

Randomized Quantum Algorithm for Statistical Phase Estimation

Kianna Wan^{1,2}, Mario Berta^{1,3,4} and Earl T. Campbell⁵

¹AWS Center for Quantum Computing, Pasadena 91125, USA

²Stanford Institute for Theoretical Physics, Stanford University, Stanford 94305, USA

³Institute for Quantum Information and Matter, California Institute of Technology, Pasadena 91125, USA

⁴Department of Computing, Imperial College London, London, SW7 2AZ, United Kingdom

⁵AWS Center for Quantum Computing, Cambridge, CB1 2GA, United Kingdom



(Received 17 December 2021; accepted 15 June 2022; published 13 July 2022)

Phase estimation is a quantum algorithm for measuring the eigenvalues of a Hamiltonian. We propose and rigorously analyze a randomized phase estimation algorithm with two distinctive features. First, our algorithm has complexity independent of the number of terms L in the Hamiltonian. Second, unlike previous L -independent approaches, such as those based on qDRIFT, all algorithmic errors in our method can be suppressed by collecting more data samples, without increasing the circuit depth.

DOI: [10.1103/PhysRevLett.129.030503](https://doi.org/10.1103/PhysRevLett.129.030503)

Quantum computers can be used to simulate dynamics and learn the spectra of quantum systems, such as interacting particles comprising complex molecules or materials, described by some Hamiltonian H . Phase estimation [1] on the unitary $U = e^{iHt}$ efficiently solves the common spectral problem of computing ground state energies, whenever we can efficiently prepare a trial state with nontrivial (not exponentially small) overlap η with the ground state [2] (see also Ref. [3]). Each run of standard phase estimation returns a single eigenvalue, with precision and success probability dependent on the number of times U is used.

Recently, statistical approaches to phase estimation have been proposed [4–6], where each run uses only a few ancillae and shorter circuits than standard phase estimation. As such, statistical phase estimation may be better suited to early fault-tolerant quantum computers that are qubit limited and depth limited. However, in these approaches, a single run gives a sample of an estimator of $\langle U^j \rangle$ for some runtime j , which alone is not enough to infer spectral properties. Multiple runs with different values of j are needed, and statistical analysis gives spectral information with a confidence that increases with the amount of data obtained. These runs could be massively parallelized across multiple quantum computers. Interestingly, the approach of Lin and Tong [6] is not only statistical in its analysis, but also generates the runtime j , and therefore the circuits, from a random ensemble.

The cost of phase estimation—statistical or standard—typically depends on the Hamiltonian sparsity L , the number of terms in the Hamiltonian when decomposed in a suitable basis, such as the Pauli basis. Simple schemes based on implementing U using Trotter formulas have $\mathcal{O}(L)$ gate complexity [7–11]. This can be prohibitive for the electronic structure problem in chemistry and materials science, where typically $L = \mathcal{O}(N^4)$ for an N -orbital

problem [12]. This increases to $L = \mathcal{O}(N^6)$ when using transcorrelated orbitals [13,14] to better resolve electron-electron interactions. Sublinear non-Clifford complexity $\mathcal{O}(\sqrt{L} + N)$ is possible by employing an efficient data-lookup oracle [15,16] in qubitization-based implementations of phase estimation [17–20]. However, these approaches require $\mathcal{O}(\sqrt{L})$ ancillae, which increases the qubit cost from $\mathcal{O}(N)$ to $\mathcal{O}(N^2)$, or even $\mathcal{O}(N^3)$ in the transcorrelated setting.

Heuristic truncation and low-rank factorizations have been proposed to decrease the sparsity L [18–20] of the electronic structure Hamiltonian. As an alternative approach, randomized compilation [21–23] has been rigorously shown to enable phase estimation with gate complexity that is independent of L for any Hamiltonian. A weakness of these randomized algorithms is a systematic error in energy estimates that can only be suppressed by increasing gate complexity, leading to high gate counts per run (cf. Ref. [20], Appendix D).

Here, we overcome this difficulty by combining the statistical approach of Lin and Tong [6] with a novel random compilation of each U^j instance that has parallels to—but is distinct from—both the QDRIFT random compiler [21] and the linear combinations of unitaries (LCU) method [24,25]. Our algorithm for phase estimation is doubly randomized in that we randomly sample j , then approximate U^j using a random gate sequence. Unlike in any previous approach, all approximation and compilation errors can be expressed in terms of statistical noise that is suppressed by collecting more data samples. This allows for a trade-off between the gate complexity per sample and the number of samples required. We explore this trade-off and show how to efficiently find the algorithmic parameters that minimize the total complexity. In contrast, QDRIFT

approximates U up to some systematic error that cannot be mitigated by increasing the number of samples.

Applied to ground state energy estimation, we can tune the gate vs sample trade-off to yield the following complexities. Given a Hamiltonian as a linear combination of Pauli operators with total weight λ , and an ansatz state with overlap at least η with the ground space, we can choose to sample from $\tilde{\mathcal{O}}(\eta^{-2})$ randomly compiled quantum circuits, where $\tilde{\mathcal{O}}(\cdot)$ hides polylogarithmic factors. Each circuit uses one ancilla and at most $\tilde{\mathcal{O}}(\lambda^2 \Delta^{-2})$ single-qubit Pauli rotations to estimate the ground state energy to within additive error Δ .

Problem setting.—We assume that the Hamiltonian H is specified as a linear combination of n -qubit Pauli operators P_ℓ :

$$H = \sum_{\ell=1}^L \alpha_\ell P_\ell, \quad \text{with} \quad \lambda := \sum_{\ell=1}^L |\alpha_\ell|. \quad (1)$$

This form can always be achieved, and is particularly natural for many physical systems of interest, such as fermionic Hamiltonians [26–30]. We consider the following problem of coarsely determining whether an ansatz state ρ has overlap with eigenstates of H with eigenvalues below some threshold: Given a threshold X , precision $\Delta > 0$, and overlap parameter $\eta > 0$, we seek to decide if (A) $\text{tr}[\rho \Pi_{\leq X-\Delta}] < \eta$ or (B) $\text{tr}[\rho \Pi_{\leq X+\Delta}] > 0$, where $\Pi_{\leq x}$ denotes the projector onto the eigenspaces of H with eigenvalues at most x . Both of these statements can simultaneously be true, in which case it suffices to output either A or B. We refer to this problem as *eigenvalue thresholding*, and its solution will allow us to estimate the ground state energy, given a suitable ansatz ρ .

Cumulative distribution function.—Similarly to Ref. [6], we define the cumulative distribution function (CDF) associated with the Hamiltonian H and ansatz state ρ as

$$C(x) := \text{tr}[\rho \Pi_{\leq x/\tau}], \quad (2)$$

where $\tau := \lceil \pi/(2\lambda + \Delta) \rceil$ is a normalization factor. The jump discontinuities in $C(x)$ occur at eigenvalues of τH , so appropriately characterising the CDF would enable us to estimate the spectrum of the Hamiltonian. We can write $C(x)$ as the convolution $(\Theta * p)(x)$ of the Heaviside function $\Theta(\cdot)$ and the probability density function $p(\cdot)$ corresponding to τH and ρ :

$$C(x) = \int_{-\pi/2}^{\pi/2} dy p(y) \Theta(x - y), \quad (3)$$

noting that $p(x)$ is supported within $x \in [-(\pi/2), (\pi/2)]$ since $\tau \|H\| \leq \tau \lambda < (\pi/2)$. This will enable us to replace $\Theta(x)$ with a periodic function that is a good approximation only within $x \in (-\pi, \pi)$. Eigenvalue thresholding then reduces to the following problem regarding the CDF.

Problem 1.—For given $x \in [-\tau\lambda, \tau\lambda]$ and $\delta > 0$, determine whether

$$C(x - \delta) < \eta \quad \text{or} \quad C(x + \delta) > 0, \quad (4)$$

outputting either statement if both are true.

In particular, solving Problem 1 for $x = \tau X$ and $\delta = \tau \Delta$ solves eigenvalue thresholding. Solving Problem 1 with these parameter values also solves the “eigenvalue threshold problem” [31,32], which, unlike eigenvalue thresholding, is a promise problem (where it is guaranteed that either $\text{tr}[\rho \Pi_{\leq X-\Delta}] \geq \eta$ or $\text{tr}[\rho \Pi_{\leq X+\Delta}] = 0$) and hence cannot be used as a subroutine for phase estimation in the same manner.

Algorithm overview.—To solve Problem 1, we will construct an approximation $\tilde{C}(\cdot)$ to the CDF $C(\cdot)$ satisfying

$$C(x - \delta) - \varepsilon \leq \tilde{C}(x) \leq C(x + \delta) + \varepsilon \quad (5)$$

for relevant values of x , δ , and ε . Observe that for $\varepsilon \in (0, \eta/2)$, $\tilde{C}(x) < \eta - \varepsilon$ would imply the first case of Eq. (4), while $\tilde{C}(x) > \varepsilon$ would imply the second case. Hence, it suffices to estimate $\tilde{C}(x)$.

Our algorithm is based on expressing $\tilde{C}(x)$ in terms of a linear combination of computationally simple unitaries, obtained via a two-step construction. First, we develop an improved Fourier series approximation to the Heaviside function (Lemma 1). Second, we combine this with a novel decomposition of the time evolution operators (Lemma 2) in the relevant Fourier series. Randomly sampling unitaries from our decomposition and estimating their expectation values using Hadamard tests [Fig. 1(a)] will give estimates for $\tilde{C}(x)$, allowing us to solve Problem 1 with high probability.

Fourier series approximation.—Following Lin and Tong [6], who use ideas similar to those in Refs. [33,34], we obtain an approximate CDF $\tilde{C}(\cdot)$ by replacing $\Theta(\cdot)$ in Eq. (3) with a finite Fourier series approximation thereof. As in Refs. [32,35] and related works, we need a Fourier series with small approximation error on $|x| \in [\delta, \pi - \delta]$ for fixed $\delta > 0$, small total weight of Fourier coefficients, and small maximal “time” parameter $|t|$ in the e^{itx} terms. We explicitly construct such a Fourier series by converting suitable Chebyshev approximations to the error function (see Supplemental Material [36], Sec. A, for proofs and explicit constants), establishing the following.

Lemma 1.—For any $\varepsilon > 0$ and $\delta \in (0, \pi/2)$, there exists a Fourier series $F(x) = \sum_{j \in S_1} F_j e^{ijx}$ with $S_1 := \{0\} \cup \{\pm(2j+1)\}_{j=0}^d$ and $d = \mathcal{O}(\delta^{-1} \log(\varepsilon^{-1}))$, that satisfies.

- (1) $|F(x) - \Theta(x)| \leq \varepsilon, \quad \forall x \in [-\pi + \delta, -\delta] \cup [\delta, \pi - \delta],$
- (2) $-\varepsilon \leq |F(x)| \leq 1 + \varepsilon, \quad \forall x \in \mathbb{R},$
- (3) $\mathcal{F} := \sum_{j \in S_1} |F_j| = \mathcal{O}(\log d).$

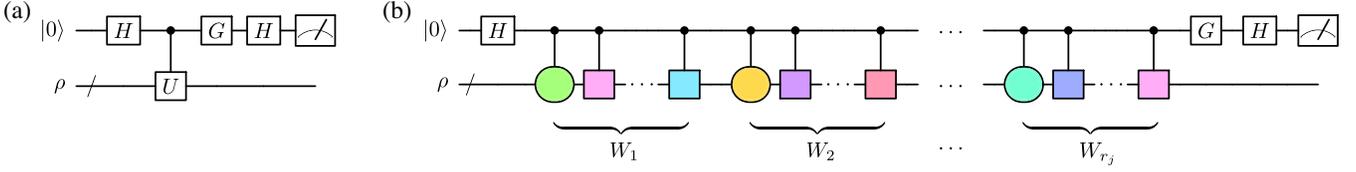


FIG. 1. (a) Hadamard test on ρ and U : setting $G = \mathbb{1}$ (respectively $G = S^\dagger := |0\rangle\langle 0| - i|1\rangle\langle 1|$) and associating the measurement outcomes ($|0\rangle, |1\rangle$) with $(+1, -1)$ produces an unbiased estimator for $\text{Re}(\text{tr}[\rho U])$ [respectively $\text{Im}(\text{tr}[\rho U])$]. (b) Schematic depiction of the randomly compiled circuits in our algorithm. For $\hat{H} = \sum_{\ell} p_{\ell} P_{\ell}$, the squares represent Pauli operators randomly drawn from $\{P_{\ell}\}_{\ell}$ according to $\{p_{\ell}\}_{\ell}$, while circles denote multiqubit Pauli rotations; see the proof of Lemma 2 for details.

This improves on the Fourier approximation of Lin and Tong, which has $d = \mathcal{O}(\delta^{-1} \log(\delta^{-1} \epsilon^{-1}))$ [6, Lemma 6]. As such, Lemma 1 also improves the asymptotic complexity of their phase estimation algorithm. Using our Fourier series $F(\cdot)$, we obtain the approximate CDF

$$\tilde{C}(x) := \int_{-\pi/2}^{\pi/2} dy p(y) F(x - y) = \sum_{j \in S_1} F_j e^{ijx} \text{tr}[\rho e^{i\hat{H}t_j}] \quad (6)$$

where $t_j := -j\tau\lambda$ and $\hat{H} := H/\lambda$, which for any $\delta \in (0, \tau\Delta]$ indeed satisfies the guarantees in Eq. (5) (see Supplemental Material [36], Sec. B).

LCU decomposition of time evolution operators.— Instead of directly implementing the time evolution operators $e^{i\hat{H}t_j}$ from Eq. (6) in Hadamard tests, as considered by Ref. [6], we further decompose each of these terms into a specific linear combination of unitaries.

Lemma 2.—Let $\hat{H} = \sum_{\ell} p_{\ell} P_{\ell}$ be a Hermitian operator that is specified as a convex combination of Pauli operators. For any $t \in \mathbb{R}$ and $r \in \mathbb{N}$, there exists a linear decomposition

$$e^{i\hat{H}t} = \sum_{k \in S_2} b_k U_k$$

for some index set S_2 , real numbers $b_k > 0$, and unitaries U_k , such that

$$\sum_{k \in S_2} b_k \leq \exp(t^2/r),$$

and for all $k \in S_2$, the non-Clifford cost of controlled- U_k is that of r controlled single-qubit Pauli rotations.

This decomposition is conceptually different from previous LCU methods (cf. [25] and references therein). The purpose of Lemma 2 is to allow for a trade-off between the sample complexity and gate complexity of our algorithm. Specifically, as shown later, the sample complexity depends on the total weight $\sum_{k \in S_2} b_k$. Since this is bounded by $\exp(t^2/r)$, we can reduce the sample complexity by increasing r , at the cost of increasing the gate complexity per sample, and vice versa.

To prove Lemma 2, we write $e^{i\hat{H}t} = (e^{i\hat{H}t/r})^r$ and Taylor expand each $e^{i\hat{H}t/r} = \mathbb{1} + i\hat{H}t/r + \mathcal{O}((t/r)^2)$. We then pair up consecutive terms in this expansion, which differ in phase by i . Since \hat{H} is a convex combination of Pauli operators, this gives rise to convex combinations of multi-qubit Pauli rotations, e.g., the leading term is

$$\mathbb{1} + i\hat{H}t/r = \sum_{\ell} p_{\ell} (\mathbb{1} + iP_{\ell}t/r) \propto \sum_{\ell} p_{\ell} e^{i\theta P_{\ell}} \quad (7)$$

with $\theta := \arccos \sqrt{1 + (t/r)^2}$. The higher-order terms contain additional Pauli operators, as illustrated in Fig. 1(b). The controlled version of each Pauli rotation can be implemented using a controlled single-qubit rotation, along with Clifford gates. Hence, each controlled- U_k requires r controlled single-qubit rotations in total. Explicit forms for the higher-order terms and proof details are given Supplemental Material [36], Sec. C.

Our algorithm for Problem 1.—Putting together the above results, we apply Lemma 2 to decompose each $e^{i\hat{H}t_j}$ in Eq. (6) as $e^{i\hat{H}t_j} = \sum_{k \in S_2} b_k^{(j)} U_k^{(j)}$. We choose a positive integer r_j for each $j \in S_1$, and define the corresponding “runtime vector” $\vec{r} = (r_j)_j \in \mathbb{N}^{S_1}$. This leads to the final decomposition

$$\tilde{C}(x) = \sum_{(j,k) \in S_1 \times S_2} \underbrace{F_j e^{ijx} b_k^{(j)}}_{=: a_{jk}} \text{tr}[\rho U_k^{(j)}] \quad (8)$$

with total weight

$$\mathcal{A}(\vec{r}) := \sum_{(j,k) \in S_1 \times S_2} |a_{jk}| \leq \sum_{j \in S_1} |F_j| \exp(t_j^2/r_j). \quad (9)$$

As a simple example,

$$r_j = \lceil 2t_j^2 \rceil \quad \forall j \in S_1 \text{ gives } \mathcal{A}(\vec{r}) \leq \sqrt{e} \mathcal{F}. \quad (10)$$

Recall that we can solve Problem 1 by determining if $\tilde{C}(x) < \eta - \epsilon$ or $\tilde{C}(x) > \epsilon$. To estimate $\tilde{C}(x)$, we sample (j, k) from $S_1 \times S_2$ with probability proportional to $|a_{jk}|$, and perform a Hadamard test on ρ and $U_k^{(j)}$, obtaining an

Algorithm 1. Algorithm for Problem 1.

Problem inputs: an n -qubit Hamiltonian $H = \sum_{\ell=1}^L \alpha_{\ell} P_{\ell}$ with $\alpha_{\ell} > 0$ and $\lambda := \sum_{\ell=1}^L \alpha_{\ell}$, an ansatz state ρ , a precision parameter $\Delta > 0$; $\tau := \lceil \pi / (2\lambda + \Delta) \rceil$.

Algorithm parameters: real numbers $x \in [-\tau\lambda, \tau\lambda]$, $\delta \in (0, \tau\Delta]$, $\eta \in (0, 1]$, and $\varepsilon \in (0, \eta/2)$, a probability ϑ .

Output: 0 if $C(x - \delta) < \eta$, 1 if $C(x + \delta) > 0$, and either 0 or 1 if both are true [where $C(\cdot)$ is the CDF defined in Eq. (2)] with probability of error at most ϑ .

1: Compute the coefficients $\{F_{jj}\}_{j \in S_1}$ of the Fourier series from Lemma 1 with approximation parameters δ and ε . Set $t_j \leftarrow -j\lambda\tau \forall j \in S_1$.

2: Choose a runtime vector $\vec{r} \in \mathbb{N}^{S_1}$ [using, e.g., Eq. (10) or Eq. (13)], and apply Lemma 2 to obtain the decomposition in Eq. (8), with total weight $\mathcal{A}(\vec{r})$ as in Eq. (9).

3:
$$C_{\text{sample}}(\vec{r}) \leftarrow \lceil \left(\frac{2\mathcal{A}(\vec{r})}{\eta/2 - \varepsilon} \right)^2 \ln \frac{1}{\vartheta} \rceil. \quad (11)$$

4: **For** $i = 1, \dots, C_{\text{sample}}(\vec{r})$:

5: Sample a unitary $U_k^{(j)}$ from Eq. (8).

6: Perform a Hadamard test with inputs ρ and $U_k^{(j)}$, obtaining an estimate m_i of $\text{tr}[\rho U_k^{(j)}]$.

7: $z_i \leftarrow \mathcal{A}(\vec{r}) e^{i \arg(F_{jj} + j\lambda)} m_i$

8: $\bar{z} \leftarrow \sum_i z_i / C_{\text{sample}}(\vec{r})$. If $\text{Re}(\bar{z}) < \eta/2$, **return** 0. Else, **return** 1.

estimate m_{jk} for $\text{tr}[\rho U_k^{(j)}]$. Then, $z_{jk} := \mathcal{A}(\vec{r}) e^{i \arg(a_{jk})} m_{jk}$ is an unbiased estimate of $\tilde{C}(x)$. Letting \bar{Z} denote the average of C_{sample} such estimates, it follows from Hoeffding's inequality that guessing $\tilde{C}(x) < \eta - \varepsilon$ if $\text{Re}[\bar{Z}] < \eta/2$ and $\tilde{C}(x) > \varepsilon$ otherwise gives a correct answer with probability at least $1 - \vartheta$ provided that $C_{\text{sample}} \geq 4\mathcal{A}(\vec{r})^2 (\eta/2 - \varepsilon)^{-2} \ln(\vartheta^{-1})$. Thus, we arrive at Algorithm 1, our algorithm for solving Problem 1, and hence eigenvalue thresholding.

Complexity.—The Hadamard test in Step 6 is the only quantum step and involves two circuits on $n + 1$ qubits. The expected number of controlled Pauli rotations per circuit is

$$C_{\text{gate}}(\vec{r}) := \frac{1}{\mathcal{A}(\vec{r})} \sum_{(j,k) \in S_1 \times S_2} |a_{jk}| r_j. \quad (12)$$

Step 6 is repeated $C_{\text{sample}}(\vec{r})$ times, so the expected total non-Clifford complexity is $2C_{\text{sample}}(\vec{r}) \cdot C_{\text{gate}}(\vec{r})$.

It remains to specify how to choose the runtime vector $\vec{r} \in \mathbb{N}^{S_1}$. For example, we could aim to minimize the total complexity

$$\underset{\vec{r}}{\text{argmin}} C_{\text{sample}}(\vec{r}) \cdot C_{\text{gate}}(\vec{r}). \quad (13)$$

Prima facie this is a high-dimensional optimization problem, as $|S_1| = \mathcal{O}(\delta^{-1} \log(\varepsilon^{-1}))$ from Lemma 1. However,

differentiating with respect to \vec{r} , one sees that the argmin is effectively described by a single free parameter. Therefore, optimizing \vec{r} is reducible to an efficiently solvable one-dimensional problem, and this further holds when minimizing C_{sample} subject to constraints on C_{gate} (see Supplemental Material [36], Sec. D). Moreover, if one is exclusively interested in asymptotic complexities, the simple choice for \vec{r} in Eq. (10) already gives

$$C_{\text{sample}}(\vec{r}) = \mathcal{O}\left(\frac{1}{\eta^2} \log^2\left(\frac{1}{\delta} \log \frac{1}{\eta}\right) \log \frac{1}{\vartheta}\right) = \tilde{\mathcal{O}}\left(\frac{1}{\eta^2}\right) \quad (14)$$

$$\text{and } C_{\text{gate}}(\vec{r}) = \mathcal{O}\left(\frac{1}{\delta^2} \log^2 \frac{1}{\eta}\right) = \tilde{\mathcal{O}}\left(\frac{1}{\delta^2}\right), \quad (15)$$

since $C_{\text{gate}}(\vec{r}) \leq \max_{j \in S_1} r_j = 2[(2d + 1)\tau\lambda]^2$ and $\mathcal{A}(\vec{r}) \leq \sqrt{e\mathcal{F}}$ for this choice, with \mathcal{F} and d given by Lemma 1 and picking $\varepsilon = \text{const} \times \eta$ in Algorithm 1. Note that the *worst-case* gate complexity thus has the same scaling as that in Eq. (15) for the *expected* gate complexity $C_{\text{gate}}(\vec{r})$. Hence, we arrive at a total complexity $\tilde{\mathcal{O}}(\delta^{-2}\eta^{-2})$. For eigenvalue thresholding, one would choose $\delta = \tau\Delta$, in which case $\delta^{-1} = \mathcal{O}(\lambda/\Delta)$.

Ground state energy estimation.—Under appropriate assumptions on the Hamiltonian H and ansatz state ρ , our method for estimating the CDF can be adapted to perform phase estimation. Specifically, eigenvalues of H coincide with the locations of discontinuities in $C(x)$, and we can estimate these locations given sufficient knowledge about the overlap of ρ with relevant eigenspaces. For simplicity, we restrict ourselves to the problem of estimating the ground state energy $[H]_{\min}$, which only requires the standard promise that $\text{tr}[\rho \Pi_{\min}] \geq \eta$ for some $\eta > 0$, where Π_{\min} denotes the projector onto the ground space of H .

The analysis in Ref. [[6], Sec. 5] shows that by solving Problem 1 for $s = \mathcal{O}(\log(\delta^{-1}))$ different values of x determined in a fashion similar to binary search, one can find an x^* such that $C(x^* - \delta) < \eta$ and $C(x^* + \delta) > 0$, which implies $|x^* - \tau[H]_{\min}| \leq \delta$. Hence, if we take $\delta = \tau\Delta$, then x^*/τ would give an estimate of the ground state energy to within additive error Δ . We use Algorithm 1 to solve Problem 1, noting that we can reuse the samples collected in Step 6 for *all* of the different x values, with only a small overhead in the sample complexity. Namely, since Algorithm 1 errs with probability at most ϑ for any x , choosing $\vartheta = \xi/s$ would ensure that the ground state is successfully estimated with probability at least $1 - \xi$.

Theorem 1.—For any n -qubit Hamiltonian H of the form in Eq. (1), let ρ be a state that has overlap $\text{tr}[\rho \Pi_{\min}] \geq \eta$ with the ground space of H . Then, the ground state energy of H can be estimated to within additive error Δ with probability at least $1 - \xi$ using $\mathcal{O}((1/\eta^2) \log^2[(\lambda/\Delta) \log(1/\eta)] \times \log[(1/\xi) \log(\lambda/\Delta)])$ quantum circuits on $n + 1$ qubits. Each circuit uses one copy of ρ and at most $\mathcal{O}((\lambda^2/\Delta^2) \log^2(1/\eta))$ single-qubit Pauli rotations.

Thus, our quantum complexities are independent of the Hamiltonian sparsity L , at the price of the quadratic dependence $\tilde{\mathcal{O}}(\lambda^2 \Delta^{-2} \eta^{-2})$ for the total gate count. This is in contrast to standard results on phase estimation (see, e.g., Ref. [20], Table I). Additionally, note that Theorem 1 is derived using the specific choice of runtime vector in Eq. (10). By tuning \vec{r} (using for instance the optimization procedures in the Supplemental Material [36], Sec. D), we can reduce the gate complexity per circuit by running more circuits, for a given set of problem parameters.

Comparisons.—Conventional phase estimation algorithms depend on the sparsity L , which is especially prohibitive for chemistry Hamiltonians. Several algorithms [20–18]] have used heuristic truncation policies to justify eliminating certain terms from the Hamiltonians, thereby reducing L . While supporting numerics were presented, these truncations are not rigorous. It was also assumed that only a single run of the algorithm suffices. In practice, a single sample might return an incorrect result due to imperfect overlap with the ground state ($\eta < 1$), inherent failure probabilities of phase estimation, or quantum error correction failure events. In contrast, our algorithm is rigorously analyzed; we use no Hamiltonian truncation, and upper bound the number of samples in terms of η and the target success probability. (We neglect quantum error correction failure events, though these can easily be suppressed to lower levels than other failure modes.)

Hydrogen chains.—As a benchmark system for assessing the scaling of quantum algorithms applied to quantum chemistry, we discuss hydrogen chains [20,37]. Using the best value $\lambda \sim \mathcal{O}(N^{1.34})$ given in Ref. [37], our algorithm scales as $\tilde{\mathcal{O}}(N^{2.68}/\Delta^2)$. For comparison, the scaling of qubitization is $\tilde{\mathcal{O}}(N^{3.34}/\Delta)$ without truncation, and with heuristic truncations, $\tilde{\mathcal{O}}(N^{2.3}/\Delta)$ for the sparse method of Ref. [18] and $\tilde{\mathcal{O}}(N^{2.1}/\Delta)$ for the tensor hypercontraction approach of Ref. [20]. Hence, for constant Δ , qubitization gives better scaling than our algorithm if the proposed truncation schemes are accurate. However, we emphasize that our rigorous analysis does not make use of heuristic strategies for truncating Hamiltonian terms [18–20] and that qubitization uses considerably more logical ancillae. Finally, if we are interested in extensive properties where $\Delta \propto N$, then our approach scales as $\tilde{\mathcal{O}}(N^{0.68})$, outperforming all qubitization algorithms.

FeMoco.—We estimate the costs of our algorithm applied to the Li *et al.* FeMoco Hamiltonian [38], another popular benchmark for which there have been several state-of-the-art resource studies [18–20]. We consider chemical accuracy $\Delta = 0.0016$ Hartree, and use $\lambda = 1511$ Hartree, obtained using the bounds in Ref. [37]. We present our results in Fig. 2, illustrating the trade-off between the expected number of gates per circuit and the number of

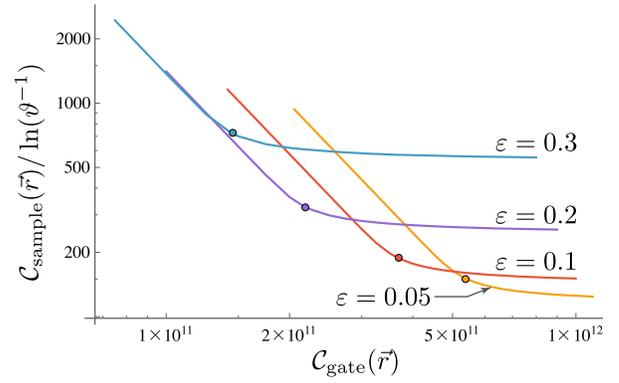


FIG. 2. Log-log plot of $C_{\text{sample}}(\vec{r})/\ln(\vartheta^{-1})$ vs $C_{\text{gate}}(\vec{r})$, for $\lambda = 1511$ (FeMoco [37,38]), $\Delta = 0.0016$ (chemical accuracy), $\eta = 1$, and $\epsilon \in \{0.05, 0.1, 0.2, 0.3\}$, for various choices of the optimized runtime vector \vec{r} . Dots indicate the values that minimize the total expected complexity $2C_{\text{sample}} \cdot C_{\text{gate}}$, while curves are obtained by fixing C_{gate} and minimizing C_{sample} . To calculate the number of samples needed to guarantee an overall failure probability $\leq \xi$ for ground state energy estimation (Theorem 1), one would multiply the y axis by $\ln(\vartheta^{-1}) = \mathcal{O}(\log(\xi^{-1}) + \log \log(\delta^{-1}))$ (see Ref. [6] for the explicit constants). As an example, $\xi = 0.1$ would give a multiplier of approximately 6.

samples required. Since the Hamiltonian from Ref. [38] has 152 spin orbitals, each circuit uses 153 qubits.

We have presented our gate counts as C_{gate} controlled Pauli rotations, but asymptotically our circuits can typically be realized using $\sim 2C_{\text{gate}}$ Toffoli gates. For modest system sizes and a modest number of logical ancilla (~ 40), the Toffoli count is $\sim 6C_{\text{gate}}$ (see Supplemental Material [36], Sec. F). The FeMoco resource estimate for the QDRIFT random compiler combined with phase estimation in Ref. [20], Appendix D] arrived at 10^{16} Toffoli gates per sample, which is $\sim 10^4$ times larger than $2C_{\text{gate}}$ from the results in Fig. 2. Moreover, our rigorous analysis will likely be loose and overestimate resources; for instance, more aggressive—though heuristic—Hamiltonian rescaling is justifiable and can further reduce costs (see Supplemental Material [36], Sec. G).

We thank Sam McArdle for helpful discussions, especially with regard to calculations for the quantum chemistry examples, and Fernando Brandão for discussions and support throughout this project.

-
- [1] A. Y. Kitaev, A. Shen, M. N. Vyalyi, and M. N. Vyalyi, *Classical and Quantum Computation* (American Mathematical Society, Providence, 2002), [10.1090/gsm/047](https://doi.org/10.1090/gsm/047).
 - [2] D. S. Abrams and S. Lloyd, Quantum Algorithm Providing Exponential Speed Increase for Finding Eigenvalues and Eigenvectors, *Phys. Rev. Lett.* **83**, 5162 (1999).

- [3] D. Poulin, A. Kitaev, D. S. Steiger, M. B. Hastings, and M. Troyer, Quantum Algorithm for Spectral Measurement with a Lower Gate Count, *Phys. Rev. Lett.* **121**, 010501 (2018).
- [4] 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).
- [5] A. Dutkiewicz, B. M. Terhal, and T. E. O'Brien, Heisenberg-limited quantum phase estimation of multiple eigenvalues with a single control qubit, [arXiv:2107.04605](https://arxiv.org/abs/2107.04605).
- [6] L. Lin and Y. Tong, Heisenberg-limited ground state energy estimation for early fault-tolerant quantum computers, *PRX Quantum* **3**, 010318 (2022).
- [7] D. Poulin, M. B. Hastings, D. Wecker, N. Wiebe, A. C. Doherty, and M. Troyer, The Trotter step size required for accurate quantum simulation of quantum chemistry, *Quantum Inf. Comput.* **15**, 361 (2015).
- [8] R. Babbush, J. McClean, D. Wecker, A. Aspuru-Guzik, and N. Wiebe, Chemical basis of Trotter-Suzuki errors in quantum chemistry simulation, *Phys. Rev. A* **91**, 022311 (2015).
- [9] I. D. Kivlichan, C. Gidney, D. W. Berry, N. Wiebe, J. McClean, W. Sun, Z. Jiang, N. Rubin, A. Fowler, A. Aspuru-Guzik *et al.*, Improved fault-tolerant quantum simulation of condensed-phase correlated electrons via trotterization, *Quantum* **4**, 296 (2020).
- [10] E. T. Campbell, Early fault-tolerant simulations of the Hubbard model, *Quantum Sci. Technol.* **7**, 015007 (2021).
- [11] S. McArdle, E. Campbell, and Y. Su, Exploiting fermion number in factorized decompositions of the electronic structure Hamiltonian, *Phys. Rev. A* **105**, 012403 (2022).
- [12] T. Helgaker, P. Jorgensen, and J. Olsen, *Molecular Electronic-Structure Theory* (John Wiley & Sons, New York, 2014), [10.1002/9781119019572](https://doi.org/10.1002/9781119019572).
- [13] M. Motta, T. P. Gujarati, J. E. Rice, A. Kumar, C. Masteran, J. A. Latone, E. Lee, E. F. Valeev, and T. Y. Takeshita, Quantum simulation of electronic structure with transcorrelated Hamiltonian: Increasing accuracy without extra quantum resources, *Phys. Chem. Chem. Phys.* **22**, 24270 (2020).
- [14] S. McArdle and D. P. Tew, Improving the accuracy of quantum computational chemistry using the transcorrelated method, [arXiv:2006.11181](https://arxiv.org/abs/2006.11181).
- [15] R. Babbush, C. Gidney, D. W. Berry, N. Wiebe, J. McClean, A. Paler, A. Fowler, and H. Neven, Encoding Electronic Spectra in Quantum Circuits with Linear T Complexity, *Phys. Rev. X* **8**, 041015 (2018).
- [16] G. H. Low, V. Kliuchnikov, and L. Schaeffer, Trading T-gates for dirty qubits in state preparation and unitary synthesis, [arXiv:1812.00954](https://arxiv.org/abs/1812.00954).
- [17] R. Babbush, D. W. Berry, J. R. McClean, and H. Neven, Quantum simulation of chemistry with sublinear scaling in basis size, *npj Quantum Inf.* **5**, 92 (2019).
- [18] 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).
- [19] 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. Research* **3**, 033055 (2021).
- [20] 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).
- [21] E. Campbell, Random Compiler for Fast Hamiltonian Simulation, *Phys. Rev. Lett.* **123**, 070503 (2019).
- [22] I. D. Kivlichan, C. E. Granade, and N. Wiebe, Phase estimation with randomized Hamiltonians, [arXiv:1907.10070](https://arxiv.org/abs/1907.10070).
- [23] Y. Ouyang, D. R. White, and E. T. Campbell, Compilation by stochastic Hamiltonian sparsification, *Quantum* **4**, 235 (2020).
- [24] A. M. Childs and N. Wiebe, Hamiltonian simulation using linear combinations of unitary operations, *Quantum Inf. Comput.* **12**, 901 (2012).
- [25] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, Simulating Hamiltonian Dynamics with a Truncated Taylor Series, *Phys. Rev. Lett.* **114**, 090502 (2015).
- [26] P. Jordan and E. Wigner, Über das Paulische Äquivalenzverbot, *Z. Phys.* **47**, 631 (1928).
- [27] F. Verstraete and J. I. Cirac, Mapping local Hamiltonians of fermions to local Hamiltonians of spins, *J. Stat. Mech.* (2005) P09012.
- [28] J. T. Seeley, M. J. Richard, and P. J. Love, The Bravyi-Kitaev transformation for quantum computation of electronic structure, *J. Chem. Phys.* **137**, 224109 (2012).
- [29] V. Havlíček, M. Troyer, and J. D. Whitfield, Operator locality in the quantum simulation of fermionic models, *Phys. Rev. A* **95**, 032332 (2017).
- [30] C. Derby, J. Klassen, J. Bausch, and T. Cubitt, Compact fermion to qubit mappings, *Phys. Rev. B* **104**, 035118 (2021).
- [31] J. M. Martyn, Z. M. Rossi, A. K. Tan, and I. L. Chuang, A grand unification of quantum algorithms, *PRX Quantum* **2**, 040203 (2021).
- [32] A. Gilyén, Y. Su, G. H. Low, and N. Wiebe, Quantum singular value transformation and beyond: Exponential improvements for quantum matrix arithmetics, in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing* (Association for Computing Machinery, New York, 2019), pp. 193–204, [10.1145/3313276.3316366](https://doi.org/10.1145/3313276.3316366).
- [33] R. Somma, G. Ortiz, J. E. Gubernatis, E. Knill, and R. Laflamme, Simulating physical phenomena by quantum networks, *Phys. Rev. A* **65**, 042323 (2002).
- [34] R. D. Somma, Quantum eigenvalue estimation via time series analysis, *New J. Phys.* **21**, 123025 (2019).
- [35] J. van Apeldoorn, A. Gilyén, S. Gribling, and R. de Wolf, Quantum SDP-solvers: Better upper and lower bounds, *Quantum* **4**, 230 (2020).
- [36] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevLett.129.030503> for full proofs of all our findings, which includes Refs. [6,10,32,37–53].
- [37] E. Koridon, S. Yalouz, B. Senjean, F. Buda, T. E. O'Brien, and L. Visscher, Orbital transformations to reduce the

- 1-norm of the electronic structure Hamiltonian for quantum computing applications, *Phys. Rev. Research* **3**, 033127 (2021).
- [38] Z. Li, J. Li, N. S. Dattani, C. J. Umrigar, and G. K.-L. Chan, The electronic complexity of the ground-state of the FeMo cofactor of nitrogenase as relevant to quantum simulations, *J. Chem. Phys.* **150**, 024302 (2019).
- [39] G. H. Low and I. L. Chuang, Hamiltonian simulation by uniform spectral amplification, [arXiv:1707.05391](https://arxiv.org/abs/1707.05391).
- [40] S. Sachdeva and N. K. Vishnoi, Faster algorithms via approximation theory, *Found. Trends Theor. Comput. Sci.* **9**, 125 (2014).
- [41] V. Makis, A. Volodin, and M. Short, Improved inequalities for the Poisson and binomial distribution and upper tail quantile functions, *ISRN Probab. Stat.* **2013**, 412958 (2013).
- [42] P. Kasperkovitz, Asymptotic approximations for modified Bessel functions, *J. Math. Phys. (N.Y.)* **21**, 6 (1980).
- [43] S. Gribling, I. Kerenidis, and D. Szilágyi, Improving quantum linear system solvers via a gradient descent perspective, [arXiv:2109.04248](https://arxiv.org/abs/2109.04248).
- [44] D. Gottesman, The Heisenberg representation of quantum computers, [arXiv:quant-ph/9807006](https://arxiv.org/abs/quant-ph/9807006).
- [45] S. Aaronson and D. Gottesman, Improved simulation of stabilizer circuits, *Phys. Rev. A* **70**, 052328 (2004).
- [46] A. Bocharov, M. Roetteler, and K. M. Svore, Efficient Synthesis of Universal Repeat-Until-Success Quantum Circuits, *Phys. Rev. Lett.* **114**, 080502 (2015).
- [47] V. Kliuchnikov, D. Maslov, and M. Mosca, Asymptotically Optimal Approximation of Single Qubit Unitaries by Clifford and T Circuits Using a Constant Number of Ancillary Qubits, *Phys. Rev. Lett.* **110**, 190502 (2013).
- [48] D. Gosset, V. Kliuchnikov, M. Mosca, and V. Russo, An algorithm for the T-count, *Quantum Inf. Comput.* **14**, 1261 (2014).
- [49] A. Bocharov, M. Roetteler, and K. M. Svore, Efficient synthesis of probabilistic quantum circuits with fallback, *Phys. Rev. A* **91**, 052317 (2015).
- [50] N. J. Ross and P. Selinger, Optimal ancilla-free Clifford + T approximation of z-rotations, *Quantum Inf. Comput.* **16**, 901 (2016).
- [51] E. Campbell, Shorter gate sequences for quantum computing by mixing unitaries, *Phys. Rev. A* **95**, 042306 (2017).
- [52] C. Gidney, Halving the cost of quantum addition, *Quantum* **2**, 74 (2018).
- [53] 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).