

State preparation and symmetries

Ivana Miháliková,^{1,2,3,4} Joseph Carlson,⁴ Duff Neill,⁴ and Ionel Stetcu⁴

¹*Institute of Physics of Materials, Czech Academy of Sciences, Brno 616 00, Czech Republic*

²*Institute of Physics, Slovak Academy of Sciences,
Bratislava 84511, Slovakia; Matej Bel University,
Národná ulica 12, Banská Bystrica, 97401, Slovakia*

³*Department of Condensed Matter Physics, Faculty of Science,
Masaryk University, Brno 611 37, Czech Republic*

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

We demonstrate the importance of symmetries in Variational Quantum Eigensolver (VQE) algorithms to prepare the ground or specific low-lying states of quantum Hamiltonians. We examine two spin problems, one with random all-to-all couplings inspired by neutrino flavor evolution in supernovae, and the standard Heisenberg spin Hamiltonian on a 4×3 lattice. The neutrino Hamiltonian has the total spin J and third component J_z as its only symmetries. The Heisenberg model has these symmetries plus translational invariance and reflection symmetry.

We demonstrate that the convergence of variational methods is dramatically improved by keeping all symmetries. In both cases a nearly exact solution can be obtained in cases where standard unconstrained variational algorithms fail. Since variational algorithms can use standard Trotter steps as part of the optimization, allowing additional correlations that obey all the symmetries of the Hamiltonian will speed convergence of variational algorithms. This will lead to faster convergence than standard projection algorithms.

I. INTRODUCTION

Quantum ground-state preparation requires good classical product-state ansatz [1–3]. Ideally, the ansatz should have large overlap with the true ground state of the system and should be simple enough to be easily prepared on a quantum computer, if possible only involving single-qubit gates. However, for a given target Hamiltonian, it may be the case that only poor guarantees can be made for the initial product state overlap, and even finding the optimal classical product state can be shown to be a hard problem, if the system is a classical spin-glass, for instance [4–8]. The guarantees are generally a function of the locality and dimensionality of the Hamiltonian, where we use locality in the computer science sense of number of sites that can interact via a single term in the Hamiltonian, and dimensionality to count the number of sites that a given site can interact with. For problems with a spatial lattice, this counting dimensionality is directly correlated with the spatial dimensionality of the lattice. Generally, as the dimensionality increases, one can expect product states to produce an ever better overlap with the ground state. The underlying intuition is the monogamy of entanglement: each site is pushed to being entangled with the sites that it can interact with, but must share that entanglement in a roughly equal fashion according to the dimensionality of the Hamiltonian.

For large lattices in low dimensions, the entanglement can be high. Variational quantum algorithms can be useful, though they often suffer from the ‘barren plateau’ issue [9–14]. In many directions the optimized variational function has very small gradients with respect to changes in the (many) variational parameters [15–19]. However, because barren plateaus describe only the average landscape, regions with non-vanishing gradients can still be found when the circuit is initialized sufficiently close to a good minimum [20–24]. Reducing the number of variational parameters can produce useful results, though sometimes at the cost of reduced accuracy.

An important consideration for many problems of physical interest is the presence of either exact or approximate symmetries in the problem [25–29]. In the limit of exact symmetries the Hilbert space will break up into non-interacting subspaces, each subspace being specified by a set of quantum numbers that characterize an irreducible representation of the symmetry group. Constructing states that belong to a specific irreducible representation in general requires entanglement. Take the case of $SU(2)$ as a concrete example of a symmetry group, and a Hilbert Space composed from N qubits. Then it is simple to see that to build a state with total spin of zero, one can pair each spin, and put the pairs into a locally spin zero state, anti-symmetric under interchange. Such symmetry-adapted product states can serve as warm starts for quantum algorithms, where the initial state is prepared with respect to the symmetry of the problem [30–34].

In this work, we examine, for small systems, how effectively a product-state ansatz can approximate the true ground state of symmetric Hamiltonians, and how variational quantum algorithms can enhance this process. We then explore the extent to which projection algorithms alone can improve the overlap of such product states with the total spin-zero subspace. Following this projection, we apply a restricted variational quantum algorithm specifically designed to preserve the underlying symmetries of the Hamiltonian. We focus on the case for an underlying $SU(2)$ symmetry, comparing an all-to-all coupled system inspired from neutrino physics [35, 36] (where the symmetry derives from the

approximate flavor symmetry of the neutrino sector of the Standard Model under neutral current interactions) to the standard magnetic Heisenberg model in 2 dimensions. For the neutrino-inspired Hamiltonian, it is simple to write down the specific product state that is supposed to dominate the ground state in the classical limit [37] as the number of neutrinos becomes large, a state called the flavor-momentum locked state. For the Heisenberg model, we use the anti-ferromagnetic Néel state as the initial product state [38].

II. MODELS AND ALGORITHMS

A. Spin models

The Heisenberg spin model [39] is a fundamental framework in statistical mechanics and condensed matter physics that describes the interactions between spins arranged on a lattice. In this work, we consider two different Hamiltonians. One is the standard spin 1/2 Heisenberg Hamiltonian on a 2D square lattice. This Hamiltonian is given by

$$H = \sum_{\langle ij \rangle} \sigma_i \cdot \sigma_j, \quad (1)$$

where the sum runs over nearest-neighbor pairs on the lattice. For illustrative purposes we use a 2D square lattice with periodic boundary conditions, but similar techniques could be used on other lattices including those that cannot be simulated exactly by classical computers with Quantum Monte Carlo methods. This Hamiltonian has several symmetries. The total spin and its projection on an axis are conserved, and there are translational symmetry and mirror symmetries in the x - and y -planes.

We also consider the Heisenberg-like spin Hamiltonian including all-to-all random interactions between pairs of spins. This model has often been used to study neutrino dynamics in supernovae [37]. In this case, the Hamiltonian is

$$H = \sum_{i < j} (1 - \cos \theta_{ij}) \sigma_i \cdot \sigma_j. \quad (2)$$

The term $(1 - \cos \theta_{ij})$ represents a coupling strength or interaction coefficient between the i -th and j -th SU(2) spin in the system, where θ_{ij} is an angle parameter that influences the interaction, and σ_i and σ_j are Pauli matrices representing the spin operators at sites i and j , respectively. In our model, $\cos \theta_{ij}$ is randomly chosen from the range $[-1, 1]$, ensuring that the coupling strength is always non-negative and varies smoothly from 0 to 2.

This Hamiltonian again conserves the total spin J and its third component J_z . However, there are no apparent additional symmetries.

B. Initial states

In our variational quantum simulations of these spin models, we start with the tensor product of individual spin coherent states for N spins [37] as an initial state. These states are very simple product states that are easy to prepare on a quantum computer. For the standard Heisenberg model we use the symmetrized Néel state

$$|\psi_i\rangle = \frac{1}{\sqrt{2}} \left[\prod_i \sigma_x(i)^{x+y} + \prod_i \sigma_x(i)^{x+y+1} \right] \prod_i |0\rangle_i. \quad (3)$$

This state is unique for a square lattice with even lengths in both directions. All nearest neighbors have expectation value -1 for that term in the Hamiltonian. For the 4×3 lattice considered here, we also use this state, although the periodic boundary conditions imply that the pairs wrapping around the lattice in one direction have a positive expectation value.

For the all-to-all coupled neutrino model, we employ a product state. In this scenario, we use the state outlined in Ref. [37], where the spin of each particle is aligned to the momentum. The state is

$$\psi_i = \prod_i R(i)|0\rangle, \quad (4)$$

with the rotation R defined so that its expectation values are parallel to the momenta used to define the Hamiltonian.

In other words, the state can be written as

$$|\psi_i\rangle = \bigotimes_{i=1}^N \exp\left(i\pi k_i \cdot \frac{\sigma_i}{2}\right) |0\rangle_i, \quad (5)$$

where k_i is parallel to σ_i , and k_j is parallel to σ_j .

The vectors k are defined as

$$k_i = \cos(\phi_i) \sin\left(\frac{\theta_i}{2}\right) e_x + \sin(\phi_i) \sin\left(\frac{\theta_i}{2}\right) e_y + \cos\left(\frac{\theta_i}{2}\right) e_z, \quad (6)$$

the azimuthal angle ϕ is randomly chosen from the range $[0, 2\pi]$. These vectors represent spin directions in spherical coordinates, capturing the randomized nature of the system's initial state.

This initial state represents a product of single-spin states, where each spin is rotated from the reference state $|0\rangle_i$ by an operator $\exp(i\pi k_i \cdot \sigma_i/2)$, where k_i is a unit vector associated with each neutrino, representing its momentum, and σ_i is the Pauli spin vector operator acting on qubit i . This tensor product state can be considered as a mean-field-like state where each spin aligns with its corresponding momentum vector. It approximates the ground state of the system in the classical limit and provides a semiclassical interpretation of the quantum dynamics.

To implement this state on quantum computer, we used the fact that any single-qubit operator can be decomposed into a sequence of three rotations about the z - and y -axes. Specifically, we used $R_z R_y R_z$ decomposition that provides a symmetric and standard way to decompose arbitrary rotations, reflecting the Euler angle decomposition

$$|\psi_i\rangle = \exp(-i\alpha_i Z) \exp(-i\beta_i Y) \exp(-i\gamma_i Z) |0\rangle_i = R_z(\alpha_i) R_y(\beta_i) R_z(\gamma_i) |0\rangle_i. \quad (7)$$

C. Spin Projection Algorithm

One of the applications of the projection algorithm introduced in Ref. [40] was to prepare states with good symmetries. However, for projection on total spin $J = 0$ it is more efficient to use a simplified algorithm [41] that requires more measurements but avoids deep quantum circuits with complex Pauli strings and Trotterization. For the Heisenberg model this approach simplifies further and consists of repeated projections on total spin projection zero using the operator J_z (which involve a measurement), followed by a single rotation by $\pi/2$ around the x axis. This significantly reduces circuit depth and makes the approach more practical for today's quantum hardware. To perform the projection on $J = 0$, the algorithm starts with an initial quantum state, which is generally a superposition of states with different values of total angular momentum J and their corresponding magnetic projections M . The goal is to remove all components that do not belong to the $J = 0$ subspace.

One of the key steps in this process is the application of unitary evolution followed by a measurement of an ancilla to eliminate unwanted components. The state evolution under J_z is given by

$$\begin{aligned} |\psi(t_i)\rangle &= \exp(-it_i J_z \otimes Y_a) |\psi\rangle \otimes |0\rangle_a \\ &= \cos(J_z t_i) |\psi\rangle \otimes |0\rangle_a + \sin(J_z t_i) |\psi\rangle \otimes |1\rangle_a. \end{aligned} \quad (8)$$

Choosing a proper series of times $t_i = \pi/2^i$ ($i > 0$), repeated measurements of the ancilla in state $|0\rangle_a$ projects on many-body states with $J_z = 0$. However, because states with $J > 0$ also have a $J_z = 0$ component, the algorithm must continue to filter out those components. A rotation around the x (or y) axis with an angle $\varphi \neq 0$ or π reshuffles contributions from $J > 0$, $J_z = 0$ states to states with $J > 0$, $|J_z| > 0$, and hence an additional projection as above reduces the weight of the $J > 0$ states, which become zero after a relatively small number of iterations that involve projection followed by rotation around the x or y axis by $\pi/2$ (the last-step rotation is not necessary) [41]. The exact number of iterations depends on the contamination with states with $J > 0$, but in practice we needed to perform about 11 iterations, somewhat better than in Ref. [41] where the number of iterations for nuclear physics problems was around 40–50. While in other cases, like in nuclear physics for even-even nuclei, this algorithm could be used to project on $J = 0$ states, increasing the gap and hence allowing for a more efficient state preparation via projections [40], in the current investigation we use it to efficiently project onto states with total $J = 0$ and $J_z = 0$ (equal numbers of spin up and down), and then continue with a variational approach.

D. Variational Ansatz

In our VQE calculations, we utilize an ansatz that incorporates the swap operator $\hat{\rho}_{ij}$, which exchanges the quantum states of two qubits. This operator has the advantage of conserving both the total and z-component of the spin of the pair, and hence the whole system. The exponentiation of this operator, $\exp(i\theta\hat{\rho}_{ij})$, generates entangling interactions that can be decomposed into a sequence of two-qubit Pauli interactions of the form $\exp(i\theta\sigma_i^a\sigma_j^a)$ for $a = X, Y, Z$, which are implemented using controlled-NOT (CNOT) gates and single-qubit rotations. The exponentiated swap operator is given by

$$\exp(i\theta\hat{\rho}_{ij}) = \begin{bmatrix} e^{i\theta} & 0 & 0 & 0 \\ 0 & \cos(\theta) & i\sin(\theta) & 0 \\ 0 & i\sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 0 & e^{i\theta} \end{bmatrix}, \quad (9)$$

where θ represents the variational parameter. The individual exponentiated Pauli interactions have the following forms

$$\exp(i\theta\sigma_i^X\sigma_j^X) = H_i H_j \text{CNOT}_{i \rightarrow j} R_{i,z}(-\theta) \text{CNOT}_{i \rightarrow j} H_i H_j. \quad (10)$$

$$\exp(i\theta\sigma_i^Y\sigma_j^Y) = H_i^Y H_j^Y \text{CNOT}_{i \rightarrow j} R_{i,z}(-\theta) \text{CNOT}_{i \rightarrow j} H_i^{Y\dagger} H_j^{Y\dagger}, \quad (11)$$

where

$$H_i^Y = S_i H_i \sigma_i^Z S_i, \quad (12)$$

and

$$H_i^{Y\dagger} = \sigma_i^Z S_i \sigma_i^Z H_i \sigma_i^Z S_i. \quad (13)$$

$$\exp(i\theta\sigma_i^Z\sigma_j^Z) = \text{CNOT}_{i \rightarrow j} R_{i,z}(-\theta) \text{CNOT}_{i \rightarrow j}. \quad (14)$$

By combining these exponentials, we obtain the full exponentiated swap operator decomposition

$$\exp(i\theta\hat{\rho}_{ij}) = e^{i\theta/2} \exp\left(i\frac{\theta}{2}\sigma_i^X\sigma_j^X\right) \exp\left(i\frac{\theta}{2}\sigma_i^Y\sigma_j^Y\right) \exp\left(i\frac{\theta}{2}\sigma_i^Z\sigma_j^Z\right). \quad (15)$$

Here, $e^{i\theta/2}$ represents an overall phase factor that does not affect the computational basis measurement results. This ansatz preserves the interaction structure while maintaining low circuit depth.

Arbitrary pair spin swaps are appropriate for the neutrino Hamiltonian as the spin symmetries are the only relevant ones. However, for the Heisenberg Hamiltonian, translational and reflection symmetry are also important.

We can enforce these symmetries approximately by first considering all nearest-neighbor pairs and assigning the same swap amplitude to each. Correlations can then be added for next-nearest-neighbor pairs, such as those separated by a distance of 2 in one direction, and so on. In each case, all pairs are equally spaced, which helps maintain approximate translational symmetry. However, the symmetry is not exact because not all pair operations commute, leading to small admixtures of states that break translational invariance.

III. NUMERICAL RESULTS

We show results for the neutrino model with 12 randomly coupled spins as described above and for the 4×3 Heisenberg model that has many additional symmetries. Our initial product states have energy distance to the ground state well within the bounds derived in Ref. [4], which are 166.7 energy units to the ground state for the neutrino

model and 42.4 for the Heisenberg model. We note that the normalization differences between our hamiltonians and that of Ref. [4], which always normalizes by the total number of spins and the degree of the interaction graph of the hamiltonian, that is, the number of sites each site can be connected to.. This indicates our initial product-states are performing much better than the worst case guarantees for on how well a product state can approximate the ground state. Specifically, these bounds quantify how well a product-state can approximate the ground of a quantum hamiltonian, given no other considerations about the underlying physics. Roughly speaking, they bound the energy per site or qubit involved in the interaction, and per number of sites/qubits that a single site/qubit can interact with. For the neutrino interactions, this becomes a more stringent bound, as the number of sites that interact with a given site scales with the system size, while for the Heisenberg model, the bound becomes a constant as the system grows in size.

The first step after creating the product state is to project onto the $J = 0$ subspace. With a finite number of steps, the convergence is not exact, but it is very rapid.

The energy versus iteration step of the projection to spin zero is shown in Fig. 1. For the neutrino case, the product state should become nearly exact for large N , projection to spin zero improves the initial state dramatically for 12 spins, improving the energy approximately 80% of the way to the exact. For the Heisenberg case, there are additional symmetries important to get the ground state correct. Projecting on good total spin (and J_z) is important, but the spin projection alone improves the energy around 50% to the ground state.




FIG. 1. Energy difference ($E - E_{\text{exact}}$) of the projected state at each step of the projection algorithm applied to the initial state of the Neutrino model, and the Heisenberg model. The unprojected initial state is shown as the first point at step 0, while the subsequent points connected by solid lines represent the energy after each of the 11 projection steps. The red line indicates the exact ground state energy.

The next step in the algorithm is the VQE convergence, with results shown in Fig. 2. For the neutrino case, see Fig. 2(a), we use random pair swaps as projection operators to construct the all-to-all connected ansatz for VQE optimization, projecting state into more entangled subspace. The swaps preserve the total $J = 0$ character of the state. This system has a fairly large gap to higher $J = 0$ states, which also improves the convergence. The final energy of the VQE projected state is $E = -47.46$ compared to the exact -47.66 (see Table I).

TABLE I. Energy values during optimization for the Neutrino and Heisenberg models.

Optimization step	Neutrino energy	Fidelity	Heisenberg energy	Fidelity
Initial product state	-24.93	0.024	-32.00	0.085
J^2 Projection 11 iterations	-44.15	0.567	-45.33	0.296
VQE convergence	-47.46	0.988	-56.66	0.897
Translation/Mirror Projection	-	-	-58.52	0.985
Reference	-47.66	1.000	-58.95	1.000




FIG. 2. Swap convergence plots using the VQE algorithm for the Neutrino model (a), and the Heisenberg model (b) starting from the projected state obtained after 11 steps of the projection algorithm represented by the orange dashed line. The red line is the exact ground state energy. The plots show the energy difference $E - E_{\text{exact}}$. In both cases, the dark green curve shows the VQE energy as a function of iteration that preserve the symmetry of the model. In the Heisenberg case (b) additional strategies are shown: the green curve uses nearest-neighbor connectivity, and the olive-green curve uses all-to-all connectivity. Final translational/reflection symmetry projected state from VQE is shown as a blue dashed line.

For Heisenberg, in Fig. 2(b), there are multiple symmetries. We compare results for three approaches: (i) a VQE with randomly ordered nearest-neighbor swaps, (ii) an all-to-all VQE, and (iii) a VQE that approximately conserves translation and mirror symmetry. For the last case, the algorithm uses the same parameters for all swaps between pairs of spins that are the same distance apart. One layer connects the nearest neighbors (vertical and horizontal), another connects the diagonal ones, and then the second nearest neighbors. We applied periodic boundary conditions to ensure that the pattern is consistent across the lattice. This setup respects the symmetries of the system and still preserve the total spin. We initialized the variational parameters (Eq. 15) in the ansatz according to the interaction type. For the first layer, which applies gates between nearest neighbors (vertical and horizontal), we set the parameter values to 0.01. For the second layer, which connects diagonal neighbors, we used slightly larger values of 0.15, and for the third layer, which connects second nearest neighbors, we again set the parameters to 0.01. Small initial parameters for the nearest and second nearest neighbor layers (0.01) keep the circuit close to the identity at the start of the optimization. The larger initial value for the diagonal layer (0.15) gives more weight to diagonal interactions, which can introduce relevant correlations in the system.

(i) The nearest-neighbor VQE, retains the $J = 0$ character of the ground state, but is difficult for the algorithm to find a translationally invariant state since it is using only local operators. Thus, the final energy does not work as well as the symmetry-preserving algorithms.

(ii) The VQE all-to-all, again preserves the total spin, in principle it can also obtain a translationally and mirror invariant state but in practice it is difficult to find that. It improves the longer distance correlations (entanglement) compared to the nearest-neighbor.

(iii) The best algorithm is the one that (approximately) preserves all symmetries. The symmetries are not exact because of non-vanishing commutators between different pairs at the same separation. This VQE yields an energy of -56.66 compared to the initial product state energy of -32.00 and to the spin projected Néel state energy of -45.33 , while the exact value of energy is -58.95 .

One can project on translational symmetry and mirror symmetries in a manner very similar to the spin projection above. In effect, we add a term to the Hamiltonian

$$\delta H = \sum_{\mathbf{x}, \mathbf{y}} c_{\mathbf{x}, \mathbf{y}} \prod_{(i, j)} T(j \rightarrow i). \quad (16)$$

This term simultaneously translates each spin to a nearest-neighbor site in the x - or y -directions. The x - and y - terms

can have different coefficients. This essentially adds a harmonic oscillator term with x - and y -masses in the center of mass. The spectrum is linear in energy (a harmonic oscillator) and all excitations can be projected out very quickly. The mirror symmetries can be treated similarly.

Projecting out exact translational and mirror symmetries from the VQE state (the k^2 projection, where $k = (k_x, k_y)$ denotes the lattice momentum quantum numbers associated with translations in the x - and y -directions and the state is restricted to the translationally invariant $k = 0$ sector) yields an energy of -58.52 compared to the exact -58.95 , and is indicated by the horizontal dark blue dash-dotted line in Fig. 2. The algorithm for such a projection is very similar to the spin projection described above.




FIG. 3. Fidelity of the prepared state with respect to the 4096 classically computed eigenstates during optimization for the Neutrino model (a), and the Heisenberg model (b). The fidelity is computed as $|\langle \psi | \phi \rangle|^2$, where $|\psi\rangle$ is the prepared state and $|\phi\rangle$ is the reference state.

Figure 3 shows the squared overlaps (fidelities) obtained by comparing each of our prepared states with all exact eigenstates of the corresponding Hamiltonian, which were obtained classically. The results are displayed for three prepared states: the initial state, the state obtained after projection onto the appropriate symmetry subspace, and the state obtained after VQE optimization. In the case of the Heisenberg model, the fourth state corresponds to the VQE optimized state after the additional symmetry projection.

In case of the neutrino model, Fig. 3(a) shows the fidelity between the initial coherent product state and the exact eigenstates of the fully connected neutrino Hamiltonian. The overlap with the exact ground state is relatively low 2.4%, and to improve it, we project the state onto the $J = 0$ subspace. The projection leads to a state whose maximum overlap matches the exact ground state 56.7%. The plot further shows only minor contributions from the exact higher-energy states, confirming that the projection procedure effectively filters out components of the studied product state outside the target symmetry subspace. The fidelities of the final state obtained after the full variational optimization with the exact eigenstates of the neutrino Hamiltonian exhibit almost perfect overlap with the true ground state 98.8%, while contributions from all higher-energy states are negligibly small.

For the Heisenberg model, the overlap squared of the superposition of two Néel states with the true ground state, that is in the total spin $J = 0$ subspace, is 8.5%, see Fig. 3(b). However, significant overlap is observed with two exact higher-energy states, which are 0.3% and 14.0%. With the projection onto the total spin $J = 0$, we reduce the size of the state space and provide a better starting point for the VQE convergence. The overlap squared of the projected of the studied product state with the exact ground state is 29.6%. Even after applying the projection algorithm for 11 iterations, the resulting state still shows a nonzero overlap with exact higher-energy eigenstates. Despite this, the projection improves the overlap with the exact ground state and helps concentrate the state within the correct symmetry sector. The state obtained after VQE optimization has a very high overlap with the exact ground state, which is 89.7%. Furthermore, this state can be exactly projected onto a state of good translational symmetry and good mirror symmetry in the x - and y - directions, yielding an overlap with the true ground state of greater than 98%, similar to the neutrino Hamiltonian.

In each case, the final VQE state has excellent overlap with the ground state, and projecting on the symmetries dramatically increases the gap. Further improvements by projection algorithms ([40]) would be very effective.

It is not entirely surprising that these symmetries are so important. The full Hilbert space for a 4×3 lattice is $2^{12} = 4096$. Of these, 924 have $J_z = 0$, restricting to $J = 0$ as well yields 132 states. Further restricting to translationally invariant states and mirror symmetry drastically reduces the number of relevant states to 9.

Similarly, the gap to the first excited state increases dramatically when all symmetries are retained. The exact ground state is at $E = -58.95$, the first excited state with $J_z = 0$ is at -51.81 , the first with $J = 0$ and $J_z = 0$ is at -46.72 . The first excited state with all the same symmetries is at an energy of -30.86 . The increased gap is extremely valuable in time projection to algorithms to the true ground state [40]. In this Hamiltonian the total J , J_z and total translation/mirror symmetry each reduce the space by a factor proportional to the size of the lattice. These impacts will remain very important for larger lattices.




FIG. 4. Fidelity of the prepared state obtained using the VQE algorithm with all-to-all connectivity, with respect to the reference ground state of the Neutrino model. The fidelity is shown on a logarithmic scale, where each total spin subspace S is indicated by a different color.

In Fig. 4 we show the probability of all eigenstates in the variational state for the neutrino case. The overlap is by far the largest with the true ground state, the other $J = 0$ states have probabilities of 10^{-5} to 10^{-4} . The spin 2 states have the next highest probabilities after 11 projection steps, around 10^{-9} . The $J = 2, 4$ states have similar amplitudes, but there are fewer of those in this finite system because we are also projecting on J_z . The $J = 1$ states and all $J = 3$ or higher odd J states have probabilities around 10^{-30} after even 11 steps. The trial state has very little $J = \text{odd}$ contamination, the states with $J = 2, 4, 6$ are higher. These probabilities could be reduced dramatically with a few more iteration steps if desired.

IV. CONCLUSIONS

We have studied two quantum spin Hamiltonians with variational quantum algorithms. These two problems have different spectral gaps, spectral densities and symmetries. In both cases projecting to the correct symmetries and then preserving them, at least approximately, with the variational ansatz is extremely effective in improving the efficiency of the variational method. Additional studies of different spin problems or continuum Fermion problems represented in a set of basis states would be very valuable.

The variational solutions obtained remain approximate and could be further refined using time-projection techniques, such as QPE and related algorithms. By constraining the variational states in this algorithm to the correct symmetries, the relevant spectral gaps are increased—sometimes substantially. Since projection techniques require time evolution proportional to the inverse of the gap, combining variational algorithms with strong symmetry enforcement and projection methods can be especially powerful. Moreover, because variational steps can be interleaved with time evolution, even more efficient algorithms may be achievable for a wide range of important quantum state preparation tasks. Further investigations into strategies for reducing circuit depth—and thereby lowering noise sensitivity—would be highly beneficial.

In the neutrino model, projection into the $J = 0$ subspace improved the energy from -24.93 to -44.15 . This

brought the state $\approx 80\%$ of the way towards the exact ground-state energy before any variational optimization. The corresponding fidelity with the exact ground state increased from 2.4% to 56.7%. In the Heisenberg model, the spin projection step led to a smaller improvement from -32.00 to -45.33 in energy, and from 8.5% to 29.6% in fidelity.

The optimization of VQE further improved both systems. The neutrino model reaching 98.8% fidelity while the Heisenberg model reaching 89.7%. By applying the additional projection to preserve translational and mirror symmetries increased the fidelity above 98%. These results confirm that symmetry preservation improves convergence and also produces states that are much closer to the true ground state.

For the 4×3 Heisenberg model, the full $2^{12} = 4096$ -dimensional space is reduced to 924 states by projecting onto $J_z = 0$, then to 132 by restricting to $J = 0$, and finally to only 9 states when translational and mirror symmetries are also preserved. This reduction in dimensionality is accompanied by the increase in the effective spectral gap, from ≈ 7 (-58.95 and -51.81) to nearly 28 (-58.95 and -46.72), making projection algorithms more efficient.

These observations suggest that the symmetry projection is applicable to other many-body systems where the symmetry group can be efficiently implemented in a quantum circuit. The combination of symmetry-preserving variational ansatz and projection steps can yield highly accurate results with shallower circuits, which makes the approach well-suited for near-term quantum hardware. Extensions of this work to larger lattices, other spin models, and fermionic problems would help to further assess the scalability and generality of this strategy.

ACKNOWLEDGEMENTS

I.M. acknowledges financial support from the Czech Academy of Sciences (the *Praemium Academiae* awarded to Martin Friák, from IPM in Brno, Czech Republic). She was also partially supported by the U.S. DOE through a quantum computing program sponsored by the LANL Information Science & Technology Institute. This work was carried out under the auspices of the National Nuclear Security Administration of the U.S. Department of Energy at Los Alamos National Laboratory under Contract No. 89233218CNA000001. IS and JC gratefully acknowledge support by the Advanced Simulation and Computing program. JC and DN also acknowledge the Quantum Science Center for partial support of their work on this project. Computational resources were provided by the Ministry of Education, Youth and Sports of the Czech Republic under the Projects e-INFRA CZ (ID:90254) at the IT4Innovations National Supercomputing Center (computational time project assigned to Martin Friák, from IPM in Brno, Czech Republic).

-
- [1] S. Fomichev, K. Hejazi, M. S. Zini, M. Kiser, J. Fraxanet, P. A. M. Casares, A. Delgado, J. Huh, A.-C. Voigt, J. E. Mueller, and J. M. Arrazola, Initial state preparation for quantum chemistry on quantum computers, [PRX Quantum](#) **5**, 040339 (2024).
 - [2] B. Bauer, S. Bravyi, M. Motta, and G. K.-L. Chan, Quantum computational chemistry in the age of quantum computing, [Chemical Reviews](#) **120**, 12685 (2020).
 - [3] D. W. Berry, Y. Tong, T. Khatyar, A. White, T. I. Kim, G. H. Low, S. Boixo, Z. Ding, L. Lin, S. Lee, G. K.-L. Chan, R. Babbush, and N. C. Rubin, Rapid initial-state preparation for the quantum simulation of strongly correlated molecules, [PRX Quantum](#) **6**, 020327 (2025).
 - [4] F. G. S. L. Brandão and A. W. Harrow, Product-State Approximations to Quantum States, [Communications in Mathematical Physics](#) **342**, 47 (2016).
 - [5] C. Krumnow, Z. Zimborás, and J. Eisert, Extendibility of fermionic states and rigorous ground state approximations of interacting fermionic systems, [arXiv](#) (2024), 2410.08322 [quant-ph].
 - [6] J. Kallaughner, O. Parekh, K. Thompson, Y. Wang, and J. Yirka, Complexity classification of product state problems for local hamiltonians (Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025).
 - [7] E. R. Anschuetz, D. Gamarnik, and B. T. Kiani, Bounds on the ground state energy of quantum p -spin hamiltonians, [arXiv](#) (2024), 2404.07231 [quant-ph].
 - [8] F. Barahona, On the computational complexity of ising spin glass models, [Journal of Physics A: Mathematical and General](#) **15**, 3241 (1982).
 - [9] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles, Cost function dependent barren plateaus in shallow parametrized quantum circuits, [Nature Communications](#) **12**, 1791 (2021).
 - [10] M. Larocca, S. Thanasilp, S. Wang, K. Sharma, J. Biamonte, P. J. Coles, L. Cincio, J. R. McClean, Z. Holmes, and M. Cerezo, Barren plateaus in variational quantum computing, [Nature Reviews Physics](#) **7**, 174 (2025).
 - [11] S. Wang, E. Fontana, M. Cerezo, K. Sharma, A. Sone, L. Cincio, and P. J. Coles, Noise-induced barren plateaus in variational quantum algorithms, [Nature Communications](#) **12**, 6961 (2021).
 - [12] J. Cunningham and J. Zhuang, Investigating and mitigating barren plateaus in variational quantum circuits: A survey, [Quantum Information Processing](#) **24**, 48 (2025).

- [13] X. Liu, G. Liu, H.-K. Zhang, J. Huang, and X. Wang, Mitigating barren plateaus of variational quantum eigensolvers, *IEEE Transactions on Quantum Engineering* **5**, 1 (2024).
- [14] C. Ortiz Marrero, M. Kieferová, and N. Wiebe, Entanglement-induced barren plateaus, *PRX Quantum* **2**, 040316 (2021).
- [15] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, Variational quantum algorithms, *Nature Reviews Physics* **3**, 625–644 (2021).
- [16] D. A. Fedorov, B. Peng, N. Govind, and Y. Alexeev, Vqe method: a short survey and recent developments, *Materials Theory* **6**, 2 (2022).
- [17] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O’Brien, A variational eigenvalue solver on a photonic quantum processor, *Nature Communications* **5**, 4213 (2014).
- [18] J. Tilly, H. Chen, S. Cao, D. Picozzi, K. Setia, Y. Li, E. Grant, L. Wossnig, I. Rungger, G. H. Booth, and J. Tennyson, The variational quantum eigensolver: A review of methods and best practices, *Physics Reports* **986**, 1 (2022).
- [19] I. Stetcu, A. Baroni, and J. Carlson, Variational approaches to constructing the many-body nuclear ground state for quantum computing, *Phys. Rev. C* **105**, 064308 (2022).
- [20] H. R. Grimsley, G. S. Barron, E. Barnes, and N. J. Mayhall, Adaptive, problem-tailored variational quantum eigensolver mitigates rough parameter landscapes and barren plateaus, *npj Quantum Information* **9**, 19 (2023).
- [21] A. A. Mele, G. B. Mbeng, G. E. Santoro, M. Collura, and P. Torta, Avoiding barren plateaus via transferability of smooth solutions in a hamiltonian variational ansatz, *Phys. Rev. A* **106**, L060401 (2022).
- [22] C.-Y. Park and N. Killoran, Hamiltonian variational ansatz without barren plateaus, *Quantum* **8**, 1239 (2024).
- [23] Y. Wang, B. Qi, C. Ferrie, and D. Dong, Trainability enhancement of parameterized quantum circuits via reduced-domain parameter initialization, *Phys. Rev. Appl.* **22**, 054005 (2024).
- [24] T. Tsuchimochi, Y. Mori, and S. L. Ten-no, Spin-projection for quantum computation: A low-depth approach to strong correlation, *Phys. Rev. Res.* **2**, 043142 (2020).
- [25] J. Gibbs, Z. Holmes, and P. Stevenson, Exploiting symmetries in nuclear hamiltonians for ground state preparation, *Quantum Machine Intelligence* **7**, 14 (2025).
- [26] C. T. Chubb and S. T. Flammia, Approximate symmetries of hamiltonians, *Journal of Mathematical Physics* **58**, 082202 (2017).
- [27] T.-C. Yen, R. A. Lang, and A. F. Izmaylov, Exact and approximate symmetry projectors for the electronic structure problem on a quantum computer., *The Journal of chemical physics* **151** **16**, 164111 (2019).
- [28] M. C. Tran, Y. Su, D. Carney, and J. M. Taylor, Faster digital quantum simulation by symmetry protection, *PRX Quantum* **2**, 010323 (2021).
- [29] D. Lacroix, E. A. Ruiz Guzman, and P. Siwach, Symmetry breaking/symmetry preserving circuits and symmetry restoration on quantum computers, *The European Physical Journal A* **59**, 3 (2023).
- [30] F. Truger, J. Barzen, F. Leymann, and J. Obst, Warm-starting the vqe with approximate complex amplitude encoding, *arXiv* (2024), 2402.17378 [quant-ph].
- [31] D. J. Egger, J. Mareček, and S. Woerner, Warm-starting quantum optimization, *Quantum* **5**, 479 (2021).
- [32] Y. Chai, K. Jansen, S. Kühn, T. Schwägerl, and T. Stollenwerk, Structure-inspired ansatz and warm start of variational quantum algorithms for quadratic unconstrained binary optimization problems, *arXiv* (2024), 2407.02569 [quant-ph].
- [33] R. Puig, M. Drudis, S. Thanasilp, and Z. Holmes, Variational quantum simulation: A case study for understanding warm starts, *PRX Quantum* **6**, 010317 (2025).
- [34] J. A. Sheikh, J. Dobaczewski, P. Ring, L. M. Robledo, and C. Yannouleas, Symmetry restoration in mean-field approaches, *Journal of Physics G: Nuclear and Particle Physics* **48**, 123001 (2021).
- [35] C. W. Bauer, Z. Davoudi, A. B. Balantekin, T. Bhattacharya, M. Carena, W. A. de Jong, P. Draper, A. El-Khadra, N. Gemelke, M. Hanada, G. K.-L. Kang, D. Rzeev, H. Lamm, Y.-Y. Li, J. Liu, M. Lukin, Y. Meurice, C. Monroe, B. Nachman, G. Pagano, J. Preskill, E. Rinaldi, A. Roggero, D. I. Santiago, M. J. Savage, I. Siddiqi, G. Siopsis, D. Van Zanten, N. Wiebe, Y. Yamauchi, K. Yeter-Aydeniz, and S. Zorzetti, Quantum simulation for high-energy physics, *PRX Quantum* **4**, 027001 (2023).
- [36] 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).
- [37] D. Neill, H. Liu, J. Martin, and A. Roggero, Scattering neutrinos, spin models, and permutations, *arXiv* (2024), 2406.18677 [hep-ph].
- [38] E. Manousakis, The spin-1/2 heisenberg antiferromagnet on a square lattice and its application to the cuprous oxides, *Rev. Mod. Phys.* **63**, 1 (1991).
- [39] W. Heisenberg, Zur Theorie des Ferromagnetismus, *Zeitschrift für Physik* **49**, 619 (1928).
- [40] I. Stetcu, A. Baroni, and J. Carlson, Projection algorithm for state preparation on quantum computers, *Phys. Rev. C* **108**, L031306 (2023).
- [41] E. Rule, I. Stetcu, and J. Carlson, Simplified projection on total spin zero for state preparation on quantum computers, *Phys. Rev. C* **110**, 064003 (2024).