



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Variational Quantum Eigensolvers for Sparse Hamiltonians

William M. Kirby and Peter J. Love

Phys. Rev. Lett. **127**, 110503 — Published 9 September 2021

DOI: [10.1103/PhysRevLett.127.110503](https://doi.org/10.1103/PhysRevLett.127.110503)

Variational quantum eigensolvers for sparse Hamiltonians

William M. Kirby¹ and Peter J. Love^{1,2}

¹*Department of Physics and Astronomy, Tufts University, Medford, MA 02155*

²*Computational Science Initiative, Brookhaven National Laboratory, Upton, NY 11973*

Hybrid quantum-classical variational algorithms such as the variational quantum eigensolver (VQE) and the quantum approximate optimization algorithm (QAOA) are promising applications for noisy, intermediate-scale quantum (NISQ) computers. Both VQE and QAOA variationally extremize the expectation value of a Hamiltonian. All work to date on VQE and QAOA has been limited to Pauli representations of Hamiltonians. However, many cases exist in which a sparse representation of the Hamiltonian is known but there is no efficient Pauli representation. We extend VQE to general sparse Hamiltonians. We provide a decomposition of a fermionic second quantized Hamiltonian into a number of one-sparse, self-inverse, Hermitian terms linear in the number of ladder operator monomials in the second quantized representation. We provide a decomposition of a general d -sparse Hamiltonian into $O(d^2)$ such terms. In both cases a single sample of any term can be obtained using two ansatz state preparations and at most six oracle queries. The number of samples required to estimate the expectation value to precision ϵ scales as ϵ^{-2} as for Pauli-based VQE. This widens the domain of applicability of VQE to systems whose Hamiltonian and other observables are most efficiently described in terms of sparse matrices.

I. INTRODUCTION

The leading applications for noisy, intermediate-scale quantum (NISQ) computers are the variational quantum eigensolver (VQE) [1] and the quantum approximate optimization algorithm (QAOA) [2]. VQE estimates the ground state energy (and other properties) of a Hamiltonian by optimizing an ansatz for an energy eigenstate [1, 3–15]. QAOA approximately optimizes a classical objective function using a parameterized quantum state [2]. Methods such as VQE and phase estimation, which compute energy eigenstates, rely on efficient representations of the Hamiltonian. Two such representations are widely used: local Hamiltonians and sparse Hamiltonians. For simplicity we will restrict our discussion henceforth to systems of qubits; generalizations to tensor factors of arbitrary dimension are straightforward.

Firstly, we describe the efficient representation of Hamiltonians based on *locality*. For this representation, a convenient basis for qubit operators is given by the Pauli operators, \mathcal{P}_i , which are tensor products of Pauli matrices and the identity. The Hamiltonian is written:

$$H = \sum_i^m \alpha_i \mathcal{P}_i, \quad (1)$$

where the α_i are real coefficients. This representation of a Hamiltonian is efficient if the number of terms m grows only polynomially with the number of qubits. The locality k of the Hamiltonian (1) is the maximum locality of any term \mathcal{P}_i , which refers to the number of non-identity tensor factors in \mathcal{P}_i . Note that locality here does not necessarily refer to geometric locality. For qubits the Hamiltonian is a sum of one-qubit terms, two-qubit terms and so on. This representation was first used for quantum simulation by Lloyd [16], and all VQE algorithms to date make use of this Hamiltonian representation.

A second efficient representation of a Hamiltonian is

based on *sparsity*, which refers to the maximum number of nonzero entries in any row or column of the Hamiltonian. For example, in a Hamiltonian of the form (1) the number of nonzero entries in any row or column in the computational basis is bounded above by d , so we refer to the Hamiltonian as d -sparse. This follows because the Pauli operators are one-sparse: they each have only one nonzero entry in each row and column. However, not every sparse Hamiltonian is local, and many sparse Hamiltonians do not admit a Pauli decomposition (1) with a polynomial number of terms. A simple example is the number operator for a bosonic mode encoded as a binary number in qubit computational basis states, which is one-sparse but has an exponential number of Pauli terms.

Quantum simulation of sparse Hamiltonians has undergone extensive study [17–27], culminating in algorithms with optimal (for time-independent Hamiltonians) [25, 26] or near-optimal (for time-dependent Hamiltonians) [27] scaling with all parameters. In these algorithms a d -sparse Hamiltonian is accessed via a pair of oracle unitaries O_F and O_H . O_F returns the location of the i th nonzero entry in a given row x . O_H returns the value of the entry in row x and column y to a given precision. The actions of O_F and O_H are given by

$$O_F|x, i\rangle = |x, y_i\rangle, \quad (2)$$

$$O_H|x, y\rangle|z\rangle = |x, y\rangle|z \oplus H_{xy}\rangle, \quad (3)$$

where x is a row-index of H (i.e., a computational basis state), and for $0 \leq i \leq d-1$, y_i is column-index of the i th nonzero entry in row x of H . The Hamiltonians are given via these oracle unitaries for the sake of modularity: the sparse Hamiltonians are very general, so oracle queries provide a standardized formalism for accessing them.

A VQE algorithm comprises two main components: a quantum subroutine for estimating the expectation value of a Hamiltonian of interest for some parametrized ansatz state, and a classical outer loop that updates the parameters of the ansatz in order to minimize the expected en-

ergy [1]. The quantum subroutine is implemented by separately estimating the expectation value of each term in the Hamiltonian under some decomposition, most commonly the Pauli decomposition (1).

In this paper, we extend VQE to sparse Hamiltonians, the possibility of which was briefly discussed in the Appendix of [1]. We decompose sparse Hamiltonians into linear combinations of self-inverse one-sparse Hermitian matrices. We then show how to estimate the expectation value of these one-sparse terms using two ansatz state preparations and calls to the oracle unitaries defining the Hamiltonian terms. We will show below that our algorithm requires at most six oracle queries per measurement circuit.

The class of sparse Hamiltonians that admit description by oracles of the forms (2) and (3) is much broader than the class of local Hamiltonians, which admit efficient Pauli decompositions and are thus simulable by standard VQE. To prove that local Hamiltonians are a subset of sparse Hamiltonians with efficient oracle descriptions, it is enough to give oracle descriptions of the Pauli operators, which we do explicitly in the Supplemental Material [28]. Hence for any local Hamiltonian, we could first decompose it into Pauli operators [29–31], and then simulate each Pauli operator using sparse VQE. All electronic structure Hamiltonians that can be simulated using standard VQE can be simulated using sparse VQE in this way. This also provides an example of oracles with simple implementations that are appropriate for NISQ devices.

However, when a Hamiltonian has an efficient Pauli decomposition, it is not a good candidate for sparse VQE because the measurement scheme in Section II requires an extra qubit and an extra ansatz preparation compared to measuring the Pauli terms directly. The cases of real interest for sparse VQE are Hamiltonians that are sparse and admit efficient oracle implementations, but do not admit efficient Pauli decompositions.

One such case is a Hamiltonian that includes bosons and is represented in a *direct encoding* [32], in which the occupation of each mode is stored in binary in its own register of qubits. Bosonic creation and annihilation operators in this encoding are naturally represented in terms of Weyl-Heisenberg shift operators but not as Pauli operators because the occupations of modes can be larger than one. In [28], we give explicit implementations of oracles for this case. These implementations can be combined with the oracles for Pauli operators to handle Hamiltonians that act on both fermions and bosons.

The class of sparse Hamiltonians is very large, and we will not attempt to give an exhaustive list of all theories that can be addressed within it. However, two more examples are as follows. The first is quantum field theory in *compact encoding* [32, 33], in which only the occupations of occupied modes are stored, providing asymptotically optimal space efficiency. Oracles for field theories in compact encoding are explicitly constructed in [33]. The second example is the CI-matrix representation of

quantum chemistry, for which oracles are explicitly constructed in [34].

The number of gates and depth of circuits required by the oracles in [33] and [34] are larger than those required for Pauli operators or for quantum field theory in the direct encoding. Implementation of these oracles in the NISQ era will require extensive error mitigation or significantly improved physical gates and qubits. However, sparse VQE will become possible before other sparsity-based simulation algorithms [19, 21–27], because these require more coherent queries to the same oracles.

In Section II we describe the basic structure of VQE for Hamiltonians that can be decomposed into self-inverse one-sparse Hermitian terms that possess efficient circuit representations. In Section III, we describe methods for obtaining such decompositions. Then in Section IV, we explain how to construct efficient circuit representations of the resulting terms. These methods permit the implementation of efficient VQE procedures for sparse Hamiltonians. We close the paper with some discussion and directions for future work in Section V.

II. SPARSE VQE

VQE was first used to estimate expectation values of the Hamiltonian [1]. However, many other quantities are of interest given an ansatz state that is a good approximation to the ground state or other energy eigenstate. For example, [15, 32, 35] study various properties of composite particles in interacting quantum field theory. Properties such as the invariant mass, mass radius, parton distribution function, and form factor are expectation values of corresponding operators, whereas quantities such as the decay constant are matrix elements between different states [35]. We will therefore consider estimation of quantities $\langle \phi | \hat{O} | \psi \rangle$ for sparse operators \hat{O} between ansatz functions $|\phi\rangle = V|0\rangle$ and $|\psi\rangle = U|0\rangle$ prepared by quantum circuits U and V .

We begin with a Hermitian operator that we assume has an efficient decomposition into a sum of Hermitian, self-inverse, one-sparse terms G_j :

$$\hat{O} = \sum_j^t \alpha_j G_j \quad (4)$$

where α_j are real coefficients and the number of terms t is polylogarithmic in the dimension of the Hilbert space on which \hat{O} acts. If \hat{O} can be efficiently decomposed into Pauli operators, then the Pauli decomposition of \hat{O} is an example of (4) because the Pauli operators are self-inverse and one-sparse. The terms G_j are both Hermitian and unitary, and we further assume that an efficient quantum circuit for each G_j is known. Circuits for sparse unitaries were studied in [36].

Any operator \hat{O} of the form (4) is sparse, and the number of one-sparse terms t is an upper bound on the sparsity. In Section III we will explicitly show how to con-




FIG. 1. Hadamard test circuit realizing estimation of the real part of the matrix element $\langle 0|V^\dagger G_j U|0\rangle$.

struct a decomposition as in (4) for any arbitrary sparse Hermitian operator. However, for the purpose of this Section it is enough to assume that this is possible, because the actual VQE implementation is agnostic to the method used to obtain the decomposition.

Given (4) we perform M Hadamard tests of the operators $V^\dagger G_j U$ via the circuit shown in Figure 1. This circuit has a state register of n qubits initialized in the all zeros state $|0^n\rangle$ and a single ancilla register initialized in the state $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$. The first operation is application of $V^\dagger G_j U$ controlled on the ancilla qubit. The second operation is a single qubit Hadamard gate applied to the ancilla qubit. This circuit is shown in Figure 1. After application of this circuit, the probability of observing zero on the ancilla qubit is:

$$p(0) = \frac{1}{2} (1 + \text{Re}\langle 0^n|V^\dagger G_j U|0^n\rangle). \quad (5)$$

To replace the real part by the imaginary part of $V^\dagger G_j U$ in (5), replace the initial $|+\rangle$ state of the ancilla by the state $| - i\rangle = (|0\rangle - i|1\rangle)/\sqrt{2}$.

After M repetitions of the circuit in Fig. 1, one obtains n_0 zeros and n_1 ones from the measurement outcomes of the ancilla bit. The quantity $(n_0 - n_1)/M$ is an estimate of $\text{Re}\langle 0^n|V^\dagger G_j U|0^n\rangle$. We can therefore interpret ancilla outcome b as determining a random variable with value $(-1)^b$. The analysis of the variance of these estimates and hence the scaling of M for given \hat{O} and precision ϵ proceeds exactly as for Pauli decompositions (given in [37, 38]), so the required M for precision ϵ scales as ϵ^{-2} .

In this Section we have given the extension of VQE to matrices that can be efficiently decomposed into self-inverse one-sparse Hermitian terms described by efficient quantum circuits. Estimation of a matrix element of a self-inverse term G_j between two ansatz states $U|0^n\rangle$ and $V|0^n\rangle$ is accomplished by controlled application of the ansatz circuits U^\dagger and V as well as G_j . For estimation of expectation values we have $U = V$ and twice as many ansatz preparations are required as for Pauli decomposition VQE. The necessity of these extra preparations is apparent when one notes the capability to also estimate matrix elements between distinct states. In the remainder of the paper, we will focus on Hamiltonians for easier comparison to prior literature on sparse Hamiltonian simulation, but all of our results will apply to general

sparse, Hermitian observables. We will discuss obtaining the necessary sparse decompositions in Section III, and efficiently applying the resulting operators in Section IV.

III. OBTAINING SPARSE DECOMPOSITIONS

Access to a d -sparse Hamiltonian is provided by the oracles O_F and O_H as defined in (2) and (3). In the case that H is one-sparse we can simplify the action of O_F :

$$O_F|x, 0\rangle = |x, y_x\rangle, \quad (6)$$

where y_x is defined to be the column-index of the single nonzero entry in row x (corresponding to $i = 0$).

Given a d -sparse Hamiltonian, we wish to decompose it into a polynomial number of one-sparse self-inverse Hermitian terms. In some cases of interest, most notably fermionic Hamiltonians in second-quantized form, we already have a decomposition into one-sparse Hermitian terms. A Hamiltonian expressed in second-quantized form is a polynomial of some set of ladder operators for various particles or modes. The basis for the Hilbert space is given by the occupation number (Fock) representation for each of the modes. Each term in the Hamiltonian is a monomial of ladder operators, which for fermions in the Fock basis is one-sparse since its action as a linear transformation is to map each single Fock state to some scaling of a single Fock state. Therefore, the fermionic Hamiltonian in the occupation number basis is at most d -sparse if it contains d terms, so assuming the number of terms is polynomial in the number of qubits, so is the sparsity.

Ladder operator monomials are in general not self-inverse, nor are they Hermitian. However, for each ladder operator monomial present in the fermionic Hamiltonian, its Hermitian conjugate must also be present, and each such pair together is one-sparse and Hermitian. One-sparseness follows because for a fermionic ladder operator monomial, at most one of the monomial and its conjugate may act on any given state; this in fact extends to any Hamiltonian that contains fermionic ladder operators with nontrivial action in every term, even if bosonic operators are also present. To obtain a decomposition into one-sparse terms that are also self-inverse, we use the following Lemma:

Lemma 1. *Any one-sparse Hamiltonian $H^{(1)}$ may be expressed up to to L bits per real and imaginary part of each entry as a linear combination of*

$$4L = 4 \left\lceil \log_2 \left(\frac{\sqrt{2} \|H^{(1)}\|_{max}}{\gamma} \right) \right\rceil \quad (7)$$

one-sparse, self-inverse Hamiltonians G_j , where γ is the resulting error in max-norm. The O_F oracles for the G_j are the same as the O_F oracle for $H^{(1)}$, and the O_H oracle for any G_j may be computed using two queries to the O_H oracle for $H^{(1)}$.

The proof may be found in the Supplemental Material [28]. Note that $\|H^{(1)}\|_{\max}$ denotes the max-norm of $H^{(1)}$, defined to be the maximum magnitude of any entry in $H^{(1)}$, which is upper-bounded by $\|H^{(1)}\|_{\infty}$ [39].

Lemma 1 is constructive, so we can use the proof to decompose each Hermitian conjugate pair of ladder operator monomials into one-sparse, self-inverse terms. If N is the number of ladder operator monomials in the second-quantized fermionic Hamiltonian, the number of conjugate pairs is $N/2$, so Lemma 1 provides a decomposition of H into a linear combination of at most

$$2N \left\lceil \log_2 \left(\frac{\sqrt{2} \|H\|_{\max}}{\gamma} \right) \right\rceil \quad (8)$$

one-sparse, self-inverse Hermitian terms, since the max-norm of each monomial is upper-bounded by $\|H\|_{\max}$.

Beyond the case of fermionic second-quantized Hamiltonians we consider an arbitrary d -sparse Hamiltonian that we only have oracle access to. This includes the case of second-quantized Hamiltonians with both fermionic and bosonic modes. In order to apply Lemma 1 we first decompose the Hamiltonian into one-sparse terms:

Lemma 2 ([40], Lemma 4.4). *If H is a d -sparse Hamiltonian, there exists a decomposition $H = \sum_{j=1}^{d^2} H_j$ where each H_j is Hermitian and one-sparse. An O_F query to any H_j can be simulated with two O_F queries to H , and an O_H query to any H_j can be simulated with one O_H query to H .*

The proof (in [40]) is again constructive, so we can use Lemma 2 to obtain Hermitian one-sparse terms H_j , and then use Lemma 1 to approximately decompose each of these into Hermitian one-sparse, self-inverse terms. The resulting total number of one-sparse, self-inverse terms in the decomposition of H is at most

$$4d^2 \left\lceil \log_2 \left(\frac{\sqrt{2} \|H\|_{\max}}{\gamma} \right) \right\rceil. \quad (9)$$

Comparing (9) and (8), we see that in cases where either decomposition could be used, which one is preferable depends on half the number of ladder operator monomials ($N/2$) versus the squared sparsity (d^2). For example, the light-front Yukawa model studied in [32] leads to a second-quantized Hamiltonian whose sparsity scales as $\Theta(N^{2/3})$. This is sublinear because each ladder operator monomial maps a large number of Fock states to zero in this model. However, even though the sparsity is asymptotically smaller than N , the squared sparsity is $d^2 = \Theta(N^{4/3})$, so for this example it is still better to separately decompose each Hermitian conjugate pair of ladder operator monomials into one-sparse, self-inverse terms using Lemma 1.

We use the decomposition provided by Lemmas 1 and 2 because it results in terms that are one-sparse and unitary, and as we will see below, have entries ± 1 or $\pm i$;

the cost is that the decomposition itself is approximate. However, the resulting terms can be implemented exactly using at most six oracle queries (see Section IV). Alternative decompositions exist that avoid approximations in the decompositions themselves (e.g., [41]), so it is possible that in future the method given above can be improved if the terms in such a decomposition can be implemented using few oracle queries.

IV. EVOLUTION UNDER ONE-SPARSE UNITARY OPERATORS

The expectation value estimation method in Section II requires controlled applications of one-sparse, self-inverse, Hermitian operators; let G be such an operator. In practice, G will be one of the operators G_j obtained from Lemma 1. There is an extensive body of methods for simulating sparse Hamiltonians [17–27], any of which could be used to implement the controlled application of G . However, because G is one-sparse and self-inverse, we can use a simpler method similar to the construction of the quantum walk operator in [21] (see the proof of Lemma 4 in [21]). The fact that methods for simulation of time evolution generated by sparse Hamiltonians can also be used for simulation of sparse unitaries was first noted in [36].

Using the oracles O_F and O_H for G , we can apply G as follows: let $|x\rangle_s$ be any input computational basis state, and let $|0\rangle_{a_1}|0\rangle_{a_2}$ be ancilla registers. The steps to apply G are: first,

$$|x\rangle_s |0\rangle_{a_1} |0\rangle_{a_2} \xrightarrow{O_F} |x\rangle_s |y_x\rangle_{a_1} |0\rangle_{a_2} \quad (10)$$

$$\xrightarrow{O_H} |x\rangle_s |y_x\rangle_{a_1} |G_{xy_x}\rangle_{a_2}, \quad (11)$$

where y_x is the column-index of the single nonzero entry in row x of G , as in (6). From the proof of Lemma 1 (in the Supplemental Material [28]), it follows that $G_{xy_x} = \pm 1 \forall x, y$ or $G_{xy_x} = \pm i \forall x, y$. Whether the entries are ± 1 or $\pm i$ is determined by j (where $G = G_j$ for some G_j resulting from Lemma 1), which is evaluated in classical preprocessing. Therefore, $|G_{xy_x}\rangle_{a_2}$ need only be a single qubit determining the sign, and as our next step we can apply entry G_{xy_x} exactly as a phase controlled by $|G_{xy_x}\rangle_{a_2}$, and then complete the implementation of G as follows:

$$\xrightarrow{\text{controlled phase}} G_{xy_x} |x\rangle_s |y_x\rangle_{a_1} |G_{xy_x}\rangle_{a_2} \quad (12)$$

$$\xrightarrow{O_H^{-1}} G_{xy_x} |x\rangle_s |y_x\rangle_{a_1} |0\rangle_{a_2} \quad (13)$$

$$\xrightarrow{\text{swap } s, a_1} G_{xy_x} |y_x\rangle_s |x\rangle_{a_1} |0\rangle_{a_2} \quad (14)$$

$$\xrightarrow{O_F^{-1}} G_{xy_x} |y_x\rangle_s |0\rangle_{a_1} |0\rangle_{a_2}, \quad (15)$$

where the last step follows because for a one-sparse, Hermitian operator, (6) implies

$$O_F |y_x, 0\rangle = |y_x, x\rangle. \quad (16)$$

The effect of these operations is to map

$$|x\rangle_s \mapsto G_{xy_x}|y_x\rangle_s = G|x\rangle_s, \quad (17)$$

i.e., we have applied G to $|x\rangle_s$. This required four queries to the oracles: one query each to O_F , O_H , and their inverses. For ancillas, we required copying the computational register s in the register a_1 to apply the O_F oracle, and one additional qubit in the register a_2 to represent the sign of $G_{xy_x} = \pm 1, \pm i$. In both queries and ancillas, these are the minimum requirements to apply the oracle unitaries at all.

Finally, recall that the one-sparse, self-inverse terms that we are estimating expectation values of were obtained via Lemma 1, above. However, the full Hamiltonian is first decomposed into one-sparse terms $H^{(1)}$, either conjugate pairs of ladder-operator monomials or via Lemma 2, which form the inputs to Lemma 1. From Lemma 1 we know that the O_F oracle for any of the G_j is identical to the O_F oracle for $H^{(1)}$. Also, from the proof of Lemma 1, we know that the O_H oracle for any of the G_j can be implemented by first applying the O_H oracle for G_j and then performing a single controlled operation (we would later undo both of these steps to apply O_H^{-1}). Hence the number of queries to each $H^{(1)}$ is still four, each of which will either be implemented directly (in the ladder operator monomial decomposition), or via one (for O_H) or two (for O_F) queries to the full Hamiltonian using Lemma 2. This gives a total of at most six oracle queries.

V. CONCLUSION

In existing studies, the only Hamiltonian input model used in VQE has been decomposition into Pauli operators. In this paper we have extended VQE to the case of sparse Hamiltonians. We accomplished this by employing a variant of techniques previously considered applicable to future fault-tolerant quantum computers [40]. For sparse Hamiltonians, we have demonstrated how VQE can be implemented via a decomposition into one-sparse, self-inverse Hermitian terms. As discussed in the introduction, simulation of second-quantized Hamiltonians in condensed matter, high energy and nuclear physics, and in compact representations of quantum chemistry are natural candidates for this sparse VQE method [15, 32–35, 42].

This paper focused on VQE, but the results may also be used in the context of QAOA [2]. QAOA to date treats classical objective functions that are a sum of local clauses, of which 3-SAT is a canonical NP-complete example. Classical objective operators are diagonal in the computational basis and hence naturally one-sparse. The techniques here would allow extension to the case where the diagonal entries are given by more complicated classical functions. This broadens the space of examples within which to search for quantum advantage and also may provide practical advantages for problems with large locality such as the travelling salesman problem. We leave the investigation of these problems for future work.

ACKNOWLEDGMENTS

The authors thank Michael Kreshchuk for helpful conversations and advice. W. M. K. acknowledges support from the National Science Foundation, Grant No. DGE-1842474. This work was supported by the NSF STAQ project (PHY-1818914).

-
- [1] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O’Brien, *Nature Communications* **5**, 4213 EP (2014).
- [2] E. Farhi, J. Goldstone, and S. Gutmann, arXiv preprint arXiv:1411.4028 (2014).
- [3] P. J. J. O’Malley, R. Babbush, I. D. Kivlichan, J. Romero, J. R. McClean, R. Barends, J. Kelly, P. Roushan, A. Tranter, N. Ding, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, E. Jeffrey, E. Lucero, A. Megrant, J. Y. Mutus, M. Neeley, C. Neill, C. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, P. V. Coveney, P. J. Love, H. Neven, A. Aspuru-Guzik, and J. M. Martinis, *Phys. Rev. X* **6**, 031007 (2016).
- [4] R. Santagati, J. Wang, A. A. Gentile, S. Paesani, N. Wiebe, J. R. McClean, S. Morley-Short, P. J. Shadbolt, D. Bonneau, J. W. Silverstone, D. P. Tew, X. Zhou, J. L. O’Brien, and M. G. Thompson, *Science Advances* **4** (2018), 10.1126/sciadv.aap9646.
- [5] Y. Shen, X. Zhang, S. Zhang, J.-N. Zhang, M.-H. Yung, and K. Kim, *Physical Review A* **95**, 020501 (2017).
- [6] S. Paesani, A. A. Gentile, R. Santagati, J. Wang, N. Wiebe, D. P. Tew, J. L. O’Brien, and M. G. Thompson, *Phys. Rev. Lett.* **118**, 100503 (2017).
- [7] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, *Nature* **549**, 242 (2017).
- [8] C. Hempel, C. Maier, J. Romero, J. McClean, T. Monz, H. Shen, P. Jurcevic, B. P. Lanyon, P. Love, R. Babbush, A. Aspuru-Guzik, R. Blatt, and C. F. Roos, *Phys. Rev. X* **8**, 031022 (2018).
- [9] 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, *Phys. Rev. Lett.* **120**, 210501 (2018).
- [10] J. I. Colless, V. V. Ramasesh, D. Dahlen, M. S. Blok, M. E. Kimchi-Schwartz, J. R. McClean, J. Carter, W. A. de Jong, and I. Siddiqi, *Phys. Rev. X* **8**, 011021 (2018).
- [11] Y. Nam, J.-S. Chen, N. C. Pienti, K. Wright, C. Delaney, D. Maslov, K. R. Brown, S. Allen, J. M. Amini, J. Apisdorf, K. M. Beck, A. Blinov, V. Chaplin, M. Chmielewski, C. Collins, S. Debnath, A. M. Ducore, K. M. Hudek, M. Keesan, S. M. Kreikemeier, J. Mizrahi, P. Solomon, M. Williams, J. D. Wong-Campos, C. Monroe, and J. Kim, arXiv preprint (2019), arXiv:1902.10171 [quant-ph].
- [12] C. Kokail, C. Maier, R. van Bijnen, T. Brydges, M. K. Joshi, P. Jurcevic, C. A. Muschik, P. Silvi, R. Blatt, C. F. Roos, and P. Zoller, *Nature* **569**, 355 (2019).
- [13] A. Kandala, K. Temme, A. D. Córcoles, A. Mezzacapo, J. M. Chow, and J. M. Gambetta, *Nature* **567**, 491 (2019).
- [14] Google AI Quantum and Collaborators, *Science* **369**, 1084 (2020).
- [15] M. Kreshchuk, S. Jia, W. M. Kirby, G. Goldstein, J. P. Vary, and P. J. Love, arXiv preprint arXiv:2009.07885 (2020).
- [16] S. Lloyd, *Science* **273**, 1073 (1996).
- [17] D. Aharonov and A. Ta-Shma, in *Proceedings of the 35th Annual ACM Symposium on Theory of Computing, STOC '03* (Association for Computing Machinery, New York, NY, USA, 2003) pp. 20–29.
- [18] A. M. Childs, R. Cleve, E. Deotto, E. Farhi, S. Gutmann, and D. A. Spielman, in *Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing, STOC '03* (Association for Computing Machinery, New York, NY, USA, 2003) pp. 59–68.
- [19] D. W. Berry, G. Ahokas, R. Cleve, and B. C. Sanders, *Communications in Mathematical Physics* **270**, 359 (2007).
- [20] A. M. Childs, *Communications in Mathematical Physics* **294**, 581 (2010).
- [21] D. W. Berry and A. M. Childs, *Quantum Information and Computation* **12**, 29 (2012).
- [22] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, 283 (2014).
- [23] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, *Phys. Rev. Lett.* **114**, 090502 (2015).
- [24] D. W. Berry, A. M. Childs, and R. Kothari, in *2015 IEEE 56th Annual Symposium on Foundations of Computer Science* (2015) pp. 792–809.
- [25] G. H. Low and I. L. Chuang, *Phys. Rev. Lett.* **118**, 010501 (2017).
- [26] G. H. Low and I. L. Chuang, *Quantum* **3**, 163 (2019).
- [27] D. W. Berry, A. M. Childs, Y. Su, X. Wang, and N. Wiebe, *Quantum* **4**, 254 (2020).
- [28] See Supplemental Material at [URL will be inserted by publisher] for explicit oracle implementations for example applications, and a proof of Lemma 1 in the main text. Additional works cited in the Supplemental Material are [43–46].
- [29] P. Jordan and E. Wigner, *Z. Phys.* **47**, 631 (1928).
- [30] S. B. Bravyi and A. Y. Kitaev, *Annals of Physics* **298**, 210 (2002).
- [31] J. T. Seeley, M. J. Richard, and P. J. Love, *Journal of Chemical Physics* **137**, 224109 (2012).
- [32] M. Kreshchuk, W. M. Kirby, G. Goldstein, H. Beauchemin, and P. J. Love, arXiv preprint (2020), arXiv:2002.04016 [quant-ph].
- [33] W. M. Kirby, S. Hadi, M. Kreshchuk, and P. J. Love, arXiv preprint (2021), arXiv:2105.10941 [quant-ph].
- [34] R. Babbush, D. W. Berry, Y. R. Sanders, I. D. Kivlichan, A. Scherer, A. Y. Wei, P. J. Love, and A. Aspuru-Guzik, *Quantum Science and Technology* **3**, 015006 (2017).
- [35] M. Kreshchuk, S. Jia, W. M. Kirby, G. Goldstein, J. P. Vary, and P. J. Love, “Simulating hadronic physics on nisq devices using basis light-front quantization,” (2020), arXiv:2011.13443 [quant-ph].
- [36] S. P. Jordan and P. Wocjan, *Physical Review A* **80**, 062301 (2009).
- [37] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, *New Journal of Physics* **18**, 023023 (2016).
- [38] N. C. Rubin, R. Babbush, and J. McClean, *New Journal of Physics* **20**, 053020 (2018).
- [39] A. M. Childs and R. Kothari, *Quantum Information & Computation* **10**, 669 (2010).
- [40] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, in *Proceedings of the 46th ACM Symposium on Theory of Computing (STOC 2014)*, (2014) pp. 283–292.
- [41] A. M. Childs and R. Kothari, in *Theory of Quantum Computation, Communication, and Cryptography*, edited by W. van Dam, V. M. Kendon, and S. Severini (Springer Berlin Heidelberg, Berlin, Heidelberg, 2011) pp. 94–103.
- [42] A. Aspuru-Guzik, A. D. Dutoi, P. J. Love, and M. Head-Gordon, *Science* **309**, 1704 (2005).
- [43] D. Gottesman, arXiv preprint (1997), arXiv:quant-ph/9705052 [quant-ph].
- [44] A. Parent, M. Roetteler, and M. Mosca, in *12th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2017)*, Leibniz International Proceedings in Informatics (LIPIcs), Vol. 73, edited by M. M. Wilde (Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2018) pp. 7:1–7:15.
- [45] S. Sultana and K. Radecka, in *2011 18th IEEE International Conference on Electronics, Circuits, and Systems* (2011) pp. 141–144.
- [46] A. JavadiAbhari, S. Patil, D. Kudrow, J. Heckey, A. Lvov, F. T. Chong, and M. Martonosi, in *Proceedings of the 11th ACM Conference on Computing Frontiers, CF '14* (Association for Computing Machinery, New York, NY, USA, 2014).