

Scaling adaptive quantum simulation algorithms via operator pool tiling

John S. Van Dyke^{1,2,*}, Karunya Shirali^{1,2,*}, George S. Barron^{1,2}, Nicholas J. Mayhall^{2,3}, Edwin Barnes^{1,2} and Sophia E. Economou^{1,2}¹Department of Physics, Virginia Tech, Blacksburg, Virginia 24061, USA²Virginia Tech Center for Quantum Information Science and Engineering, Blacksburg, Virginia 24061, USA³Department of Chemistry, Virginia Tech, Blacksburg, Virginia 24061, USA

(Received 27 June 2022; revised 27 October 2023; accepted 14 December 2023; published 16 February 2024)

Adaptive variational quantum simulation algorithms use information from a quantum computer to dynamically create optimal trial wave functions for a given problem Hamiltonian. A key ingredient in these algorithms is a predefined operator pool from which trial wave functions are constructed. Finding suitable pools is critical for the efficiency of the algorithm as the problem size increases. Here, we present a technique called operator pool tiling that facilitates the construction of problem-tailored pools for arbitrarily large problem instances. By first performing an Adaptive Derivative-Assembled Problem-Tailored Ansatz Variational Quantum Eigensolver (ADAPT-VQE) calculation on a smaller instance of the problem using a large, but computationally inefficient, operator pool, we extract the most relevant operators and use them to design more efficient pools for larger instances. We demonstrate the method here on strongly correlated quantum spin models in one and two dimensions, finding that ADAPT automatically finds a highly effective ansatz for these systems. Given that many problems, such as those arising in condensed matter physics, have a naturally repeating lattice structure, we expect the pool tiling method to be a widely applicable technique apt for such systems.

DOI: [10.1103/PhysRevResearch.6.L012030](https://doi.org/10.1103/PhysRevResearch.6.L012030)

Variational quantum eigensolvers (VQEs) comprise an important and much studied class of algorithms for near-term quantum computers [1–4]. As examples of hybrid quantum-classical algorithms, they divide the computational task of optimizing an objective function $f(\theta)$ (for instance, the ground state energy of a physical system) between a quantum processor and a classical processor, in which the former efficiently queries the objective function while the latter determines updated guesses for the variational parameters θ . In a VQE, the objective function takes the form of an expectation value $f(\theta) = \langle H \rangle_\theta$ of an observable H in the state $|\psi(\theta)\rangle$. The latter is prepared by a parameterized quantum circuit and is subsequently measured to determine the relevant expectation value.

While the choice of H is typically fixed by the problem at hand, the success of the algorithm depends strongly on the form of the circuit that prepares $|\psi(\theta)\rangle$ (also known as the “ansatz”) and the optimization method used. Many different ansätze have been proposed, including hardware-efficient [5] and physics- or chemistry-inspired [1,6–8] ones. In general, a good ansatz ought to be expressive enough to closely approximate the solution for suitable values of θ , while maintaining a low circuit depth, for compatibility with the limited coherence times of noisy quantum devices.

An important development in this regard was the introduction of the Adaptive Derivative-Assembled Problem-Tailored Ansatz Variational Quantum Eigensolver (ADAPT-VQE) method [9], which put forward the idea of adaptive VQEs. This approach iteratively constructs an ansatz

$$|\psi_{n+1}(\theta^{(n+1)})\rangle = e^{-i\theta_{n+1}A_{n+1}}|\psi_n(\theta^{(n)})\rangle, \quad (1)$$

where $\theta^{(n)} = (\theta_1, \dots, \theta_n)$ are the variational parameters at step n in the algorithm. In adaptive VQEs, the operators A_n are selected from a fixed, predefined “operator pool” according to a selection criterion. In the ADAPT-VQE approach, the operator chosen at each step is the one that maximizes the gradient of the objective function, i.e., the A_k for which $|\partial \langle H \rangle / \partial \theta_{n+1}|_{\theta_{n+1}=0}$ is maximized, where H is the objective function operator. Intuitively, this criterion seeks to identify the operators A_k which are most effective at decreasing the objective function at a given step.

The ADAPT-VQE algorithm was first applied to finding the ground state energies of small molecules, where it obtained high accuracy with shallower circuit depths compared to other commonly used ansätze. The approach was subsequently generalized to a number of other applications, including optimization problems [10,11], real and imaginary time evolution [12,13], the preparation of excited states [14], and strongly correlated lattice models [15]. The adaptive philosophy towards ansatz construction has also found other expressions, as in evolutionary algorithms and related methods [16,17].

As one might expect, the success of ADAPT is crucially dependent on the choice of operator pool. The first application of the method used pools consisting of fermionic operators that correspond to single and double excitations of electrons. It was then discovered that qubit-based pools, consisting of

*These authors contributed equally to this work.

[†]Present address: Johns Hopkins University Applied Physics Laboratory, 11100 Johns Hopkins Road, Laurel, Maryland 20723, USA.

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

individual Pauli strings, can also reach the ground states of molecular systems but with considerably reduced counts of two-qubit gates [18]. Subsequent work [19] demonstrated that even more savings can be achieved by symmetrizing the pool of Ref. [18].

In general, the optimal choice of pool for a given problem is subject to a trade-off: On the one hand, pools with more operators are more likely to lead to convergence (and more quickly) due to the greater flexibility of operator selection at each step in the algorithm. On the other hand, since $|\partial \langle H \rangle / \partial \theta|$ is calculated for each operator in the pool at every step, this can lead to considerable measurement overhead when estimating the expectation values on a quantum computer. In seeking to reduce this, it was proven that “minimally complete” pools can be constructed that contain only $2L - 2$ operators (where L is the number of qubits) [18]. These pools have the property that all states in the Hilbert space can be obtained from an ansatz using only the operators of the given pool. However, due to the local nature of the ansatz updates, it is not guaranteed that the sequence of operators required can be found by the algorithm. One method for alleviating some of these issues is to explicitly encode information about symmetries of the problem into the choice of operator pool [20]. More recently, it was shown [21] that certain hardware-efficient pools containing up to $O(L^4)$ operators incur only an $O(L)$ measurement overhead on top of what is needed for standard, nonadaptive VQEs; this is the same measurement overhead scaling as for minimal complete pools, but without a reduction in operator selection flexibility.

The presence of symmetries in some systems leads to degeneracies in the gradient of the objective function among various operators in the pool. One strategy to address this involved adding multiple disjoint operators to the ansatz at each step, which was found to lead to significantly shallower circuits [22] and, consequently, a reduction in the overall measurement count as well.

Despite these advances, the space of possible choices for operator pools remains staggeringly large, growing exponentially in the number of qubits. At present, identifying ultimately successful pools is somewhat of an art, often based on a trial-and-error approach in which various candidates are investigated and compared, slowing the pace of discovery.

In this work, we present a general systematic procedure for constructing pools that are both small in size and also facilitate rapid convergence to the ground state. This is achieved by tiling successful pools from small problem instances to construct pools that are effective for larger problems. We then illustrate the method with the example of the spin-1/2 XXZ model, a prototypical correlated spin model. While we find that our approach leads to shorter or similar-depth state preparation ansätze compared to previous methods, our approach is a more general technique to systematically solve lattice problems rather than relying on *ad hoc* solutions. Finally, we discuss other variations of the underlying idea and their connection to many-body expansions used in quantum chemistry.

ADAPT instances that use larger pools are expected to converge more often and with fewer iterations, owing to the greater number of states that can be accessed at each step. Among qubit-based pools, the largest possible one is the “full Pauli pool,” consisting of all Pauli strings on L qubits, excluding the identity. As this pool includes all other qubit pools as

proper subsets, it performs at least as well as each of them, in the sense of making the locally optimal choice. However, the exponential scaling of the size, $4^L - 1$, clearly prevents its application beyond small problem instances.

Many problems, such as those arising in condensed matter physics, have a simple repeated (lattice) structure that relates instances defined for different numbers of qubits, up to finite-size effects. Exploiting the lattice structure of these systems on near-term quantum devices has attracted increasing attention [23–25]. Example problems include spin models (as used in the field of quantum magnetism), as well as fermionic ones like the Hubbard model, well known from the study of high-temperature superconductivity. The lattice structure of these models suggests the following approach to designing effective pools: (1) Start with a small instance of the particular model or problem of interest on L_1 qubits and run ADAPT-VQE with a large, highly expressive pool that can quickly converge to the desired solution. To address the degenerate gradients that typically arise when using a sizable operator pool, run a large number of trials, growing the ansatz by picking randomly from the degenerate operators, thus sampling a variety of paths to convergence. (2) Take the set of operators chosen by ADAPT over all trials as the basis for a new pool defined on a larger instance of the same problem on $L_2 > L_1$ qubits. This requires taking the tensor product of each operator with the single-qubit identity $L_2 - L_1$ times, such that operators in the new pool have the correct dimension. Since the products can be applied either on the left or the right, a single operator of the original pool yields $L_2 - L_1 + 1$ operators for the new one, corresponding to “tiling” the operator across the system in the one-dimensional case (with straightforward generalizations to higher dimensions). Importantly, while the original pool size can be exponential in the number of qubits L_1 (as with the full Pauli pool), the number of operators in the new pool scales linearly with the system size L_2 .

We now demonstrate the viability of this strategy through numerical simulations on a concrete example. We consider both one-dimensional and two-dimensional spin-1/2 XXZ models. The first case consists of a chain of L sites with open boundary conditions, whose Hamiltonian is given by

$$H_{\text{XXZ}}^{(1D)} = \sum_{i=1}^{L-1} J_{xy}(X_i X_{i+1} + Y_i Y_{i+1}) + J_z Z_i Z_{i+1}, \quad (2)$$

where X_i , Y_i , and Z_i are Pauli operators acting on site i . We set $J_{xy} = 1$ throughout this work. For the first part of the tiling method, we solve the $L = 3$ case of the XXZ chain using the full Pauli pool. ADAPT always converges in three steps for this simple problem, and running trials to accommodate degenerate gradients as described above reveals that it chooses a set of operators that preserve time-reversal symmetry and simultaneously commute with the operator $Z_1 Z_2 Z_3$. The complete list of chosen operators is available in Appendix B. Convergence of the algorithm is defined to occur when the maximum gradient for a given step drops below a threshold ε [9]. We then tile these operators to create pools for systems with $L > 3$ and run ADAPT. Figure 1(a) shows the relative error in the ground state energy compared to the exact solution for this model for system sizes $4 \leq L \leq 16$ and several values of the interaction strength J_z . In each case, the error is less




FIG. 1. (a) Relative error in the ground state energy versus system size L for the XXZ chain. E is the energy found by ADAPT after convergence, while E_0 is the exact ground state energy. (b) Number of ADAPT steps required to converge the algorithm versus system size L . The initial state is the Néel state $|\uparrow\downarrow\uparrow\downarrow\cdots\rangle$, and we set $\varepsilon = 0.01$. Variational parameters are optimized using the LBFGS algorithm, with $\theta_{n+1} = 0$ and $\{\theta_j\}_{j=1,\dots,n}$ set to the optimal values from the previous ADAPT iteration.

than 0.06%, indicating good convergence across a range of J_z that includes both the paramagnetic and ordered regimes.

In Fig. 1(b) we present the number of ADAPT iterations required to find the ground state as a function of L . We note that this value is also equal to the number of optimization parameters in the final ansatz. The behavior over the simulated range appears to be close to polynomial. To our knowledge, previous works that applied VQE to the XXZ model focused on only the case of closed boundary conditions. Reference [26] used a fixed-circuit ansatz to study the $J_z = 0.5$ XXZ model for $L = 12$ sites, finding that 168 parameters were required to obtain an energy accuracy of $\log_{10}(1/|E - E_0|) \approx 1.0$, whereas in our calculations with the tiled pool, ADAPT achieves a higher accuracy of $\log_{10}(1/|E - E_0|) \approx 1.9$, with an ansatz containing 132 variational parameters. On the other hand, for $J_z = 1$ we find that ADAPT converges with 122 parameters, which is comparable to what Ref. [27] found using the Hamiltonian variational ansatz. These examples suggest that ADAPT with tiling can perform at least as well as previous approaches, if not better. Further evidence of this comes from Ref. [28], which examined various VQE ansätze for the XXZ model and found that the best-performing ansatz, which they termed the XY-ansatz, is constructed from operators with support on two or three qubits, similar to those in the tiled pool discovered by ADAPT.

We now turn to the details of how ADAPT-VQE approaches convergence in the present example. The energy and the absolute value of the maximum gradient at each iteration of the algorithm are shown in Figs. 2(a) and 2(b), respectively. Three distinct regimes are apparent in both quantities. The energy decreases sharply during the first $L/2$ iterations, while the maximum gradient remains at its initial value for $L/2$




FIG. 2. (a) Energy and (b) absolute value of the maximum gradient as a function of ADAPT iteration number for the XXZ chain. Simulation parameters are the same as for Fig. 1 with $J_z = 1$.

iterations. This is followed by a second region in which the energy decreases roughly linearly with iteration number, while the gradient shows a plateau with oscillations between several values. The final regime shows a slower decay in both the energy and maximum gradient. One can obtain some physical insight into this behavior by inspecting the operators chosen by ADAPT and the corresponding optimal parameter values. The first $L/2$ operators in the ansatz are of the form $e^{-i\theta_k X_{2k-1} Y_{2k}}$, with $k = 1, \dots, L/2$ and $\theta_k \approx -\pi/4$. Acting on the initial Néel state, this sequence produces a superposition of basis states in which pairs of up and down spins have been swapped, with a nontrivial relative sign structure between them. The other two regimes in the ADAPT convergence process involve the three-local operators and possess different optimal parameter values which do not form an easily discernible pattern. Nevertheless, it is clear that the operations performed during the first and second regimes are more effective for reducing the energy overall, while the final regime involves finer tuning of the wave function.

Importantly, the operators chosen by ADAPT for the *small* problem instance span the symmetry subsector that respects time-reversal symmetry, conserves the expectation value of $Z_1 Z_2 \cdots Z_L$, and contains the reference state and the target state, implying that the tiled operators constructed from these operators span the relevant symmetry sector for the *large* problem instance. This in turn demonstrates that the tiled pool is sufficiently expressible to represent the target state. More detailed proofs are available in Appendixes A and B.

For the two-dimensional XXZ model we study the case of open boundary conditions and nearest-neighbor interactions on rectangular lattices of $N = L_x L_y$ sites:

$$H_{\text{XXZ}}^{(2D)} = \sum_{\langle i,j \rangle} J_{xy} (X_i X_j + Y_i Y_j) + J_z Z_i Z_j, \quad (3)$$




FIG. 3. Illustration of the tiling procedure for the two-dimensional XXZ model.

where $\langle i, j \rangle$ denotes nearest neighbors. In generating the ADAPT operator pool, it is natural to consider tiles consisting of 2×2 blocks, given the two-dimensional nature of the problem. Inspired by the successful strategy of collecting operators from a number of ADAPT runs on the tile in the one-dimensional case, we run several trials, randomly picking operators in the case of degenerate gradients. In this case, ADAPT always converges in either five or six steps, choosing a set of operators that simultaneously satisfy time-reversal symmetry and commute with $Z_1 Z_2 Z_3 Z_4$ and, importantly, form a complete basis for the symmetry subsector conserving the two. These blocks are then tiled across the ladder (in an overlapping fashion) to create the operator pool. The procedure is illustrated in Fig. 3.

In Fig. 4 we present the energy error [Fig. 4(a)] and number of steps to convergence [Fig. 4(b)] for this case. While the ground state is accurately approximated here as well, the computational difficulty grows rapidly with increasing L_x and L_y . Although a quantitative extrapolation to large systems would not be warranted, it is clear that the required number of ADAPT steps grows more quickly with system size than in the one-dimensional (1D) case. Notably, this nonlinear growth remains when the number of steps is plotted as a function of $L_x \times L_y$, suggesting that the computational difficulty is at least partially due to the intrinsic connectivity of the two-dimensional (2D) lattice and not just the total number of sites. In any case, it appears that the 2D XXZ model presents a significant challenge for VQE methods in the quantum advantage regime due to the large number of optimization steps needed. Nevertheless, the rapid progress in variational quantum algorithms lends hope that such difficulties may be surmounted.

Having demonstrated the power of the “tiling” approach to designing operator pools for ADAPT, we turn to possible generalizations of the method. While the above examples focused on 1D and 2D systems, it is clear how to generalize to higher dimensions (although it is likely to be computationally challenging). Generalizations to other popular models in condensed matter physics, such as the Hubbard model, are also straightforward. While such constructions are quite natural for lattice systems with short-range interactions, it is less clear that a simple tiling procedure will work for more complex geometries and/or long-range interactions (as found in quantum chemistry applications, for instance). In this case, one can consider the more general approach of starting from the full problem, selecting various subsets of the degrees of




FIG. 4. (a) Relative error in the ground state energy and (b) number of ADAPT steps required for convergence for the 2D XXZ model with a square geometry. The simulation parameter is $J_z = 1$.

freedom, and solving these reduced problems using ADAPT with a large operator pool. The resulting operators chosen from each subproblem can then be collected to generate a pool for executing ADAPT on the original problem. This method bears a close resemblance to many-body expansion techniques used in quantum chemistry [29,30]. We leave further development of these ideas for future work.

To conclude, we have proposed and demonstrated a systematic operator tiling method for designing small but accurate operator pools for use with the ADAPT-VQE algorithm. The strength of this approach lies in its versatility in producing efficient representations of the ground state for lattice models in general and is a significant advantage for using ADAPT in noisy quantum devices. We have illustrated the approach using the open XXZ spin chain, and notably, we found that the tiling method automatically found a highly effective ansatz for this system [28]. Confirming the usefulness of the operator tiling method for other models in strongly correlated physics is an interesting direction for future work.

Acknowledgments. Numerical calculations of the exact ground state energies were performed with the QUSPIN PYTHON library [31]. This work was supported by the Department of Energy. E.B. and N.J.M. acknowledge Award No. DE-SC0019199, and S.E.E. acknowledges the DOE Office of Science, National Quantum Information Science Research Centers, Co-design Center for Quantum Advantage (C2QA), Contract No. DE-SC0012704.

Appendix A: Outline of the proof of completeness of the tiled pool. Here, we outline how to prove that for a translation-symmetric Hamiltonian on a lattice, given reference states in the appropriate symmetry sectors, if we start from a complete operator pool (i.e., the operators span the entire Lie algebra in the relevant symmetry sector) on a small problem instance (the “tile”) on L_1 qubits, then the tiled pool for a larger problem instance on $L_2 > L_1$ qubits, generated using the operators chosen by ADAPT-VQE [9] for the tile, will also be complete. Completeness of the pool in turn implies that the pool has sufficient expressibility to represent any state in the relevant symmetry sector of the L_2 -qubit system [18,20], including the target state.

We begin by summarizing our procedure for constructing the pool on L_1 qubits as described in the main text. First, consider a Hamiltonian H for a small problem instance on L_1 qubits along with a highly expressive operator pool of Pauli string-type operators and a given reference state. At each step, ADAPT adds to the ansatz the operator with the highest gradient among the pool operators, growing the ansatz until convergence is achieved. To account for the possibility of multiple operators having identical gradients, several ADAPT-VQE calculations are run in which the algorithm randomly picks operators from those with degenerate gradients to add to the ansatz at each step. Thus, various routes to convergence are sampled. The chosen operators obtained from these runs are collected and stored. Note that if the system has a conserved quantity S that commutes with the Hamiltonian and if we start from a reference state $|\psi_{\text{init}}\rangle$ which lies in the same symmetry subsector as the target state, then ADAPT will pick only operators that respect this symmetry [20].

For all the examples considered in the main text, we find numerically that ADAPT chooses a set of operators that generates the *entire* Lie algebra of the symmetry subsector; that is, these operators constitute a complete pool for the small problem on L_1 qubits. For other lattice problems, one can, of course, easily check whether the set of operators collected in this way forms a complete pool on L_1 qubits. Assuming this is the case, we will now outline how to prove that the set of tiled pool operators on L_2 qubits, which are obtained by padding the original set with $L_2 - L_1$ identity operators in all possible ways without breaking the tile, necessarily forms a complete pool on L_2 qubits.

To start, consider the concrete example of a linear qubit array with $L_1 = 3$ and $L_2 = 4$. There are then two types of tiled pool operators that we construct for the four-qubit system by adding identities to the three-qubit pool obtained from ADAPT: $I \otimes \sigma_a \otimes \sigma_b \otimes \sigma_c$ and $\sigma_d \otimes \sigma_e \otimes \sigma_f \otimes I$, where $a, b, c, d, e, f \in \{0, 1, 2, 3\}$, with $\sigma_0 = I$, $\sigma_1 = X$, $\sigma_2 = Y$, and $\sigma_3 = Z$. Here, we assume that the sets of operators $\sigma_a \otimes \sigma_b \otimes \sigma_c$ and $\sigma_d \otimes \sigma_e \otimes \sigma_f$ generate the full Lie subalgebra of Pauli strings on three qubits that commute with S . If the Hamiltonian also possesses time-reversal symmetry (i.e., the Hamiltonian is real), then the subalgebra will be further restricted to contain only odd Pauli strings (operators that contain an odd number of Y 's) that commute with S since in this case the eigenstate wave functions can always be chosen to be real. The question is then whether we can generate any four-qubit (odd) Pauli string that commutes with S from commutators of tiled operators like $I \otimes \sigma_a \otimes \sigma_b \otimes \sigma_c$

and $\sigma_d \otimes \sigma_e \otimes \sigma_f \otimes I$. (We assume that S can be readily extended to the larger four-qubit system in accordance with the translation symmetry of the Hamiltonian.) In order to prove that this is indeed the case, we need to first specify a particular Hamiltonian H and symmetry S . We prove this for the XXZ models considered in this work in Appendix B. We then use this completeness result for two overlapping tiles (which is the $L_2 = 4$ case in our linear $L_1 = 3$ example) as the base case for an inductive proof that extends the same result to any L_2 .

Let us illustrate how the inductive proof works by continuing on to the case of $L_2 = 5$ for our linear $L_1 = 3$ example. Assuming that it is possible to generate any symmetry-preserving four-qubit Pauli string from the $L_1 = 3$ tiled pool, the next question is whether we can generate any five-qubit symmetry-preserving Pauli string in a similar fashion, where now $L_2 = 5$. We could try to form five-qubit strings by taking commutators of operators like $I \otimes I \otimes \sigma_a \otimes \sigma_b \otimes \sigma_c$ and $\sigma_d \otimes \sigma_e \otimes \sigma_f \otimes \sigma_g \otimes I$. If the $L_2 = 4$ case holds, then it immediately follows that five-qubit strings for which the four-qubit substring on qubits 2 through 5 is symmetry preserving can be generated via such commutators. However, if the four-qubit substring is not symmetry preserving, then further analysis is needed. In particular, we would need to show that any symmetry-violating four-qubit Pauli string can be produced from nested commutators of the form

$$[\cdots [[V, P_1], P_2] \cdots], \quad (\text{A1})$$

where $V = \sigma_e \otimes \sigma_f \otimes \sigma_g \otimes I$ violates either time-reversal symmetry or S symmetry or both, while each of the $P_i = \sigma_{a_i} \otimes \sigma_{b_i} \otimes \sigma_{c_i} \otimes \sigma_{d_i}$ preserves both symmetries, where at least one of σ_{a_i} and σ_{d_i} is the identity operator. Note that the nested commutator above has the same parity as V and it commutes with S if and only if V commutes with S . Thus, showing that any symmetry-preserving five-qubit Pauli string can be generated from commutators of tiled pool operators is tantamount to showing that, given a V of a certain parity that commutes or anticommutes with S , the above nested commutators span the entire set of operators that share the same parity and commutativity relation with S .

The above argument can, in fact, be extended to symmetry-preserving Pauli strings on any number of qubits L_2 via induction. If we can generate any symmetry-preserving string on L_2 qubits, then we can obtain any symmetry-preserving string on $L_2 + 1$ qubits by taking nested commutators of an L_2 -qubit string with tiled pool operators that have support only on the last four qubits, $L_2 - 2, \dots, L_2 + 1$. Thus, in order to prove that the tiled pool is complete, it suffices to show that any four-qubit Pauli string can be generated from nested commutators as in Eq. (A1) for some choice of V and the P_i .

Although here we focused on a linear array of qubits with $L_1 = 3$, this argument works regardless of the dimension or connectivity of the lattice and regardless of the size L_1 of the tiles, as long as the requisite conditions on overlapping tiles are satisfied. More generally, one would need to show that any Pauli string with support on the union of two overlapping tiles can be generated by taking nested commutators of some Pauli string V that has support on one tile with tiled pool operators P_i that each have support on one of the two tiles. Since this proof does not require assumptions about the form of the Hamiltonian, we expect it to be applicable to a wide

variety of models. In Appendix B, we show explicitly that the above conditions hold for the tiled pool for the XXZ model and that this pool is thus complete for any number of qubits L_2 .

Appendix B: Proof of tiled pool completeness for the XXZ model. For the spin models considered in the main text, the operator $P_Z = Z_1 Z_2 \cdots Z_n$ (where n is the length of the spin system) commutes with the Hamiltonian. In addition to P_Z , the Hamiltonian has time-reversal symmetry \mathcal{T} .

Upon running ADAPT calculations on this system using the Néel antiferromagnetic state as a reference state and the full Pauli pool as the operator pool, we find that ADAPT chooses a set of operators that generate the Lie algebra for the subspace that simultaneously preserves the expectation value of P_Z and respects \mathcal{T} . For a one-dimensional tile of size $L_1 = 3$ with the Néel antiferromagnetic state as the reference state, ADAPT chooses the operators $\{IXY, IYX, XYI, YXI, ZXY, YXZ, ZYX, XYZ\}$ for $0.0 \leq J_z < 1.0$ and $\{IXY, IYX, XYI, XIY, YXI, YIX, ZXY, YZX, YXZ, XZY, ZYX, XYZ\}$ for $J_z \geq 1.0$. Note that both these sets are complete in the relevant symmetry sector. For two-dimensional systems, we use a tile of size 2×2 and find that, for $J_z = 1.0$, ADAPT chooses the following set of operators: $IIXY, IYYX, IXYZ, IXZY, IYXZ, IYZX, IZXY, IZYX, XIYZ, XIZY, XYII, XYIZ, XYZI, XZZZ, XZII, XZII, YIXZ, YIZX, YXII, YXIZ, YXZI, YXZZ, YZIX, YZII, ZIIX, ZIYX, ZXYI, ZYIX, ZYXI, ZZZY, and ZZYX$. These operations, in one and two dimensions, respectively, generate through their Lie algebras all Pauli string operators for the respective system sizes that preserve P_Z and \mathcal{T} .

We next explicitly demonstrate the ability of the tiled operator pool to generate symmetry-preserving Pauli strings of any size via commutators of the tiled operators, focusing on the one-dimensional system. We do this by following the general strategy outlined in Appendix A. That is, we first need to show that we can produce any symmetry-preserving four-qubit Pauli string from commutators of tiled operators. We then need to show that we can transform a given symmetry-violating four-qubit substring V into any other symmetry-violating substring via such commutators as in Eq. (A1), as this ensures that we can generate symmetric Pauli strings on more than four qubits. For example, we could use $[XZYYI, IIXYI] \propto XZZYI$ as our initial Pauli string on $L_2 = 5$ qubits, in which case the four-qubit substring $V = ZZYI$ violates P_Z symmetry but not time-reversal symmetry. We then need to show that this V can be transformed into any other four-qubit Pauli string that violates P_Z but preserves time reversal using commutators with tiled operators [Eq. (A1)] so that we obtain all symmetry-preserving five-qubit Pauli strings that start with X on the first qubit.

Consider the commutator of a pair of four-qubit tiled operators constructed from three-qubit operator tiles with overlapping support,

$$[I \otimes \sigma_a \otimes \sigma_b \otimes \sigma_c, \sigma_d \otimes \sigma_e \otimes \sigma_f \otimes I]. \quad (\text{B1})$$

Here, $a, b, c, d, e, f \in \{0, 1, 2, 3\}$ such that $\sigma_a \otimes \sigma_b \otimes \sigma_c$ and $\sigma_d \otimes \sigma_e \otimes \sigma_f$ are odd Pauli strings (containing one or three Y 's) and commuting with $Z \otimes Z \otimes Z$. Since the only possible three-qubit Pauli string containing three Y 's fails to commute with $Z \otimes Z \otimes Z$, the strings necessarily contain a single Y .

Thus, the indices $\{a, b, c\}$ and $\{d, e, f\}$ are permutations of either (0,1,2) or (1,2,3). Next, if the two operators fail to commute, their commutator is proportional to their product, $\sigma_d \otimes (\sigma_a \sigma_e) \otimes (\sigma_b \sigma_f) \otimes \sigma_c$. Furthermore, their product commutes with $Z \otimes Z \otimes Z \otimes Z$ since they individually commute with it. Letting $\{a, b, c\}$ and $\{d, e, f\}$ in Eq. (B1) take on all possible combinations of the values discussed above, as well as each of the operators in the commutator, and including the higher-order commutators, it is straightforward to show that one obtains exactly the set of four-qubit Pauli strings that commute with $Z \otimes Z \otimes Z \otimes Z$ and contain an odd number of Y 's. Thus, we have shown that the tiled operator pool can represent any state in the relevant symmetry sector for a system of $L_2 = 4$ qubits.

Next, consider the commutator of a pair of five-qubit tiled operators—one constructed from a three-qubit operator tile and the other from a four-qubit operator obtained in the previous step. This can be represented as

$$[I \otimes I \otimes \sigma_a \otimes \sigma_b \otimes \sigma_c, \sigma_d \otimes \sigma_e \otimes \sigma_f \otimes \sigma_g \otimes I]. \quad (\text{B2})$$

If the commutator is nonzero, then it is proportional to the product of the two operators. Now, $\{a, b, c\}$ must be a permutation of either (0,1,2) or (1,2,3) as in the earlier step. We will consider three separate scenarios for $\sigma_d \otimes \sigma_e \otimes \sigma_f \otimes \sigma_g \otimes I$, ensuring that it is overall odd and commutes with $Z \otimes Z \otimes Z \otimes Z \otimes Z$: (1) First, let $d = 0$ or 3. In this case, $\sigma_e \otimes \sigma_f \otimes \sigma_g$ must be an odd Pauli string commuting with $Z \otimes Z \otimes Z$; that is, it reduces to the four-qubit case, and $\{e, f, g\}$ can take on the same values as $\{a, b, c\}$. (2) Next, suppose $d = 1$; now, $\sigma_e \otimes \sigma_f \otimes \sigma_g$ must be an odd Pauli string *anticommuting* with $Z \otimes Z \otimes Z$. This restricts $\{e, f, g\}$ to permutations of (2,0,0), (2,1,1), (2,3,3), (2,0,3), and (2,2,2). (3) Last, let $d = 2$: Here, $\sigma_e \otimes \sigma_f \otimes \sigma_g$ must be an even Pauli string *anticommuting* with $Z \otimes Z \otimes Z$. $\{e, f, g\}$ must then take on permutations of the values (0,0,1), (1,3,3), (0,1,3), (1,1,1), and (1,2,2). Letting $\{a, b, c\}$ and $\{d, e, f, g\}$ take on the appropriate values as detailed in each scenario and including higher-order commutators, it is straightforward to show that one obtains all five-qubit operators conserving P_Z and \mathcal{T} .

Finally, consider the commutator of a pair of six-qubit tiled operators,

$$[I \otimes I \otimes I \otimes \sigma_a \otimes \sigma_b \otimes \sigma_c, \sigma_d \otimes \sigma_e \otimes \sigma_f \otimes \sigma_g \otimes \sigma_h \otimes I]. \quad (\text{B3})$$

We previously showed that we obtain all five-qubit operators in the relevant symmetry sector for $\sigma_e \otimes \sigma_f \otimes \sigma_g$ in Eq. (B2) (1) odd and commuting with $Z \otimes Z \otimes Z$, (2) odd and anti-commuting with $Z \otimes Z \otimes Z$, and (3) even and anticommuting with $Z \otimes Z \otimes Z$. In Eq. (B3), we can use the same argument to show that for all choices of $\{d, e\}$ such that $\sigma_f \otimes \sigma_g \otimes \sigma_h$ falls into one of these three cases, the result will be a six-qubit operator in the appropriate symmetry sector. This leaves us with one case—that of $\sigma_f \otimes \sigma_g \otimes \sigma_h$ in Eq. (B3) being even and commuting with $Z \otimes Z \otimes Z$, which restricts $\{f, g, h\}$ to the values (1,1,0), (1,1,3), (0,0,0), (0,0,3), (0,3,3), (3,3,3), (0,2,2), and (2,2,3). Allowing $\{a, b, c\}$ and $\{d, e, f, g, h\}$ to take on the corresponding values and including higher-order commutators, it is again straightforward to show that one obtains all six-qubit operators conserving P_Z and \mathcal{T} .

By induction, it follows that symmetry-preserving Pauli strings of any size can therefore be obtained from commutators of the tiled operators. The resultant operator pool spans the Lie algebra for the Hilbert space sector in which P_Z and

\mathcal{T} are preserved and in which the target state lies. Thus, we have shown how to construct an efficient operator pool for ADAPT-VQE calculations on translation-symmetric systems that is sufficiently expressive to reach the target state.

-
- [1] 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, *Nat. Commun.* **5**, 4213 (2014).
 - [2] 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, *Nat. Rev. Phys.* **3**, 625 (2021).
 - [3] K. Bharti, A. Cervera-Lierta, T. H. Kyaw, T. Haug, S. Alperin-Lea, A. Anand, M. Degroote, H. Heimonen, J. S. Kottmann, T. Menke, W.-K. Mok, S. Sim, L.-C. Kwek, and A. Aspuru-Guzik, Noisy intermediate-scale quantum algorithms, *Rev. Mod. Phys.* **94**, 015004 (2022).
 - [4] 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, *Phys. Rep.* **986**, 1 (2022).
 - [5] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets, *Nature (London)* **549**, 242 (2017).
 - [6] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, The theory of variational hybrid quantum-classical algorithms, *New J. Phys.* **18**, 023023 (2016).
 - [7] D. Wecker, M. B. Hastings, and M. Troyer, Progress towards practical quantum variational algorithms, *Phys. Rev. A* **92**, 042303 (2015).
 - [8] H. R. Grimsley, D. Claudino, S. E. Economou, E. Barnes, and N. J. Mayhall, Is the trotterized UCCSD ansatz chemically well-defined? *J. Chem. Theory Comput.* **16**, 1 (2020).
 - [9] H. R. Grimsley, S. E. Economou, E. Barnes, and N. J. Mayhall, An adaptive variational algorithm for exact molecular simulations on a quantum computer, *Nat. Commun.* **10**, 3007 (2019).
 - [10] L. Zhu, H. L. Tang, G. S. Barron, F. A. Calderon-Vargas, N. J. Mayhall, E. Barnes, and S. E. Economou, Adaptive quantum approximate optimization algorithm for solving combinatorial problems on a quantum computer, *Phys. Rev. Res.* **4**, 033029 (2022).
 - [11] Y. Chen, L. Zhu, C. Liu, N. J. Mayhall, E. Barnes, and S. E. Economou, How much entanglement do quantum optimization algorithms require? [arXiv:2205.12283](https://arxiv.org/abs/2205.12283).
 - [12] Y.-X. Yao, N. Gomes, F. Zhang, C.-Z. Wang, K.-M. Ho, T. Iadecola, and P. P. Orth, Adaptive variational quantum dynamics simulations, *PRX Quantum* **2**, 030307 (2021).
 - [13] N. Gomes, A. Mukherjee, F. Zhang, T. Iadecola, C.-Z. Wang, K.-M. Ho, P. P. Orth, and Y.-X. Yao, Adaptive variational quantum imaginary time evolution approach for ground state preparation, *Adv. Quantum Technol.* **4**, 2100114 (2021).
 - [14] F. Zhang, N. Gomes, Y. Yao, P. P. Orth, and T. Iadecola, Adaptive variational quantum eigensolvers for highly excited states, *Phys. Rev. B* **104**, 075159 (2021).
 - [15] G. Gyawali and M. J. Lawler, Adaptive variational preparation of the Fermi-Hubbard eigenstates, *Phys. Rev. A* **105**, 012413 (2022).
 - [16] A. G. Rattew, S. Hu, M. Pistoia, R. Chen, and S. Wood, A domain-agnostic, noise-resistant, hardware-efficient evolutionary variational quantum eigensolver, [arXiv:1910.09694](https://arxiv.org/abs/1910.09694).
 - [17] D. Chivilikhin, A. Samarin, V. Ulyantsev, I. Iorsh, A. R. Oganov, and O. Kyriienko, MoG-VQE: Multiobjective genetic variational quantum eigensolver, [arXiv:2007.04424](https://arxiv.org/abs/2007.04424).
 - [18] H. L. Tang, V. O. Shkolnikov, G. S. Barron, H. R. Grimsley, N. J. Mayhall, E. Barnes, and S. E. Economou, Qubit-ADAPT-VQE: An adaptive algorithm for constructing hardware-efficient ansätze on a quantum processor, *PRX Quantum* **2**, 020310 (2021).
 - [19] Y. S. Yordanov, V. Armaos, C. H. W. Barnes, and D. R. M. Arvidsson-Shukur, Qubit-excitation-based adaptive variational quantum eigensolver, *Commun. Phys.* **4**, 228 (2021).
 - [20] V. O. Shkolnikov, N. J. Mayhall, S. E. Economou, and E. Barnes, Avoiding symmetry roadblocks and minimizing the measurement overhead of adaptive variational quantum eigensolvers, *Quantum* **7**, 1040 (2023).
 - [21] P. G. Anastasiou, N. J. Mayhall, E. Barnes, and S. E. Economou, How to really measure operator gradients in ADAPT-VQE, [arXiv:2306.03227](https://arxiv.org/abs/2306.03227).
 - [22] P. G. Anastasiou, Y. Chen, N. J. Mayhall, E. Barnes, and S. E. Economou, Tetris-ADAPT-VQE: An adaptive algorithm that yields shallower, denser circuit ansätze, [arXiv:2209.10562](https://arxiv.org/abs/2209.10562).
 - [23] J.-G. Liu, Y.-H. Zhang, Y. Wan, and L. Wang, Variational quantum eigensolver with fewer qubits, *Phys. Rev. Res.* **1**, 023025 (2019).
 - [24] A. Eddins, M. Motta, T. P. Gujarati, S. Bravyi, A. Mezzacapo, C. Hadfield, and S. Sheldon, Doubling the size of quantum simulators by entanglement forging, *PRX Quantum* **3**, 010309 (2022).
 - [25] K. Fujii, K. Mizuta, H. Ueda, K. Mitarai, W. Mizukami, and Y. O. Nakagawa, Deep variational quantum eigensolver: A divide-and-conquer method for solving a larger problem with smaller size quantum computers, *PRX Quantum* **3**, 010346 (2022).
 - [26] C. Bravo-Prieto, J. Lumbrales-Zarapico, L. Tagliacozzo, and J. I. Latorre, Scaling of variational quantum circuit depth for condensed matter systems, *Quantum* **4**, 272 (2020).
 - [27] R. Wiersema, C. Zhou, Y. de Sereville, J. F. Carrasquilla, Y. B. Kim, and H. Yuen, Exploring entanglement and optimization within the hamiltonian variational ansatz, *PRX Quantum* **1**, 020319 (2020).

- [28] M. S. Jattana, F. Jin, H. De Raedt, and K. Michielsen, Assessment of the variational quantum eigensolver: Application to the Heisenberg model, *Front. Phys.* **10**, 907160 (2022).
- [29] C. Gros and R. Valentí, Cluster expansion for the self-energy: A simple many-body method for interpreting the photoemission spectra of correlated Fermi systems, *Phys. Rev. B* **48**, 418 (1993).
- [30] D. Sénéchal, D. Perez, and M. Pioro-Ladrière, Spectral weight of the Hubbard model through cluster perturbation theory, *Phys. Rev. Lett.* **84**, 522 (2000).
- [31] P. Weinberg and M. Bukov, QuSpin: A Python package for dynamics and exact diagonalisation of quantum many body systems part I: Spin chains, *SciPost Phys.* **2**, 003 (2017).