

Simulation of collective neutrino oscillations on a quantum computer

Benjamin Hall,¹ Alessandro Roggero,^{2,3} Alessandro Baroni,⁴ and Joseph Carlson⁴

¹Facility for Rare Isotope Beams (FRIB), Michigan State University, East Lansing, Michigan 48824, USA

²Institute for Nuclear Theory, University of Washington, Seattle, Washington 98195, USA

³InQuBator for Quantum Simulation (IQuS), Department of Physics, University of Washington, Seattle, Washington 98195, USA

⁴Theoretical Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA



(Received 24 March 2021; accepted 21 July 2021; published 3 September 2021)

In astrophysical scenarios with large neutrino density, like supernovae and the early universe, the presence of neutrino-neutrino interactions can give rise to collective flavor oscillations in the out-of-equilibrium collective dynamics of a neutrino cloud. The role of quantum correlations in these phenomena is not yet well understood, in large part due to complications in solving for the real-time evolution of the strongly coupled many-body system. Future fault-tolerant quantum computers hold the promise to overcome much of these limitations and provide direct access to the correlated neutrino dynamic. In this work, we present the first simulation of a small system of interacting neutrinos using current generation quantum devices. We introduce a strategy to overcome limitations in the natural connectivity of the qubits and use it to track the evolution of entanglement in real-time. The results show the critical importance of error-mitigation techniques to extract meaningful results for entanglement measures using noisy, near term, quantum devices.

DOI: 10.1103/PhysRevD.104.063009

I. INTRODUCTION

The flavor evolution of neutrinos in dense astrophysical environments have, by now, a long history. It has been pointed out by Pantelone, Raffelt, and Sigl [1,2] and others that, through forward scattering, neutrinos can exchange their flavors. Given an anisotropic initial distribution in energy and/or angle as found in supernovae, neutron star mergers, or the early universe, the neutrino energy flux versus energy and flavor may be impacted by this nontrivial quantum many-body evolution. This can in turn affect the dynamics of these environments and other observable signatures, including nucleosynthesis in the ejected material (see [3,4] for recent reviews).

The Hamiltonian for neutrino flavor evolution in a dense neutrino environment includes three terms: the vacuum mixing that has been determined from solar and accelerator neutrino experiments [5], the forward scattering in matter leading to the well-known Mikheyev–Smirnov–Wolfenstein (MSW) effect [6,7], and neutrino-neutrino forward scattering.

In the neutrino flavor basis, the vacuum term includes diagonal contributions describing the mass differences between different neutrino flavors and an off-diagonal term characterized by a mixing angle θ_v . The interaction describing forward scattering in matter is diagonal in the flavor basis. The neutrino-neutrino interaction can exchange flavors of two neutrinos and has a forward scattering amplitude that depends on the angle between

their momenta. For the two-flavor case considered here, this interaction is proportional to the dot product $\vec{\sigma}_i \cdot \vec{\sigma}_j$ of the SU(2) matrices describing the different flavor amplitudes of the two neutrinos

$$V_{ij} \propto \left(1 - \frac{\vec{q}_i \cdot \vec{q}_j}{\|\vec{q}_i\| \|\vec{q}_j\|} \right) \vec{\sigma}_i \cdot \vec{\sigma}_j. \quad (1)$$

Here we denoted by \vec{q}_k the momentum of the k th neutrino and with $\vec{\sigma}_k = (\sigma_k^x, \sigma_k^y, \sigma_k^z)$ the vector of Pauli operators acting on its amplitude. Generalization to the three-flavor case is straightforward in principle; here we assume the μ and τ flavors evolve similarly.

For this simplified two-flavor case, we seek to understand the time and space evolution of the set of amplitudes from a Schrödinger equation:

$$|\Phi(t)\rangle = \exp[-iHt]|\Phi_0\rangle, \quad (2)$$

with H the total Hamiltonian including both the vacuum and forward-scattering interaction contributions. For simplicity here we consider $|\Phi_0\rangle$ to be a product state, but generalizations to arbitrary states are straightforward.

Most often these quantum equations have been treated on the mean-field level by replacing one of the spin operators in Eq. (1) by its expectation value, yielding a set of nonlinear coupled differential equations. This makes the calculations tractable for several hundred energies and

angles on modern computers (see e.g., [8]). More recently, studies of neutrino propagation as a quantum many-body problem have appeared, including for example [9–16]. These works highlight the importance of understanding the role of quantum correlations, such as entanglement, in order to quantify beyond mean-field effects in out-of-equilibrium neutrino simulations. A direct solution of the Schrödinger equation in Eq. (2), for a system of N configurations in energy and angle, incurs a computational cost that is exponential in N . This has limited early explorations of the problem to systems with $N = \mathcal{O}(10)$ neutrinos. An alternative to reach larger system sizes, explored recently by one of us in Refs. [15,16], employs a matrix product state representation for $|\Phi(t)\rangle$ which allows one to track the exact time evolution in situations where entanglement never grows too much. For conditions leading to strong entanglement instead, simulations on digital/analog quantum computers have the potential to tackle the full neutrino dynamics while still enjoying a polynomial computational cost in system size N [17].

In this work we explore the time-dependent many-body evolution of the neutrinos on a current-generation digital quantum computer. In Sec. II we introduce in more detail the SU(2) spin model used to describe collective neutrino oscillations and describe an implementation of the time evolution operator appearing in Eq. (2) suitable for an array of qubits with linear connectivity. We present the results obtained for a small system with $N = 4$ neutrino amplitudes in Sec. III and provide a summary and conclusions in Sec. IV.

II. SPIN MODEL FOR NEUTRINO OSCILLATIONS

For the simplified two-flavor case studied here, the state of the system can be described as an amplitude for a neutrino of each energy E_i (equal to the magnitude of momentum $\|\vec{q}_i\|$) and direction of momentum (denoted by \hat{q}_i), with α_{\uparrow} and α_{\downarrow} describing the amplitude of being in the electron flavor or in a heavy (μ or τ) flavor respectively. These two amplitudes can be encoded in an SU(2) spinor basis. In this basis, the Hamiltonian can be written in terms of Pauli operators as the sum of a one-body term, describing both vacuum oscillations and forward scattering in matter,

$$H_1 = \frac{1}{2} \sum_i [(-\Delta_i \cos 2\theta_v + A) \sigma_i^z + \Delta_i \sin 2\theta_v \sigma_i^x], \quad (3)$$

and a two-body term, coming from the neutrino-neutrino forward-scattering potential V_{ij} from Eq. (1), which takes the following form [12]

$$H_2 = \sum_{i < j} \eta [1 - \hat{q}_i \cdot \hat{q}_j] \vec{\sigma}_i \cdot \vec{\sigma}_j. \quad (4)$$

In the one-body term, θ_v represents the vacuum mixing angle, while the strength is given by $\Delta_i = \delta m^2 / (2E_i)$

with δm^2 the mass squared difference for neutrinos of different flavor. The matter potential enters as the diagonal contribution in the one-body term through the constant $A = \sqrt{2}G_F n_e$, with G_F the Fermi coupling constant and n_e the electron density.

As described in the introduction, the two-body term is a sum over spin-spin interactions with a coupling depending upon the relative angle between them. The overall strength depends on the neutrino density as

$$\eta = \frac{G_F}{\sqrt{2}V} = \frac{G_F n_\nu}{\sqrt{2}N}, \quad (5)$$

with N the number of neutrino momenta considered, given by the neutrino density n_ν times the quantization volume V . The Hamiltonian is similar to a Heisenberg model, but the two-body term is all-to-all rather than nearest neighbor. Its coupling strength $\eta \propto 1/N$ assures that the energy of the system is extensive. This allows us to obtain a well-defined many-body solution, in the limit of large numbers of neutrino momenta by extrapolating in system size N .

Currently available quantum devices are able to perform only a relatively limited number of operations while maintaining a high fidelity [18], this in turn poses limits on the maximum time that could be reached in the simulation of neutrino dynamics. Given this practical constraint, it is then useful to consider a test case where the one- and two-body interaction terms are similar in magnitude and the evolution can occur rapidly. An example is the environment of order ≈ 100 km from the surface of a proto-neutron star in a core collapse supernovae. Here the background matter density has decreased to a point where its contribution to the Hamiltonian is similar in magnitude to the neutrino-neutrino forward scattering. The relative angles of neutrino propagation are fairly small as neutrinos are emitted from a typical proto-neutron star radius of order 10 km. In the neutrino bulb model [8] one further assumes the evolution in a supernovae depends only on the energy and the angle from the normal. Averaging over the azimuthal angles results in an average coupling $\langle 1 - \hat{q}_i \cdot \hat{q}_j \rangle = 1 - \cos(\theta_i) \cos(\theta_j)$. Often a further simplification, usually called single-angle approximation, is made where an average coupling is taken between all pairs of neutrinos, resulting in a two-body term simply related to the square of the total spin S of the many-body state.

For our test case, we take a monochromatic neutrino beam with energy $E_\nu = \delta m^2 / (4\eta)$ and measure energies in units of the two-body coupling η . In order to avoid the symmetries introduced by the single angle approximation, we employ an anisotropic distribution of momentum directions using a simple grid of angles with

$$\theta_{pq} = \arccos(0.9) \frac{|p - q|}{N - 1}. \quad (6)$$

This is similar to the standard bulb model as the relative couplings $1 - \cos(\theta_{pq})$ are small.

The final Hamiltonian for the simple model we implement here can be written compactly, in units of η , as

$$H = \sum_{k=1}^N \vec{b} \cdot \vec{\sigma}_k + \sum_{p < q} J_{pq} \vec{\sigma}_p \cdot \vec{\sigma}_q, \quad (7)$$

with the external field $\vec{b} = (\sqrt{1 - 0.925^2}, 0, -0.925)$ obtained by choosing the mixing angle $\theta_v = 0.195$ and pair coupling matrix $J_{pq} = (1 - \cos(\theta_{pq}))$. Note that in this model we set the matter potential A in the one-body contribution to the Hamiltonian Eq. (3) to zero.

A. Real time evolution

The major challenge in implementing the time evolution in Eq. (2) in a quantum simulation is to find an accurate approximation to the evolution operator $U(t) = \exp[-iHt]$ that can also be decomposed efficiently into local unitary operations [17]. A simple and popular approach is to use a first-order Trotter-Suzuki decomposition [19] of the propagator leading to the approximation

$$U_1(t) = \prod_{j=1}^N e^{-it\vec{b} \cdot \vec{\sigma}_j} \prod_{p < q} e^{-itJ_{pq} \vec{\sigma}_p \cdot \vec{\sigma}_q}, \quad (8)$$

which is correct up to an additive error $\epsilon = \mathcal{O}(t^2)$. Past experience with the Euclidean version of this evolution operator in quantum Monte Carlo suggests that a better approximation to the full propagator $U(t)$ can be obtained by using the exact propagators for pairs (see e.g., [20,21]). In order to construct this alternative approximation, we first rewrite the Hamiltonian in Eq. (7) manifestly as a sum of $\binom{N}{2}$ two-body Hamiltonians acting on each pair of qubits

$$H = \sum_{p < q}^N \left(\frac{\vec{b} \cdot (\vec{\sigma}_p + \vec{\sigma}_q)}{N-1} + J_{pq} \vec{\sigma}_p \cdot \vec{\sigma}_q \right) := \sum_{p < q}^N h_{pq}. \quad (9)$$

We can then define an approximate propagator U_2 using the exact pair propagator as follows

$$U_2(t) = \prod_{p < q}^N e^{-ith_{pq}} := \prod_{p < q}^N u_{pq}. \quad (10)$$

Note that the implementation of this operator is efficient since each pair Hamiltonian acts nontrivially only on a 4×4 subset of the total Hilbert space and therefore, as shown for instance in Refs. [22,23], can be implemented exactly using at most 3 entangling operations. Note that the error in this approximation still scales as $\mathcal{O}(t^2)$ but now with a possibly reduced prefactor. In Appendix A we present a direct comparison between the two approximations.

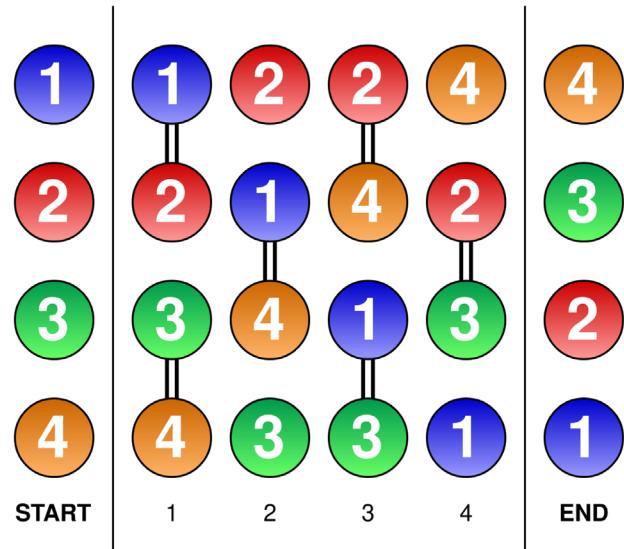


FIG. 1. Pictorial representation of the swap network used in our simulation in the case of $N = 4$ neutrinos.

Finally, the approximation order could also be improved by symmetrizing over the ordering of operators or by applying symmetry transformations (see e.g., [24]).

Owing to the long range of the interactions, a naive implementation of this scheme will require either a device with all-to-all connectivity (like trapped ion systems [25]) or an extensive use of the SWAP operation, represented in matrix form as

$$\text{SWAP} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \quad (11)$$

The effect of this operation is to exchange the state of two qubits. One can then use this operation to bring a pair of qubits that we want to interact close to each other by applying a sequence of SWAP gates of order N . Since we need to apply all possible pair interactions, we will show that it is actually possible to carry out a complete step, under the unitary in Eq. (10), without incurring any overhead due to the application of the SWAP operations. The scheme is inspired by the more general fermionic swap network construction presented in Ref. [26].

We illustrate this idea using the diagram shown in Fig. 1 for a simple case with $N = 4$ neutrinos. Starting from the initial state on the left, we first apply the unitaries u_{pq} from Eq. (10) to the odd bonds: for the $N = 4$ case, these are the bonds between the (1,2) and (3,4) pairs of qubits. Before moving to the next pairs, we also apply a SWAP operation to the same pairs we just acted upon. The resulting unitary operation is denoted as a double line joining qubits in Fig. 1 and the net effect is that at the next step the qubits that have

interacted get interchanged. Given the discussion following Eq. (10) above, this modified two-qubit unitary still requires at most 3 entangling operations. At the end of a sequence of N such combined operations we will have implemented the full unitary in Eq. (10) while, at the same, we inverted the ordering of qubits, as shown in Fig. 1. This approach requires exactly the minimum number $\binom{N}{2}$ of nearest-neighbor pair operations, while the shifted ordering can be controlled completely, and efficiently, by classical means. Note that if we were to repeat at this point the same swap network in reverse order, the full unitary will correspond to a second order step for time $2t$ and the final ordering of qubits will be restored to its original one. This is the strategy used in Refs. [15,16] to study the neutrino Hamiltonian with matrix product states. In this first implementation on quantum hardware, we focus instead on a single, linear-order, time step.

Note that since we are only using nearest neighbor two-qubit gates, the total number of entangling gates required for a full time evolution step is bounded from above by $3\binom{N}{2}$ while the maximum number of single qubit operations is bounded by $15\binom{N}{2}$. As we will see in the results presented below, the presence of a large number of arbitrary single qubit rotations seems to be the limiting factor in implementing this scheme on the quantum device we used in this first exploration.

III. RESULTS WITH FOUR NEUTRINOS

In order to study the build up of correlations and entanglement generated by the time-evolution under the Hamiltonian in Eq. (7), we first initialize a system of $N = 4$ qubits in the following product state

$$|\Phi_0\rangle = |e\rangle \otimes |e\rangle \otimes |x\rangle \otimes |x\rangle = |\uparrow\uparrow\downarrow\downarrow\rangle. \quad (12)$$

We then perform one step of time evolution for time t by applying the N layers of nearest-neighbor gates as described in the previous section. This corresponds to a single Trotter-Suzuki step for different values of the time step. The four SU(2) spins representing the neutrinos are mapped to qubits (2,1,3,4) on the IBMQ Vigo quantum processor [27], whose connectivity is schematically depicted in Fig. 2. The resulting qubits are linearly connected, allowing us to carry out natively the complete simulation scheme depicted in Fig. 1 above.

The first observable we compute is the flavor polarization of individual neutrinos as a function of time. Since the spin Hamiltonian in Eq. (7) is invariant under the simultaneous exchanges $1 \leftrightarrow 4$ and $2 \leftrightarrow 3$, while the flavor content of the initial state $|\Phi_0\rangle$ gets reversed by it, we show directly the probability $P_{\text{inv}}(t)$ to find a neutrino in the opposite flavor to the starting one it had at $t = 0$. In the limit of no error, $P_{\text{inv}}(t)$ should then be the same for the pair of neutrinos (1,4) and (2,3). The errors in the approximation of the propagator in Eq. (10) do not exactly

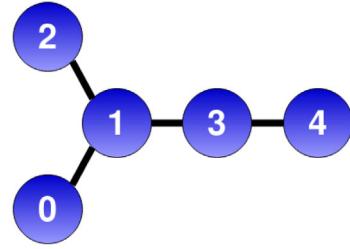


FIG. 2. Layout of the IBM Quantum Canary Processor Vigo [27]. Shown are the five qubits, labeled from 0 to 4, and their connectivity denoted as solid black lines.

follow this symmetry, with deviations in the range 3–7%. We show the results for P_{inv} obtained with the approximate evolution operator $U_2(t)$ as solid black lines in Fig. 3, for the pair (1,4), and in Fig. 4 for the pair (2,3). The ideal, and symmetric, result is shown instead as a purple dashed line. We see that the approximation error is very small up to relatively large time $\eta t \approx 6$. As we discuss more in detail in Appendix A, this is in large part an effect of using the pair propagator $U_2(t)$ instead of the naive first order formula in Eq. (8).

The results shown in Fig. 3 and Fig. 4 were obtained using either the real quantum device (right panels denoted QPU) or a local virtual machine simulation employing the noise model implemented in Qiskit [28] (left panels denoted VM) initialized with calibration data from the device. In both plots we report the results (denoted by [bare]) obtained directly from the simulation and including only statistical errors coming from a finite sample size (here and in the rest of the paper we use 8192 repetition, or “shots,” for every data point), as well as results obtained

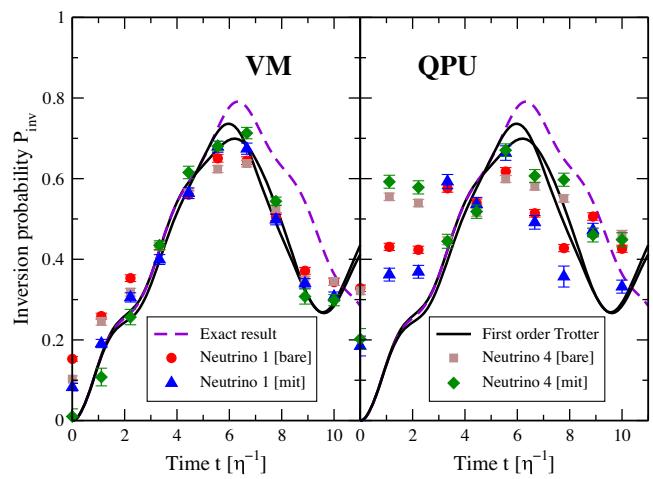


FIG. 3. Inversion probability $P_{\text{inv}}(t)$ for neutrinos 1 and 4: the red circle and brown square correspond to the bare results, the blue triangle and the green diamond are obtained after error mitigation (see text). The left panel (VM) are virtual machine results while the right panel (QPU) are results obtained on the Vigo [27] quantum device.

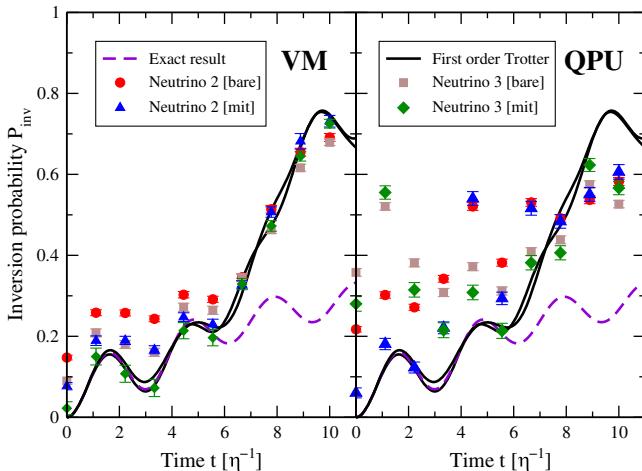


FIG. 4. Inversion probability $P_{\text{inv}}(t)$ for neutrinos 2 and 3. The notation is the same as for Fig. 3.

after performing error mitigation (denoted by [mit]). This corresponds to a final post-processing step that attempts to reduce the influence of the two main sources of errors: the read-out errors associated with the imperfect measurement apparatus and the gate error associated with the application of entangling gates. The latter error is dealt with using a zero noise extrapolation strategy (see [29,30] and Appendix B for additional details).

As seen also in previous similar calculations (see for instance [31,32]), the VM results obtained using the simulated noise are much closer to the ideal result than those obtained with the real device. This is also reflected in the fact that the error mitigation protocol is not as successful with the real QPU data as it is with the simulated VM data. This behavior is possibly linked to the substantial noise caused by the presence of a large number of single qubit operations (up to 90 rotations for time evolution +2 for state preparation) together with the relatively large CNOT count of 18. In fact, the performance of error mitigation for the results with the largest state preparation circuits presented in [32] is superior to the one obtained here, despite the use of the same device, the same error mitigation strategy and a comparable number of entangling gates (15 CNOT in that case) while the number of rotations was only 14. This suggests coherent errors constitute a considerable fraction of the overall error seen in the results above.

In order to highlight the difficulties encountered when performing noise extrapolation for this data, we plot in Fig. 5 the results obtained from both the QPU (black circles) and the VM (red squares) for the inversion probability of the first neutrino at the initial time $t = 0$ together with a linear extrapolation using the first two points for the QPU (green line) and the first three points for the VM (blue line). The exact result is of course $P_{\text{inv}}(0) = 0$ and we see that neither strategy is able to predict the correct value. The horizontal dashed line is the value

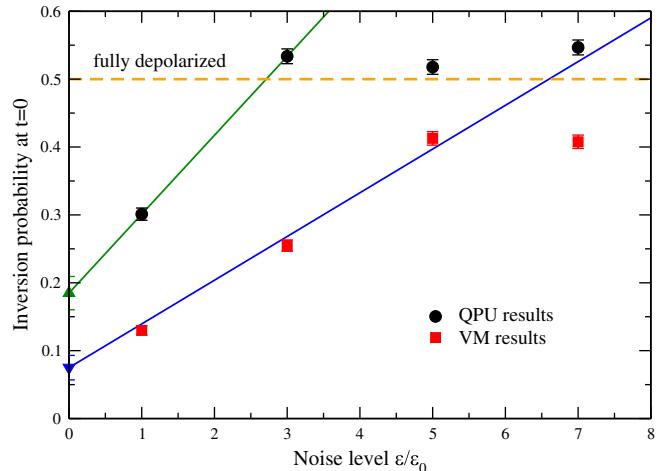


FIG. 5. Inversion probability P_{inv} at the initial time $t = 0$ for the first neutrino. Black solid circles are results from the Vigo QPU [27] while the red squares correspond to results obtained using the VM with simulated noise. Also shown are extrapolations to the zero noise limit, for both the QPU (green line) and the VM (blue line), together with the extrapolated value (green triangle up and blue triangle down respectively). The dashed orange line denotes the result for a maximally mixed state.

expected when the system is in the maximally mixed state, corresponding to full depolarization. As shown in the data, for the real QPU results, only the first level of noise extrapolation contains useful information and a more gentle noise amplification strategy, like the one proposed in Ref. [33], could provide a substantial advantage over the strategy adopted here.

A. Dynamics of entanglement

In order to track the evolution of entanglement in the system we perform complete state tomography for each of the 6 possible qubit pairs in our system by estimating, for each pair (k, q) , the 16 expectation values

$$M_{\alpha, \beta}^{k, q}(t) = \langle \Phi(t) | P_k^\alpha \otimes P_q^\beta | \Phi(t) \rangle, \quad (13)$$

with $P_k = \{1, X, Y, Z\}$ the basis for $U(2)$ and $|\Phi(t)\rangle$ the state obtained from $|\Phi_0\rangle$ by applying the time-evolution operator as in Eq. (2). In principle, we might reconstruct the density matrix for the pair of qubits (k, q) directly from these expectation values as

$$\rho_{kq}^D(t) = \sum_{\alpha=1}^4 \sum_{\beta=1}^4 M_{\alpha, \beta}^{k, q}(t) P_k^\alpha \otimes P_q^\beta. \quad (14)$$

In practice however, we can only estimate the matrix elements $M_{\alpha, \beta}^{k, q}(t)$ to some finite additive precision, and the approximation in Eq. (14) is not guaranteed to be a physical density matrix (positive definite and with trace equal to 1). In this work we use the common approach (see

eg. [34]) of performing a maximum-likelihood (ML) optimization, while enforcing the reconstructed density matrix $\rho_{kq}^{ML}(t)$ to be physical. We note in passing that it is possible to devise operator basis that are more robust than the choice used in Eq. (13) (see e.g., [35]) but we did not explore this further in our work.

In order to propagate the effect of statistical errors into the final estimator for $\rho_{kq}^{ML}(t)$, we use a resampling strategy similar to what was introduced in [32] but using a Bayesian approach to determine the empirical posterior distribution. We provide a detailed description of the adopted protocol in Appendix B 1.

1. Entanglement entropies

As we mentioned in the introduction, one of the main differences between a mean field description and the full many-body description of the dynamics of the neutrino cloud is the absence of quantum correlations, or entanglement, in the former. Past work on the subject [13,14] looked at the single spin entanglement entropy defined as

$$S_k(t) = -\text{Tr}[\rho_k(t) \log_2 (\rho_k(t))], \quad (15)$$

with $\rho_k(t)$ the reduced density matrix of the k th spin. A value of the entropy $S_k(t)$ different from zero indicates the presence of entanglement between the k th neutrino and the rest of the system.

In our setup, we compute the one-body reduced density matrix from the maximum-likelihood estimator of the pair density matrix defined above, explicitly

$$S_{k;q}^{ML}(t) = -\text{Tr}_q[\rho_{k;q}^{ML}(t) \log_2 (\rho_{k;q}^{ML}(t))], \quad (16)$$

where the reduced density matrices are computed from

$$\rho_{k;q}^{ML}(t) = \text{Tr}_q[\rho_{kq}^{ML}(t)], \quad (17)$$

and Tr_q denotes the trace over the states of the q th qubit. We combine the 3 values obtained in this way for each neutrinos as follows: the estimator for the single-spin entanglement entropy is obtained from the average

$$S_k^{\text{avg}}(t) = \frac{1}{3} \sum_q S_{k;q}^{ML}(t), \quad (18)$$

summing over pairs containing the k th spin, while as an error estimate we use the average of the 3 errors.

As for the case of the inversion probability $P_{\text{inv}}(t)$ studied in the previous section, the substantial noise present in the QPU data prevents us from using the full set of results at the 4 effective noise levels. In order to overcome this difficulty, we have performed zero noise extrapolations using only results for effective noise levels $r = \epsilon/\epsilon_0 = (1, 3)$ and performed a Richardson extrapolation (in this

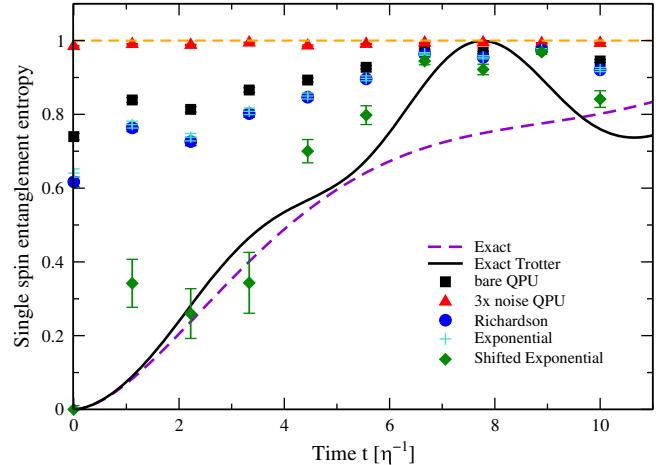


FIG. 6. Single spin entanglement entropy for neutrino 2. Black square are bare results obtained from the QPU, red triangles are results obtained by amplifying the noise to $\epsilon/\epsilon_0 = 3$, the blue circles are obtained using Richardson extrapolation, the turquoise plus symbols indicate results obtained using the standard exponential extrapolation and the green diamonds correspond to the results obtained from a shifted exponential extrapolation using the maximum value of the entropy (indicated as a dashed orange line).

case equivalent to a simple linear fit as done in Ref. [30]), a two point exponential extrapolation [29], and an exponential extrapolation with shifted data. The latter technique consists in shifting the data for the entropy by -1 (its maximum value) so that the result, in the limit of large noise, tends to 0 instead of $\log_2(2) = 1$. We then shift back the result obtained after extrapolation. The exponential extrapolation method is well suited for situations where expectation values decay to zero as a function of the noise strength ϵ , while maintaining a consistent sign, and this shift allows us to make the data conform to this ideal situation (see Appendix B for more details on the method). The impact on the efficacy of the error mitigation is dramatic as can be seen in the results presented in Fig. 6 for the entropy of the second neutrino (the entropies for the other neutrinos follow a similar pattern; see Appendix C for all four results). The results with the standard exponential extrapolation are presented as the turquoise plus symbols, they are almost the same as those obtained using Richardson extrapolation (blue circles) and show a significant systematic error. On the contrary, the results obtained with the shifted exponential extrapolation (green diamonds) are much more close to the expected results with our pair propagator (solid black curve). We expect more general multiexponential extrapolation schemes, like those proposed in Refs. [36,37], to enjoy a similar efficiency boost in the large noise limit achieved with deep circuits.

Using the reconstructed pair density matrix $\rho_{kq}^{ML}(t)$, we can clearly also evaluate directly the entanglement entropy of the pair

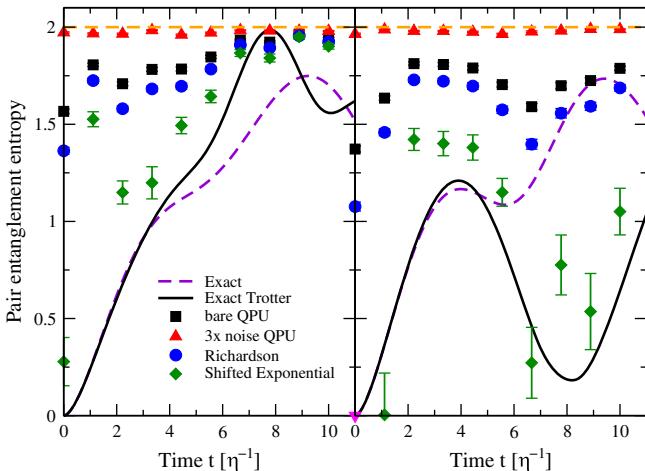


FIG. 7. Pair entanglement entropy for the neutrino pair (1,2) starting as $|e\rangle \otimes |e\rangle$ (left panel) and pair (2,4) which starts as the flavor state $|e\rangle \otimes |x\rangle$ (right panel). Results obtained directly from the QPU are shown as black squares ($r = 1$) and red triangles ($r = 3$) while blue circles and green diamonds indicate mitigated results using Richardson and the shifted exponential extrapolations respectively. For the shifted exponential ansatz we use the maximum value of the entropy (indicated as a dashed orange line). The magenta triangle indicates a mitigated result with shifted exponential extrapolation below zero within error bars.

$$S_{kq}^{ML}(t) = -\text{Tr}[\rho_{kq}^{ML}(t) \log_2 (\rho_{kq}^{ML}(t))]. \quad (19)$$

In Fig. 7 we show the result of this calculation for the pair (1,2), which started as electron flavor at $t = 0$, and the pair (2,4) which started instead as heavy flavor states.

2. Concurrence

In order to better understand these quantum correlations, we also compute the concurrence [38] for all the pair states. This measure of entanglement is defined for a 2 qubit density matrix as

$$C(\rho) = \max \{0, \lambda_0 - \lambda_1 - \lambda_2 - \lambda_3\}, \quad (20)$$

where λ_i are the square roots of the eigenvalues, in decreasing order, of the non-Hermitian matrix

$$M = \rho(Y \otimes Y)\rho^*(Y \otimes Y), \quad (21)$$

with the star symbol indicating complex conjugation. The usefulness of this measure is its relation with the entanglement of formation [38,39], which is the minimum number of maximally-entangled pairs needed to represent ρ with an ensemble of pure states [39].

The definition of concurrence in Eq. (20) does not lend itself as easily to be adapted in an error extrapolation procedure as the one we used to obtain the mitigated results in the previous sections. This is due to the presence of the *max* function in the definition of the concurrence: when the

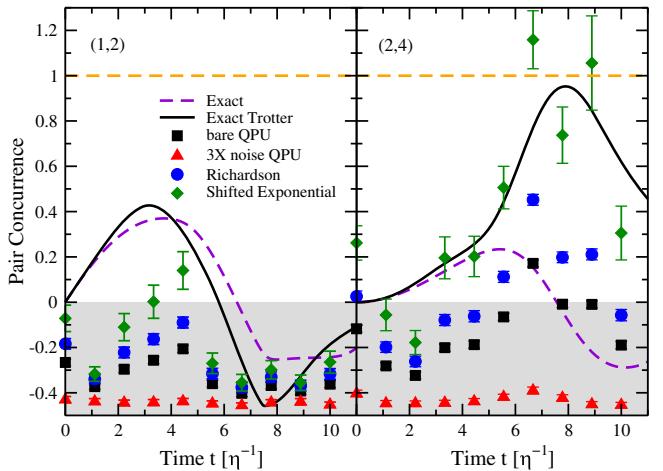


FIG. 8. Extended concurrence \tilde{C} for two pairs of neutrinos, (1,2) in the left and (2,4) in the right panel. The convention for the curves and date point used here is the same as in Fig. 7. The gray area indicates the region where the concurrence $C(\rho)$ is zero. The maximum value for the concurrence is shown as a dashed orange line.

error is sufficiently strong to make the difference in eigenvalues

$$\tilde{C}(\rho) = \lambda_0 - \lambda_1 - \lambda_2 - \lambda_3 \quad (22)$$

negative, the concurrence in Eq. (20) ceases to carry information about the error free result. For this reason, we will regard \tilde{C} as an “extended concurrence” which varies smoothly for large error levels and perform the truncation to positive values only after the zero noise extrapolation. The results obtained from the simulation on the Vigo QPU are shown in Fig. 8 for two pairs of neutrinos: pair (1,2) starting as like spin at $t = 0$ and pair (2,4) which started as opposite flavors. The complete set of results for all pairs can be found in Figs. 10–12 in Appendix C.

The bare results are shown as black squares and we can immediately notice why the definition of \tilde{C} is so important in our case: the only bare data point with a measurable concurrence $C(\rho)$ is at $t \approx 6.7\eta^{-1}$ for pair (2,4) (the right panel in Fig. 8) while all the other results, including those obtained with a larger noise level (red triangles), are compatible with zero. In this situation, no mitigation of $C(\rho)$ would be possible.

By keeping the negative contributions, we see that the bare results often contain a substantial signal, while those at a higher error rate are already almost at the asymptotic value $\tilde{C} = -0.5$ expected for a completely depolarized system [40]. This allowed us to perform error extrapolation using both the Richardson and shifted exponential ansatz. Similarly to what we observed for the entanglement entropies in the previous section, the shifted exponential ansatz (with shift -0.5) produces consistently better results

than Richardson extrapolation. This indicates that we are more close to the asymptotic large error regime than the small error limit used to motivate a polynomial expansion. The resilience of the exponential extrapolations to large errors, especially augmented by an appropriate shift, is seen here to be critical in extracting physical information from quantum simulations carried out near the coherence limit of the device used for the implementation.

IV. CONCLUSIONS

In this work, we presented the first digital quantum simulation of the flavor dynamics in collective neutrino oscillations using current quantum technology. The results reported here for the evolution of flavor and entanglement properties of a system with $N = 4$ neutrino amplitudes are not directly indicative to the behavior at large $N \gg 1$ but nevertheless show that current quantum devices based on superconducting qubits are starting to become a viable option for studying out-of-equilibrium dynamics of interacting many-body systems. The reduced fidelity in the results obtained here, compared to the simulations reported previously in Ref. [32] employing the same quantum processor and a comparable number of entangling gates, points to the importance of controlling unitary errors associated with the imperfect implementation of arbitrary single-qubit rotations (on average $< 1\%$ for the device used in both works). In future work we plan to explore the use of more advanced error mitigation strategies, such as Pauli twirling [41] or symmetry protection [24], to achieve a better overall fidelity. A complementary strategy to the use of digital quantum devices as discussed in the present work would be to adopt quantum platforms, like trapped ion quantum devices [25], that are able to physically implement all-to-all interactions similar to those needed for the neutrino Hamiltonian in Eq. (7). This analog simulation strategy can possibly be implemented in the near term on systems with 50–100 neutrino amplitudes (see e.g., Ref. [42]).

We showed the zero-noise error extrapolation using a shifted Gaussian ansatz to be remarkably efficient in predicting the expected error-free estimator of observables. Given the large circuits employed in this work, past experience with zero-noise extrapolations (see e.g. [31,32]) suggest the exponential ansatz to be appropriate due to the large noise rates, and we find it to indeed outperforms Richardson extrapolation in this regime. The current results highlight the importance of using alternative measures of entanglement to the entropy in order to extract reliable information about quantum correlations in the states generated on the quantum device. Using the pair concurrence together with the entropy provides a robust way to detect entanglement even in the presence of substantial noise, like in the results shown here. We expect these insights, and the mapping of the neutrino evolution problem into a swap network, to prove very valuable in

future explorations of out-of-equilibrium neutrino dynamics with near-term, noisy, quantum devices.

After the completion of the present work we have become aware of Ref [43], which also present results obtained from simulating collective neutrino oscillations on current generation quantum devices.

ACKNOWLEDGMENTS

This work was supported by the InQubator for Quantum Simulation under U.S. DOE Grant No. DE-SC0020970, by the Quantum Science Center (QSC), a National Quantum Information Science Research Center of the U.S. Department of Energy (DOE), by the U.S. Department of Energy under Grants No. DE-FG02-00ER41132, No. DE-SC0021152 and by the U.S. National Science Foundation under Grants No. PHY-1404159 and No. PHY-2013047. B. H. acknowledges support from the U.S. Department of Energy (DOE) through a quantum computing program sponsored by the Los Alamos National Laboratory (LANL) Information Science & Technology Institute. We acknowledge use of the IBM Q for this work. The views expressed are those of the authors and do not reflect the official policy or position of IBM or the IBM Q team.

APPENDIX A: PAIR PRODUCT PROPAGATOR

In this appendix, we provide additional details about the pair approximation of the time-evolution operator $U_2(t)$ introduced in Eq. (10) of the main text. In particular, we will present a direct comparison with the simpler approximation $U_1(t)$ corresponding to the canonical first order Trotter-Suzuki step.

As mentioned in the main text, the asymptotic scaling of the approximation error ϵ is quadratic in the time step t for both approximations [19]. The pair approximation is expected, however, to perform better in practice for cases where an accurate description of pair evolution is important due, for instance, to strong cancellations between the one-body and two-body contributions in the Hamiltonian. In the neutrino case, these situations can occur with appropriate initial conditions so that, for typical states in the evolution, we have for most pairs

$$|\langle K_{pq} \rangle + \langle V_{pq} \rangle| \ll |\langle K_{pq} \rangle| + |\langle V_{pq} \rangle|, \quad (A1)$$

where we used the short-hand (cf. Eq. (9) in main text)

$$\langle K_{pq} \rangle = \frac{\vec{b}}{N-1} \cdot \langle \vec{\sigma}_k + \vec{\sigma}_q \rangle \langle V_{pq} \rangle = J_{kq} \langle \vec{\sigma}_k \cdot \vec{\sigma}_q \rangle. \quad (A2)$$

Since the difference between the two approximation is not expected to hold for a generic initial state, standard error measures like the matrix norm of the difference with the exact propagator

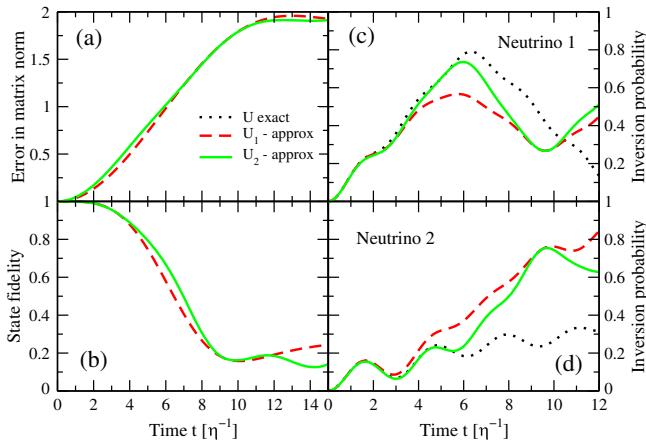


FIG. 9. Panel (a) shows the error in matrix 2-norm Eq. (A3) of the two approximations U_1 and U_2 described in the text. Panel (b) shows the state fidelity and the right panels show results for the inversion probability $P_{\text{inv}}(t)$. Panel (c) is for neutrino 1 while panel (d) is for neutrino 2.

$$\|\exp(-itH) - U_{1/2}(t)\| \quad (\text{A3})$$

are not expected to capture the effect. This is in fact found in practice for our system. In panel (a) of Fig. 9 we show the estimate from Eq. (A3) for the $N = 4$ neutrino model considered in this work. This error estimate indicates that the U_1 approximation has a smaller maximum error than U_2 up to long times. We can look at a more direct measure of accuracy for our specific setup by considering instead the state fidelity

$$f(t) = |\langle \Psi_{1/2}(t) | \Psi(t) \rangle|^2 \quad (\text{A4})$$

between the exact state $|\Psi(t)\rangle$ at time t and one of its approximations $|\Psi_{1/2}(t)\rangle$ obtained using either U_1 or U_2 . We show $f(t)$ for both approximations in panel (b) of Fig. 9. The result here suggest that instead the pair approximation produces a state with a higher fidelity than the simple linear propagator U_1 , especially at relatively long time steps $t \in [4, 8]$.

Finally, since we are mostly interested in flavor observables diagonal in the computational basis, we also show a direct comparison of the inversion probability for two out of the $N = 4$ neutrinos using both approximations and the exact propagator [panels (c) and (d)]. These results show more clearly that the pair approximation allows us to correctly describe the evolution of flavor for substantially longer times than the canonical U_1 approximation. The results reported here do depend on the specific choice of ordering of qubits in the time evolution layers shown in Fig. 1. In both the present analysis and the simulation results in the main text we used the best ordering which we empirically found to be (1,3,2,4) as one would have expected based on the initial state and the criterion Eq. (A1) above.

A more rigorous discussion of the relative accuracy between the canonical first order and the pair approximation, together with the effect of ordering choices, will be explored in future work.

APPENDIX B: ERROR MITIGATION

In the following subsections we describe in more detail the error mitigation techniques used in this work.

1. Propagation of statistical uncertainties

In this section we describe the procedure we have adopted for propagating statistical errors in the results reported in the main text. We found that careful treatment of statistical errors was important for non linear functions of the expectation values like entropy and concurrence of a reconstructed density matrix.

In the following, we will symbolically denote as $\langle O \rangle$, expectation values of Pauli operators which can be measured directly on the device. These are, for instance, the expectation values $\langle XX \rangle$, $\langle XY \rangle$, etc. needed to reconstruct a two-qubit density matrix.

We use a Bayesian approach to perform inference from the bare counts obtained from the device. The idea is best described initially for the simple case of a single qubit measurement. The probability of obtaining m measurements of the state $|1\rangle$ out of a total of M trials can be modelled as a binomial distribution

$$P_b(m; p) = \binom{M}{m} p^m (1-p)^{M-m}, \quad (\text{B1})$$

with p the probability of a $|1\rangle$ measurement. In order to infer the parameter p from a given sample m_i of measurement outcomes, we use Bayes theorem

$$P(p|m_i) = \frac{P(m_i|p)P(p)}{\int dq P(m_i|q)P(q)}. \quad (\text{B2})$$

For the single qubit measurement, we use the binomial distribution as likelihood $P(m_i|p)$ and, in order to obtain a posterior $P(p|m_i)$ in closed form, we use the conjugate prior of the binomial: the beta distribution

$$P_\beta(p; \alpha, \beta) = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} p^{\alpha-1} (1-p)^{\beta-1}. \quad (\text{B3})$$

Here $\alpha, \beta > 0$ are the parameters defining the distribution and with $\alpha = \beta = 1$ we obtain a uniform distribution. The advantage of using the Beta distribution as a prior is that, after a measurement m_i of the system is available, the parameters (α_0, β_0) of the prior distribution get updated as

$$\alpha_i = \alpha_0 + m_i \quad \beta_i = \beta_0 + M - m_i. \quad (\text{B4})$$

Intuitively we can interpret the parameters (α_0, β_0) of the prior as assigning an *a priori* number of measurements to the measurement outcomes, which are then updated as more measurements are performed. In this work we used a simple uniform prior corresponding to the choice $\alpha_0 = \beta_0 = 1$ for the prior parameters.

After the inference step described above, we calculate the expectation value of a generic nonlinear function $\langle F[O] \rangle$ by sampling new outcomes m'_k using the posterior distribution. More in detail, we generate a new artificial measurement m'_k after the measured m_i by the following procedure

- (i) sample a value p'_k from the posterior $P(p'_k|m_i)$
- (ii) sample a new measurement outcome m'_k from the likelihood $P_b(m'_k; p'_k)$

The new measurements m'_k obtained in this way are then samples from the predictive posterior distribution.

Using an ensemble of size L obtained in this way, we compute $\langle F[O] \rangle$ by taking an average of the results obtained for each individual sample

$$\langle F[O] \rangle \approx \frac{1}{L} \sum_{k=1}^L F[O_k]. \quad (\text{B5})$$

The error bars reported in the main text are 68% confidence intervals which we found in most cases where well approximated by a Gaussian approximation.

This scheme is complete only for single qubit measurements but a generalization to generic multiqubit observables can be obtained in a straightforward way. In the situation where we are estimating expectation values over N qubits, the probability of measuring a specific collection of N bit strings m_i in M repeated trials can be described with a multinomial distribution with N probabilities. We use this distribution as the likelihood $P(m_i|\vec{p})$ in Bayes theorem and, for similar reasons as above, we take its conjugate prior distribution: the Dirichlet distribution (also initialized as uniform as for the Beta above). The procedure we follow is otherwise exactly equivalent to what we described above.

2. Read-out mitigation

The qubit measurements on a real device are not perfect and it is therefore important to understand the associated systematic errors. We refer the reader to Appendix H.1 of Ref. [32] for a more detailed derivation of the exact procedure we employ and the motivations behind it. Here, we instead describe the main difference with the scheme described there which comes from the use of the Bayesian inference scheme described in the previous subsection.

In the calculations presented here, we work under the assumption that read-out errors are independent on each qubit and perform a set of $2N$ calibration measurements c_i

(requiring two separate executions) to extract the parameters (\vec{e}_0, \vec{e}_1) of the noise model (see Eq. (H1) of Ref. [32]). In order to consistently propagate the statistical uncertainties associated from the finite sample statistic used to estimate the noise parameters, we use an additional layer of Bayesian sampling using a binomial prior for the two error probabilities (e_0^n, e_1^n) associated to each qubit n .

Using a single pair of error probability vectors $\epsilon_i = (\vec{e}_0, \vec{e}_1)_i$, obtained either by direct measurement or by sampling from the posterior, we can generate a linear transformation \mathbf{C}_i that maps a set of (in general multiqubit) measurements m_i to a new set \tilde{m}_i with reduced read-out errors (see Ref. [32] for more details).

The complete procedure that we use to generate an ensemble of measurements $\{\tilde{m}_i\}$ with read-out mitigation starting from a single calibration measurement c_i and Pauli operator measurement m_i is as follows

- (i) sample a value p'_k from the posterior $P(p'_k|m_i)$
- (ii) sample a new measurement outcome m'_k from the likelihood $P_b(m'_k; p'_k)$
- (iii) for each qubit $n = \{1, \dots, N\}$
sample a pair (e_0^n, e_1^n) of error probabilities from the posterior $P(e_0^n, e_1^n|c_i)$
- (iv) use the sampled error probabilities (\vec{e}'_0, \vec{e}'_1) to generate the linear transformation \mathbf{C}'_k
- (v) apply the sampled correction matrix \mathbf{C}'_k to m'_k to obtain the read-out mitigated estimator \tilde{m}'_k

The resulting ensemble of measurements can be used directly to estimate expectation values and confidence intervals as described above. In this way, we avoid having to explicitly construct the variance of the correction matrix \mathbf{C}'_k using maximum likelihood estimation and then propagating the error perturbatively to arbitrary observables as done in Ref. [32].

3. Zero-noise-extrapolation

For observables like the inversion probability, we adopt the procedure developed in Ref. [32]. For entanglement observables we adopt a two point shifted exponential extrapolation that we briefly describe here. We denote the entanglement observable as $\langle F[O] \rangle^{(L)}(r)$ where L is the number of samples used and r denotes the noise level of the circuit, proportional to the number of CNOT gates in the circuit. We first note that in the case of very high noise levels, denoted here with $\langle F[O] \rangle(r \rightarrow \infty)$ the density matrix corresponds to the maximally mixed state given by $1/4$. Therefore, the concurrence in this case is $-1/2$ and the pair entanglement saturates to 2 .

Using an estimate for the large noise expected value $\langle F[O] \rangle(r \rightarrow \infty)$, we can then consider a simple exponential extrapolation of the form

$$\langle F[O] \rangle^{(L)}(r) - \langle F[O] \rangle(r \rightarrow \infty) = A_F^{(L)} e^{-ar}, \quad (\text{B6})$$

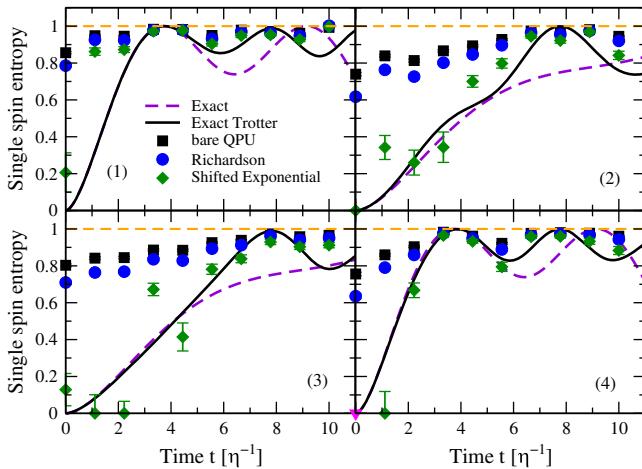


FIG. 10. Single spin entanglement entropy for all four neutrinos. Black square are bare results obtained from the QPU, the blue circles are obtained using Richardson extrapolation and the green diamonds correspond to the results obtained from a shifted exponential extrapolation using the maximum value of the entropy (dashed orange line). The magenta triangle indicates a mitigated result with shifted exponential extrapolation below zero within errorbars.

with α and $A_F^{(L)}$ the parameters of the model which can be obtain using results at two different noise levels r and r' . The zero-noise extrapolated result in this model corresponds to the limit $r \rightarrow 0$ and is given simply by the estimated $A_F^{(L)}$. More explicitly this becomes

$$A_F^{(L)} = \langle F[O] \rangle^{(L)}(r) \left(\frac{\langle F[O] \rangle^{(L)}(r')}{\langle F[O] \rangle^{(L)}(r)} \right)^{r/(r-r')}, \quad (B7)$$

and the zero noise extrapolated observable is

$$\langle F[O] \rangle^{(L)}(0) = A_F^{(L)} + \langle F[O] \rangle(r \rightarrow \infty). \quad (B8)$$

Finally, the estimated statistical error is obtained by calculating the standard deviation of the L copies as above.

APPENDIX C: ADDITIONAL DATA FOR CONCURRENCE AND ENTANGLEMENT ENTROPY

Here we show the full set of results for both entanglement entropy and concurrence for all the other pairs of qubits not shown in the main text. We denote with a

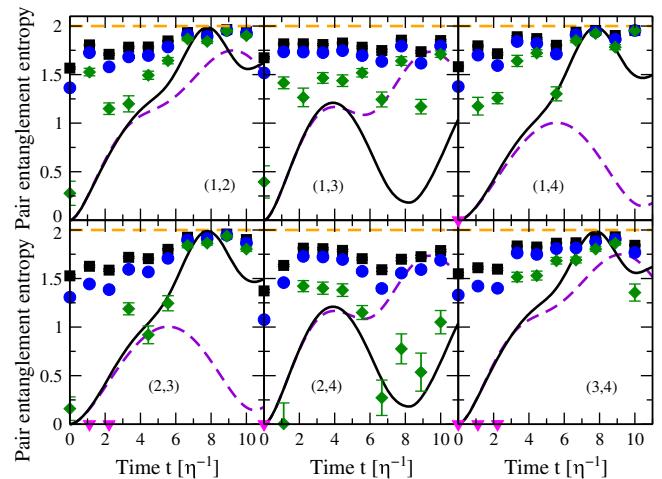


FIG. 11. Pair entanglement entropy for all pair of neutrinos. Black square are bare results obtained from the QPU, red triangles are results obtained by amplifying the noise to $\epsilon/\epsilon_0 = 3$, the blue circles are obtained using Richardson extrapolation and the green diamonds correspond to the results obtained from a shifted exponential extrapolation using the maximum value of the entropy (indicated as a dashed orange line). The magenta triangle points are mitigated results with shifted exponential extrapolation below zero within errorbars.

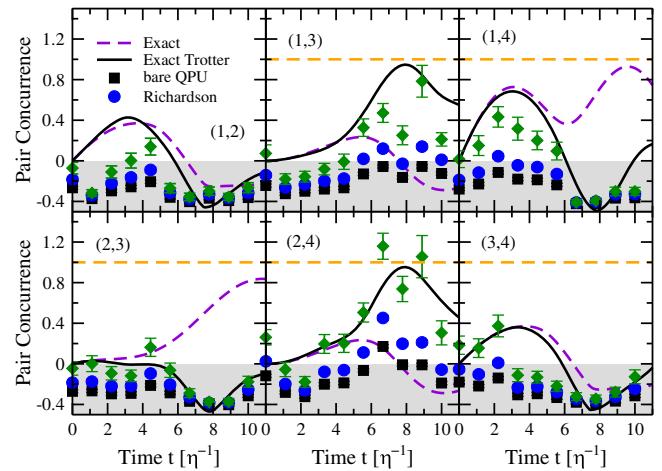


FIG. 12. Entanglement concurrence for all the pairs of qubits. The maximum value for the concurrence is shown as a dashed orange line.

magenta triangle, data points that fall below zero for the entropy as in Fig. 8 of the main text.

[1] J. Pantaleone, Neutrino oscillations at high densities, *Phys. Lett. B* **287**, 128 (1992).

[2] G. Sigl and G. Raffelt, General kinetic description of relativistic mixed neutrinos, *Nucl. Phys.* **B406**, 423 (1993).

[3] H. Duan, G. M. Fuller, and Y.-Z. Qian, Collective neutrino oscillations, *Annu. Rev. Nucl. Part. Sci.* **60**, 569 (2010).

[4] S. Chakraborty, R. Hansen, I. Izaguirre, and G. Raffelt, Collective neutrino flavor conversion: Recent developments, *Nucl. Phys.* **B908**, 366 (2016); Neutrino oscillations: Celebrating the Nobel Prize in Physics 2015.

[5] M. C. Gonzalez-Garcia and Y. Nir, Neutrino masses and mixing: Evidence and implications, *Rev. Mod. Phys.* **75**, 345 (2003).

[6] L. Wolfenstein, Neutrino oscillations in matter, *Phys. Rev. D* **17**, 2369 (1978).

[7] S. Mikheyev and A. Smirnov, Resonance amplification of oscillations in matter and spectroscopy of solar neutrinos, *Sov. J. Nucl. Phys.* **42**, 913 (1985).

[8] H. Duan, G. M. Fuller, J. Carlson, and Y.-Z. Qian, Simulation of coherent nonlinear neutrino flavor transformation in the supernova environment: Correlated neutrino trajectories, *Phys. Rev. D* **74**, 105014 (2006).

[9] N. F. Bell, A. A. Rawlinson, and R. Sawyer, Speed-up through entanglement many-body effects in neutrino processes, *Phys. Lett. B* **573**, 86 (2003).

[10] A. Friedland and C. Lunardini, Do many-particle neutrino interactions cause a novel coherent effect? *J. High Energy Phys.* **10** (2003) 043.

[11] R. F. Sawyer, “Classical” instabilities and “quantum” speed-up in the evolution of neutrino clouds, [arXiv:hep-ph/0408265](https://arxiv.org/abs/hep-ph/0408265).

[12] Y. Pehlivan, A. B. Balantekin, T. Kajino, and T. Yoshida, Invariants of collective neutrino oscillations, *Phys. Rev. D* **84**, 065008 (2011).

[13] E. Rrapaj, Exact solution of multiangle quantum many-body collective neutrino-flavor oscillations, *Phys. Rev. C* **101**, 065805 (2020).

[14] 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).

[15] A. Roggero, Entanglement and Many-body Effects in Collective Neutrino Oscillations (to be published).

[16] A. Roggero, Dynamical phase transitions in models of collective neutrino oscillations, [arXiv:2103.11497](https://arxiv.org/abs/2103.11497).

[17] S. Lloyd, Universal quantum simulators, *Science* **273**, 1073 (1996).

[18] J. Preskill, Quantum computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).

[19] M. Suzuki, General theory of fractal path integrals with applications to manybody theories and statistical physics, *J. Math. Phys. (N.Y.)* **32**, 400 (1991).

[20] D. M. Ceperley, Path integrals in the theory of condensed helium, *Rev. Mod. Phys.* **67**, 279 (1995).

[21] J. Carlson, S. Gandolfi, F. Pederiva, S. C. Pieper, R. Schiavilla, K. E. Schmidt, and R. B. Wiringa, Quantum monte carlo methods for nuclear physics, *Rev. Mod. Phys.* **87**, 1067 (2015).

[22] G. Vidal and C. M. Dawson, Universal quantum circuit for two-qubit transformations with three controlled-not gates, *Phys. Rev. A* **69**, 010301 (2004).

[23] F. Vatan and C. Williams, Optimal quantum circuits for general two-qubit gates, *Phys. Rev. A* **69**, 032315 (2004).

[24] M. C. Tran, Y. Su, D. Carney, and J. M. Taylor, Faster digital quantum simulation by symmetry protection, *PRX Quantum* **2**, 010323 (2021).

[25] C. Monroe, W. C. Campbell, L. M. Duan, Z. X. Gong, A. V. Gorshkov, P. Hess, R. Islam, K. Kim, N. Linke, G. Pagano, P. Richerme, C. Senko, and N. Y. Yao, Programmable quantum simulations of spin systems with trapped ions, *Rev. Mod. Phys.* **93**, 025001 (2021).

[26] I. D. Kivlichan, J. McClean, N. Wiebe, C. Gidney, A. Aspuru-Guzik, G. K.-L. Chan, and R. Babbush, Quantum Simulation of Electronic Structure with Linear Depth and Connectivity, *Phys. Rev. Lett.* **120**, 110501 (2018).

[27] 5qubit backed: IBM Q team, IBM Vigo backend specification v1.0.2 (2020), retrieved from <https://quantum-computing.ibm.com>.

[28] H. A. *et al.*, Qiskit: An open-source framework for quantum computing (2019).

[29] S. Endo, S. C. Benjamin, and Y. Li, Practical Quantum Error Mitigation for Near-Future Applications, *Phys. Rev. X* **8**, 031027 (2018).

[30] E. F. Dumitrescu, A. J. McCaskey, G. Hagen, G. R. Jansen, T. D. Morris, T. Papenbrock, R. C. Pooser, D. J. Dean, and P. Lougovski, Cloud Quantum Computing of an Atomic Nucleus, *Phys. Rev. Lett.* **120**, 210501 (2018).

[31] A. Roggero, A. C. Y. Li, J. Carlson, R. Gupta, and G. N. Perdue, Quantum computing for neutrino-nucleus scattering, *Phys. Rev. D* **101**, 074038 (2020).

[32] A. Roggero, C. Gu, A. Baroni, and T. Papenbrock, Preparation of excited states for nuclear dynamics on a quantum computer, *Phys. Rev. C* **102**, 064624 (2020).

[33] A. He, B. Nachman, W. A. de Jong, and C. W. Bauer, Zero-noise extrapolation for quantum-gate error mitigation with identity insertions, *Phys. Rev. A* **102**, 012426 (2020).

[34] K. Banaszek, G. M. D’Ariano, M. G. A. Paris, and M. F. Sacchi, Maximum-likelihood estimation of the density matrix, *Phys. Rev. A* **61**, 010304 (1999).

[35] J. Czartowski, D. Goyeneche, M. Grassl, and K. Życzkowski, Isoentangled Mutually Unbiased Bases, Symmetric Quantum Measurements, and Mixed-State Designs, *Phys. Rev. Lett.* **124**, 090503 (2020).

[36] T. Giurgica-Tiron, Y. Hindy, R. LaRose, A. Mari, and W. J. Zeng, Digital zero noise extrapolation for quantum error mitigation, in *2020 IEEE International Conference on Quantum Computing and Engineering (QCE)* (2020), pp. 306–316.

[37] Z. Cai, Multi-exponential error extrapolation and combining error mitigation techniques for nisq applications, [arXiv: 2007.01265](https://arxiv.org/abs/2007.01265).

[38] W. K. Wootters, Entanglement of Formation of an Arbitrary State of Two Qubits, *Phys. Rev. Lett.* **80**, 2245 (1998).

[39] S. Hill and W. K. Wootters, Entanglement of a Pair of Quantum Bits, *Phys. Rev. Lett.* **78**, 5022 (1997).

[40] Note that our results seem to converge to a larger asymptotic value of $\tilde{C} \approx -0.44$ instead of $\tilde{C} = -0.5$. We can empirically explain this difference as the effect of statistical fluctuations.

[41] J. J. Wallman and J. Emerson, Noise tailoring for scalable quantum computation via randomized compiling, *Phys. Rev. A* **94**, 052325 (2016).

[42] J. Zhang, G. Pagano, P. W. Hess, A. Kyprianidis, P. Becker, H. Kaplan, A. V. Gorshkov, Z.-X. Gong, and C. Monroe, Observation of a many-body dynamical phase transition with a 53-qubit quantum simulator, *Nature (London)* **551**, 601 (2017).

[43] K. Yeter-Aydeniz, S. Bangar, G. Siopsis, and R. C. Pooser, Collective neutrino oscillations on a quantum computer, [arXiv:2104.03273](https://arxiv.org/abs/2104.03273).