

Ground-State Preparation and Energy Estimation on Early Fault-Tolerant Quantum Computers via Quantum Eigenvalue Transformation of Unitary Matrices

Yulong Dong^{1,4}, Lin Lin^{1,2,3,*} and Yu Tong¹

¹*Department of Mathematics, University of California, Berkeley, California 94720, USA*

²*Applied Mathematics and Computational Research Division, Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA*

³*Challenge Institute of Quantum Computation, University of California, Berkeley, California 94720, USA*

⁴*Berkeley Center for Quantum Information and Computation, Berkeley, California 94720, USA*



(Received 24 April 2022; accepted 25 August 2022; published 12 October 2022)

Under suitable assumptions, some recently developed quantum algorithms can estimate the ground-state energy and prepare the ground state of a quantum Hamiltonian with near-optimal query complexities. However, this is based on a block-encoding input model of the Hamiltonian, the implementation of which is known to require a large resource overhead. We develop a tool called quantum eigenvalue transformation of unitary matrices with real polynomials (QETU), which uses a controlled Hamiltonian evolution as the input model, a single ancilla qubit, and no multiqubit control operations and is thus suitable for early fault-tolerant quantum devices. This leads to a simple quantum algorithm that outperforms all previous algorithms with a comparable circuit structure for estimating the ground-state energy. For a class of quantum spin Hamiltonians, we propose a new method that exploits certain anticommutation relations and further removes the need to implement the controlled Hamiltonian evolution. Coupled with a Trotter-based approximation of the Hamiltonian evolution, the resulting algorithm can be very suitable for early fault-tolerant quantum devices. We demonstrate the performance of the algorithm using IBM QISKIT for the transverse-field Ising model. If we are further allowed to use multiqubit Toffoli gates, we can then implement amplitude amplification and a new binary amplitude-estimation algorithm, which increases the circuit depth but decreases the total query complexity. The resulting algorithm saturates the near-optimal complexity for ground-state preparation and energy estimation using a constant number of ancilla qubits (no more than three).

DOI: [10.1103/PRXQuantum.3.040305](https://doi.org/10.1103/PRXQuantum.3.040305)

I. INTRODUCTION

Preparation of the ground state and estimation of the ground-state energy of a quantum Hamiltonian have a wide range of applications in condensed-matter physics, quantum chemistry, and quantum information. To solve such problems, quantum computers promise to deliver a new level of computational power that can be significantly beyond the boundaries set by classical computers. Recently developed quantum algorithms in [1] can perform these tasks with near-optimal query complexity under suitable conditions. Despite exciting early progress on noisy intermediate-scale quantum (NISQ) devices [2],

it is widely believed that most scientific advances in quantum sciences require some version of fault-tolerant quantum computers, which are expected to be able to accomplish much more complicated tasks. On the other hand, the fabrication of full-scale fault-tolerant quantum computers remains a formidable technical challenge for the foreseeable future and it is reasonable to expect that early fault-tolerant quantum computers share the following characteristics: (1) the number of logical qubits is limited; and (2) it can be difficult to execute certain controlled operations (e.g., multiqubit control gates), the implementation of which requires a large number of non-Clifford gates. Besides these, the maximum circuit depth of early fault-tolerant quantum computers, which is determined by the maximum coherence time of the devices, may still be limited. Therefore it is still important to reduce the circuit depth, sometimes even at the expense of a larger total run time (via a larger number of repetitions). Quantum algorithms tailored for early fault-tolerant quantum computers [3–10] need to properly

*linlin@math.berkeley.edu

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

take these limitations into account and the resulting algorithmic structure can be different from those designed for fully fault-tolerant quantum computers.

To gain access to the quantum Hamiltonian H , a standard input model is the *block-encoding* (BE) model, which directly encodes the matrix H (after proper rescaling) as a submatrix block of a larger unitary matrix U_H [11,12]. Combined with techniques such as linear combination of unitaries (LCU) [13], quantum signal processing [14], or quantum singular-value transformation (QSVT) [15], one can implement a large class of matrix functions of H on a quantum computer. This leads to quantum algorithms for ground-state preparation and ground-state energy estimation with near-optimal query complexities to U_H [1]. The block-encoding technique is also very useful in many other tasks such as Hamiltonian simulation, the solution of linear systems, preparation of the Gibbs state, and computation of Green's function and the correlation functions [11,14,16–18]. However, the block encoding of a quantum Hamiltonian (e.g., a sparse matrix) often involves a relatively large number of ancilla qubits, as well as multiqubit controlled operations that lead to a large number of two-qubit gates and long circuit depths [15], and is therefore not suitable in the early fault-tolerant setting.

A widely used alternative approach for accessing the information in H is the time-evolution operator $U = \exp(-i\tau H)$ for some time τ . This input model is referred to as the *Hamiltonian-evolution* (HE) model. While Hamiltonian simulation can be performed using quantum signal processing for sparse Hamiltonians with optimal query complexity [14], such an algorithm queries a block encoding of H , which defeats the purpose of employing the HE model. On the other hand, when H can be efficiently decomposed into a linear combination of Pauli operators, the time-evolution operator can be efficiently implemented using, e.g., the Trotter product formula [19,20] without using any ancilla qubit. This remarkable feature has inspired quantum algorithms for performing a variety of tasks using controlled time evolution and one ancilla qubit. A textbook example of such an algorithm is the Hadamard test. It uses one ancilla qubit and the controlled Hamiltonian evolution to estimate the average value $\text{Re} \langle \psi | U | \psi \rangle$, which is encoded by the probability of measuring the ancilla qubit with outcome 0 [see Fig. 1(a)]. The number of repeated measurements of this procedure is $\mathcal{O}(\epsilon^{-2})$, where ϵ is the desired precision. Assume that the spectrum of the Hamiltonian H is contained in $[\eta, \pi - \eta]$ for some $\eta > 0$. If $|\psi\rangle$ is the exact ground state of H , we can retrieve the eigenvalue as $\lambda = \arccos(\text{Re} \langle \psi | U | \psi \rangle)$. By allowing a series of longer simulation times of $t = d$ for some integer d , this leads to Kitaev's algorithm, which uses only $\log \epsilon^{-1}$ measurements, at the expense of increasing the circuit depth to $\mathcal{O}(\epsilon^{-1})$ [see Fig. 1(b)]. The total simulation time is therefore $\mathcal{O}(\epsilon^{-1} \log \epsilon^{-1})$, which reaches the Heisenberg limit [21–24] up to a logarithmic factor.

When the input state $|\phi_0\rangle$ (prepared by an oracle U_I) is different from the exact ground state of H , denoted by $|\psi_0\rangle$, there have been multiple quantum algorithms using the circuit of Fig. 1(b) or its variant to estimate the ground-state energy [7,25–27]. Let γ be a lower bound of the initial overlap, i.e., $|\langle \phi_0 | \psi_0 \rangle| \geq \gamma$. It is worth noting that *all* quantum algorithms with provable performance guarantees require *a priori* knowledge that γ is reasonably large (assuming black-box access to the Hamiltonian). Without such an assumption, this problem is QMA hard [28–31]. Candidates for such $|\phi_0\rangle$ include the Hartree-Fock state in quantum chemistry [32,33] and quantum states prepared using the variational quantum eigensolver [34–36]. Some techniques can be used to boost the overlap using low-depth circuits [9]. Furthermore, algorithms using the circuit of Fig. 1(b) typically cannot be used to prepare the ground state. With the time-evolution operator as the input, one can use the LCU algorithm to prepare the ground state [37,38], thus reducing the number of ancilla qubits needed to implement the block encoding of the Hamiltonian. Note that LCU requires additional ancilla qubits to store the coefficients and, as a result, cannot be implemented using $\mathcal{O}(1)$ qubits.

We show that both the ground-state preparation and the energy estimation can be solved by (repeatedly) preparing a quantum state of the form $|\psi_f\rangle \propto f(H) |\phi_0\rangle$, where f is a real polynomial approximating a shifted sign function. A main technical tool developed in this paper is called *quantum eigenvalue transformation of unitary matrices with real polynomials* (QETU), which allows us to prepare such a state $|\psi_f\rangle$ by querying $U = e^{-iH}$ using only one ancilla qubit and does not require any multiqubit control operation (Theorem 1). The circuit structure [Fig. 1(c)] is only slightly different from that in Fig. 1(b). The QETU technique is closely related to concepts such as quantum signal processing (QSP), quantum eigenvalue transformation, and QSVT. The relations among these techniques are detailed in Appendix A. The information of the function f of interest is stored in the adjustable parameters $\{\varphi_i\}$, called *phase factors*. To find such parameters, we need to identify a polynomial approximation to the shifted sign function and then evaluate the phase factors corresponding to the approximate polynomial. Most QSP-based applications construct such a polynomial approximation analytically, which can sometimes lead to cumbersome expressions and suboptimal approximation results. We provide a convex-optimization-based procedure to streamline this process and yield the near-optimal approximation (see Sec. IV). Both the QETU technique and the convex-optimization method can be useful in applications beyond ground-state preparation and energy estimation.

The computational cost is primarily measured in terms of the query complexity, i.e., how many times we need to query U and U_I in total, and we also analyze the additional one- and two-qubit gates needed, such as the

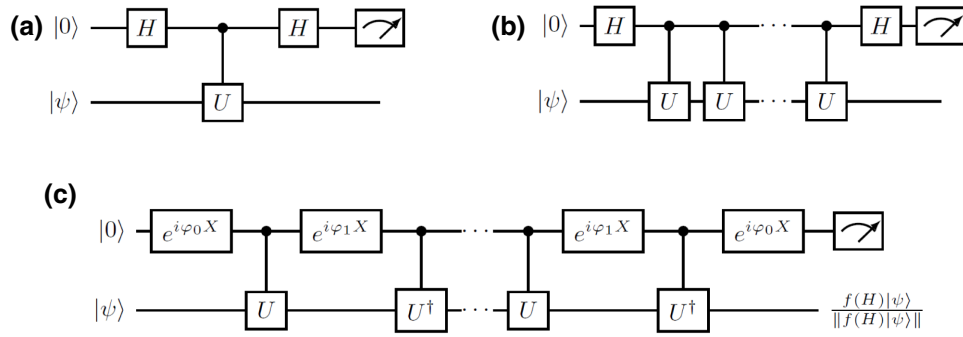


FIG. 1. After a proper rescaling, the n -qubit circuit U implements e^{-iH} and the Hadamard test circuit [(a), short time evolution] estimates $\text{Re} \langle \psi | e^{-iH} | \psi \rangle$. Repeating the controlled evolution d times, the circuit shown in (b) estimates the average value of a long time evolution $\text{Re} \langle \psi | e^{-idH} | \psi \rangle$. For a very general class of functions f , the QETU circuit shown in (c) can approximately prepare a normalized quantum state $f(H)|\psi\rangle / \|f(H)|\psi\rangle\|$ with approximate success probability $p = \|f(H)|\psi\rangle\|^2$, by interleaving the forward (U) and backward (U^\dagger) time evolution with some properly chosen X rotations in the ancilla qubit.

single-qubit-rotation gates, and the two-qubit gates needed to implement the n -qubit reflection operator. In our algorithms, the number of additional gates has the same scaling as the query complexity or involves an n factor, where n is the system size. We measure the circuit-depth requirement in terms of *query depth*: the number of times we need to query U in one coherent run of the circuit. Note that this term is not to be confused with the circuit depth for implementing the oracle U , which we do not consider in this work. This metric reflects the circuit-depth requirement faithfully because in our algorithm, in one coherent run of the circuit, the number of queries to U_i is also upper bounded by this metric and the additional circuit depth needed for additional gates is upper bounded by this metric up to a factor of $\mathcal{O}(n)$.

Besides the query depth, we also focus on whether multiqubit control needs to be implemented. Algorithms such as amplitude amplification and amplitude estimation [39] can be used to reduce the total query complexity but they also need to use $(n + 1)$ -bit Toffoli gates (specifically, an n -qubit reflection operator with respect to the zero state $|0^n\rangle$), which can be implemented using $\mathcal{O}(n)$ two-qubit gates and one ancilla qubit [40]. These operations are referred to as “low-level” multiqubit control gates. Some other quantum algorithms may require more complex multiqubit control operations as well as more ancilla qubits. For instance, the high-confidence quantum-phase-estimation (QPE) algorithm [41–43] requires a circuit to carry out the arithmetic operation of taking the median of multiple energy-measurement results, which can require $\text{poly}(n)$ two-qubit gates and ancilla qubits. Such operations are referred to as “high-level” multiqubit control gates.

To solve the ground-state-preparation and energy-estimation problem, we propose two different types of algorithm, with two different goals in mind. For the first type of algorithm, which we call the *short-query-depth algorithm*, we only use QETU to prioritize reducing the

quantum resources needed. No multiqubit controlled operation is involved. For the second type of algorithm, which we call the *near-optimal algorithm*, we optimize the total query complexity by using amplitude amplification and a new binary amplitude-estimation algorithm (Lemma 12). Such algorithms only use low-level multiqubit control operations. Both types of algorithms only use a small number of ancilla qubits (no more than two or three).

For ground-state energy estimation, surprisingly, even though the total query complexity of the short-query-depth algorithm does not have the optimal asymptotic scaling, it still outperforms *all* previous algorithms with the same ancilla qubit number constraint [7,44–46], in terms of total query complexity (see Table I). Most notably, in this setting, we achieve a quadratic improvement on the γ dependence, from $\tilde{\mathcal{O}}(\gamma^{-4})$ to $\tilde{\mathcal{O}}(\gamma^{-2})$ [the notation $\tilde{\mathcal{O}}(g)$ means $\mathcal{O}[g \text{ poly log}(g)]$ unless otherwise stated]. Moreover, the circuit depth from the previous state-of-the-art result is preserved in our algorithm. Numerical comparison (see Fig. 4) demonstrates that our algorithm outperforms QPE, not only in terms of the asymptotic scaling but also the exact nonasymptotic number of queries for moderately small values of γ . Our near-optimal algorithm takes this advantage even further, matching the best known query-complexity scaling in Ref. [1] (which saturates the query-complexity lower bound).

For ground-state preparation, the only other algorithm that can use at most constantly many ancilla qubits is the QPE algorithm with a semiclassical Fourier transform [47]. Compared to this algorithm, our short-query-depth algorithm has an exponentially improved precision dependence and a quadratically improved γ dependence, from $\tilde{\mathcal{O}}(\gamma^{-4})$ to $\tilde{\mathcal{O}}(\gamma^{-2})$, while maintaining the same circuit depth. The near-optimal algorithm further improves the dependence to $\tilde{\mathcal{O}}(\gamma^{-1})$. A comparison of the algorithms for ground-state preparation can be found in Table II. We remark that here we consider the case where we know a

TABLE I. A comparison of the performance of quantum algorithms for ground-state energy estimation in terms of the query complexity, the query depth, the number of ancilla qubits, and the level of multiqubit control (abbreviated as “MQC” in the table) operations needed. γ is the overlap between the initial guess $|\phi_0\rangle$ and the ground state and ϵ is the allowed error. “HE” denotes the Hamiltonian-evolution model (assuming no ancilla qubits) and “BE” for the block-encoding model. The sharper estimate for estimating the ground-state energy using the quantum-eigenvalue-estimation algorithm (QEEA) is given in Ref. [7, Appendix C]. We assume that Ref. [1] uses m ancilla qubits and that the high-level MQC operation is due to the block encoding of H .

	Query depth	Query complexity	Number of ancilla qubits	Need MQC?	Input model
This work (Theorem 4)	$\tilde{\mathcal{O}}(\epsilon^{-1})$	$\tilde{\mathcal{O}}(\epsilon^{-1}\gamma^{-2})$	$\mathcal{O}(1)$	No	HE
This work (Theorem 5)	$\tilde{\mathcal{O}}(\epsilon^{-1}\gamma^{-1})$	$\tilde{\mathcal{O}}(\epsilon^{-1}\gamma^{-1})$	$\mathcal{O}(1)$	Low	HE
QPE (high confidence) [41–43]	$\tilde{\mathcal{O}}(\epsilon^{-1})$	$\tilde{\mathcal{O}}(\epsilon^{-1}\gamma^{-2})$	$\mathcal{O}(\text{poly} \log(\gamma^{-1}\epsilon^{-1}))$	High	HE
QPE (semiclassical) [44,45]	$\tilde{\mathcal{O}}(\epsilon^{-1}\gamma^{-2})$	$\tilde{\mathcal{O}}(\epsilon^{-1}\gamma^{-4})$	$\mathcal{O}(1)$	No	HE
QEEA [7,46]	$\tilde{\mathcal{O}}(\epsilon^{-1})$	$\tilde{\mathcal{O}}(\epsilon^{-4}\gamma^{-4})$	$\mathcal{O}(1)$	No	HE
GTC19 (Theorem 4) [37]	$\tilde{\mathcal{O}}(\epsilon^{-3/2}\gamma^{-1})$	$\tilde{\mathcal{O}}(\epsilon^{-3/2}\gamma^{-1})$	$\mathcal{O}(\log(\epsilon^{-1}))$	High	HE
LT20 [1]	$\tilde{\mathcal{O}}(\epsilon^{-1}\gamma^{-1})$	$\tilde{\mathcal{O}}(\epsilon^{-1}\gamma^{-1})$	$m + \mathcal{O}(\log(\epsilon^{-1}))$	High	BE
LT22 [7]	$\tilde{\mathcal{O}}(\epsilon^{-1})$	$\tilde{\mathcal{O}}(\epsilon^{-1}\gamma^{-4})$	$\mathcal{O}(1)$	No	HE

parameter μ such that $\lambda_0 \leq \mu - \Delta/2 < \mu + \Delta/2 \leq \lambda_1$, as in Theorem 6. If no such μ is known, we first need to estimate the ground-state energy to precision $\mathcal{O}(\Delta)$ and the resulting algorithm is discussed in Theorem 13. If the ground-state energy is known *a priori*, then the algorithm in Ref. [48] may yield a similar speed-up for preparing the ground state, but such knowledge is generally not available.

In the above analysis, specifically in Tables I and II, we compare with algorithms the complexity of which can be rigorously analyzed under the assumptions of a good initial overlap (and spectral gap for ground-state preparation). We do not compare with heuristic algorithms such as the variational quantum eigensolver [34–36]. There are also algorithms that are designed with different but similar goals in mind, such as the quantum algorithmic cooling technique in Ref. [49], which can estimate an eigenvalue λ_j belonging to a given range $[\lambda_j^L, \lambda_j^R]$ (assuming that all other eigenvalues are away from this range). Then, the

total run-time scaling is $\tilde{\mathcal{O}}(\gamma^{-4})$, where γ is the overlap between the initial guess and the target eigenstate [49, Theorem 2]. The same technique can also be used to estimate the expectation value of an observable associated with the target eigenstate, without coherently preparing the target eigenstate.

For certain Hamiltonians, QETU can be implemented with the standard Hamiltonian evolution rather than the controlled version. Note that “control-free” only means that the Hamiltonian evolution is not controlled by one or more qubits but control gates that are independent of the Hamiltonian can still be used. In Ref. [50], the control-free setting for an n -qubit time evolution is achieved by introducing an n -qubit reference state, on which the time evolution acts trivially. The algorithm also requires the implementation of the controlled n -qubit SWAP gate and therefore has a relatively large overhead. There are other control-free algorithms proposed in Refs. [7,51,52] for energy and phase estimation via the measurement of

TABLE II. A comparison of the performance of quantum algorithms for ground-state preparation in terms of the query complexity, the query depth, the number of ancilla qubits, and the level of multiqubit control (abbreviated as “MQC” in the table) operations needed. γ is the overlap between the initial guess $|\phi_0\rangle$ and the ground state, Δ is a lower bound of the spectral gap, and $1 - \epsilon$ is the target fidelity. “HE” denotes the Hamiltonian-evolution model (assuming no ancilla qubits) and “BE” the block-encoding model. Here, we assume that an upper bound of the ground-state energy is known (μ in Theorem 6). The algorithm in Ref. [37] (GTC19 in the table) requires precise knowledge of the ground-state energy. We assume that Ref. [1] uses m ancilla qubits and that the high-level MQC operation is due to the block encoding of H .

	Query depth	Query complexity	Number of ancilla qubits	Need MQC?	Input model
This work (Theorem 6)	$\tilde{\mathcal{O}}(\Delta^{-1})$	$\tilde{\mathcal{O}}(\Delta^{-1}\gamma^{-2})$	$\mathcal{O}(1)$	No	HE
This work (Theorem 11)	$\tilde{\mathcal{O}}(\Delta^{-1}\gamma^{-1})$	$\tilde{\mathcal{O}}(\Delta^{-1}\gamma^{-1})$	$\mathcal{O}(1)$	Low	HE
QPE (high confidence) [41–43]	$\tilde{\mathcal{O}}(\Delta^{-1})$	$\tilde{\mathcal{O}}(\Delta^{-1}\gamma^{-2})$	$\mathcal{O}(\text{poly} \log(\Delta^{-1}\gamma^{-1}\epsilon^{-1}))$	High	HE
QPE (semiclassical) [44,45]	$\tilde{\mathcal{O}}(\Delta^{-1}\gamma^{-2})$	$\tilde{\mathcal{O}}(\Delta^{-1}\gamma^{-4})$	$\mathcal{O}(1)$	No	HE
GTC19 (Theorem 1) [37]	$\tilde{\mathcal{O}}(\Delta^{-1}\gamma^{-1})$	$\tilde{\mathcal{O}}(\Delta^{-1}\gamma^{-1})$	$\mathcal{O}(\log(\Delta^{-1}) + \log \log(\epsilon^{-1}))$	High	HE
LT20 [1]	$\tilde{\mathcal{O}}(\Delta^{-1}\gamma^{-1})$	$\tilde{\mathcal{O}}(\Delta^{-1}\gamma^{-1})$	m	High	BE

certain scalar expectation values as the output. In particular, such algorithms cannot coherently implement a controlled time evolution and are therefore not compatible with the implementation of QETU. In this paper, we exploit certain anticommutation relations and structures of the Hamiltonian to propose a new control-free implementation. In the context of QETU, the algorithm does not introduce any ancilla qubit and requires a small number of two-qubit gates that scales linearly in n . We demonstrate the optimized circuit implementation of the transverse-field Ising model under the control-free setting. To the extent of our knowledge, this circuit is significantly simpler than all previous QSP-type circuits for simulating a physical Hamiltonian. We show the numerical performance of our algorithm for estimating the ground-state energy estimation in the presence of tunable quantum error using IBM QISKIT.

II. QUANTUM EIGENVALUE TRANSFORMATION OF UNITARY MATRICES

Given the Hamiltonian-evolution input model $U = e^{-iH}$, we first demonstrate that by slightly modifying the circuit for the Hadamard test in Fig. 1(b), we can approximately prepare a target state $|\psi_f\rangle = f(H)|\psi\rangle / \|f(H)|\psi\rangle\|$ efficiently and with controlled accuracy for a large class of real functions f . Specifically, this requires alternately applying the controlled forward time-evolution operator U , a single qubit X rotation in the ancilla qubit, and the controlled backward time-evolution operator U^\dagger [see Fig. 1(c)]. This circuit does not store the eigenvalues of H either in a classical or a quantum register and the information of the function f of interest is entirely stored in the adjustable parameters $\varphi_0, \varphi_1, \varphi_2, \dots, \varphi_{d/2}$. These parameters form a set of symmetric phase factors $(\varphi_0, \varphi_1, \varphi_2, \dots, \varphi_2, \varphi_1, \varphi_0) \in \mathbb{R}^{d+1}$ used in the circuit Fig. 1(c). The symmetry of the phase factors is the key to attaining the reality of the function f of interest.

Theorem 1: (QETU). *Let $U = e^{-iH}$ with an n -qubit Hermitian matrix H . For any even real polynomial $F(x)$ of degree d satisfying $|F(x)| \leq 1, \forall x \in [-1, 1]$, we can find a sequence of symmetric phase factors $\Phi_z := (\varphi_0, \varphi_1, \dots, \varphi_1, \varphi_0) \in \mathbb{R}^{d+1}$, such that the circuit in Fig. 1(c) denoted by \mathcal{U} satisfies $\langle (|0\rangle \otimes I_n) \mathcal{U} (|0\rangle \otimes I_n) = F(\cos(H/2))$.*

The proof of Theorem 1 is given in Appendix B. It is worth mentioning that the concept of ‘‘qubitization’’ [12,15] appears very straightforwardly in QETU. Let the matrix function of interest be expressed as $f(H) = (f \circ g)(\cos(H/2))$, where $g(x) = 2 \arccos(x)$. Therefore, we can find a polynomial approximation $F(x)$ so that

$$\sup_{x \in [\sigma_{\min}, \sigma_{\max}]} |(f \circ g)(x) - F(x)| \leq \epsilon. \quad (1)$$

Here, $\sigma_{\min} = \cos(\lambda_{\max}/2)$ and $\sigma_{\max} = \cos(\lambda_{\min}/2)$, respectively [note that $\cos(x/2)$ is a monotonically decreasing function on $[0, \pi]$]. This ensures that the operator norm error satisfies

$$\| \langle (|0\rangle \otimes I_n) \mathcal{U} (|0\rangle \otimes I_n) - f(H) \| \leq \epsilon.$$

The implementation of $U = e^{-iH}$ corresponds to a Hamiltonian simulation problem of H at time $t = 1$. In practice, we can use the Trotter decomposition to obtain an approximate implementation of U without ancilla qubits, i.e., we can partition the time interval into r steps with $\tau = r^{-1}$ and use a low-order Trotter method to implement an approximation to $U_\tau \approx e^{-iH\tau}$. Then,

$$U = e^{-iH} \approx (U_\tau)^r. \quad (2)$$

In line with other works in analyzing the performance of quantum algorithms using the HE input model [7,44–46], in the discussion below, unless otherwise specified, we assume that U is implemented exactly and that the errors are due to other sources such as polynomial approximation, the binary search process, etc. We refer readers to Appendix C for the complexity analysis of QETU when U is implemented using a p th-order Trotter formula, as well as its implication in the ground-state energy estimation.

III. GROUND-STATE ENERGY ESTIMATION AND GROUND-STATE PREPARATION

In this section, we discuss how to estimate the ground-state energy and to prepare the ground state within the QETU framework. The setup of the problems is as follows. We assume that the Hamiltonian H can be accessed through its time-evolution operator e^{-iH} . The goal is (1) to estimate the ground-state energy and (2) to prepare the ground state. For the first task, we assume that we have access to a good initial guess $|\phi_0\rangle$ of the ground state, i.e., $|\langle \phi_0 | \psi_0 \rangle| \geq \gamma$, where $|\psi_0\rangle$ is the ground state. For the second task, we need the additional assumption that the ground-state energy λ_0 is separated from the rest of the spectrum by a gap Δ . These assumptions are stated more formally in the definitions below.

Definition 2: (ground-state energy estimation). Suppose that we are given a Hamiltonian H on n qubits, the spectrum of which is contained in $[\eta, \pi - \eta]$ for some $\eta > 0$. The Hamiltonian can be accessed through a unitary $U = e^{-iH}$. Also suppose that we have an initial guess $|\phi_0\rangle$ of the ground state $|\psi_0\rangle$ satisfying $|\langle \phi_0 | \psi_0 \rangle| \geq \gamma$. This initial guess can be prepared by U_I . The oracles U and U_I are provided as black-box oracles. The goal is to estimate the ground-state energy λ_0 to within additive error ϵ .

Definition 3: (ground-state preparation). Under the same assumptions as in Definition 2 and with the additional

assumption that there is a spectral gap of at least Δ separating the ground-state energy λ_0 from the rest of the spectrum, the goal is to prepare a quantum state $|\tilde{\psi}_0\rangle$ such that $|\langle\psi_0|\tilde{\psi}_0\rangle| \geq 1 - \epsilon$.

We primarily focus on ground-state energy estimation. This is because once we have the ground-state energy, preparing the ground state can be done by applying an approximate projection, which can be directly performed using QETU. We consider two settings: the short-query-depth setting and the near-optimal setting. In the first setting we prioritize lowering the query depth (and hence the circuit depth) and in the second setting we prioritize lowering the query complexity (and hence the total run time). Our results for the two settings are stated in the following theorems.

Theorem 4: (ground-state energy estimation using QETU). *Under the assumptions stated in Definition 2, we can estimate the ground-state energy to within additive error ϵ , with probability at least $1 - \vartheta$, with the following cost:*

- (1) $\tilde{\mathcal{O}}(\epsilon^{-1}\gamma^{-2} \log(\vartheta^{-1}))$ queries to (controlled-) U and $\mathcal{O}(\gamma^{-2} \text{poly} \log(\epsilon^{-1}\vartheta^{-1}))$ queries to U_I
- (2) One ancilla qubit
- (3) $\tilde{\mathcal{O}}(\epsilon^{-1}\gamma^{-2} \log(\vartheta^{-1}))$ additional one-qubit quantum gates
- (4) $\mathcal{O}(\epsilon^{-1} \log(\gamma^{-1}))$ query depth of U

Note that here, using the short-query-depth algorithm, we do not need to use extra two-qubit gates beyond what is needed in (controlled-) U .

Theorem 5: (near-optimal ground-state energy estimation with QETU and improved binary amplitude estimation). *Under the assumptions stated in Definition 2, we can estimate the ground-state energy to within additive error ϵ , with probability at least $1 - \vartheta$, with the following cost:*

- (1) $\tilde{\mathcal{O}}(\epsilon^{-1}\gamma^{-1} \log(\vartheta^{-1}))$ queries to (controlled-) U and $\mathcal{O}(\gamma^{-1} \text{poly} \log(\epsilon^{-1}\vartheta^{-1}))$ queries to U_I
- (2) Three ancilla qubits
- (3) $\tilde{\mathcal{O}}(n\gamma^{-1} \log(\epsilon^{-1}\vartheta^{-1}) + \epsilon^{-1}\gamma^{-1} \log(\vartheta^{-1}))$ additional one- and two-qubit quantum gates
- (4) $\tilde{\mathcal{O}}(\epsilon^{-1}\gamma^{-1} \log(\vartheta^{-1}))$ query depth of U

To the best of our knowledge, this is also the first algorithm that can estimate the ground-state energy with $\tilde{\mathcal{O}}(\gamma^{-1}\epsilon^{-1})$ query complexity using only a constant number of ancilla qubits.

We can see from the two theorems stated above that the trade-off between the query depth and the query complexity, which is also shown in Table I: the short-query-depth algorithm has a $\tilde{\mathcal{O}}(\gamma^{-2})$ dependence on γ , which is suboptimal. This is compensated by the fact that the query depth

is only logarithmic in γ , which can be significantly smaller than that required in the near-optimal algorithm. A similar trade-off exists for the ground-state-preparation algorithms (Theorems 6 and 11), as shown in Table II.

We first discuss in Sec. III A the quantum algorithms to solve these tasks with short query depth. As a side note, when the initial state is indeed an eigenstate of H , Theorem 4 also directly gives rise to a new algorithm for performing QPE using QETU that achieves the Heisenberg-limited precision scaling (see Sec. III B). Finally, assuming access to $(n + 1)$ -bit Toffoli gates, Sec. III C describes the quantum algorithms for solving the ground-state-preparation and energy-estimation problems with near-optimal complexity.

In Secs. III A–III C, we mostly describe the algorithms to solve these tasks and state the results as lemmas and theorems along the way. We believe that this can help readers to grasp the whole picture better. An exception is the proofs of Theorems 4 and 5, which are presented as formal proofs.

A. Algorithms with short query depths

Let us first focus on the ground-state-preparation problem. We first consider a simple setting in which we assume knowledge of a parameter μ such that

$$\lambda_0 \leq \mu - \Delta/2 < \mu + \Delta/2 \leq \lambda_1, \quad (3)$$

where λ_1 is the first-excited-state energy. We need to find a polynomial approximation to the shifted sign function

$$\theta(x - \mu) = \begin{cases} 1, & x \leq \mu, \\ 0, & x > \mu \end{cases}$$

and the polynomial should satisfy the requirement in Theorem 1. To this end, given a number $0 < c < 1$, we would like to find a real polynomial $f(x)$ satisfying

$$\begin{aligned} |f(x) - c| &\leq \epsilon, \quad \forall x \in [\eta, \mu - \Delta/2]; \quad |f(x)| \leq \epsilon, \\ \forall x &\in [\mu + \Delta/2, \pi - \eta]. \end{aligned} \quad (4)$$

As is discussed in Sec. IV, it is preferable to choose c to be slightly smaller than 1 to avoid numerical overshooting. Compared to $c = 1$, this has a negligible effect in practice and does not affect the asymptotic scaling of the algorithm. Taking the cosine transformation in Theorem 1 into account, we need to find a real even polynomial satisfying

$$\begin{aligned} |F(x) - c| &\leq \epsilon, \quad x \in [\sigma_+, \sigma_{\max}]; \quad |F(x)| \leq \epsilon, \\ x &\in [\sigma_{\min}, \sigma_-]; \quad |F(x)| \leq 1, \quad x \in [-1, 1], \end{aligned} \quad (5)$$

where

$$\sigma_{\pm} = \cos \frac{\mu \mp \Delta/2}{2}, \quad \sigma_{\min} = \cos \frac{\pi - \eta}{2}, \quad \sigma_{\max} = \cos \frac{\eta}{2}. \quad (6)$$

Here, we use the fact that $\cos(\cdot)$ is a monotonically decreasing function on $[0, \pi/2]$.

To find such a polynomial $F(x)$, we may use the result in Ref. [53, Corollary 7], which constructs a polynomial of degree $\mathcal{O}(\Delta^{-1} \log \epsilon^{-1})$ for $c = 1$ and any $\mu \in [\eta, \pi - \eta]$. This algorithm first replaces the discontinuous shifted sign function by a continuous approximation using error functions (we need to shift both horizontally and vertically and symmetrize to obtain an even polynomial) and then truncates a polynomial expansion of the resulting smooth function. The construction is specific to the shifted sign function. Its implementation relies on modified Bessel functions of the first kind, which should be carefully treated to ensure numerical stability especially when Δ is small. In Sec. IV, we introduce a simple convex-optimization-based method for generating a near-optimal approximation, which does not rely on any analytic computation. The convex-optimization procedure can be used not only to approximate the shifted sign function but also to find polynomial approximations in a wide range of settings. The process of obtaining the phase factors can also be streamlined using QSPACK [54]. The details of this procedure are described in Sec. IV and an example of the optimal approximate polynomial is given in Fig. 2.

We can then run the QETU circuit to apply $f(H) = F(\cos(H/2))$ to an initial guess $|\phi_0\rangle$. If $|\phi_0\rangle$ has a nonzero component in the direction of the ground state $|\psi_0\rangle$, then $f(H)$ will preserve this component up to a factor $c \approx 1$ but will suppress the orthogonal component by a factor $\epsilon \approx 0$, thus giving us a quantum state close to the ground state. This procedure does not always succeed due to the nonunitary nature of $f(H)$ and consequently we need to repeat it multiple times until we get a success. The number of repetitions needed is $\mathcal{O}(\gamma^{-2} \log(\vartheta^{-1}))$ to guarantee a success probability of at least $1 - \vartheta$. The result is summarized in Theorem 6.

Theorem 6: (ground-state preparation using QETU). *Under the same assumptions as in Definition 3, with the additional assumption that we have μ satisfying Eq. (3), we can prepare the ground state up to fidelity $1 - \epsilon$, with probability at least $2/3$, with the following cost:*

- (1) $\tilde{\mathcal{O}}(\gamma^{-2} \Delta^{-1} \log(\epsilon^{-1}))$ queries to (controlled-) U and $\mathcal{O}(\gamma^{-2})$ queries to U_I
- (2) One ancilla qubit
- (3) $\tilde{\mathcal{O}}(\gamma^{-2} \Delta^{-1} \log(\epsilon^{-1}))$ additional one-qubit quantum gates
- (4) $\mathcal{O}(\Delta^{-1} \log(\epsilon^{-1} \gamma^{-1}))$ maximal query depth of U

Again, using the short-query-depth algorithm, we do not need to use extra two-qubit gates beyond what is needed in (controlled-) U . We can repeat the procedure multiple times to make the success probability exponentially close

to 1. In this algorithm, to be more precise than the $\tilde{\mathcal{O}}$ notation used in Theorem 6, we need $\mathcal{O}(\gamma^{-2} \Delta^{-1} \log(\gamma^{-1} \epsilon^{-1}))$ queries to U . There is a logarithmic dependence on γ^{-1} because we need to account for subnormalization that comes from postselecting measurement results when analyzing the error. The success of the above procedure is flagged by the measurement outcome of the ancilla qubit.

For ground-state energy estimation, our strategy is to adapt the binary search algorithm in Ref. [1, Theorem 8] to the current setting. In order to estimate the ground-state energy with increasing precision, we need to repeatedly solve a decision problem.

Definition 7: (the fuzzy-bisection problem). Under the same assumptions as in Definition 2, we are asked to solve the following problem: output 0 when $\lambda_0 \leq x - h$ and output 1 when $\lambda_0 \geq x + h$.

Here, the fuzziness is in the fact that when $x - h < \lambda_0 < x + h$, we are allowed to output either 0 or 1. This is, in fact, essential for making this problem efficiently solvable. Solution of the fuzzy-bisection problem will enable us to find the ground-state energy through binary search. We discuss the details in the proof of Theorem 4.

To solve the fuzzy-bisection problem, we need a real even polynomial $F(x)$ satisfying the following:

$$\begin{aligned} c - \epsilon' &\leq F(x) \leq c + \epsilon', & x &\in [\cos((x - h)/2), 1] \\ |F(x)| &\leq \epsilon', & x &\in [0, \cos((x + h)/2)]. \end{aligned} \quad (7)$$

For asymptotic analysis, we can use the approximate sign function from Ref. [53, Corollary 6] and the degree of $F(x)$ is $\mathcal{O}(h^{-1} \log(\epsilon'^{-1}))$. With a choice of $F(x)$ that satisfies the above requirements, if $\lambda_0 \geq x + h$, then $\|F(\cos(H/2)) |\phi_0\rangle\| \leq \epsilon'$; if $\lambda_0 \leq x - h$, then

$$\|F(\cos(H/2)) |\phi_0\rangle\| \geq \|F(\cos(\lambda_0/2)) |\phi_0\rangle\| \geq (c - \epsilon')\gamma.$$

Therefore, after choosing $\epsilon' = \gamma c / [2(\gamma + 1)]$, to solve the fuzzy-bisection problem, we only need to distinguish between the following two cases: $\|F(\cos(H/2)) |\phi_0\rangle\| \leq \epsilon' = \gamma c / [2(\gamma + 1)]$ or $\|F(\cos(H/2)) |\phi_0\rangle\| \geq (c - \epsilon')\gamma = (\gamma + 2)\gamma c / [2(\gamma + 1)]$. These two cases are well separated, because

$$\frac{(\gamma + 2)\gamma c}{2(\gamma + 1)} - \frac{\gamma c}{2(\gamma + 1)} = \frac{\gamma c}{2}.$$

Hence these two quantities are separated by a gap of order $\Omega(\gamma)$, which enables us to distinguish between them using a modified version of amplitude estimation, as is discussed later. A block encoding of $F(\cos(H/2))$ can be constructed using QETU, which we denote by U_{proj} :

$$(|0\rangle \otimes I_n) U_{\text{proj}} (|0\rangle \otimes I_n) = F(\cos(H/2)). \quad (8)$$

Because of the estimate of the degree of $F(x)$, U_{proj} here uses $\mathcal{O}(h^{-1} \log(\gamma^{-1}))$ queries to $U = e^{-iH}$. All we need to

do is to distinguish between the following two cases:

$$\begin{aligned} \|(\langle 0| \otimes I)U_{\text{proj}}(I \otimes U_I)(|0\rangle |0^n)\| &\leq \frac{\gamma c}{2(\gamma + 1)} \text{ or} \\ \|(\langle 0| \otimes I)U_{\text{proj}}(I \otimes U_I)(|0\rangle |0^n)\| &\geq \frac{(\gamma + 2)\gamma c}{2(\gamma + 1)}. \end{aligned}$$

This problem can be generalized into the following binary amplitude-estimation problem.

Definition 8: (binary amplitude estimation). Let W be a unitary acting on two registers (one with one qubit and the other with n qubits), with the first register indicating success or failure. Let $A = \|(\langle 0| \otimes I_n)W(|0\rangle |0^n)\|$ be the success amplitude. Given $0 \leq \gamma_1 < \gamma_2$, provided that A is either smaller than γ_1 or greater than γ_2 , we want to correctly distinguish between the two cases, i.e., output 0 for the former and 1 for the latter.

In the context of the fuzzy-bisection problem in Definition 7, we need to choose $W = U_{\text{proj}}(I \otimes U_I)$, $\gamma_1 = \gamma c/[2(\gamma + 1)]$, $\gamma_2 = (\gamma + 2)\gamma c/[2(\gamma + 1)]$. Note that $\gamma_2/\gamma_1 = \gamma + 2 \geq 2$ and therefore henceforth we only consider the case where for some constant c' we have $\gamma_2/\gamma_1 \geq c'$.

Now we can use Monte Carlo sampling to estimate $A = \|(\langle 0| \otimes I_n)W(|0\rangle |0^n)\|$. We estimate how many samples are needed to distinguish whether $A \geq \gamma_2$ or $A \leq \gamma_1$. We implement $W|0\rangle |0^n\rangle$ and measure the first qubit and the output is a random variable, taking a value in $\{0, 1\}$, following the Bernoulli distribution and its expectation value is $1 - A^2$. We denote $p_1 = 1 - \gamma_1^2$ and $p_2 = 1 - \gamma_2^2$. We generate N_s samples and check whether the average is larger than $p_{1/2} = (p_1 + p_2)/2$ (in which case we choose to believe that $A \leq \gamma_1$) or smaller than $p_{1/2}$ (in which case we choose to believe that $A \geq \gamma_2$). By the Chernoff-Hoeffding theorem, the error probability is upper bounded by

$$\max \{e^{-D(p_{1/2}||p_1)N_s}, e^{-D(p_{1/2}||p_2)N_s}\}, \quad (9)$$

where

$$D(x||y) = x \log(x/y) + (1 - x) \log[(1 - x)/(1 - y)]$$

is the Kullback-Leibler divergence between Bernoulli distributions. Direct calculation, using the fact that $\gamma_2 \geq c'\gamma_1$, shows that $D(p_{1/2}||p_1), D(p_{1/2}||p_2) \geq \Omega(\gamma_1^2)$. Therefore, to ensure that the error probability is below ϑ' , we only need to choose N_s such that $e^{-\Omega(\gamma_1^2 N_s)} \leq \vartheta'$. Thus we obtain the scaling of $N_s = \mathcal{O}(\gamma_1^{-2} \log(\vartheta'^{-1}))$. From the above analysis, we have the following lemma.

Lemma 9: (Monte Carlo method for solving binary estimation). *The binary amplitude-estimation problem in Definition 8, with the additional assumption that there exists constant $c' > 0$ such that $\gamma_2/\gamma_1 \geq c'$, can be solved*

correctly with probability at least $1 - \vartheta'$ by querying W $\mathcal{O}(\gamma_1^{-2} \log(\vartheta'^{-1}))$ times and this procedure does not require additional ancilla qubits besides the ancilla qubits already required in W . The maximal query depth of W is $\mathcal{O}(1)$.

Lemma 9 enables us to solve the fuzzy-bisection problem stated in Definition 7. With the tools introduced above, we can now prove Theorem 4.

Proof of Theorem 4. We solve the ground-state energy-estimation problem by performing a binary search and at each search step we need to solve a fuzzy-bisection problem, which we know can be done from the above discussion. Below, we discuss how the binary search works, i.e., why repeatedly solving the fuzzy-bisection problem can help us find the ground-state energy. For simplicity of the discussion, we set $\eta = \pi/4$ in Definition 2, i.e., the spectrum of the Hamiltonian is contained in the interval $[\pi/4, 3\pi/4]$. In each iteration of the binary search, we have l and r such that $l \leq \lambda_0 \leq r$. In the first iteration, we choose $l = \pi/4$ and $r = 3\pi/4$. With l and r , we want to solve the following fuzzy-bisection problem: output 0 when $\lambda_0 \leq (2l + r)/3$ and output 1 when $\lambda_0 \geq (l + 2r)/3$. In other words, we let $x = (l + r)/2$ and $h = (r - l)/6$ in Definition 7. After solving this fuzzy-bisection problem, if the output is 0, then we know that $l \leq \lambda_0 \leq (l + 2r)/3$ and therefore we can update r to be $(l + 2r)/3$. Similarly, if the output is 1, we can update l to be $(2l + r)/3$. In this way, we obtain a new pair of l and r such that $l \leq \lambda_0 \leq r$ and $r - l$ shrinks by $2/3$.

The values of l and r will converge to λ_0 from both sides and therefore when $r - l \leq 2\epsilon$, λ_0 will be within ϵ distance from $(l + r)/2$, thus giving us the ground-state energy estimate that we want. This will take $\lceil \log_{3/2}(\pi\epsilon^{-1}/2) \rceil$ iterations, because $r - l$ shrinks by a factor $2/3$ in each iteration and the initial value is $\pi/2$.

In our context, we need to perform a binary amplitude estimation at each of the $\mathcal{O}(\log(\epsilon^{-1}))$ steps to solve the fuzzy-bisection problem and therefore to ensure a final success probability of at least ϑ we need to choose $\vartheta' = \Theta(\vartheta/\log(\epsilon^{-1}))$ in Lemma 9. Since $\gamma_1 = \gamma c/[2(\gamma + 1)] = \Omega(\gamma)$, as discussed immediately after we introduce Definition 8, each time we solve the binary amplitude-estimation problem we need to use $W = U_{\text{proj}}(I \otimes U_I)$ for $\mathcal{O}(\gamma^{-2}(\log(\vartheta'^{-1}) + \log \log(\epsilon^{-1})))$ times. Note that each U_{proj} requires using $U = e^{-iH}$ for $\mathcal{O}(h^{-1} \log(\gamma^{-1}))$ times. Adding up for h that decreases exponentially until it is of order ϵ , we can obtain the estimate for the cost of estimating the ground-state energy, as stated in Theorem 4. ■

B. QPE revisited

As an application of the ground-state energy-estimation algorithm using QETU, let us revisit the task of

performing QPE. Assuming access to a unitary $U = e^{-iH}$ and an eigenstate $|\psi\rangle$ such that $U|\psi\rangle = e^{-i\lambda}|\psi\rangle$, the goal of the phase estimation is to estimate λ to precision ϵ . This can be viewed as the ground-state energy-estimation problem with an initial overlap $\gamma = 1$. Although λ may not be the ground-state energy of H , other eigenvalues of H do not matter because the initial state has zero overlap with other eigenstates.

In order to estimate λ , we can repeatedly solve the fuzzy-bisection problem in Definition 7 and this gives us an algorithm that is essentially identical to the one described in the proof of Theorem 4. As a corollary to Theorem 4, the phase-estimation problem can be solved with the following cost that achieves the Heisenberg limit.

Corollary 10: (phase estimation). *Suppose that we are given $U = e^{-iH}$, a quantum state $|\psi\rangle$ such that $U|\psi\rangle = e^{-i\lambda}|\psi\rangle$, where $\lambda \in [\eta, \pi - \eta]$ for some constant $\eta > 0$. We can estimate λ to precision ϵ with probability at least $1 - \vartheta$ using $\tilde{\mathcal{O}}(\epsilon^{-1} \log(\vartheta^{-1}))$ applications to (controlled) U and its inverse, a single copy of $|\psi\rangle$, and $\mathcal{O}(\epsilon^{-1}(\log(\vartheta^{-1}) + \log \log(\epsilon^{-1})))$ additional one qubit gates.*

It may require some explanation as to why we only need a single copy of $|\psi\rangle$ rather than repeatedly apply a circuit that prepares $|\psi\rangle$. This is because $|\psi\rangle$ is an eigenstate and consequently will be preserved up to a phase factor in the circuit depicted in Figure 1(c). Therefore, it can be reused throughout the algorithm and there is no need to prepare it more than once.

C. Algorithms with near-optimal query complexities

With a block-encoding input model, Theorems 6 and 8 in Ref. [1] are near-optimal algorithms for preparing the ground state and for estimating the ground-state energy, respectively. In this section, we combine QETU with amplitude amplification and a new binary amplitude-estimation method to yield quantum algorithms with the same near-optimal query complexities. For amplitude estimation, we avoid using quantum Fourier transform, as it would require an additional register of qubits. Instead, we use a procedure described in Appendix D based on QETU. Unlike previous near-term methods for amplitude estimation [26,55] that typically rely on Bayesian inference techniques and thus require knowledge of a prior distribution, our method does not require such prior knowledge. One could also adapt the QFT-free approximate counting algorithms in Refs. [56,57] to the amplitude-estimation problem but our approach in Appendix D is better tailored for the QETU framework.

We need to use amplitude amplification to quadratically improve the γ dependence in Theorem 6 and to achieve the near-optimal query complexity for preparing the ground

state in Definition 3 (assuming knowledge of μ). Let us first study the number of ancilla qubits needed for this task. For amplitude amplification, we need to construct a reflection operator around the initial guess $|\phi_0\rangle$. This requires implementing $2|0^n\rangle\langle 0^n| - I$, which is equivalent, using phase kickback and Pauli- X gates, to implementing an $(n + 1)$ -bit Toffoli gate:

$$|1^n\rangle\langle 1^n| \otimes \sigma_x + (I - |1^n\rangle\langle 1^n|) \otimes I.$$

This $(n + 1)$ -bit Toffoli gate can be implemented using $\mathcal{O}(n)$ elementary one- or two-qubit gates, on $n + 2$ qubits [40, Corollary 7.4]. Note that this can be a relatively costly operation on early fault-tolerant quantum devices. We need two ancilla qubits to implement the reflection operator but one of them can be reused for other purposes. The reason is as follows: one ancilla qubit is the one that σ_x acts on conditionally in the $(n + 1)$ -bit Toffoli gate. This one cannot be reused because it needs to start from $|0\rangle$ and will be returned to $|0\rangle$. The other qubit, however, can start from any state and will be returned to the original state, as discussed in Ref. [40, Corollary 7.4], and therefore we can use any qubit in the circuit for this task, except for the $n + 1$ qubits already involved in the $(n + 1)$ -bit Toffoli gate. Note that in Theorem 6 we have one ancilla qubit that is used for QETU. This qubit can therefore serve as the ancilla qubit needed in implementing the $(n + 1)$ -bit Toffoli gate. Thus we only need two ancilla qubits in the whole procedure.

To summarize the cost, we have the following theorem.

Theorem 11: (near-optimal ground-state preparation with QETU and amplitude amplification). *Under the same assumptions as in Definition 3, with the additional assumption that we have μ satisfying Eq. (3), we can prepare the ground state, with probability $2/3$, up to fidelity $1 - \epsilon$ with the following cost:*

- (1) $\tilde{\mathcal{O}}(\gamma^{-1} \Delta^{-1} \log(\epsilon^{-1}))$ queries to (controlled-) U and $\mathcal{O}(\gamma^{-1})$ queries to U_I
- (2) Two ancilla qubits
- (3) $\tilde{\mathcal{O}}(n\gamma^{-1} \Delta^{-1} \log(\epsilon^{-1}))$ additional one- and two-qubit quantum gates
- (4) $\tilde{\mathcal{O}}(\gamma^{-1} \Delta^{-1} \log(\epsilon^{-1}))$ query depth for U

Note that we can repeat this procedure multiple times to make the success probability exponentially close to 1.

For ground-state energy estimation, in the algorithm described in the proof of Theorem 4, our short-query-depth algorithm has a $\tilde{\mathcal{O}}(\gamma^{-2})$ scaling because of the Monte Carlo sampling in the binary amplitude estimation (Definition 8) procedure. Here, we improve the scaling to $\tilde{\mathcal{O}}(\gamma^{-1})$ using the technique developed in Ref. [1, Lemma 7]. The technique in Ref. [1, Lemma 7] uses phase estimation and requires $\mathcal{O}(\log((\gamma_2 - \gamma_1)^{-1})) = \mathcal{O}(\log(\gamma^{-1}))$ ancilla qubits. In Appendix D, we propose a method

to solve the binary amplitude-estimation problem using QETU, which reduces the number of additional ancilla qubits down from $\mathcal{O}(\log(\gamma^{-1}))$ to two. The result is summarized here.

Lemma 12: (binary amplitude estimation). *The binary amplitude-estimation problem in Definition 8 can be solved correctly with probability at least $1 - \vartheta'$ by querying $W \mathcal{O}((\gamma_2 - \gamma_1)^{-1} \log(\vartheta'^{-1}))$ times and this procedure requires one additional ancilla qubit (besides the ancilla qubits already required in W).*

The key idea is to treat the walk operator in amplitude estimation as a time-evolution operator corresponding to a Hamiltonian and this allows us to apply QETU to extract information about that Hamiltonian.

As a result of this new method to solve the binary amplitude-estimation problem, the ground-state energy-estimation problem can now be solved using only three ancilla qubits: one for QETU and two others for binary amplitude estimation.

With these results, we can now analyze the cost of ground-state energy estimation in the near-optimal setting and thereby prove Theorem 5.

Proof of Theorem 5. We adopt the same strategy of performing a binary search to locate the ground-state energy, as used in Theorem 4. The main difference is that instead of using Monte Carlo sampling to solve the binary amplitude-estimation problem, we now use QETU to do so, with the complexity stated in Lemma 12.

We now count how many times we need to query $U = e^{-iH}$ in this approach. Each time we perform binary amplitude estimation, we need to use U_{proj} for $\mathcal{O}(\gamma^{-1} \log(\vartheta'^{-1}))$ times to have at least $1 - \vartheta'$ success probability each time. We need to perform binary amplitude estimation for each step of the binary search and there are in total $\mathcal{O}(\log(\epsilon^{-1}))$ steps. Therefore, to ensure a final success probability of at least $1 - \vartheta$, we need to choose $\vartheta' = \Theta(\vartheta / \log(\epsilon^{-1}))$. At the k th binary search step, U_{proj} uses $U = e^{-iH}$ for $\mathcal{O}((3/2)^k \log(\gamma^{-1}))$ times. Therefore, in total we need to query U , up to a constant factor

$$\begin{aligned} & \sum_{k=0}^{\lceil \log_{3/2}(\pi \epsilon^{-1/2}) \rceil} (3/2)^k \log(\gamma^{-1}) \gamma^{-1} \log(\vartheta'^{-1}) \\ &= \mathcal{O}\left(\epsilon^{-1} \gamma^{-1} \log(\gamma^{-1}) (\log(\vartheta^{-1}) + \log \log(\epsilon^{-1}))\right) \end{aligned}$$

times. This query complexity agrees with that in Ref. [1, Theorem 8] up to a logarithmic factor.

When we count the number of additional quantum gates needed, there will be an n dependence, which comes from the fact that we need to implement a reflection operator $2|0^{n+1}\rangle\langle 0^{n+1}| - I$ each time we implement U_{proj} in the binary amplitude-estimation procedure (see Appendix D).

These reflection operators require $\tilde{\mathcal{O}}(n\gamma^{-1} \log(\epsilon^{-1} \vartheta^{-1}))$ gates. We also need $\mathcal{O}(\epsilon^{-1} \gamma^{-1} \log(\vartheta^{-1}))$ additional quantum gates that come from implementing QETU. Combining the two numbers, we obtain the number of gates as shown in the theorem. ■

When preparing the ground state, the parameter μ in Theorem 11 is generally not known *a priori*. For exactly solvable models or small quantum systems, despite the ability to simulate them classically, one might still want to prepare the ground state and in such cases μ is available. In the general case, to prepare the ground state without knowing a parameter μ as in Theorem 11, we can first estimate the ground-state energy to within additive error $\mathcal{O}(\Delta)$, and then run the algorithm in Theorem 11. This results in an algorithm with the following costs.

Theorem 13: *Under the assumptions stated in Definition 3, we can prepare the ground state to fidelity at least $1 - \epsilon$, with probability at least $1 - \vartheta$, with the following cost:*

- (1) $\tilde{\mathcal{O}}(\Delta^{-1} \gamma^{-1} \text{poly log}(\epsilon^{-1} \vartheta^{-1}))$ queries to (controlled-) U and $\mathcal{O}(\gamma^{-1} \text{poly log}(\Delta^{-1} \epsilon^{-1} \vartheta^{-1}))$ queries to U_I
- (2) Three ancilla qubits
- (3) $\tilde{\mathcal{O}}(n\gamma^{-1} \log(\Delta^{-1} \epsilon^{-1} \vartheta^{-1}) + \Delta^{-1} \gamma^{-1} \log(\epsilon^{-1} \vartheta^{-1}))$ additional one- and two-qubit quantum gates
- (4) $\tilde{\mathcal{O}}(\Delta^{-1} \gamma^{-1} \text{poly log}(\epsilon^{-1} \vartheta^{-1}))$ query depth of U

IV. CONVEX-OPTIMIZATION-BASED METHOD FOR CONSTRUCTING APPROXIMATING POLYNOMIALS

To approximate an even target function using an even polynomial of degree d , we can express the target polynomial as the linear combination of Chebyshev polynomials with some unknown coefficients $\{c_k\}$:

$$F(x) = \sum_{k=0}^{d/2} T_{2k}(x) c_k. \quad (10)$$

To formulate this as a discrete optimization problem, we first discretize $[-1, 1]$ using M grid points [e.g., roots of Chebyshev polynomials $\{x_j = -\cos(j\pi/M - 1)\}_{j=0}^{M-1}$]. One can also restrict them to $[0, 1]$ due to symmetry. We define the coefficient matrix, $A_{jk} = T_{2k}(x_j)$, $k = 0, \dots, d/2$. Then, the coefficients for approximating the shifted sign function can be found by solving the following optimization problem:

$$\begin{aligned} & \min_{\{c_k\}} \max \left\{ \max_{x_j \in [\sigma_{\max}, \sigma_+]} |F(x_j) - c|, \max_{x_j \in [\sigma_{\min}, \sigma_-]} |F(x_j)| \right\} \\ & \text{s.t. } F(x_j) = \sum_k A_{jk} c_k, \quad |F(x_j)| \leq c, \\ & \forall j = 0, \dots, M-1. \end{aligned} \quad (11)$$

This is a convex-optimization problem and it can be solved using software packages such as CVX [58]. The norm constraint $|F(x)| \leq 1$ is relaxed to $|F(x_j)| \leq c$ to take into account that the constraint can only be imposed on the sampled points and the values of $|F(x)|$ may slightly overshoot on $[-1, 1] \setminus \{x_j\}_{j=0}^{M-1}$. The effect of this relaxation is negligible in practice and we can choose c to be sufficiently close to 1 (for instance, c can be 0.999). Since Eq. (11) approximately solves a min-max problem, it achieves the near-optimal solution (in the sense of the L^∞ norm) by definition both in the asymptotic and preasymptotic regimes.

Once the polynomial $F(x)$ is given, the Chebyshev coefficients can be used as the input to find the symmetric phase factors using an optimization-based method. Due to the parity constraint, the number of degrees of freedom in the target polynomial $F(x)$ is $\tilde{d} := \lceil (d+1)/2 \rceil$. Hence $F(x)$ is entirely determined by its values on \tilde{d} distinct points. We may choose these points to be $x_k = \cos((2k-1)/4\tilde{d}\pi)$, $k = 1, \dots, \tilde{d}$, which are the positive nodes of the Chebyshev polynomial $T_{2\tilde{d}}(x)$. The problem of finding the symmetric phase factors can be equivalently solved via the following optimization problem:

$$\begin{aligned} \Phi^* &= \arg \min_{\substack{\Phi \in [-\pi, \pi]^{d+1}, \\ \text{symmetric.}}} \mathcal{F}(\Phi), \\ \mathcal{F}(\Phi) &:= \frac{1}{\tilde{d}} \sum_{k=1}^{\tilde{d}} |g(x_k, \Phi) - F(x_k)|^2, \end{aligned} \quad (12)$$

where

$$\begin{aligned} g(x, \Phi) &:= \text{Re}[\langle 0 | e^{i\phi_0 Z} e^{i \arccos(x) X} e^{i\phi_1 Z} e^{i \arccos(x) X} \\ &\dots e^{i\phi_{d-1} Z} e^{i \arccos(x) X} e^{i\phi_d Z} | 0 \rangle]. \end{aligned}$$

The desired phase factor achieves the global minimum of the cost function with $\mathcal{F}(\Phi^*) = 0$. It has been found that a quasi-Newton method to solve Eq. (12) with a particular symmetric initial guess

$$\Phi^0 = (\pi/4, 0, 0, \dots, 0, 0, \pi/4) \quad (13)$$

can robustly find the symmetric phase factors. Although the optimization problem is highly nonlinear, the success of the optimization-based algorithm can be explained in terms of the strongly convex energy landscape near Φ^0 [59]. The numerical results indicate that on a laptop computer, CVX can find the near-optimal polynomials for $d \sim 5000$. Given the target polynomial, the optimization-based algorithm can find phase factors for $d \sim 10\,000$ [54]. This should be more than sufficient for most QSP-based applications on early fault-tolerant quantum computers. The streamlined process of finding near-optimal polynomials and the associated phase factors has been implemented in QSPPACK [60].

As an illustrative example of the numerically optimized min-max polynomials, we set $\eta = 0.1, \mu = 1.0, \Delta = 0.4, M = 400, c = 0.999$. This corresponds to $\sigma_{\min} = 0.0500, \sigma_- = 0.8253, \sigma_+ = 0.9211, \sigma_{\max} = 0.9997$. The resulting polynomial and the pointwise errors with $d = 20$ and $d = 80$ are shown in Fig. 2. We remark that the polynomial is reconstructed by means of the numerically optimized phase factors using QSPPACK and hence takes the error in the entire process into account. We find that the pointwise error of the polynomial approximation satisfies the equioscillation property on each of the intervals $[\sigma_{\min}, \sigma_-], [\sigma_+, \sigma_{\max}]$. This resembles the Chebyshev equioscillation theorem of the best polynomial approximation on a single interval (see, e.g., Ref. [61, Chapter 10]). Figure 3 shows that the maximum pointwise error on the desired intervals converges exponentially with the increase of the polynomial degree.

More generally, to find a min-max polynomial approximation to a general even target function $h(x)$ on a set $\mathcal{I} \subseteq [-1, 1]$ satisfying $|h(x)| \leq c < 1, x \in \mathcal{I}$, we may solve the optimization problem

$$\begin{aligned} \min_{\{c_k\}} \max_{x_j \in \mathcal{I}} |F(x_j) - h(x_j)| \\ \text{s.t. } F(x_j) &= \sum_k A_{jk} c_k, \quad |F(x_j)| \leq c, \\ &\forall j = 0, \dots, M-1. \end{aligned} \quad (14)$$

We also remark that even though QETU only concerns even polynomials, the same strategy can be applied if the target function h is odd or does not have a definite parity.

V. NUMERICAL COMPARISON WITH QPE FOR GROUND-STATE ENERGY ESTIMATION

In this section, we compare the numerical performance of our ground-state energy-estimation algorithm in Theorem 4 with the QPE algorithm implemented using semiclassical Fourier transform [47] to save the number of ancilla qubits, as done in Refs. [44,45]. We evaluate how many queries to $U = e^{-iH}$ are needed in both algorithms to reach the target accuracy $\epsilon \leq 10^{-3}$. The Hamiltonian H used here has a randomly generated spectrum and is a 200×200 matrix. The initial state $|\phi_0\rangle$ is guaranteed to satisfy $|\langle \phi_0 | \psi_0 \rangle| \geq \gamma$ with a tunable value of γ . In our algorithm, the number of queries is counted by adding up the degrees of all the polynomials we need to implement using QETU.

Figure 4 shows that to achieve comparable accuracy, our algorithm uses significantly fewer queries than quantum phase estimation, in terms of both the asymptotic scaling (improves from γ^{-4} to γ^{-2}) as well as the actual number of queries for moderately small values of γ . In Fig. 4, the error of our method is computed by running the algorithm in Theorem 4 on a classical computer and comparing the

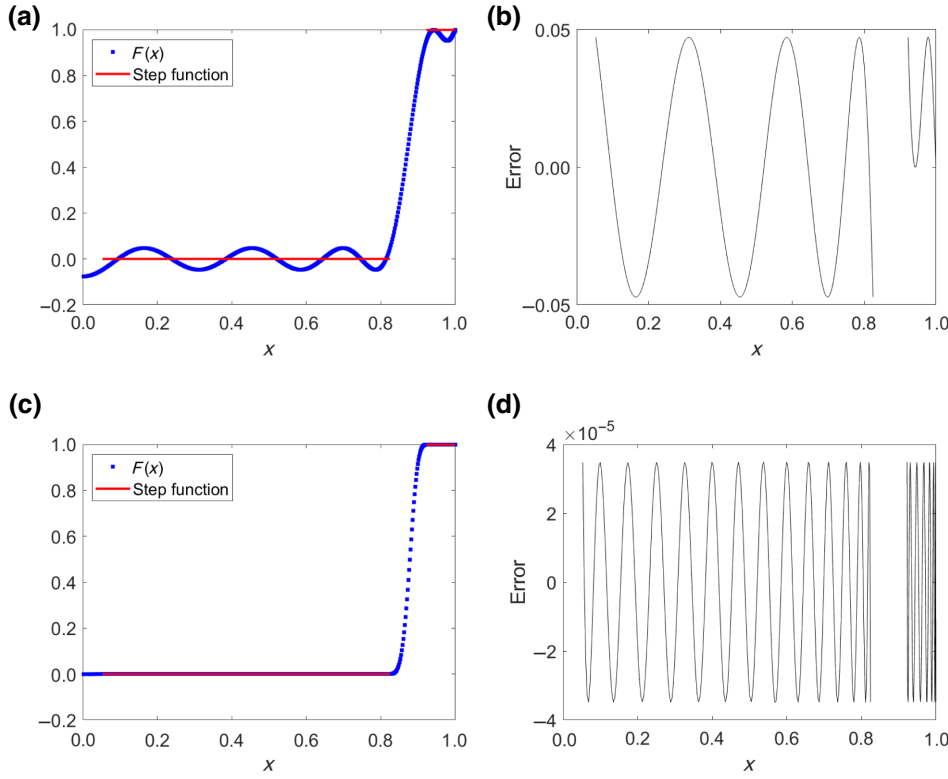


FIG. 2. The polynomial obtained by convex optimization approximating the shifted sign function and the pointwise error on $[\sigma_{\min}, \sigma_-] \cup [\sigma_+, \sigma_{\max}]$ with $d = 20$ (a),(b) and $d = 80$ (c),(d).

output with the exact ground-state energy and we show in the figure the mean of the absolute error in multiple trials. In our method, we need to determine the polynomial degree needed for each binary search step (or each time we solve the fuzzy-bisection problem in Definition 7). This polynomial degree is determined by running the algorithm in Sec. IV and selecting the smallest degree that provides an error below the target accuracy. In our numerical tests, we require the approximation error to be below 10^{-3} , so




FIG. 3. The exponential convergence of the maximum pointwise error on $[\sigma_{\min}, \sigma_-] \cup [\sigma_+, \sigma_{\max}]$ with respect to the increase of the polynomial degree obtained by the convex-optimization method.

that it is much smaller than the squared overlap. The error of QPE is computed by sampling from the exact energy-measurement output distribution, which is again simulated on a classical computer, and comparing the output with the exact ground-state energy. We also compute the absolute error for QPE in multiple trials and take the mean in Figure 4. The mean absolute errors in Fig. 4(b) show that the advantage of our algorithm does not come from a loose error estimate for QPE, since our algorithm reaches the target precision ($\epsilon = 5 \times 10^{-4}$ in this case) consistently and QPE does not achieve a higher precision than our algorithm.

VI. CONTROL-FREE IMPLEMENTATION OF QUANTUM SPIN MODELS

In this section, we demonstrate that for certain quantum spin models, the QETU circuit can be simplified without the need of accessing the controlled Hamiltonian evolution.

Consider a Hamiltonian H that is a linear combination of $\text{poly}(n)$ terms of Pauli operators. Note that two Pauli operators either commute or anticommute. Hence for each term in the Hamiltonian, we can easily find another Pauli operator K that anticommutes with this term. More generally, we assume that H admits a grouping

$$H = \sum_{j=1}^{\ell} H^{(j)}, \quad H^{(j)} = \sum_{s=1}^{d_j} h_s^{(j)}, \quad (15)$$

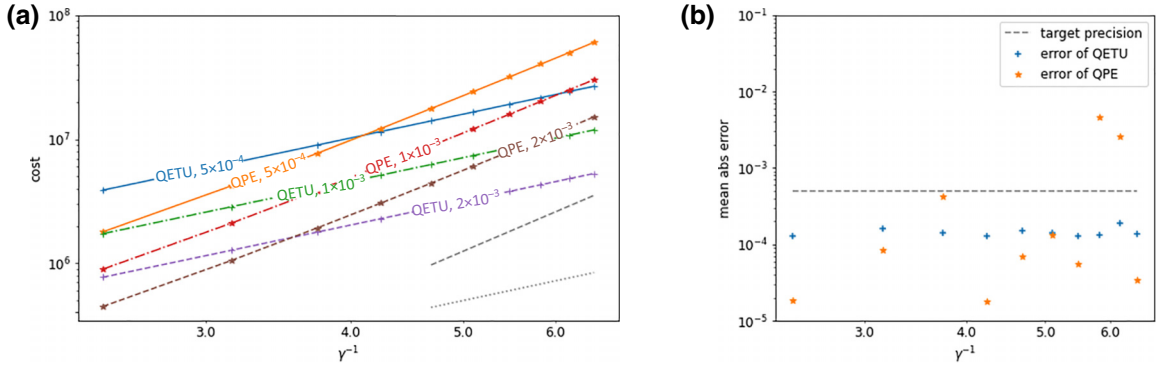


FIG. 4. A comparison of the performance of the algorithm in Theorem 4 and the single-ancilla qubit quantum phase estimation using a semiclassical Fourier transform [44,45,47]. (a) The number of queries of U needed to reach target precision $\epsilon = 5 \times 10^{-4}$, 10^{-3} , and 2×10^{-3} for different values of γ . The gray dashed line and dotted line show γ^{-4} scaling and γ^{-2} scaling, respectively. Both axes are in logarithmic scale. (b) The mean absolute error achieved by the two algorithms for target accuracy $\epsilon = 5 \times 10^{-4}$.

where each $h_s^{(j)}$ is a weighted Pauli operator. For each j , we assume that there exists a single Pauli operator K_j that anticommutes with $H^{(j)}$, i.e., $K_j H^{(j)} K_j = -H^{(j)}$. The number of groups ℓ is poly(n) in the worst case but for many Hamiltonians in practice, ℓ may be much smaller. For example, it may be upper bounded by a constant (see examples below).

Conjugation of the Pauli string on the time-evolution operator flips the sign of the evolution time, i.e., $K_j e^{-i\tau H^{(j)}} K_j = e^{i\tau H^{(j)}}$. Since the time evolution of each Hamiltonian component is the building block of the Trotter splitting algorithm, the time flipping gives a simple implementation of the controlled time evolution without controlling the Hamiltonian components or their time evolution. To implement the controlled time evolution, it suffices to conjugate the circuit implementing the time evolution with the corresponding controlled Pauli string. Suppose that $W_j(\tau) \approx U_j(\tau) := e^{-i\tau H^{(j)}}$ is the quantum circuit approximately implementing the time evolution using Trotter splitting. Then, $K_j W_j(\tau) K_j = W_j(-\tau)$ approximates the reversed time evolution using the same splitting algorithm. This allows us to use Corollary 17 (see Appendix B), which simplifies the circuit implementation of QETU.

To illustrate the control-free implementation, let us consider the transverse-field Ising model (TFIM), for instance, the Hamiltonian of which takes the form

$$H_{\text{TFIM}} = - \underbrace{\sum_{j=1}^{n-1} Z_j Z_{j+1}}_{H_{\text{TFIM}}^{(1)}} - g \underbrace{\sum_{j=1}^n X_j}_{H_{\text{TFIM}}^{(2)}}. \quad (16)$$

Here, $g > 0$ is the coupling constant. Note that a Pauli string

$$K := Y_1 \otimes Z_2 \otimes Y_3 \otimes Z_4 \otimes \dots \quad (17)$$

anticommutes with both components of the Hamiltonian, namely $K H_{\text{TFIM}}^{(j)} K = -H_{\text{TFIM}}^{(j)}$ for $j = 1$ and 2 . Therefore, conjugation of the Pauli string K on the time-evolution operator flips the sign of the evolution time, i.e., $K e^{-i\tau H_{\text{TFIM}}^{(j)}} K = e^{i\tau H_{\text{TFIM}}^{(j)}}$ for $j = 1$ and 2 . In the sense of Eq. (15), we have $\ell = 1$. As a consequence, for the TFIM, Fig. 1(c) is equivalent to the circuit in Fig. 5 in which the controlled time evolution is implicitly implemented by inserting controlled Pauli strings.

It should be noted that the controlled Pauli string only requires implementation of controlled single-qubit gates, rather than the controlled two-qubit gates of the form $e^{-i\tau Z_j Z_{j+1}}$ (note that the Hamiltonian involves two-qubit terms of the form $Z_j Z_{j+1}$). If the quantum circuit conceptually queries the controlled time evolution d times, the simplified circuit only inserts $2d$ controlled Pauli strings in the circuit. In this case, when implementing $W(\frac{1}{2})$ using several Trotter steps, the controlled Pauli strings only need to be inserted before and after each $W(\frac{1}{2})$ but not between the Trotter layers. This simplified implementation gives the quantum circuit in Fig. 5(b). Note that in the general case where controlled Pauli strings are inserted in between Trotter layers, the number of controlled Pauli strings required for the implementation is $\mathcal{O}(d\ell r)$, where r is the number of Trotter steps to implement each time-evolution operator. The simplified quantum circuit in Fig. 5(b) only uses $2d$ controlled Pauli strings. Therefore, this simplified implementation significantly reduces the cost when the number of Trotter step r is large.

In contrast to the TFIM, in which all Hamiltonian components share the same anticommuting Pauli

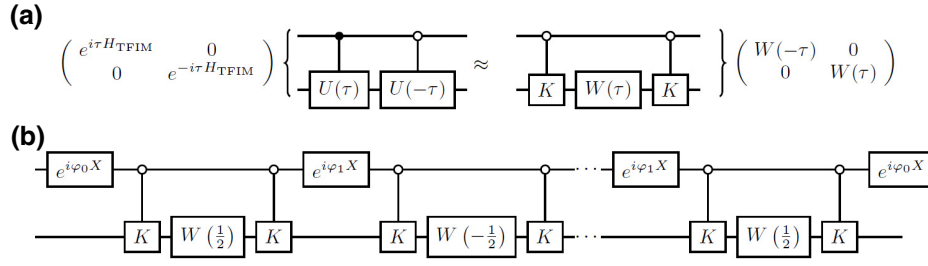


FIG. 5. The simplified quantum circuit for simulating the TFIM Hamiltonian using QETU. (a) The controlled time evolution of the TFIM Hamiltonian without directly controlling the Hamiltonian. (b) The simplified circuit for implementing QETU. The evolution operator $W(-\frac{1}{2})$ can also be implemented as $W(\frac{1}{2})$ conjugated by Pauli X operators (see Fig. 7).

string, the control-free implementation of a general spin Hamiltonian may require more Pauli strings. For example, the Hamiltonian of the Heisenberg model takes the form

$$H_{\text{Heisenberg}} = - \underbrace{\sum_{j=1}^{n-1} J_x X_j X_{j+1} - \sum_{j=1}^{n-1} J_y Y_j Y_{j+1}}_{H_{\text{Heisenberg}}^{(1)}} - \underbrace{\sum_{j=1}^{n-1} J_z Z_j Z_{j+1}}_{H_{\text{Heisenberg}}^{(2)}}.$$

Let us consider two Pauli strings, $K_1 = Z_1 \otimes I_2 \otimes Z_3 \otimes I_4 \otimes \dots$ and $K_2 = X_1 \otimes I_2 \otimes X_3 \otimes I_4 \otimes \dots$. Then, we have the anticommutation relations $K_1 H_{\text{Heisenberg}}^{(1)} K_1 = -H_{\text{Heisenberg}}^{(1)}$ and $K_2 H_{\text{Heisenberg}}^{(2)} K_2 = -H_{\text{Heisenberg}}^{(2)}$. Therefore,

conjugating each basic time evolution $e^{-itH_{\text{Heisenberg}}^{(j)}}$ $j = 1, 2$ by controlled K_1 or K_2 , respectively, we can implement the controlled time evolution without directly controlling Hamiltonians or the corresponding time-evolution operators. This corresponds to $\ell = 2$ in Eq. (15). Unlike the implementation for the TFIM, the controlled Pauli strings cannot be canceled between each Trotter layer. Therefore, the simulation of a Heisenberg model requires additional controlled Pauli gates compared to that in the TFIM simulation.

Other types of quantum Hamiltonians may also be mapped to spin Hamiltonians to perform control-free time evolution using the anticommutation relation. Consider the one-dimensional Fermi-Hubbard model of interacting fermions in a lattice:

$$H_{\text{FH}} = -\mu \sum_{j=1}^n \sum_{\sigma \in \{\uparrow, \downarrow\}} c_{j,\sigma}^\dagger c_{j,\sigma} + u \sum_{j=1}^n c_{j,\uparrow}^\dagger c_{j,\uparrow} c_{j,\downarrow}^\dagger c_{j,\downarrow} - t \sum_{j=1}^{n-1} \sum_{\sigma \in \{\uparrow, \downarrow\}} \left(c_{j,\sigma}^\dagger c_{j+1,\sigma} + c_{j+1,\sigma}^\dagger c_{j,\sigma} \right),$$

where $c_{j,\sigma}^\dagger$ and $c_{j,\sigma}$ ($\sigma \in \{\uparrow, \downarrow\} = \{0, 1\}$) are creation and annihilation operators for different fermionic modes, μ is the chemical potential, u is the on-site Coulomb repulsion energy, and t is the hopping energy. The equivalent spin Hamiltonian can be derived by applying a Jordan-Wigner transformation (see, e.g., Ref. [62]), which gives (up to a global constant)

$$H_{\text{FH,qubits}} = \underbrace{\frac{1}{2} \left(\frac{1}{2} u - \mu \right) \sum_{j=1}^n \sum_{\sigma \in \{0,1\}} Z_{j,\sigma}}_{H_{\text{FH,qubits}}^{(1)}} + \underbrace{\frac{1}{4} u \sum_{j=1}^n Z_{j,0} Z_{j,1}}_{H_{\text{FH,qubits}}^{(2)}} - t \underbrace{\sum_{j=1}^{n-1} \sum_{\sigma \in \{0,1\}} \left(\Sigma_{j,\sigma}^+ \Sigma_{j+1,\sigma}^- + \Sigma_{j+1,\sigma}^+ \Sigma_{j,\sigma}^- \right)}_{H_{\text{FH,qubits}}^{(3)}},$$

where the subscript (j, σ) denotes the j th qubit on the σ th chain, and $\Sigma_{j,\sigma}^\pm := X_{j,\sigma} \pm iY_{j,\sigma}$. Let $K_1 = \otimes_{j,\sigma} X_{j,\sigma}$, $K_2 = \left(\otimes_{j=1}^n X_{j,0} \right) \otimes \left(\otimes_{j=1}^n I_{j,1} \right)$, and $K_3 = \left(\otimes_{j=\text{even}} (Z_{j,0} \otimes Z_{j,1}) \right) \otimes \left(\otimes_{j=\text{odd}} (I_{j,0} \otimes I_{j,1}) \right)$ be three Pauli strings. Then, we have the anticommutation relations $K_j H_{\text{FH,qubits}}^{(j)} K_j = -H_{\text{FH,qubits}}^{(j)}$ for $j = 1, 2, 3$. Thus, control of the time evolution of the spin-1/2 Fermi-Hubbard model can be implemented by controlling these Pauli strings. The above construction can also be generalized to two-dimensional (2D) Fermi-Hubbard models. Note that direct Trotterization of the 2D Fermi-Hubbard model following the Jordan-Wigner transformation leads to nonoptimal complexities and the complexity can be improved via fermionic swap networks [63]. The control-free implementation of these more complex instances will be our future work.

For simplicity of implementation, the energy estimation can be derived from the measurement frequencies of bit strings using the standard variational quantum eigensolver (VQE) algorithm (see, e.g., Ref. [34]). We state the algorithm for deriving energy estimation from measurement results in Appendix E for completeness. Using the

TABLE III. The system-dependent parameters for different numbers of system qubits n .

n	μ	Δ	σ_+	σ_-	c_1	c_2	γ
2	0.7442	1.2884	0.9988	0.7686	0.1824	1.5708	0.5301
4	0.3926	0.5851	0.9988	0.9419	0.0909	1.5708	0.3003
6	0.2887	0.3773	0.9988	0.9717	0.0605	1.5708	0.1703
8	0.2394	0.2788	0.9988	0.9821	0.0453	1.5708	0.0965

control-free implementation and VQE-type energy estimation, the implementation of the ground-state preparation using QETU can be carried out efficiently on quantum hardware.

VII. NUMERICAL RESULTS FOR THE TFIM

Despite the potentially wide range of applications of QSP and QSVT, their implementation has been limited by the large resource overhead needed to implement the block encoding of the input matrix. To our knowledge, QSP-based quantum simulation has only been implemented for matrices encoded by random circuits [64–66]. Using QETU and control-free implementation in previous section, we show that the short-depth version of our algorithm for ground-state preparation and energy estimation can be readily implemented for certain physical Hamiltonians. Our implementation has a very small overhead compared to the Trotter-based Hamiltonian simulation and the circuit uses only one- and two-qubit gate operations. We demonstrate this for the TFIM via

IBM QISKIT. To demonstrate the algorithm, we prepare the ground state of the Ising model with a varying number of qubits n and the coupling strength is set to $g = 4$. In the quantum circuit, we set the initial state to $|0\rangle |\psi_{\text{in}}\rangle$, where $|\psi_{\text{in}}\rangle = |0^n\rangle$ and the additional one qubit is the ancilla qubit on which X rotations are applied. To simplify the numerical test, we compute the value of μ and Δ by explicitly diagonalizing the Hamiltonian. For each Hamiltonian evolution U in the QETU circuit, the number of Trotter steps is set to $r = 3$. We list system-dependent parameters and the initial overlap $\gamma = |\langle \psi_{\text{in}} | \psi_0 \rangle|$ for different numbers of system qubits n in Table III. According to the analysis in Sec. VII, it is sufficient to measure two quantum circuits to estimate the ground-state energy of the TFIM. Each quantum circuit in the numerical experiment is measured with 10^5 measurement shots and we independently repeat the numerical test 30 times to estimate the statistical fluctuation. To emulate the noisy quantum operation, we add a depolarizing error channel to each gate operation in the quantum circuit, where the error of single-qubit gate operation and that of the two-qubit gate operation are set to $r_{\text{dplz}}/10$ and r_{dplz} respectively. Assuming the digital error model (DEM), the total effect of the noise can be written as

$$\rho_{\text{exp}} = \alpha_{\text{DEM}} \rho_{\text{exact}} + (1 - \alpha_{\text{DEM}}) \mathcal{E}(\rho_{\text{input}})$$

and the noise channel $\mathcal{E}(\rho_{\text{input}})$ can be modeled as a global depolarized error channel [67] with circuit fidelity α_{DEM} .

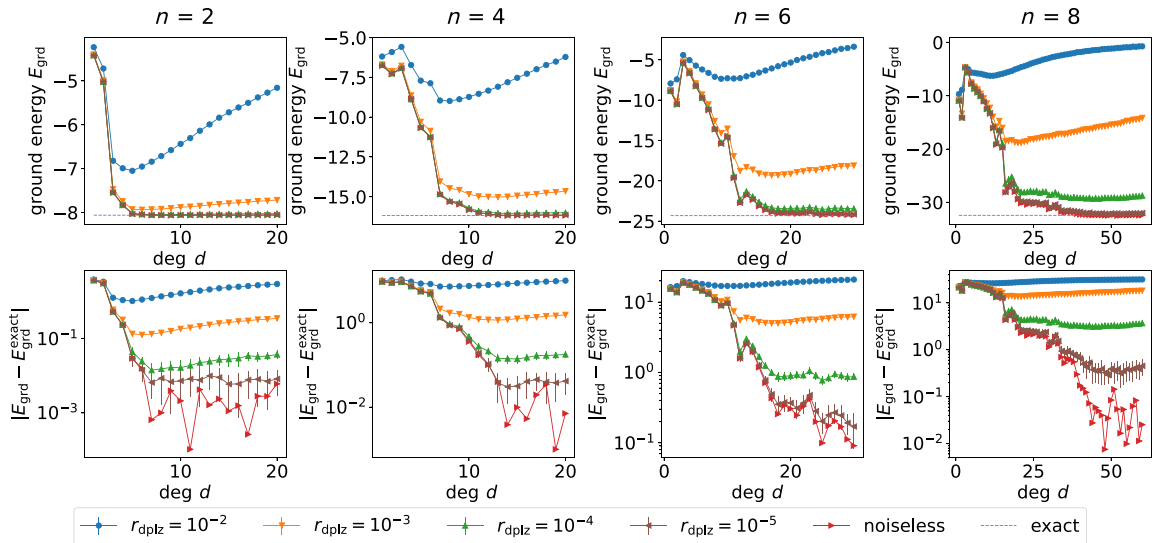


FIG. 6. Estimation of the ground-state energy of the TFIM model using QETU. Each marker labels the data simulated with a given depolarizing error rate r_{dplz} . The dashed line is the exact ground-state energy $E_{\text{grd}}^{\text{exact}}$ of the spin system. In each column, the number of qubits n is shown. The (red) right triangles denote the data computed from the best polynomial approximation by the convex-optimization solver, which only include the approximation error and are independent of quantum noise. The error bar denotes the standard deviation estimated from 30 repetitions.

Given n and d , the numbers of single- and two-qubit gates involved in the quantum circuit are

$$n_{g,1} = d(nr + 1) + 1 \quad \text{and} \quad n_{g,2} = d((n - 1)r + 2n). \quad (18)$$

Therefore, the circuit fidelity can be modeled as

$$\alpha_{\text{DEM}} = \left(1 - \frac{r_{\text{dplz}}}{10}\right)^{n_{g,1}} (1 - r_{\text{dplz}})^{n_{g,2}}.$$

The numerical result is presented in Fig. 6. The convergence of the noiseless data to the exact ground-state energy suggests that the energy can be computed accurately when d is modest ($10 \sim 30$). The statistical fluctuation, quantified by the standard deviation derived from 30 repetitions, is on the order of 10^{-2} and is not visible in the top panels. When simulating in the presence of the depolarizing noise, the numerical results suggest that accurate estimation of the energy requires r_{dplz} to be 10^{-4} or less. This requirement is beyond the noise level that can be achieved by current NISQ devices. Therefore we expect that QETU-based algorithms are more suited for early fault-tolerant quantum devices. For the TFIM, the spectral gap Δ decreases as the number of qubits increases. Therefore, the degree of the polynomial also needs to be increased to approximate the shifted sign function and to prepare the ground state to a fixed precision (see Fig. 6).

VIII. CONCLUSIONS

In this work, we develop algorithms for preparing the ground state and for estimating the ground-state energy of a quantum Hamiltonian suitable on early fault-tolerant quantum computers. The early fault-tolerant setting limits the number of qubits, the circuit depth, and the type of multiqubit control operations that can be employed. While block encoding is an elegant technique for abstractly encoding the information of an input Hamiltonian, existing block-encoding strategies (such as those for s -sparse matrices [15,68]) can lead to a large resource overhead and cannot meet the stringent requirements of early fault-tolerant devices. The resource overhead for approximately implementing a Hamiltonian-evolution input model is much lower and can be a suitable starting point for constructing more complex quantum algorithms.

Many computational tasks can be expressed in the form of applying a matrix function $f(H)$ to a quantum state $|\psi\rangle$. We develop a tool called quantum eigenvalue transformation of unitary matrices with real polynomials (QETU), which performs this task using the controlled Hamiltonian evolution as the input model (similar to that in quantum phase estimation), only one ancilla qubit, and no multi-qubit control operations. Combined with a fuzzy-bisection procedure, the total query complexity of the resulting algorithm to estimate the ground-state energy scales as

$\tilde{\mathcal{O}}(\epsilon^{-1}\gamma^{-2})$, which saturates the Heisenberg limit with target precision ϵ . The scaling with the initial overlap γ is not optimal but this result already outperforms all previous quantum algorithms for estimating the ground-state energy using a comparable circuit structure (see Table I).

The QETU technique and the new convex-optimization-based technique for streamlining the process of finding phase factors could readily be useful in many other contexts, such as preparing the Gibbs state. It is worth mentioning that other than using shifted sign functions, one can also use the exponential function $e^{-\beta(H-\mu I)}$ (the same as that needed for preparing Gibbs states, with an appropriate choice of β, μ) to approximately prepare the ground state. This gives rise to the imaginary-time-evolution method. Unlike the quantum imaginary-time-evolution (QITE) method [69], which performs both real-time evolution and a certain quantum state tomography procedure, QETU only queries the time evolution with performance guarantees and therefore can be significantly more advantageous in the early fault-tolerant regime.

If we are further allowed to use the $(n+1)$ -bit Toffoli gates (which is a relatively low-level multiqubit operation, as the additional two-qubit operations scale linearly in n), we can develop a new binary amplitude-estimation algorithm that is also based on QETU. The total query complexity for estimating the ground-state energy can be improved to the near-optimal scaling of $\tilde{\mathcal{O}}(\epsilon^{-1}\gamma^{-1})$, at the expense of increasing the circuit depth from $\tilde{\mathcal{O}}(\epsilon^{-1})$ to $\tilde{\mathcal{O}}(\epsilon^{-1}\gamma^{-1})$. This matches the results in Ref. [1] with a block-encoding input model. This also provides an answer to a question raised in Ref. [7], i.e., whether it is possible to have a quantum algorithm that does not use techniques such as LCU or block encoding, with a short query depth that scales as $\tilde{\mathcal{O}}(\epsilon^{-1})$, and with a total query complexity that scales better than $\mathcal{O}(\gamma^{-4})$. Our short-query-depth algorithm shows that it is possible to improve the total query complexity to $\mathcal{O}(\gamma^{-2})$ while satisfying all other constraints. The construction of our near-optimal algorithm (using binary amplitude estimation) indicates that it is unlikely that one can improve the total query complexity to $\mathcal{O}(\gamma^{-1})$ without introducing a factor that scales with γ^{-1} in the circuit depth.

The improvements in circuit depth and query complexity for preparing the ground state are similar to that of the ground-state energy estimation (see Table II). It is worth mentioning that many previous works using a single ancilla qubit cannot be easily modified to prepare the ground state. It is currently an open question whether the query complexity can be reduced to the near-optimal scaling without using any multiqubit controlled operation (specifically, whether the additional one- and two-qubit quantum gates can be independent of the system size n).

In practice, the cost of implementing the controlled Hamiltonian evolution can still be high. By exploiting certain anticommutation relations, we develop a new control-free implementation of QETU for a class of quantum spin Hamiltonians. The results on quantum simulators using IBM QISKIT indicate that relatively accurate estimates for the ground-state energy can already be obtained with a modest polynomial degree ($10 \sim 30$). However, the results of QETU can be sensitive to quantum noises (such as gate-wise depolarizing noises). On the one hand, while the QETU circuit (especially, the control-free variant) may be simple enough to fit on a NISQ device, the error on the NISQ devices may be too large to obtain meaningful results. On the other, it may be possible to combine QETU with randomized compilation [70] and/or error-mitigation techniques [71] to significantly reduce the impact of the noise, which may then enable us to obtain qualitatively meaningful results on near-term devices [72]. These will be our future works.

The source code for our use of IBM QISKIT is available in the GitHub repository [73].

ACKNOWLEDGMENTS

This work was supported by the U.S. Department of Energy under the Quantum Systems Accelerator program, under Grant No. DE-AC02-05CH11231 (Y.D.), by the NSF Quantum Leap Challenge Institute (QLCI) program, under Grant No. OMA-2016245 (Y.T.), and by the Google Quantum Research Award (L.L.). L.L. is a Simons Investigator. We thank Zhiyan Ding and Subhayan Roy Moulik for discussions and the anonymous referees for helpful suggestions.

APPENDIX A: BRIEF SUMMARY OF POLYNOMIAL MATRIX TRANSFORMATIONS

In the past few years, there have been significant algorithmic advancements in efficient representation of certain polynomial matrix transformations on quantum computers [12,14,15], which have found applications in Hamiltonian simulation, solving linear systems of equations, and eigenvalue problems, to name a few. The commonality of these approaches is to (1) encode a certain polynomial using a product of parametrized $SU(2)$ matrices and (2) lift the $SU(2)$ representation to matrices of arbitrary dimensions (a procedure called “qubitization” [12], which is related to quantum walks [74,75]). This framework often leads to a very concise quantum circuit and can unify a large class of quantum algorithms that have been developed in the literature [15,76]. For clarity of the presentation, the term QSP specifically refers to the $SU(2)$ representation. It is worth noting that depending on the structure of the matrix and the input model, the resulting quantum circuits can be different. Block encoding [12,15] is a commonly used input

model for representing nonunitary matrices on a quantum computer.

Definition 14: (block encoding). Given an n -qubit matrix A ($N = 2^n$), if we can find $\alpha, \epsilon \in \mathbb{R}_+$, and an $(m + n)$ -qubit unitary matrix U_A so that

$$\|A - \alpha (|0^m\rangle \otimes I_n) U_A (|0^m\rangle \otimes I_n)\|_2 \leq \epsilon, \quad (A1)$$

then U_A is called an (α, m, ϵ) -block-encoding of A .

When a polynomial of interest is represented by QSP, we can use the block-encoding input model to implement the polynomial transformation of a Hermitian matrix, which gives the quantum eigenvalue transformation (QET) [12]. Similarly, the polynomial transformation of a general matrix (called singular-value transformation) gives the QSVT [15]. In fact, for a Hermitian matrix with a block-encoding input model, the quantum circuits of QET and QSVT can be the same.

It is worth noting that the original presentation of QSP [14] combines the $SU(2)$ representation and a trigonometric polynomial transformation of a Hermitian matrix H and the input model is provided by a quantum walk operator [74]. If H is a s -sparse matrix, the use of a walk operator is actually not necessary and QET or QSVT gives a more concise algorithm than that in Ref. [14].

Using the Hamiltonian-evolution input model, our QETU algorithm provides a circuit structure that is similar to that in Ref. [14, Figure 1] and the derivation of QETU is both simpler and more constructive. Note that Ref. [14] only states the existence of the parametrization without providing an algorithm to evaluate the phase factors and the connection with the more explicit parametrization such as those in Refs. [15,77] has not been shown in the literature. Our QETU algorithm in Theorem 1 directly connects to the parametrization in Ref. [15] and, in particular, QSP with symmetric phase factors [54,59]. This gives rise to a concise way to represent real polynomial transformations that is encountered in most applications.

The QETU technique is also related to QSVT. From the Hamiltonian-evolution input model $U = e^{-iH}$, we can first use one ancilla qubit and linear combination of unitaries to implement a block encoding of $\cos(H) = (U + U^\dagger)/2$. Using another ancilla qubit, we can use QET or QSVT to implement $H = \arccos[\cos(H)]$ approximately. In other words, from the Hamiltonian evolution U , we can implement the matrix logarithm of U to approximately block encode H . Then, we can implement a matrix function $f(H)$ using the above block encoding and another layer of QSVT. QETU simplifies the above procedure by directly querying U . The concept of “qubitization” [12] appears very straightforwardly in QETU (see Appendix B). It also saves one ancilla qubit and perhaps gives a slightly smaller circuit depth.

APPENDIX B: QUANTUM EIGENVALUE TRANSFORMATION FOR UNITARY MATRICES

Let

$$W(x) = e^{i \arccos(x)X} = \begin{pmatrix} x & i\sqrt{1-x^2} \\ i\sqrt{1-x^2} & x \end{pmatrix}, \quad x \in [-1, 1]. \quad (\text{B1})$$

We first state the result of quantum signal processing for real polynomials [16, Corollary 10] and specifically the symmetric quantum signal processing [59, Theorem 10] in Theorem 15.

Theorem 15: (symmetric quantum signal processing, W -convention). *Given a real polynomial $F(x) \in \mathbb{R}[x]$, and $\deg F = d$, satisfying*

- (1) F has parity $d \bmod 2$
- (2) $|F(x)| \leq 1, \forall x \in [-1, 1]$

then there exist polynomials $G(x), Q(x) \in \mathbb{R}[x]$ and a set of symmetric phase factors $\Phi := (\phi_0, \phi_1, \dots, \phi_1, \phi_0) \in \mathbb{R}^{d+1}$ such that the following QSP representation holds:

$$e^{i\phi_0 Z} \prod_{j=1}^d [W(x)e^{i\phi_j Z}] = \begin{pmatrix} F(x) + iG(x) & iQ(x)\sqrt{1-x^2} \\ iQ(x)\sqrt{1-x^2} & F(x) - iG(x) \end{pmatrix}. \quad (\text{B2})$$

In order to derive QETU, we define

$$W_z(x) = e^{i \arccos(x)Z} = \begin{pmatrix} e^{i \arccos(x)} & 0 \\ 0 & e^{-i \arccos(x)} \end{pmatrix}, \quad x \in [-1, 1]. \quad (\text{B3})$$

Then Theorem 16 is equivalent to Theorem 15 but uses the variable x as encoded in the W_z matrix instead of the W matrix.

Theorem 16: (symmetric quantum signal processing, W_z -convention). *Given a real even polynomial $F(x) \in \mathbb{R}[x]$, and $\deg F = d$, satisfying $|F(x)| \leq 1, \forall x \in [-1, 1]$, then there exist polynomials $G(x), Q(x) \in \mathbb{R}[x]$ and symmetric phase factors $\Phi_z := (\varphi_0, \varphi_1, \dots, \varphi_1, \varphi_0) \in \mathbb{R}^{d+1}$ such that the following QSP representation holds:*

$$\begin{aligned} U_{\Phi_z}(x) &= e^{i\varphi_0 X} W_z^*(x) e^{i\varphi_1 X} W_z(x) e^{i\varphi_2 X} \dots e^{i\varphi_2 X} W_z^*(x) e^{i\varphi_1 X} W_z(x) e^{i\varphi_0 X} \\ &= \begin{pmatrix} F(x) & -Q(x)\sqrt{1-x^2} + iG(x) \\ Q(x)\sqrt{1-x^2} + iG(x) & F(x) \end{pmatrix}. \end{aligned} \quad (\text{B4})$$

Proof. Using

$$e^{i\varphi X} = H e^{i\varphi Z} H \quad (\text{B5})$$

and

$$H W_z(x) H = W(x), \quad H W_z^*(x) H = -e^{-i\pi/2Z} W(x) e^{-i\pi/2Z}, \quad (\text{B6})$$

we have

$$\begin{aligned} U_{\Phi_z}(x) &= (-1)^{d/2} H \left\{ e^{i(\varphi_0 - \pi/2)Z} W(x) e^{i(\varphi_1 - \pi/2)Z} W(x) e^{i(\varphi_2 - \pi/2)Z} \right. \\ &\quad \left. \dots e^{i(\varphi_2 - \pi/2)Z} W(x) e^{i(\varphi_1 - \pi/2)Z} W(x) e^{i\varphi_0 Z} \right\} H \\ &= (-1)^{d/2} H e^{-i\pi/4Z} \left\{ e^{i(\varphi_0 - \pi/4)Z} W(x) e^{i(\varphi_1 - \pi/2)Z} W(x) e^{i(\varphi_2 - \pi/2)Z} \right. \\ &\quad \left. \dots e^{i(\varphi_2 - \pi/2)Z} W(x) e^{i(\varphi_1 - \pi/2)Z} W(x) e^{i(\varphi_0 - \pi/4)Z} \right\} e^{i\pi/4Z} H. \end{aligned} \quad (\text{B7})$$

The term in the parentheses satisfies the condition of Theorem 15. We may choose a symmetric phase factor $(\phi_0, \phi_1, \dots, \phi_1, \phi_0)$, so that

$$e^{i\phi_0 Z} W(x) e^{i\phi_1 Z} W(x) e^{i\phi_2 Z} \dots e^{i\phi_2 Z} W(x) e^{i\phi_1 Z} W(x) e^{i\phi_0 Z} = (-1)^{d/2} \begin{pmatrix} F(x) + iG(x) & iQ(x)\sqrt{1-x^2} \\ iQ(x)\sqrt{1-x^2} & F(x) - iG(x) \end{pmatrix}. \quad (\text{B8})$$

We then define $\varphi_j = \phi_j + (2 - \delta_{j0})\pi/4$ for $j = 0, \dots, d/2$ and direct computation shows that

$$U_{\Phi_z}(x) = \begin{pmatrix} F(x) & -Q(x)\sqrt{1-x^2} + iG(x) \\ Q(x)\sqrt{1-x^2} + iG(x) & F(x) \end{pmatrix}, \quad (\text{B9})$$

which proves the theorem. ■

Proof of Theorem 1. For any eigenstate $|v_j\rangle$ of H with eigenvalue λ_j , note that $\text{span}\{|0\rangle|v_j\rangle, |1\rangle|v_j\rangle\}$ is an invariant subspace of $U, U^\dagger, X \otimes I_n$ and hence of \mathcal{U} . Together with the fact that for any phase factors φ and φ' ,

$$e^{i\varphi X} W_z^*(x) e^{i\varphi' X} W_z(x) = e^{i\varphi X} \begin{pmatrix} 1 & 0 \\ 0 & e^{2i \arccos(x)} \end{pmatrix} e^{i\varphi' X} \begin{pmatrix} 1 & 0 \\ 0 & e^{-2i \arccos(x)} \end{pmatrix}, \quad (\text{B10})$$

we have

$$\mathcal{U}|0\rangle|v_j\rangle = (U_{\Phi_z}(\cos(\lambda_j/2))|0\rangle)|v_j\rangle = F(\cos(\lambda_j/2))|0\rangle|v_j\rangle + \alpha_j|1\rangle|v_j\rangle. \quad (\text{B11})$$

Here, we use $x = \cos(\lambda_j/2)$ and

$$\alpha_j = Q(\cos(\lambda_j/2)) \sin(\lambda_j/2) + iG(\cos(\lambda_j/2)) \quad (\text{B12})$$

is an irrelevant constant according to Eq. (B4).

Since any state $|\psi\rangle$ can be expanded as the linear combination of eigenstates $|v_j\rangle$ as

$$|\psi\rangle = \sum_j c_j |v_j\rangle, \quad (\text{B13})$$

we have

$$\begin{aligned} \mathcal{U}|0\rangle|\psi\rangle &= \sum_j c_j \mathcal{U}|0\rangle|v_j\rangle = |0\rangle \sum_j c_j F(\cos(\lambda_j/2)) |v_j\rangle + |1\rangle|\perp\rangle \\ &= |0\rangle F(\cos(H/2)) |\psi\rangle + |1\rangle|\perp\rangle, \end{aligned} \quad (\text{B14})$$

where $|\perp\rangle$ is some unnormalized quantum state. This proves the theorem. ■

Sometimes, instead of controlled U , we have direct access to an oracle that simultaneously implements a controlled forward and backward time evolution:

$$V = \begin{pmatrix} e^{iH} & 0 \\ 0 & e^{-iH} \end{pmatrix}. \quad (\text{B15})$$

This is the case, for instance, in a certain implementation of QETU in a control-free setting. Corollary 17 describes this version of QETU.

Corollary 17: (QETU with forward and backward time evolution). *Let V be the unitary matrix given in Eq. (B15), corresponding to an n -qubit Hermitian matrix H . For any even real polynomial $F(x)$ of degree d satisfying $|F(x)| \leq 1, \forall x \in [-1, 1]$, we can find a sequence of symmetric phase factors $\Phi_z := (\varphi_0, \varphi_1, \dots, \varphi_1, \varphi_0) \in \mathbb{R}^{d+1}$, such that the circuit in Fig. 7 denoted by \mathcal{U} satisfies $\langle 0| \otimes I_n \mathcal{U}(|0\rangle \otimes I_n) = F(\cos H)$.*




FIG. 7. The variant of QETU with an oracle implementing the controlled forward and backward time evolution. The implementation of V^\dagger can be carried out by conjugating V with Pauli X gates acting on the first qubit and the Pauli X gate can be first combined with the phase rotation as $e^{i\varphi X} X = ie^{i(\varphi+\pi/2)X}$.

Proof. Let $|v_j\rangle$ be an eigenstate of H with eigenvalue λ_j . For any phase factors φ and φ' , using

$$\begin{aligned} e^{i\varphi X} W_z^*(x) e^{i\varphi' X} W_z(x) &= e^{i\varphi X} \begin{pmatrix} e^{-i\arccos(x)} & 0 \\ 0 & e^{i\arccos(x)} \end{pmatrix} \\ &\times e^{i\varphi' X} \begin{pmatrix} e^{i\arccos(x)} & 0 \\ 0 & e^{-i\arccos(x)} \end{pmatrix}, \end{aligned} \quad (\text{B16})$$

we have

$$\begin{aligned} \mathcal{U} |0\rangle |v_j\rangle &= (U_{\Phi_z}(\cos \lambda_j) |0\rangle) |v_j\rangle \\ &= F(\cos \lambda_j) |0\rangle |v_j\rangle + \alpha_j |1\rangle |v_j\rangle. \end{aligned} \quad (\text{B17})$$

Here, $x = \cos \lambda_j$, and $\alpha_j = Q(\cos \lambda_j) \sin \lambda_j + iG(\cos \lambda_j)$. The rest of the proof follows that of Theorem 1. ■

APPENDIX C: COST OF QETU USING TROTTER FORMULAS

If the time-evolution operator $U = e^{-iH}$ is implemented using Trotter formulas, we can directly analyze the circuit depth and gate complexity of estimating the ground-state energy in the setting of Theorems 4 and 5.

We suppose that the Hamiltonian H can be decomposed into a sum of terms $H = \sum_{\gamma=1}^L H_\gamma$, where each term H_γ can be efficiently exponentiated, i.e., with a gate complexity that is independent of time. In other words, we assume that each H_γ can be fast forwarded [78–80]. We assume that the gate complexity for implementing a single Trotter step is G_{Trotter} and the circuit depth required is D_{Trotter} . For initial state preparation, we assume that we need gate complexity G_{initial} and circuit depth D_{initial} . A p th-order Trotter formula applied to $U = e^{-iH}$ with r Trotter steps gives us a unitary operator U_{HS} with error

$$\|U_{HS} - U\| \leq C_{\text{Trotter}} r^{-p},$$

where C_{Trotter} is a prefactor, for which the simplest bound is $C_{\text{Trotter}} = \mathcal{O}((\sum_\gamma \|H_\gamma\|)^{p+1})$. Tighter bounds in the form

of a sum of commutators are proved in Refs. [19,81] and there are many works on how to decompose the Hamiltonian to reduce the resource requirement [82–85]. If the circuit queries U, U^\dagger for d times and the desired precision is δ , then we can choose

$$d \times C_{\text{Trotter}} r^{-p} = \delta$$

or, equivalently,

$$r = \mathcal{O}\left(\max\{d^{1/p} C_{\text{Trotter}}^{1/p} \delta^{-1/p}, 1\}\right). \quad (\text{C1})$$

As an example, let us now analyze the number of Trotter steps needed in the context of estimating the ground-state energy in Theorem 4 without amplitude amplification. When replacing the exact U with U_{HS} , we only need to ensure that the resulting error in U_{proj} defined in Eq. (8) is $\mathcal{O}(\gamma)$. This will enable us to solve the binary amplitude-estimation problem (see Definition 8) with the same asymptotic query complexity. Since U_{proj} uses U (and therefore U_{HS}) at most $\tilde{\mathcal{O}}(\epsilon^{-1} \log(\gamma^{-1}))$ times, we only need to ensure

$$\tilde{\mathcal{O}}(\epsilon^{-1} \log(\gamma^{-1})) \times C_{\text{Trotter}} r^{-p} = \mathcal{O}(\gamma). \quad (\text{C2})$$

Consequently, we need to choose

$$r = \tilde{\mathcal{O}}\left(\max\{C_{\text{Trotter}}^{1/p} \epsilon^{-1/p} \gamma^{-1/p}, 1\}\right). \quad (\text{C3})$$

The query depth of U is $\tilde{\mathcal{O}}(\epsilon^{-1} \log(\gamma^{-1}))$ and therefore the circuit depth is

$$\tilde{\mathcal{O}}\left(\max\{C_{\text{Trotter}}^{1/p} \epsilon^{-1/p} \gamma^{-1/p}, \log(\gamma^{-1})\} \epsilon^{-1} D_{\text{Trotter}} + D_{\text{initial}}\right).$$

Similarly, the total number of queries to U is $\tilde{\mathcal{O}}[\epsilon^{-1} \gamma^{-2} \log(\vartheta^{-1})]$ times, the total number of queries to U_I is $\mathcal{O}[\gamma^{-2} \text{poly} \log(\epsilon^{-1} \vartheta^{-1})]$, and the resulting total gate complexity is

$$\begin{aligned} &\tilde{\mathcal{O}}\left(\max\{C_{\text{Trotter}}^{1/p} \epsilon^{-1/p} \gamma^{-1/p}, 1\} \epsilon^{-1} \gamma^{-2} G_{\text{Trotter}} \log(\vartheta^{-1})\right) \\ &+ \gamma^{-2} G_{\text{initial}} \log(\vartheta^{-1}). \end{aligned}$$

The analysis of the number of Trotter steps in the setting of Theorem 5 is similar. We still want to ensure Eq. (C2), which results in the same choice of the number of Trotter steps r as in Eq. (C3). Combined with the query complexity in Theorem 5, the total gate complexity is

$$\begin{aligned} &\tilde{\mathcal{O}}\left(\max\{C_{\text{Trotter}}^{1/p} \epsilon^{-1/p} \gamma^{-1/p}, 1\} \epsilon^{-1} \gamma^{-1} G_{\text{Trotter}} \log(\vartheta^{-1})\right) \\ &+ \gamma^{-1} G_{\text{initial}} \log(\vartheta^{-1}). \end{aligned} \quad (\text{C4})$$

APPENDIX D: BINARY AMPLITUDE ESTIMATION WITH A SINGLE ANCILLA QUBIT AND QETU

In this appendix, we discuss how to solve the binary amplitude-estimation problem in Definition 8 using a single ancilla qubit. We restate the problem here.

Definition 18: (binary amplitude estimation). Let W be a unitary acting on two registers, with the first register indicating success or failure. Let $A = \|(\langle 0| \otimes I_n)W(|0\rangle |0^n\rangle)\|$ be the success amplitude. Given $0 \leq \gamma_1 < \gamma_2$, provided that A is either smaller than γ_1 or greater than γ_2 , we want to correctly distinguish between the two cases, i.e., output 0 for the former and 1 for the latter.

In the following, we describe an algorithm to solve this problem and thereby prove Lemma 12. We can find quantum states $|\Phi\rangle$ and $|\perp\rangle$ such that

$$W(|0\rangle |0^n\rangle) = A |0\rangle |\Phi\rangle + \sqrt{1 - A^2} |\perp\rangle$$

and $(\langle 0| \otimes I) |\perp\rangle = 0$. We also define

$$|\perp'\rangle = -\sqrt{1 - A^2} |0\rangle |\Phi\rangle + A |\perp\rangle.$$

As in amplitude amplification, we define two reflection operators:

$$R_0 = (2 |0\rangle \langle 0| - I) \otimes I_n, \quad R_1 = W(2 |0^{n+1}\rangle \langle 0^{n+1}| - I_{n+1})W^\dagger.$$

Relative to the basis $\{W(|0\rangle |0^n\rangle), |\perp'\rangle\}$ the two reflection operators can be represented by the matrices

$$\begin{pmatrix} 2A^2 - 1 & -2A\sqrt{1 - A^2} \\ -2A\sqrt{1 - A^2} & 1 - 2A^2 \end{pmatrix}, \quad \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

Therefore, we can verify that $|\Psi_\pm\rangle = (W|0\rangle |0^n\rangle \pm i|\perp'\rangle)/\sqrt{2}$ are eigenvectors of R_0R_1 :

$$R_0R_1 |\Psi_\pm\rangle = e^{\mp i2 \arccos(A)} |\Psi_\pm\rangle.$$

If we use the usual amplitude-estimation algorithm to estimate A , we can simply perform phase estimation with R_0R_1 on the quantum state $W(|0\rangle |0^n\rangle)$, which is an equal superposition of $|\Psi_\pm\rangle$:

$$W(|0\rangle |0^n\rangle) = \frac{1}{\sqrt{2}}(|\Psi_+\rangle + |\Psi_-\rangle).$$

However, here we do something different. We view R_0R_1 has a time-evolution operator corresponding to some

Hamiltonian L :

$$R_0R_1 = e^{-iL},$$

where, in the subspace spanned by $\{W(|0\rangle |0^n\rangle), |\perp'\rangle\}$, we have

$$L = 2 \arccos(A) |\Psi_+\rangle \langle \Psi_+| - 2 \arccos(A) |\Psi_-\rangle \langle \Psi_-|.$$

Then, using QETU in Theorem 1, we can implement a block encoding, which we denote by \mathcal{U} , of $P[\cos(L/2)]$ for any suitable polynomial P and in the same subspace we have

$$\begin{aligned} P[\cos(L/2)] &= P(A) (|\Psi_+\rangle \langle \Psi_+| + |\Psi_-\rangle \langle \Psi_-|) \\ &= P(A) (W|0\rangle |0^n\rangle \langle 0| \langle 0^n| W^\dagger + |\perp'\rangle \langle \perp'|). \end{aligned}$$

Using QETU, we can use Monte Carlo sampling to estimate the quantity

$$\|(\langle 0| \otimes I_{n+1})\mathcal{U}(|0\rangle \otimes (W|0\rangle |0^n\rangle))\|^2 = |P(A)|^2. \quad (\text{D1})$$

To be more precise, we can start from the state $|0\rangle |0\rangle |0^n\rangle$ on $n + 2$ qubits, apply W to the last $n + 1$ qubits, then \mathcal{U} to all $n + 2$ qubits, and in the end measure the first qubit. The probability of obtaining 0 in the measurement outcome is exactly as described in Eq. (D1).

Now let us consider an even polynomial $P(x)$ such that $|P(x)| \leq 1$ for $x \in [-1, 1]$ and

$$P(x) \geq 1 - \delta, \quad x \in [\gamma_2, 1], \quad |P(x)| \leq \delta, \quad x \in [0, \gamma_1].$$

Such a polynomial of degree $\mathcal{O}((\gamma_2 - \gamma_1)^{-1} \log(\delta^{-1}))$ can be constructed using the approximate sign function in Ref. [53] (if we take this approach, we need to symmetrize the polynomial through $P(x) = [Q(x) + Q(-x)]/2$) or the optimization procedure described in Sec. IV. Using this polynomial, we can then employ Monte Carlo sampling to distinguish two cases, which will solve the binary amplitude-estimation problem:

$$\begin{aligned} \|(\langle 0| \otimes I_{n+1})\mathcal{U}(|0\rangle \otimes (W|0\rangle |0^n\rangle))\|^2 &\geq (1 - \delta)^2, \\ \text{or } \|(\langle 0| \otimes I_{n+1})\mathcal{U}(|0\rangle \otimes (W|0\rangle |0^n\rangle))\|^2 &\leq \delta^2. \end{aligned}$$

We can choose $\delta = 1/4$ and it takes running W and \mathcal{U} and measuring the first qubit each $\mathcal{O}(\log(\vartheta^{-1}))$ times to successfully distinguish between the above two cases with probability at least $1 - \vartheta$. In this, we use the standard majority-voting procedure to boost the success probability. Each single run of \mathcal{U} requires $\mathcal{O}((\gamma_2 - \gamma_1)^{-1})$ applications of W , which corresponds to the polynomial degree. Therefore, in total we need to apply W $\mathcal{O}((\gamma_2 - \gamma_1)^{-1} \log(\vartheta^{-1}))$ times.

In this whole procedure, we need one additional ancilla qubit for QETU. Note that the $(n + 1)$ -qubit reflection

operator in R_1 can be implemented using the $(n + 2)$ -qubit Toffoli gate and phase kickback. Using Ref. [40, Corollary 7.4], we can implement the $(n + 2)$ -qubit Toffoli gate on $(n + 3)$ qubits. As a result, another ancilla qubit is needed. We have proved Lemma 12, which we restate here.

Lemma 19: *The binary amplitude-estimation problem in Definition 8 can be solved correctly with probability at least $1 - \vartheta$ by querying $W \mathcal{O}((\gamma_2 - \gamma_1)^{-1} \log(\vartheta^{-1}))$ times and this procedure requires two additional ancilla qubits (besides the ancilla qubits already required in W).*

APPENDIX E: DETAILS OF NUMERICAL SIMULATION OF THE TFIM

In the numerical simulation for estimating the ground-state energy of the TFIM, we explicitly diagonalize the Hamiltonian to obtain the exact ground state $|\psi_0\rangle$, the ground energy E_0 , the first excited energy E_1 , and the highest excited energy E_{n-1} . We then perform an affine transformation to the shifted Hamiltonian

$$H^{\text{sh}} = c_1 H + c_2 I_n, \quad c_1 = \frac{\pi - 2\eta}{E_{n-1} - E_0} \text{ and } c_2 = \eta - c_1 E_0. \quad (\text{E1})$$

Consequently, the eigenvalues of the shifted Hamiltonian are exactly in the interval $[\eta, \pi - \eta]$, i.e., $E_0^{\text{sh}} = \eta$ and $E_{n-1}^{\text{sh}} = \pi - \eta$. The time evolution is then $e^{-itH^{\text{sh}}} = e^{-itc_2} e^{-itc_1 H}$, which means that the evolution time is scaled to $\tau^{\text{sh}} = \tau c_1$ with an additional phase shift $\phi^{\text{sh}} = \tau c_2$. The system-dependent parameters are then given by

$$\begin{aligned} \mu &= \frac{1}{2} (E_0^{\text{sh}} + E_1^{\text{sh}}), \quad \Delta = E_1^{\text{sh}} - E_0^{\text{sh}}, \text{ and } \sigma_{\pm} \\ &= \cos \frac{\mu \mp \Delta/2}{2}. \end{aligned} \quad (\text{E2})$$

We set the input quantum state to $|0\rangle |\psi_{\text{in}}\rangle$, where $|\psi_{\text{in}}\rangle = |0^n\rangle$ and the additional one qubit is the ancilla qubit for performing X rotations in QETU. The initial overlap is $\gamma = |\langle \psi_{\text{in}} | \psi_0 \rangle|$. We list the system-dependent parameters used in the numerical experiments in Fig. 6 in Table III.

For completeness, we briefly introduce the algorithm for deriving the energy estimation from the measurement of bit-string frequencies. The energy-estimation process can be optimized so that it is sufficient to measure a few quantum circuits to compute the energy. For the TFIM, if the ground state is $|\psi_0\rangle$, its ground-state energy is

$$\begin{aligned} E_0 &= \langle \psi_0 | H_{\text{TFIM}} | \psi_0 \rangle = - \sum_{j=1}^{n-1} \langle \psi_0 | Z_j Z_{j+1} | \psi_0 \rangle \\ &\quad - g \sum_{j=1}^n \langle \psi_0 | X_j | \psi_0 \rangle =: - \sum_{j=1}^{n-1} \Psi_{j,j+1} - g \sum_{j=1}^n \Psi_j^H. \end{aligned}$$

We show that the energy component $\Psi_{j,j+1}$ and Ψ_j^H can be exactly expressed as the marginal probabilities readable from measurements. Decomposing $Z_j Z_{j+1}$ with respect to eigenvectors, we have

$$\begin{aligned} \Psi_{j,j+1} &= \langle \psi_0 | Z_j Z_{j+1} | \psi_0 \rangle = \sum_{z_j=0}^1 \sum_{z_{j+1}=0}^1 (-1)^{z_j+z_{j+1}} \\ &\quad |\langle \psi_0 | z_j, z_{j+1} \rangle|^2 = \sum_{z_j=0}^1 \sum_{z_{j+1}=0}^1 (-1)^{z_j+z_{j+1}} \\ &\quad \mathbb{P}(z_j, z_{j+1} | \psi_0). \end{aligned}$$

Here, $\mathbb{P}(z_j, z_{j+1} | \psi_0)$ is the marginal probability measuring the j th qubit with z_j and the $(j + 1)$ th qubit with z_{j+1} under computational basis when the quantum circuit for preparing the ground state $|\psi_0\rangle$ is given. Similarly, the other quantity involved in the energy is

$$\begin{aligned} \Psi_j^H &= \langle \psi_0 | X_j | \psi_0 \rangle = \langle \psi_0 | H^{\otimes n} Z_j H^{\otimes n} | \psi_0 \rangle = \langle \psi_0^H | Z_j | \psi_0^H \rangle \\ &= \sum_{z_j=0}^1 (-1)^{z_j} \mathbb{P}(z_j | \psi_0^H). \end{aligned}$$

Here, $\mathbb{P}(z_j | \psi_0^H)$ is the marginal probability measuring the j th qubit with z_j under computational basis when the quantum circuit for preparing the ground state $|\psi_0\rangle$ following a Hadamard transformation, which is denoted as $|\psi_0^H\rangle := H^{\otimes n} |\psi_0\rangle$, is given.

In order to estimate the ground-state energy of the TFIM, it suffices to measure all qubits in two circuits: the circuit in Fig. 5(b) and that following a Hadamard transformation on all system qubits. The measurement results estimate the marginal probabilities up to the Monte Carlo measurement error. Furthermore, their linear combination with signs gives the ground-state energy estimate based on the previous analysis.

The procedure for estimating the energy can readily be generalized to other models. Consider a Hamiltonian

$$H = \sum_{k=1}^L H_k, \quad H_k = \sum_{j=1}^{v_k} h_{k,j}. \quad (\text{E3})$$

Here, we group the components of the Hamiltonian into L classes and, for a fixed k , the components $h_{k,j}$ can be simultaneously diagonalized by an efficiently implementable unitary V_k . The strategies of Hamiltonian grouping have also been used in, e.g., Refs. [86,87]. We want to estimate the expectation $\langle \psi_0 | H | \psi_0 \rangle$, where $|\psi_0\rangle$ is the quantum state prepared by some quantum circuit. Then, it suffices to measure L different quantum circuits $\{V_k | \psi_0\rangle : k = 1, \dots, L\}$ and to compute the expectation from the

measurement data by some signed linear combination. For example, to estimate the ground-state energy of the Heisenberg model, we can let $L = 3$ and $V_1 = I^{\otimes n}$, $V_2 = H^{\otimes n}$ and $V_3 = (HS^\dagger)^{\otimes n}$, where H and S are the Hadamard gate and the phase gate, respectively.

-
- [1] L. Lin and Y. Tong, Near-optimal ground state preparation, *Quantum* **4**, 372 (2020).
- [2] J. Preskill, Quantum computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
- [3] R. Babbush, J. R. McClean, M. Newman, C. Gidney, S. Boixo, and H. Neven, Focus beyond Quadratic Speedups for Error-Corrected Quantum Advantage, *PRX Quantum* **2**, 010103 (2021).
- [4] K. E. Booth, B. O’Gorman, J. Marshall, S. Hadfield, and E. Rieffel, Quantum-accelerated constraint programming (2021), *arXiv preprint ArXiv:2103.04502*.
- [5] E. T. Campbell, Early fault-tolerant simulations of the Hubbard model, *Quantum Sci. Technol.* **7**, 015007 (2021).
- [6] D. Layden, First-order Trotter error from a second-order perspective (2021), *arXiv preprint ArXiv:2107.08032*.
- [7] L. Lin and Y. Tong, Heisenberg-Limited Ground State Energy Estimation for Early Fault-Tolerant Quantum Computers, *PRX Quantum* **3**, 010318 (2022).
- [8] K. Wan, M. Berta, and E. T. Campbell, A randomized quantum algorithm for statistical phase estimation (2021), *arXiv preprint ArXiv:2110.12071*.
- [9] G. Wang, S. Sim, and P. D. Johnson, State preparation boosters for early fault-tolerant quantum computation (2022), *arXiv preprint ArXiv:2202.06978*.
- [10] R. Zhang, G. Wang, and P. Johnson, Computing ground state properties with early fault-tolerant quantum computers (2021), *arXiv preprint ArXiv:2109.13957*.
- [11] S. Chakraborty, A. Gilyén, and S. Jeffery, The power of block-encoded matrix powers: Improved regression techniques via faster Hamiltonian simulation (2018), *ArXiv:1804.01973*.
- [12] G. H. Low and I. L. Chuang, Hamiltonian simulation by qubitization, *Quantum* **3**, 163 (2019).
- [13] D. W. Berry, A. M. Childs, and R. Kothari, Hamiltonian simulation with nearly optimal dependence on all parameters. *Proc. 56th IEEE Symp. Found. Comput. Sci.*, p. 792, 2015.
- [14] G. H. Low and I. L. Chuang, Optimal Hamiltonian Simulation by Quantum Signal Processing, *Phys. Rev. Lett.* **118**, 010501 (2017).
- [15] A. Gilyén, Y. Su, G. H. Low, and N. Wiebe, in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, p. 193, 2019.
- [16] A. Gilyén, Y. Su, G. H. Low, and N. Wiebe, Quantum singular value transformation and beyond: Exponential improvements for quantum matrix arithmetics (2018), *ArXiv:1806.01838*.
- [17] P. Rall, Quantum algorithms for estimating physical quantities using block encodings, *Phys. Rev. A* **102**, 022408 (2020).
- [18] Y. Tong, D. An, N. Wiebe, and L. Lin, Fast inversion, preconditioned quantum linear system solvers, and fast evaluation of matrix functions (2020), *ArXiv:2008.13295*.
- [19] A. M. Childs, Y. Su, M. C. Tran, N. Wiebe, and S. Zhu, Theory of Trotter Error with Commutator Scaling, *Phys. Rev. X* **11**, 011020 (2021).
- [20] S. Lloyd, Universal quantum simulators. *Science*, p. 1073, 1996.
- [21] V. Giovannetti, S. Lloyd, and L. Maccone, Quantum Metrology, *Phys. Rev. Lett.* **96**, 010401 (2006).
- [22] V. Giovannetti, S. Lloyd, and L. Maccone, Advances in quantum metrology, *Nat. Photon.* **5**, 222 (2011).
- [23] M. Zwierz, C. A. Pérez-Delgado, and P. Kok, General Optimality of the Heisenberg Limit for Quantum Metrology, *Phys. Rev. Lett.* **105**, 180402 (2010).
- [24] M. Zwierz, C. A. Pérez-Delgado, and P. Kok, Ultimate limits to quantum metrology and the meaning of the Heisenberg limit, *Phys. Rev. A* **85**, 042112 (2012).
- [25] T. E. O’Brien, B. Tarasinski, and B. M. Terhal, Quantum phase estimation of multiple eigenvalues for small-scale (noisy) experiments, *New J. Phys.* **21**, 023022 (2019).
- [26] D. Wang, O. Higgott, and S. Brierley, Accelerated Variational Quantum Eigensolver, *Phys. Rev. Lett.* **122**, 140504 (2019).
- [27] N. Wiebe, C. Granade, A. Kapoor, and K. M. Svore, Bayesian inference via rejection filtering (2015), *arXiv preprint ArXiv:1511.06458*.
- [28] D. Aharonov, D. Gottesman, S. Irani, and J. Kempe, The power of quantum systems on a line, *Comm. Math. Phys.* **287**, 41 (2009).
- [29] J. Kempe, A. Kitaev, and O. Regev, The complexity of the local Hamiltonian problem, *SIAM J. Comput.* **35**, 1070 (2006).
- [30] A. Y. Kitaev, A. Shen, and M. N. Vyalyi, *Classical and Quantum Computation*. Number 47. American Mathematical Soc., 2002.
- [31] R. Oliveira and B. M. Terhal, The complexity of quantum spin systems on a two-dimensional square lattice (2005), *arXiv preprint ArXiv:quant-ph/0504050*.
- [32] I. D. Kivlichan, J. McClean, N. Wiebe, C. Gidney, A. Aspuru-Guzik, G. K.-L. Chan, and R. Babbush, Quantum Simulation of Electronic Structure with Linear Depth and Connectivity, *Phys. Rev. Lett.* **120**, 110501 (2018).
- [33] N. M. Tubman, C. Mejuto-Zaera, J. M. Epstein, D. Hait, D. S. Levine, W. Huggins, Z. Jiang, J. R. McClean, R. Babbush, and M. Head-Gordon, *et al.*, Postponing the orthogonality catastrophe: Efficient state preparation for electronic structure simulations on quantum devices (2018), *arXiv preprint ArXiv:1809.05523*.
- [34] 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).
- [35] P. J. O’Malley, R. Babbush, I. D. Kivlichan, J. Romero, J. R. McClean, R. Barends, J. Kelly, P. Roushan, A. Tranter, and N. Ding, *et al.*, Scalable Quantum Simulation of Molecular Energies, *Phys. Rev. X* **6**, 031007 (2016).
- [36] 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).

- [37] Y. Ge, J. Tura, and J. I. Cirac, Faster ground state preparation and high-precision ground energy estimation with fewer qubits, *J. Math. Phys.* **60**, 022202 (2019).
- [38] T. Keen, E. Dumitrescu, and Y. Wang, Quantum algorithms for ground-state preparation and Green's function calculation (2021), *arXiv preprint ArXiv:2112.05731*.
- [39] G. Brassard, P. Hoyer, M. Mosca, and A. Tapp, Quantum amplitude amplification and estimation, *Contemp. Math.* **305**, 53 (2002).
- [40] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter, Elementary gates for quantum computation, *Phys. Rev. A* **52**, 3457 (1995).
- [41] E. Knill, G. Ortiz, and R. D. Somma, Optimal quantum measurements of expectation values of observables, *Phys. Rev. A* **75**, 012328 (2007).
- [42] D. Nagaj, P. Wocjan, and Y. Zhang, Fast amplification of QMA, *Quantum Inf. Comput.* **9**, 1053 (2009).
- [43] D. Poulin and P. Wocjan, Sampling from the Thermal Quantum Gibbs State and Evaluating Partition Functions with a Quantum Computer, *Phys. Rev. Lett.* **103**, 220502 (2009).
- [44] D. W. Berry, B. L. Higgins, S. D. Bartlett, M. W. Mitchell, G. J. Pryde, and H. M. Wiseman, How to perform the most accurate possible phase measurements, *Phys. Rev. A* **80**, 052114 (2009).
- [45] B. L. Higgins, D. W. Berry, S. D. Bartlett, H. M. Wiseman, and G. J. Pryde, Entanglement-free Heisenberg-limited phase estimation, *Nature* **450**, 393 (2007).
- [46] R. D. Somma, Quantum eigenvalue estimation via time series analysis, *New J. Phys.* **21**, 123025 (2019).
- [47] R. B. Griffiths and C.-S. Niu, Semiclassical Fourier Transform for Quantum Computation, *Phys. Rev. Lett.* **76**, 3228 (1996).
- [48] K. Choi, D. Lee, J. Bonitati, Z. Qian, and J. Watkins, Rodeo Algorithm for Quantum Computing, *Phys. Rev. Lett.* **127**, 040505 (2021).
- [49] P. Zeng, J. Sun, and X. Yuan, Universal quantum algorithmic cooling on a quantum computer (2021), *arXiv preprint ArXiv:2109.15304*.
- [50] W. J. Huggins, J. Lee, U. Baek, B. O'Gorman, and K. B. Whaley, A non-orthogonal variational quantum eigensolver, *New J. Phys.* **22**, 073009 (2020).
- [51] S. Lu, M. C. Bañuls, and J. I. Cirac, Algorithms for quantum simulation at finite energies (2020), *arXiv preprint ArXiv:2006.03032*.
- [52] T. E. O'Brien, S. Polla, N. C. Rubin, W. J. Huggins, S. McArdle, S. Boixo, J. R. McClean, and R. Babbush, Error mitigation via verified phase estimation (2020), *arXiv preprint ArXiv:2010.02538*.
- [53] G. H. Low and I. L. Chuang, Hamiltonian simulation by uniform spectral amplification (2017), *ArXiv:1707.05391*.
- [54] Y. Dong, X. Meng, K. B. Whaley, and L. Lin, Efficient phase factor evaluation in quantum signal processing, *Phys. Rev. A* **103**, 042419 (2021).
- [55] G. Wang, D. E. Koh, P. D. Johnson, and Y. Cao, Minimizing Estimation Runtime on Noisy Quantum Computers, *PRX Quantum* **2**, 010346 (2021).
- [56] S. Aaronson and P. Rall, in *Symposium on Simplicity in Algorithms* (SIAM, 2020), p. 24.
- [57] C.-R. Wie, Simpler quantum counting (2019), *arXiv preprint ArXiv:1907.08119*.
- [58] M. Grant and S. Boyd, CVX: MATLAB software for disciplined convex programming, version 2.1. <http://cvxr.com/cvx>, Mar. 2014.
- [59] J. Wang, Y. Dong, and L. Lin, On the energy landscape of symmetric quantum signal processing (2021), *arXiv preprint ArXiv:2110.04993*.
- [60] <https://github.com/qsppack/QSPPACK>.
- [61] L. N. Trefethen, *Approximation Theory and Approximation Practice*, volume 164. SIAM, 2019.
- [62] J.-M. Reiner, M. Marthaler, J. Braumüller, M. Weides, and G. Schön, Emulating the one-dimensional Fermi-Hubbard model by a double chain of qubits, *Phys. Rev. A* **94**, 032338 (2016).
- [63] I. Kivlichan, J. McClean, N. Wiebe, C. Gidney, A. Aspuru-Guzik, G.-L. Chan, and R. Babbush, Quantum Simulation of Electronic Structure with Linear Depth and Connectivity, *Phys. Rev. Lett.* **120**, 110501 (2018).
- [64] A. Cornelissen, J. Bausch, and A. Gilyén, Scalable benchmarks for gate-based quantum computers (2021), *arXiv preprint ArXiv:2104.10698*.
- [65] Y. Dong and L. Lin, Random circuit block-encoded matrix and a proposal of quantum LINPACK benchmark, *Phys. Rev. A* **103**, 062412 (2021).
- [66] Y. Dong, K. B. Whaley, and L. Lin, A quantum Hamiltonian simulation benchmark (2021), *arXiv preprint ArXiv:2108.03747*.
- [67] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, Characterizing quantum supremacy in near-term devices, *Nat. Phys.* **14**, 595 (2018).
- [68] A. M. Childs, R. Kothari, and R. D. Somma, Quantum algorithm for systems of linear equations with exponentially improved dependence on precision, *SIAM J. Comput.* **46**, 1920 (2017).
- [69] M. Motta, C. Sun, A. T. Tan, M. J. O'Rourke, E. Ye, A. J. Minnich, F. G. Brandão, and G. K.-L. Chan, Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution, *Nat. Phys.* **16**, 205 (2020).
- [70] J. J. Wallman and J. Emerson, Noise tailoring for scalable quantum computation via randomized compiling, *Phys. Rev. A* **94**, 052325 (2016).
- [71] Z. Cai, X. Xu, and S. C. Benjamin, Mitigating coherent noise using Pauli conjugation, *npj Quantum Inf.* **6**, 1 (2020).
- [72] R. N. Tazhigulov, S.-N. Sun, R. Haghshenas, H. Zhai, A. T. Tan, N. C. Rubin, R. Babbush, A. J. Minnich, and G. K. Chan, Simulating challenging correlated molecules and materials on the Sycamore quantum processor (2022), *arXiv preprint ArXiv:2203.15291*.
- [73] <https://github.com/qsppack/QETU>.
- [74] A. M. Childs, On the relationship between continuous- and discrete-time quantum walk, *Commun. Math. Phys.* **294**, 581 (2010).
- [75] M. Szegedy, in *45th Annual IEEE Symposium on Foundations of Computer Science*, p. 32, 2004.
- [76] J. M. Martyn, Z. M. Rossi, A. K. Tan, and I. L. Chuang, A grand unification of quantum algorithms (2021), *ArXiv:2105.02859*.

- [77] J. Haah, Product decomposition of periodic functions in quantum signal processing, *Quantum* **3**, 190 (2019).
- [78] Y. Atia and D. Aharonov, Fast-forwarding of Hamiltonians and exponentially precise measurements, *Nat. Commun.* **8**, 1572 (2017).
- [79] S. Gu, R. D. Somma, and B. Şahinoğlu, Fast-forwarding quantum evolution, *Quantum* **5**, 577 (2021).
- [80] Y. Su, Fast-forwardable quantum evolution and where to find them, *Quantum Views* **5**, 62 (2021).
- [81] Y. Su, H.-Y. Huang, and E. T. Campbell, Nearly tight Trotterization of interacting electrons (2020), *arXiv preprint ArXiv:2012.09194*.
- [82] D. W. Berry, C. Gidney, M. Motta, J. R. McClean, and R. Babbush, Qubitization of arbitrary basis quantum chemistry leveraging sparsity and low rank factorization, *Quantum* **3**, 208 (2019).
- [83] J. Lee, D. W. Berry, C. Gidney, W. J. Huggins, J. R. McClean, N. Wiebe, and R. Babbush, Even More Efficient Quantum Computations of Chemistry through Tensor Hypercontraction, *PRX Quantum* **2**, 030305 (2021).
- [84] S. McArdle, E. Campbell, and Y. Su, Exploiting fermion number in factorized decompositions of the electronic structure Hamiltonian, *Phys. Rev. A* **105**, 012403 (2022).
- [85] V. von Burg, G. H. Low, T. Häner, D. S. Steiger, M. Reiher, M. Roetteler, and M. Troyer, Quantum computing enhanced computational catalysis, *Phys. Rev. Res.* **3**, 033055 (2021).
- [86] A. F. Izmaylov, T.-C. Yen, R. A. Lang, and V. Verteletskyi, Unitary partitioning approach to the measurement problem in the variational quantum eigensolver method, *J. Chem. Theory Comput.* **16**, 190 (2019).
- [87] V. Verteletskyi, T.-C. Yen, and A. F. Izmaylov, Measurement optimization in the variational quantum eigensolver using a minimum clique cover, *J. Chem. Phys.* **152**, 124114 (2020).