Z. Qian, Collective neutrino flavor transformation in supernovae, Phys. Rev. D **74**, 123004 (2006), arXiv:astro-ph/0511275.

[12] H. Duan, G. M. Fuller, J. Carlson, and Y.-Z. Qian, Coherent development of neutrino flavor in the supernova environment, Physical Review Letters **97**, 10.1103/physrevlett.97.241101 (2006).

[13] H. Duan, G. M. Fuller, and Y.-Z. Qian, Collective neutrino oscillations, Annual Review of Nuclear and Particle Science **60**, 569 (2010).

[14] S. Chakraborty, R. Hansen, I. Izaguirre, and G. Raffelt, Collective neutrino flavor conversion: Recent developments, Nuclear Physics B **908**, 366–381 (2016).

[15] L. Johns, S. Richers, and M.-R. Wu, Neutrino oscillations in core-collapse supernovae and neutron star mergers, Annual Review of Nuclear and Particle Science https://doi.org/10.1146/annurev-nucl-121423-100853 (2025).

[16] A. V. Patwardhan, M. J. Cervia, E. Rrapaj, P. Siwach, and A. B. Balantekin, Many-body collective neutrino oscillations: Recent developments, in *Handbook of Nuclear Physics*, edited by I. Tanihata, H. Toki, and T. Kajino (Springer Nature Singapore, Singapore, 2020) pp. 1–16.

[17] A. B. Balantekin, M. J. Cervia, A. V. Patwardhan, E. Rrapaj, and P. Siwach, Quantum information and quantum simulation of neutrino physics, The European Physical Journal A **59**, 10.1140/epja/s10050-023-01092-7 (2023).

[18] M. C. Volpe, Neutrinos from dense environments: Flavor mechanisms, theoretical approaches, observations, and new directions, Rev. Mod. Phys. **96**, 025004 (2024).

[19] J. D. Martin, D. Neill, A. Roggero, H. Duan, and J. Carlson, Equilibration of quantum many-body fast neutrino flavor oscillations, Phys. Rev. D **108**, 123010 (2023).

[20] R. F. Sawyer, "classical" instabilities and "quantum" speed-up in the evolution of neutrino clouds (2004).

[21] 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** (03), 042.

[22] B. Dasgupta, A. Mirizzi, and M. Sen, Fast neutrino flavor conversions near the supernova core with realistic flavor-dependent angular distributions, Journal of Cosmology and Astroparticle Physics **2017** (02), 019.

[23] J. D. Martin, J. Carlson, and H. Duan, Spectral swaps in a two-dimensional neutrino ring model, Phys. Rev. **D101**, 023007 (2020), arXiv:1911.09772 [hep-ph].

[24] A. Roggero, E. Rrapaj, and Z. Xiong, Entanglement and correlations in fast collective neutrino flavor oscillations, Phys. Rev. D **106**, 043022 (2022).

[25] Z. Xiong and Y.-Z. Qian, Stationary solutions for fast flavor oscillations of a homogeneous dense neutrino gas, Physics Letters B **820**, 136550 (2021).

[26] V. Amitrano, A. Roggero, P. Luchi, F. Turro, L. Vespucci, and F. Pederiva, Trapped-ion quantum simulation of collective neutrino oscillations, Phys. Rev. D **107**, 023007 (2023).

[27] B. Hall, A. Roggero, A. Baroni, and J. Carlson, Simulation of collective neutrino oscillations on a quantum computer, Phys. Rev. D **104**, 063009 (2021).

[28] F. Tacchino, A. Chiesa, S. Carretta, and D. Gerace, Quantum computers as universal quantum simulators: State-of-the-art and perspectives, Advanced Quantum Technologies **3**, 1900052 (2020).

[29] A. Miessen, P. J. Ollitrault, F. Tacchino, and I. Tavernelli, Quantum algorithms for quantum dynamics, Nature Computational Science **3**, 25 (2023).

[30] Y. Kim, A. Eddins, S. Anand, K. X. Wei, E. Van Den Berg, S. Rosenblatt, H. Nayfeh, Y. Wu, M. Zaletel, K. Temme, *et al.*, Evidence for the utility of quantum computing before fault tolerance, Nature **618**, 500 (2023).

[31] L. E. Fischer, M. Leahy, A. Eddins, N. Keenan, D. Ferracin, M. A. Rossi, Y. Kim, A. He, F. Pietracaprina, B. Sokolov, *et al.*, Dynamical simulations of many-body quantum chaos on a quantum computer, arXiv preprint arXiv:2411.00765 (2024).

[32] A. Di Meglio, K. Jansen, I. Tavernelli, C. Alexandrou, S. Arunachalam, C. W. Bauer, K. Borras, S. Carrazza, A. Crippa, V. Croft, R. de Putter, A. Delgado, V. Dunjko, D. J. Egger, E. Fernández-Combarro, E. Fuchs, L. Funcke, D. González-Cuadra, M. Grossi, J. C. Halimeh, Z. Holmes, S. Kühn, D. Lacroix, R. Lewis, D. Lucchesi, M. L. Martinez, F. Meloni, A. Mezzacapo, S. Montangero, L. Nagano, V. R. Pascuzzi, V. Radescu, E. R. Ortega, A. Roggero, J. Schuhmacher, J. Seixas, P. Silvi, P. Spentzouris, F. Tacchino, K. Temme, K. Terashi, J. Tura, C. Tüysüz, S. Vallecorsa, U.-J. Wiese, S. Yoo, and J. Zhang, Quantum computing for high-energy physics: State of the art and challenges, PRX Quantum **5**, 037001 (2024).

[33] J. C. Halimeh, N. Mueller, J. Knolle, Z. Papić, and Z. Davoudi, Quantum simulation of out-of-equilibrium dynamics in gauge theories, arXiv preprint arXiv:2509.03586 https://doi.org/10.48550/arXiv.2509.03586 (2025).

[34] Y. Chai, A. Crippa, K. Jansen, S. Kühn, V. R. Pascuzzi, F. Tacchino, and I. Tavernelli, Fermionic wave packet scattering: a quantum computing approach, Quantum **9**, 1638 (2025).

[35] J. Schuhmacher, G.-X. Su, J. J. Osborne, A. Gandon, J. C. Halimeh, and I. Tavernelli, Observation of hadron scattering in a lattice gauge theory on a quantum computer, arXiv preprint arXiv:2505.20387 https://doi.org/10.48550/arXiv.2505.20387 (2025).

[36] Y. Chai, J. Gibbs, V. R. Pascuzzi, Z. Holmes, S. Kühn, F. Tacchino, and I. Tavernelli, Resource-efficient simulations of particle scattering on a digital quantum computer (2025).

[37] R. C. Farrell, M. Illa, A. N. Ciavarella, and M. J. Savage, Quantum simulations of hadron dynamics in the schwinger model using 112 qubits, Phys. Rev. D **109**,

114510 (2024).
[38] N. A. Zemlevskiy, Scalable quantum simulations of scattering in scalar field theory on 120 qubits, arXiv:2411.02486 , (2024).
[39] R. C. Farrell, N. A. Zemlevskiy, M. Illa, and J. Preskill, Digital quantum simulations of scattering in quantum field theories using w states, arXiv preprint arXiv:2505.03111 https://doi.org/10.48550/arXiv.2505.03111 (2025).
[40] R. C. Farrell, M. Illa, A. N. Ciavarella, and M. J. Savage, Scalable circuits for preparing ground states on digital quantum computers: The schwinger model vacuum on 100 qubits, PRX Quantum **5**, 020315 (2024).
[41] J. Cobos, J. Fraxanet, C. Benito, F. di Marcantonio, P. Rivero, K. Kapás, M. A. Werner, Ö. Legeza, A. Bermudez, and E. Rico, Real-time dynamics in a (2+1)-d gauge theory: The stringy nature on a superconducting quantum simulator, arXiv preprint arXiv:2507.08088 (2025).
[42] T. A. Cochran, B. Jobst, E. Rosenberg, Y. D. Lensky, G. Gyawali, N. Eassa, M. Will, A. Szasz, D. Abanin, R. Acharya, L. Aghababaie Beni, T. I. Andersen, M. Ansmann, F. Arute, K. Arya, A. Asfaw, J. Atalaya, R. Babbush, B. Ballard, J. C. Bardin, A. Bengtsson, A. Bilmes, A. Bourassa, J. Bovaird, M. Broughton, D. A. Browne, B. Buchea, B. B. Buckley, T. Burger, B. Burkett, N. Bushnell, A. Cabrera, J. Campero, H.-S. Chang, Z. Chen, B. Chiaro, J. Claes, A. Y. Cleland, J. Cogan, R. Collins, P. Conner, W. Courtney, A. L. Crook, B. Curtin, S. Das, S. Demura, L. De Lorenzo, A. Di Paolo, P. Donohoe, I. Drozdov, A. Dunsworth, A. Eickbusch, A. M. Elbag, M. Elzouka, C. Erickson, V. S. Ferreira, L. F. Burgos, E. Forati, A. G. Fowler, B. Foxen, S. Ganjam, R. Gasca, É. Genois, W. Giang, D. Gilboa, R. Gosula, A. Grajales Dau, D. Graumann, A. Greene, J. A. Gross, S. Habegger, M. Hansen, M. P. Harrigan, S. D. Harrington, P. Heu, O. Higgott, J. Hilton, H.-Y. Huang, A. Huff, W. Huggins, E. Jeffrey, Z. Jiang, C. Jones, C. Joshi, P. Juhas, D. Kafri, H. Kang, A. H. Karamlou, K. Kechedzhi, T. Khaire, T. Khattar, M. Khezri, S. Kim, P. Klimov, B. Kobrin, A. Korotkov, F. Kostritsa, J. Kreikebaum, V. Kurilovich, D. Landhuis, T. Lange-Dei, B. Langley, K.-M. Lau, J. Ledford, K. Lee, B. Lester, L. Le Guével, W. Li, A. T. Lill, W. Livingston, A. Locharla, D. Lundahl, A. Lunt, S. Madhuk, A. Maloney, S. Mandrà, L. Martin, O. Martin, C. Maxfield, J. McClean, M. McEwen, S. Meeks, A. Megrant, K. Miao, R. Molavi, S. Molina, S. Montazeri, R. Movassagh, C. Neill, M. Newman, A. Nguyen, M. Nguyen, C.-H. Ni, K. Ottosson, A. Pizzuto, R. Potter, O. Pritchard, C. Quintana, G. Ramachandran, M. Reagor, D. Rhodes, G. Roberts, K. Sankaragomathi, K. Satzinger, H. Schurkus, M. Shearn, A. Shorter, N. Shutty, V. Shvarts, V. Sivak, S. Small, W. C. Smith, S. Springer, G. Sterling, J. Suchard, A. Sztein, D. Thor, M. Torunbalci, A. Vaishnav, J. Vargas, S. Vdovichev, G. Vidal, C. Vollgraff Heidweiller, S. Waltman, S. X. Wang, B. Ware, T. White, K. Wong, B. W. K. Woo, C. Xing, Z. J. Yao, P. Yeh, B. Ying, J. Yoo, N. Yosri, G. Young, A. Zalcman, Y. Zhang, N. Zhu, N. Zobrist, S. Boixo, J. Kelly, E. Lucero, Y. Chen, V. Smelyanskiy, H. Neven, A. Gammon-Smith, F. Pollmann, M. Knap, and P. Roushan, Visualizing dynamics of charges and strings in (2 + 1)d lattice gauge theories, Nature **642**, 315–320 (2025).
[43] A. B. Balantekin and Y. Pehlivan, Neutrino–neutrino interactions and flavour mixing in dense matter, Journal of Physics G: Nuclear and Particle Physics **34**, 47 (2006).
[44] Y. Pehlivan, A. B. Balantekin, T. Kajino, and T. Yoshida, Invariants of collective neutrino oscillations, Phys. Rev. D **84**, 065008 (2011).
[45] D. F. G. Fiorillo and G. G. Raffelt, Slow and fast collective neutrino oscillations: Invariants and reciprocity, Phys. Rev. D **107**, 043024 (2023).
[46] D. Neill, H. Liu, J. Martin, and A. Roggero, Scattering neutrinos, spin models, and permutations, Phys. Rev. Res. **7**, 023157 (2025).
[47] D. Neill and H. Liu, Universality of $su(\infty)$ relaxation dynamics for $su(n_f)$-symmetric spin-models (2025), arXiv:2504.08925 [hep-th].
[48] R. Bhaskar, A. Roggero, and M. J. Savage, Timescales in many-body fast-neutrino-flavor conversion, Phys. Rev. C **110**, 045801 (2024).
[49] L. D'Alessio, Y. Kafri, A. Polkovnikov, and M. Rigol, From quantum chaos and eigenstate thermalization to statistical mechanics and thermodynamics, Advances in Physics **65**, 239 (2016).
[50] J. D. Martin, A. Roggero, H. Duan, and J. Carlson, Many-body neutrino flavor entanglement in a simple dynamic model (2023), arXiv:2301.07049 [hep-ph].
[51] The first local minimum or maximum.
[52] A. Gilyén, Y. Su, G. H. Low, and N. Wiebe, Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics, in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing* (ACM, 2019) p. 193–204.
[53] V. Amitrano, A. Roggero, P. Luchi, F. Turro, L. Vespucci, and F. Pederiva, Trapped-ion quantum simulation of collective neutrino oscillations, Phys. Rev. D **107**, 023007 (2023).
[54] A. Roggero, Entanglement and many-body effects in collective neutrino oscillations, Phys. Rev. D **104**, 103016 (2021).
[55] M. Illa and M. J. Savage, Multi-neutrino entanglement and correlations in dense neutrino systems, Phys. Rev. Lett. **130**, 221003 (2023).
[56] D. Lacroix, A. Balantekin, M. J. Cervia, A. V. Patwardhan, and P. Siwach, Role of non-gaussian quantum fluctuations in neutrino entanglement, Physical Review D **106**, 10.1103/physrevd.106.123006 (2022).
[57] P. Siwach, A. M. Suliga, and A. B. Balantekin, Entanglement in three-flavor collective neutrino oscillations, Phys. Rev. D **107**, 023019 (2023).
[58] I. Chernyshev, C. E. P. Robin, and M. J. Savage, Quantum magic and computational complexity in the neutrino sector, Phys. Rev. Res. **7**, 023228 (2025).
[59] By intermediate time, we mean a timescale sufficiently large that it cannot be accurately captured by a low-order Taylor expansion.
[60] Note that g is invariant for permutation of qubits on the initial state and we can just use the value in $|\phi_0\rangle$.
[61] E. Campbell, Random compiler for fast hamiltonian simulation, Phys. Rev. Lett. **123**, 070503 (2019).
[62] O. Kiss, M. Grossi, and A. Roggero, Importance sampling for stochastic quantum simulations, Quantum **7**, 977 (2023).
[63] C.-F. Chen, H.-Y. Huang, R. Kueng, and J. A. Tropp,

Concentration for random product formulas, PRX Quantum **2**, 040305 (2021).
[64] E. H. Lieb and D. W. Robinson, The finite group velocity of quantum spin systems, Communications in Mathematical Physics **28**, 251 (1972).
[65] D. Perez-Garcia, F. Verstraete, M. M. Wolf, and J. I. Cirac, Matrix product state representations, Quantum Info. Comput. **7**, 401–430 (2007).
[66] M. S. Rudolph, T. Jones, Y. Teng, A. Angrisani, and Z. Holmes, Pauli propagation: A computational framework for simulating quantum systems (2025), arXiv:2505.21606 [quant-ph].
[67] J. Gray, quimb: a python library for quantum information and many-body calculations, Journal of Open Source Software **3**, 819 (2018).
[68] S. Ayik, A stochastic mean-field approach for nuclear dynamics, Phys. Lett. B **658**, 174 (2008).
[69] D. Lacroix, S. Ayik, and B. Yilmaz, Symmetry breaking and fluctuations within stochastic mean-field dynamics: Importance of initial quantum fluctuations, Phys. Rev. C **85**, 041602 (2012).
[70] D. Lacroix and S. Ayik, Stochastic quantum dynamics beyond mean field, Eur. Phys. J. A **50**, 95 (2014).
[71] C. Gardiner and P. Zoller, *Quantum noise: a handbook of Markovian and non-Markovian quantum stochastic methods with applications to quantum optics* (Springer Science & Business Media, 2004).
[72] D. Lacroix, A. Bauge, B. Yilmaz, M. Mangin-Brinet, A. Roggero, and A. B. Balantekin, Phase-space methods for neutrino oscillations: Extension to multibeams, Phys. Rev. D **110**, 103027 (2024).
[73] D. Lacroix, D. Gambacurta, and S. Ayik, Quantal corrections to mean-field dynamics including pairing, Phys. Rev. C **87**, 061302 (2013).
[74] D. Lacroix, S. Hermanns, C. M. Hinz, and M. Bonitz, Ultrafast dynamics of finite hubbard clusters: A stochastic mean-field approach, Phys. Rev. B **90**, 125112 (2014).
[75] D. Lacroix, A. B. Balantekin, M. J. Cervia, A. V. Patwardhan, and P. Siwach, Role of non-gaussian quantum fluctuations in neutrino entanglement, Phys. Rev. D **106**, 123006 (2022).
[76] M. Mangin-Brinet, A. Bauge, and D. Lacroix, Three-flavor neutrino oscillations using the phase space approach (2025), arXiv:2507.18482 [hep-ph].
[77] O. Kiss, D. Teplitskiy, M. Grossi, and A. Mandarino, Statistics of topological defects across a phase transition in a superconducting quantum processor, Quantum Sci. Technol. , 035037 (2025).
[78] M. Urbanek, B. Nachman, V. R. Pascuzzi, A. He, C. W. Bauer, and W. A. de Jong, Mitigating depolarizing noise on quantum computers with noise-estimation circuits, Phys. Rev. Lett. **127**, 270502 (2021).
[79] O. Kiss, M. Grossi, and A. Roggero, Quantum error mitigation for fourier moment computation, Phys. Rev. D **111**, 034504 (2025).
[80] Z. Cai and S. Benjamin, Constructing smaller pauli twirling sets for arbitrary error channels, Sci Rep **9**, 1128 (2019).
[81] P. D. Nation, H. Kang, N. Sundaresan, and J. M. Gambetta, Scalable mitigation of measurement errors on quantum computers, PRX Quantum **2**, 040326 (2021).
[82] E. van den Berg, Z. K. Minev, and K. Temme, Model-free readout-error mitigation for quantum expectation values, Phys. Rev. A **105**, 032620 (2022).
[83] P. Siwach, A. B. Balantekin, A. V. Patwardhan, and A. M. Suliga, Exploring entanglement and spectral split correlations in three-flavor collective neutrino oscillations, Phys. Rev. D **111**, 063038 (2025).
[84] L. Spagnoli, N. Goss, A. Roggero, E. Rrapaj, M. J. Cervia, A. V. Patwardhan, R. K. Naik, A. B. Balantekin, E. Younis, D. I. Santiago, I. Siddiqi, and S. Aldaihan, Collective neutrino oscillations in three flavors on qubit and qutrit processors, Phys. Rev. D **111**, 103054 (2025).
[85] F. Turro, I. A. Chernyshev, R. Bhaskar, and M. Illa, Qutrit and qubit circuits for three-flavor collective neutrino oscillations, Phys. Rev. D **111**, 043038 (2025).
[86] I. A. Chernyshev, Three-flavor collective neutrino oscillation simulations on a qubit quantum annealer, Phys. Rev. D **111**, 043017 (2025).
[87] M. Mangin-Brinet, A. Bauge, and D. Lacroix, Three-flavor neutrino oscillations using the phase space approach (2025), arXiv:2507.18482 [hep-ph].
[88] V. Cirigliano, S. Sen, and Y. Yamauchi, Neutrino many-body flavor evolution: The full hamiltonian, Phys. Rev. D **110**, 123028 (2024).
[89] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao, D. A. Buell, B. Burkett, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, A. Dunsworth, E. Farhi, B. Foxen, A. Fowler, C. Gidney, M. Giustina, R. Graff, K. Guerin, S. Habegger, M. P. Harrigan, M. J. Hartmann, A. Ho, M. Hoffmann, T. Huang, T. S. Humble, S. V. Isakov, E. Jeffrey, Z. Jiang, D. Kafri, K. Kechedzhi, J. Kelly, P. V. Klimov, S. Knysh, A. Korotkov, F. Kostritsa, D. Landhuis, M. Lindmark, E. Lucero, D. Lyakh, S. Mandrà, J. R. McClean, M. McEwen, A. Megrant, X. Mi, K. Michielsen, M. Mohseni, J. Mutus, O. Naaman, M. Neeley, C. Neill, M. Y. Niu, E. Ostby, A. Petukhov, J. C. Platt, C. Quintana, E. G. Rieffel, P. Roushan, N. C. Rubin, D. Sank, K. J. Satzinger, V. Smelyanskiy, K. J. Sung, M. D. Trevithick, A. Vainsencher, B. Villalonga, T. White, Z. J. Yao, P. Yeh, A. Zalcman, H. Neven, and J. M. Martinis, Quantum supremacy using a programmable superconducting processor, Nature **574**, 505–510 (2019).
[90] A. Baroni, J. Carlson, R. Gupta, A. C. Y. Li, G. N. Perdue, and A. Roggero, Nuclear two point correlation functions on a quantum computer, Phys. Rev. D **105**, 074503 (2022).

Appendix A: Computing the rescaling factor

In this section, we provide additional details about the calculation of the short time approximation for the variance.We will start the discussion with the full neutrino Hamiltonian in Eq. (4), following the derivation in Ref. [19], and comment on the extensions needed to deal with the randomized circuit in second moment. We will express the Hamiltonian as

$$H = \sum_{i<j} \mu_{ij} P_{ij} \equiv \sum_{i<j} \nu_{ij} \vec{\sigma}_i \cdot \vec{\sigma}_j, \tag{A1}$$

with $\vec{\sigma}$ the vector of Pauli matrices, $\nu_{ij} = \mu_{ij}/2$ and $\mu_{ij} = \mu_{ji}$. The flavor variance is

$$V(t) = \frac{1}{N}\sum_{j=0}^{N-1}\langle Z_j(t)\rangle^2. \tag{A2}$$

We are interested in performing a Taylor expansion

$$V(t) = 1 + \frac{d}{dt}V(t)\Big|_{t=0} t + \frac{d^2}{dt^2}V(t)\Big|_{t=0}\frac{t^2}{2} + \mathcal{O}(\tau^3)\ . \tag{A3}$$

We can compute the first derivatives as follows

$$\begin{aligned}\frac{d}{dt}V(t)\Big|_{t=0} &= \frac{d}{dt}\frac{1}{N}\sum_{j=0}^{N-1}\langle Z_j(t)\rangle^2\Big|_{t=0} \\ &= -\frac{2i}{N}\sum_{j=0}^{N-1}\langle Z_j\rangle\langle[H, Z_j]\rangle,\end{aligned} \tag{A4}$$

where the commutator can be expanded as

$$\begin{aligned}&[H, Z_k] \\ =&\sum_{i<j}\left(\nu_{ij}[X_iX_j, Z_k] + \nu_{ij}[Y_iY_j, Z_k] + \nu_{ij}[Z_iZ_j, Z_k]\right) \\ =&i\sum_{i<j}\nu_{ij}\left[\delta_{ik}\left(X_iY_j - Y_iX_j\right) + \delta_{jk}\left(Y_iX_j - X_iY_j\right)\right]\ .\end{aligned} \tag{A5}$$

The expectation value of this operator on a product state in the Z basis is then identically zero. For the second derivative we instead obtain [19]

$$\begin{aligned}\langle\phi_0|[H,[H,Z_k]]|\phi_0\rangle &= 4m_k\sum_{j,m_j=-m_k}\mu_{kj}^2 \\ &= 2m_k\sum_{j=0}^{N-1}\mu_{kj}^2(1 - m_jm_k)\end{aligned} \tag{A6}$$

where $|\phi_0\rangle = |0\rangle^{\otimes N/2}\otimes|1\rangle^{\otimes N/2}$ is the state at $t = 0$, the factors m_i are the eigenvalues of Z_i on the initial state

$$Z_i|\phi_0\rangle = m_i|\phi_0\rangle\ , \tag{A7}$$

and the sum is over all spins k with opposite polarization as the k-th. Using these result we finally get

$$\begin{aligned}\frac{d^2}{dt^2}V(t)\Big|_{t=0} &= -\frac{2}{N}\sum_{j=0}^{N-1}\langle Z_j\rangle\langle[H,[H,Z_j]]\rangle \\ &= -\frac{4}{N}\sum_{j=0}^{N-1}\sum_{k=0}^{N-1}\mu_{jk}^2(1 - m_jm_k)\ ,\end{aligned} \tag{A8}$$

where we used the symmetry $\mu_{ij} = \mu_{ji}$ of the coupling matrix. For the specific initial sate $|\phi_0\rangle$ we use here this simplifies to

$$\begin{aligned}\frac{d^2}{dt^2}V(t)\Big|_{t=0} &= -\frac{8}{N}\sum_{j=0}^{N-1}m_j^2\sum_{k,m_k=-m_j}\mu_{jk}^2 \\ &= -\frac{8}{N}\left(\sum_{j=0}^{N/2-1}\sum_{k=N/2}^{N-1}\left(\mu_{jk}^2 + \mu_{kj}^2\right)\right) \\ &= -\frac{16}{N}\sum_{j=0}^{N/2-1}\sum_{k=N/2}^{N-1}\mu_{jk}^2\ ,\end{aligned} \tag{A9}$$

where we used the symmetry $\mu_{ij} = \mu_{ji}$ of the coupling matrix. For the full neutrino Hamiltonian we then expect on average that

$$\mathbb{E}\left[\frac{d^2}{dt^2}V(t)\Big|_{t=0}\right] = -\frac{4\mu^2}{N}\mathbb{E}\left[1 - \mathbf{v}_1\cdot\mathbf{v}_2\right]\ , \tag{A10}$$

where the vectors $\mathbf{v}_1$ and $\mathbf{v}_2$ are sampled from the distribution in Eq. (3) of the main text.

When instead we apply a random channel $\mathcal{U}_R^{(k)}(\tau, L)$, we will write for small times τ that

$$\mathcal{U}_R^{(k)}(\tau, L) \approx \exp\left(-i\tau H_{\text{eff}}(k, L)\right)\ , \tag{A11}$$

where the effective Hamiltonian is given by

$$H_{\text{eff}}(k, L) = \sum_{l=1}^{L}\sum_{n=0}^{N-2}\widetilde{h}_n^{(k,l)}P_{i(n,l),j(n+1,l)}\ , \tag{A12}$$

where the indices $i(n, l)$ and $j(n+1, l)$ correspond to the neutrino labels occurring at layer l for the n-the term in the even/odd evolutions implement in W from Eq. (10) in the main text. These labels can be reconstructed efficiently once a given circuit layout has been generated by following the SWAP network and keeping track of the indices. The effective Hamiltonian $H_{\text{eff}}(k, L)$ can then be represented exactly as the generic H in Eq. (A1) with

$$\mu_{ij} = \sum_{l=1}^{L}\sum_{n=0}^{N-2}\widetilde{h}_n^{(k,l)}\delta_{i,i(n,l)}\delta_{j,j(n,l)}\ . \tag{A13}$$

We will then expect that for the randomized unitary dynamics, the second derivative of the flavor variance should scale as

$$\frac{d^2}{dt^2}V_R(t)\Big|_{t=0} = \mathcal{O}\left(\frac{\mu^2 L}{N^2}\right)\ , \tag{A14}$$

which will then lead to a rescaling factor

$$\alpha = \sqrt{\frac{\frac{d^2}{dt^2}V(t)\big|_{t=0}}{\frac{d^2}{dt^2}V_R(t)\big|_{t=0}}} = \mathcal{O}\left(\sqrt{\frac{N}{L}}\right)\ , \tag{A15}$$

which is consistent with numerical results shown in Fig. 2.