Quantifying math xmlns="http://www.w3.org/1998/Math/MathML">mi>T/mi> /math>-gate-count improvements for ground-state-energy estimation with near-optimal state preparation

S. Pathak, A. E. Russo, S. K. Seritan, and A. D. Baczewski

Phys. Rev. A **107**, L040601 — Published 17 April 2023

DOI: 10.1103/PhysRevA.107.L040601

Quantifying T-count improvements for ground state energy estimation with near-optimal state preparation

S. Pathak,[1] A.E. Russo,[1] S.K. Seritan,[2] and A.D. Baczewski[1]

[1] *Quantum Algorithms and Applications Collaboratory, Sandia National Laboratories, Albuquerque NM, USA*
[2] *Quantum Algorithms and Applications Collaboratory, Sandia National Laboratories, Livermore CA, USA*

We study the question of when investing additional quantum resources in preparing a ground state will improve the aggregate runtime associated with estimating its energy. We analyze Lin and Tong's near-optimal state preparation algorithm and show that it can reduce a proxy for the runtime, the T-gate count, of ground state energy estimation near quadratically. Resource estimates are provided that specify the conditions under which the added cost of state preparation is worthwhile.

Introduction.— A key task in all quantum simulation algorithms is the preparation of a state that encodes the observables of a physical system of interest [1–3]. This is often the ground state $|\Psi_0\rangle$ of a Hamiltonian $\mathcal{H}$ on an n-qubit Hilbert space [4–8]. Independent of whether this represents interacting electrons [9–14], spins [15–17], or quantum fields [18–20], we generally produce an approximation to the desired state $|\Phi_0\rangle$ with overlap $\gamma = |\langle\Phi_0|\Psi_0\rangle|$. Then the probability of successfully processing $|\Phi_0\rangle$ to estimate some observable of interest (e.g., the ground state energy E_0) is generally upper bounded by γ^2 [21], which would ideally be 1.

While it is likely not possible to efficiently prepare ground states of *generic* local Hamiltonians on quantum computers [22, 23] physical arguments suggest that the *specific* instances for which nature can efficiently find the ground state will be efficiently preparable on a quantum computer [24]. Though the question of which instances these are remains an active area of research [25], it is generally of interest to develop algorithms that increase γ_i from some easy-to-prepare initial approximation, $|\Phi_{0,i}\rangle$, to γ_f with some associated final approximation $|\Phi_{0,f}\rangle$. It might be that $|\Phi_{0,i}\rangle$ comes from the outcome of a classical calculation (e.g., an approximate solution to a mean-field theory) or a hybrid quantum-classical approach like the variational quantum eigensolver [26, 27]. Regardless of the source of $|\Phi_{0,i}\rangle$, it is often assumed that the cost of preparing it is negligible relative to the cost of boosting $\gamma_i \to \gamma_f$, making use of some unitary $\mathcal{U}_{SP}(\mathcal{H})\,|\Phi_{0,i}\rangle = |\Phi_{0,f}\rangle$ [28].

The question that we answer in this Letter is "when does the added cost of implementing $\mathcal{U}_{SP}(\mathcal{H})$ outweigh the cost of repeated trials with a lower probability of success?" We are specifically concerned with the potential benefit to estimating E_0, for which conventional approaches that apply quantum phase estimation (QPE) [29] to $|\Phi_{0,i}\rangle$ will project onto $|\Psi_0\rangle$ after a single round with probability γ_i^2 [21] (see Fig. 1(a)). One can use more elaborate strategies in which repeated rounds of QPE can iteratively improve our knowledge of E_0 [2, 30–35], though we assume a simple strategy of repeating the circuit $\mathcal{O}(\gamma_i^{-2})$ times for ease of analysis [36].

Broadly speaking there are two classes of state preparation algorithms, those that make use of the adiabatic theorem [37] and those that apply a filter in the eigenbasis of $\mathcal{H}$ [4]. In this Letter, we consider a $\mathcal{U}_{SP}(\mathcal{H})$ in the latter category, derived from a near-optimal approach of Lin and Tong [7]. While adiabatic state preparation is conceptually straightforward, it generally requires time-dependent Hamiltonian simulation, the analysis of a family of Hamiltonians along the entire adiabatic pathway, and potentially many different initial Hamiltonians, rather than a single instance. It also has worse scaling with the minimum spectral gap [38]. However, there are some important exceptions [39] and it may be the case that adiabatic algorithms are found to be the optimal solution in some cases, so we make no claim to the optimality of our results. An added benefit of analyzing filter-based state preparation is that it relies on a block encoding [40] of $\mathcal{H}$ that could be identical to one used in QPE, making it straightforward to compare costs.

Whether exponential, polynomial, or nonexistent, the advantages realized by quantum computers in physical simulation are likely to be problem-specific and depend critically on the cost of implementing $\mathcal{U}_{SP}(\mathcal{H})$ [41]. We provide estimates for the cost of state preparation based on the T-gate count of implementations of Lin and Tong's $\mathcal{U}_{SP}(\mathcal{H})$ [7]. We choose the total number of T gates as a simple-to-compute proxy for an actual runtime estimate because their implementation will dominate the runtime in T-factory-limited surface code architectures [42, 43] and a more precise analysis would involve detailed scheduling of the algorithm's implementation. Actual runtimes could also be reduced relative to this proxy in contexts in which the availability of magic states is not a limiting factor, in which case the T depth would be the more appropriate quantifier.

We propose a ratio ι that compares the T-gate count for $\mathcal{U}_{SP}(\mathcal{H})$ and subsequent QPE to the count for "trivial" state prep and QPE repeated until success. This quantifies the improvement in runtime associated with "better" state preparation and allows us to answer the titular question. Even in a scenario where γ_i vanishes exponentially with increasing n, there are parameter regimes where there is a robust near-quadratic speedup for better state preparation. This is consistent with

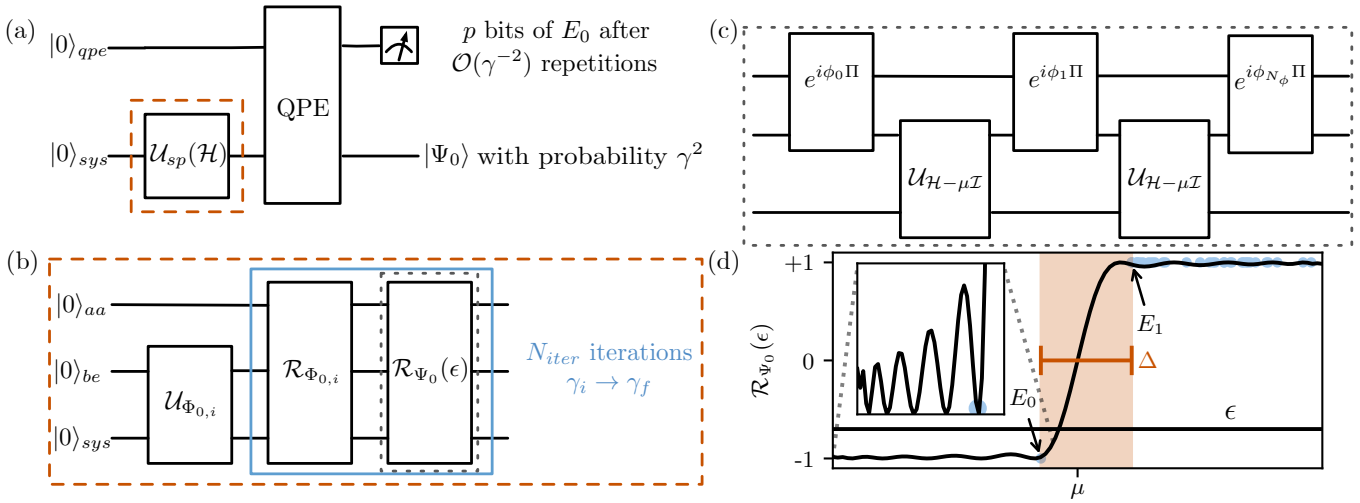

FIG. 1. Circuit diagrams and schematic visualizations for estimating the ground state energy E_0 of a Hamiltonian $\mathcal{H}$ with a near-optimal ground state preparation technique using amplitude amplification. (a) An energy estimation circuit like the one in the reference implementation under consideration [10] in which the probability of correctly measuring E_0 in a single round is γ^2. The value of γ is determined by details of $\mathcal{U}_{sp}(\mathcal{H})$. (b) A circuit that implements $\mathcal{U}_{sp}$ by first preparing an initial guess $|\Phi_{0,i}\rangle$ and then running N_{iter} rounds of AA in which the reflector about $|\Phi_{0,i}\rangle$ ($\mathcal{R}_{\Phi_{0,i}}$) is straightforward to implement and the reflector about $|\Psi_0\rangle$ ($\mathcal{R}_{\Psi_0}(\epsilon)$) requires the use of quantum signal processing, as proposed in Ref. [7], to produce $|\Phi_{0,f}\rangle$. N_{iter} is determined by γ_i and the desired final overlap γ_f. The former overlap might vanish exponentially in the size of *sys*. (c) A circuit that uses quantum signal processing to implement an ϵ-approximation to a degree-N_ϕ polynomial that suppresses support on energy eigenstates above a known lower bound (μ), in favor of support on energy eigenstates below μ. The phases are chosen using the method in Ref. [44]. (d) An illustration of the ϵ-approximate reflector, $\mathcal{R}_{\Psi_0}(\epsilon)$, and how it acts on eigenstates above and below the energy gap Δ. Here, the calculated phases realize an approximation that clearly outperforms the target $\epsilon = 0.3$.

the Grover-like speedup [45] that one would expect [46], though we note that we are exploring this in terms of T-gate counts as a proxy for runtime, instead of query complexity or some other more abstract quantifier.

Methods.— Circuit diagrams that illustrate the implementation costs being studied are provided in Fig. 1. The quantum computer that runs these circuits consists of four registers: (*sys*) the n-qubit system register that encodes $|\Phi_{0,i}\rangle$, (*qpe*) the auxiliary register that encodes p bits of an estimate for E_0, (*aa*) which implement amplitude amplification (AA), and (*be*) which block encode $\mathcal{H}$. In what follows, we describe the structure of the circuits in Fig. 1 and any attendant assumptions. Many details that were originally elaborated elsewhere in the literature are summarized in the Supplemental Materials (SM) [47].

For QPE, we use the highly optimized implementation of Babbush *et al.* [10]. We indicate the T count associated with estimating E_0 with Holevo variance ΔE as $T_{QPE}(\Delta E)$ and note that more relevant details can be found in the SM [47]. We assume a particularly simple approach to estimating E_0. Each repetition of the circuit will sample an eigenvalue from the spectrum of $\mathcal{H}$ with probability proportional to the overlap of the input state state with the associated eigenstate. So after $\mathcal{O}(\gamma^{-2})$ repetitions the smallest observed eigenvalue will likely be E_0. We now derive conditions under which boosting γ^2 with AA will reduce the expected total runtime for estimating E_0 relative to only relying on $\mathcal{U}_{\Phi_{0,i}}$, the cost of which we will denote $T_{\Phi_{0,i}}$ (see Fig. 1(b)).

AA consists of N_{iter} applications of a product of two reflections ($\mathcal{R}_{\Phi_{0,i}}$, $\mathcal{R}_{\Psi_0}$) that boosts the overlap of the state in *sys* to γ_f. N_{iter} is determined by γ_i,

$$N_{iter} = \left\lceil \frac{1}{2}\left(\frac{\sin^{-1}\gamma_f}{\sin^{-1}\gamma_i} - 1\right)\right\rceil. \tag{1}$$

While implementing $\mathcal{R}_{\Phi_{0,i}}$ only requires controlled applications of $\mathcal{U}_{\Phi_{0,i}}$, the implementation of $\mathcal{R}_{\Psi_0}$ is complicated by $|\Psi_0\rangle$ generally being unknown. We will follow Ref. [7] and construct an ϵ-approximation to this reflector using quantum signal processing (QSP) [48], $\mathcal{R}_{\Psi_0}(\epsilon)$ (see Fig. 1(c)). This is a degree-N_ϕ polynomial in $(\mathcal{H} - \mu\mathcal{I})$ that approximates a function that is ideally -1 for states with energy less than $\mu - \Delta/2$ and ideally 1 for states with energy greater than $\mu + \Delta/2$. Details pertaining to the calculation of μ and Δ are included in the SM [47].

However, in practice each eigenvalue of $\mathcal{R}_{\Psi_0}(\epsilon)$ is only ϵ-close to $\{-1, +1\}$ and $\mathcal{R}_{\Psi_0}(\epsilon)$ is not an exact reflector. One of the technical advances in this Letter is a bound on ϵ such that the QSP circuit produces $|\Phi_{0,f}\rangle$ with $\gamma_f \geq \gamma_i$,

$$\epsilon \leq \left(1 - \gamma_f^2\right)/6N_{iter}^2. \tag{2}$$

A proof can be found in the SM [47]. The cost of implementing the entire AA circuit is denoted $T_{AA}(\gamma_i, \gamma_f)$.

The only remaining details for the implementation under consideration pertain to the block encoding of $\mathcal{H}$ [40].

Block encoding is used both to encode the eigenspectrum of $\mathcal{H}$, as sampled in QPE, and to implement a polynomial in $(\mathcal{H} - \mu\mathcal{I})$ in QSP. As many of the details are model-specific and extensively developed in other work, we relegate a detailed discussion of block encoding to the SM [47]. While there are alternatives to block encoding, we leave it to future work to consider variants on the approach in Fig. 1 making use of, e.g., Trotterized Hamiltonian evolution, either for encoding the eigenspectrum of $\mathcal{H}$ for QPE or for implementing time-dependent Hamiltonian evolution in adiabatic state preparation [49].

With all of the components of our implementation specified, we can prepare T counts for the circuits with and without AA. The ratio of these counts defines the improvement,

$$\iota = \frac{\gamma_i^{-2}(T_{\Phi_{0,i}} + T_{QPE}(\Delta E))}{\gamma_f^{-2}(T_{\Phi_{0,i}} + T_{QPE}(\Delta E) + T_{AA}(\gamma_i, \gamma_f))}. \quad (3)$$

Here the cost of AA is

$$T_{AA} = N_{iter}\left(T_{\mathcal{R}_{\Phi_{0,i}}} + N_\phi(T_{\mathcal{U}_\mathcal{H}} + T_{e^{i\phi\Pi}})\right). \quad (4)$$

$T_{\mathcal{R}_{\Phi_{0,i}}}$ involves two applications of $\mathcal{U}_{\Phi_{0,i}}$ and a single multi-controlled X gate. $T_{\mathcal{U}_\mathcal{H}}$ is the cost of a single application of the block-encoded Hamiltonian. $T_{e^{i\phi\Pi}}$ involves two multi-controlled X gates and a single-qubit rotation with angles determined using the protocol in Ref. [44]. N_ϕ is the number of phases used to implement $\mathcal{R}_{\Psi_0}(\epsilon)$ and all parameter values will be chosen to saturate their bounds. We note that Eq. 4 depends on the rotation synthesis error incurred in implementing the Hamiltonian block encoding and the controlled rotations in QSP, and the error analysis used in deriving our results is considered in the SM [47]. We consider better state preparation as being worthwhile when $\iota > 1$.

Results.— We first consider resource estimates for the 1D transverse field Ising model (TFIM) [50] with periodic boundary conditions [51]. *sys* is encoded such that each of the n qubits represents one of the L sites. We consider a simple form for $\mathcal{U}_{\Phi_{0,i}}$ in which R_y rotations are applied to each qubit to generate a product state. We tune the rotation angles to construct a $|\Phi_{0,i}\rangle$ with a target value of γ_i. While not a particularly sophisticated choice for initializing *sys*, it suffices for our purposes. We select ϵ to saturate the bound in Eq. 2 with a target $\gamma_f^2 = 0.75$, and assume that the true final overlap is also $\gamma_f^2 = 0.75$; a presentation of the difference between the target and true overlaps can be found in Fig. 3.

In Fig. 2 we present ι and T_{AA} for the TFIM as a function of L, γ_i, and ΔE. We find $\iota > 1$ for all instances with $\gamma_i^2 \leq 10^{-3}$, and even for larger γ_i for the higher accuracy calculation. For $\gamma_i \ll 1$, the costs of QPE and AA are both dominated by repeated applications of the block-encoded Hamiltonian, leading to a simple approximate form of ι

$$\iota \sim \frac{\gamma_f^2}{\gamma_i^2}\left(\frac{\Delta}{\Delta E}\frac{1}{\sin^{-1}\gamma_f}\frac{\gamma_i}{\log\ \gamma_i^{-2}}\right), \quad (5)$$

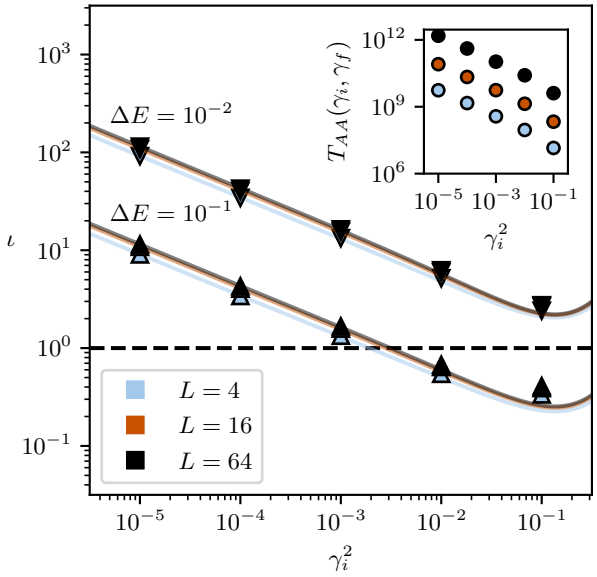

FIG. 2. T counts and improvement ι for TFIM with $L = 4$, 16, and 64 sites, with ϵ saturating the bound in Eq. 2 with $\gamma_f^2 = 0.75$, γ_i^2 ranging from 10^{-5} to 10^{-1}, and two different values of ΔE. The inset shows T_{AA} and the main figure is a plot of the improvement ι with solid lines fit to the asymptotic form Eq. 5. The dashed line is a guide for the eye at $\iota = 1$, above which improvement is seen when conducting state preparation.

consistent with a near-quadratic Grover-like speedup in γ_i from AA. Importantly, the asymptotic improvement has no explicit dependence on the system size, with the only system size dependence in the finite-size error of Δ. We find that our computed ι follows this asymptotic trend closely for $\iota \leq 10^{-2}$. We note that 88 logical qubits are required for the $L = 64$ calculation with $\Delta E = 10^{-2}$, with a detailed analysis of qubit counts in the SM [47].

To test the performance of our $\mathcal{U}_{sp}(\mathcal{H})$ implementation, we explicitly simulate circuits for $L = 2, 4, 6$ site TFIM using the upper bound on ϵ in Eq. 2. For each L we run nine simulations, with $\gamma_f^2 = 0.9$, 0.99 and 0.999, and γ_i chosen such that $N_{iter} = 4$, 6, 10. We are then able to compare the values of γ_f^2 actually realized in the simulation to the specified value of γ_f^2 in Fig. 3. Noise-free simulations were carried out using the PyTKET package [52] with the full source available in the SM [47].

We find that the simulated infidelity, $1 - \gamma_f^2$, is consistently one or two orders of smaller than the target, indicating that our bound for ϵ in Eq. 2 could be adjusted to potentially realize further savings in implementing this approach to state preparation. However, the impact of ϵ on ι is only logarithmic, so removing this looseness will only affect our results by a small constant factor. As such, we are confident that our conclusions are not skewed by a loose bound in the state preparation T counts.

Next, we consider resource estimates for a more realistic and technologically important Hamiltonian. The solid

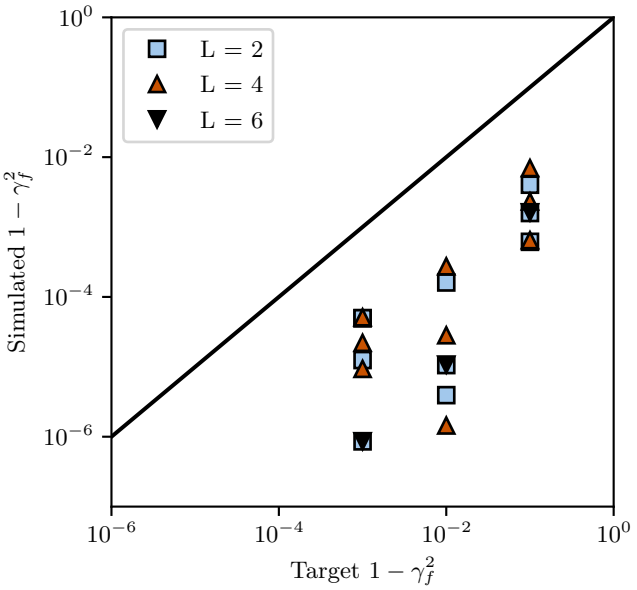

FIG. 3. Comparison of the target γ_f and simulated γ_f for the transverse field Ising model with $L = 2$, 4, and 6 sites. For each value of L, state-vector simulations of the full $\mathcal{U}_{sp}(\mathcal{H})$ circuit were carried out for three target values, $1-\gamma_f^2 = 10^{-1}, 10^{-2}, 10^{-3}$, and three values of γ_i satisfying $N_{iter} = 4, 6, 10$. The black line is the $x = y$ reference and our results are consistent with the validity of the bound in Eq. 2.

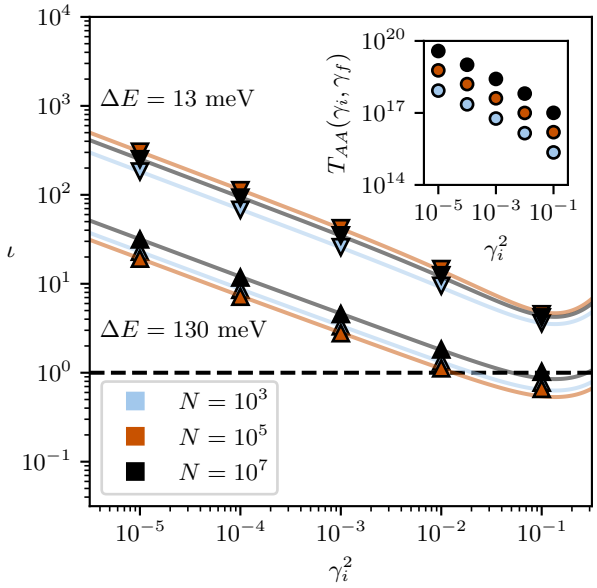

FIG. 4. T counts and improvement ι for β-alumina structure $Na_4Al_{22}O_{35}$ with $\eta = 610$, $N = 10^3$, 10^5, 10^7, with ϵ saturating the bound in Eq. 2 with $\gamma_f^2 = 0.75$, γ_i ranging from 10^{-5} to 10^{-1} and two different values of ΔE. The inset shows T_{AA} for various system sizes, with the main figure presenting ι with solid lines fit to the asymptotic form Eq. 5. The dashed line is a guide for the eye at $\iota = 1$.

electrolyte β-alumina is known for its high ionic conductivity [53, 54] and there is broad general interest in using it for low-carbon energy storage [55]. Of particular interest for these applications is the accurate calculation of equilibrium voltage, ionic mobility, and thermal stability, which are all directly related to accurate ground state energies of the battery as outlined by Delgado *et al.* in their work on the lithium-ion battery Li_2FeSiO_4 [56].

Classical computation of accurate ground state energies of the β-aluminas is challenged by the non-stoichiometric chemical composition, $Na_{1+x}Al_{11}O_{17+x/2}$, which requires large supercells to resolve finite-size effects. This is a larger supercell than has been considered in other resource estimates of materials. We choose this particular example because it is large enough that mean-field classical heuristics (e.g., density functional theory) are likely to be the only methods that are viable, potentially leading to small values of γ_i. However, an analysis of the precise value of γ_i that is classically achievable with these heuristics is beyond the scope of this Letter and thus we leave it as a free parameter.

In Fig. 4 we present resource estimates for a $Na_4Al_{22}O_{35}$ supercell with 610 electrons using a first quantized representation [12, 57] of the electronic structure Hamiltonian in a plane-wave basis set with cardinality N. We find that $\iota > 1$ in all cases where $\gamma_i^2 \leq 10^{-2}$, from small to large basis sets, and for ΔE below chemical accuracy to different extents. We also see that the T counts for AA are such that it is plausible to imagine implementing calculations like this on fault-tolerant quantum computers that are perhaps a generation beyond the ones considered in Refs. [10, 11]. Even when starting with a fairly large $\gamma_i^2 = 0.1$, we still find an order of magnitude improvement in T counts for a high accuracy calculation, as a result of investing resources in better state preparation. We note that 33,275 logical qubits are required for the $N = 10^7$ calculation with $\Delta E = 13$ meV; 18,635 logical qubits for anti-symmetrization and 14,640 logical qubits for all other computations. A detailed analysis of qubit counts is relegated to the SM [47], as well as the full source code for computing resource estimates.

Conclusions.— We have developed resource estimates for end-to-end ground state energy determination using near-optimal state preparation [7]. The ratio of the T count for successful ground state energy estimation, without and with this state preparation, define an improvement factor that is related to likely runtime reductions. This improvement is near-quadratic in γ_i and demonstrated credible multiple-order-of-magnitude speedups for a toy problem and a highly realistic electronic structure problem.

Future work will involve determining more realistic estimates for scenarios under which these types of speedups will be realized. In particular, categorizing the values of γ_i typical of classical heuristics that are efficiently implementable as $\mathcal{U}_{\Phi_{0,i}}$ is an open research area. It also remains unclear whether efficient implementations

of adiabatic state preparation or other variants on filter-based state preparation are more or less efficient than the one examined in this Letter. Finally, whether quantum phase estimation protocols with built-in tolerance to state preparation errors [33, 58] can be exploited to achieve better improvements is a topic for future work.

Note added.— Between uploading the first and second versions of this Letter to the arXiv, we became aware of another manuscript considering similar aspects of ground state preparation [59].

ACKNOWLEDGMENTS

We gratefully acknowledge useful conversations with Ryan Babbush, Anand Ganti, Lucas Kocia, Alina Kononov, Michael Kreshchuk, Andrew Landahl, Lin Lin, Alicia Magann, Jonathan Moussa, Setso Metodi, Mason Rhodes, and Norm Tubman. All authors were supported by the National Nuclear Security Administration's Advanced Simulation and Computing Program. AER was partially supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Quantum Computing Application Teams program. ADB was partially supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers program and Sandia National Laboratories' Laboratory Directed Research and Development program (Project 222396).

This article has been co-authored by employees of National Technology & Engineering Solutions of Sandia, LLC under Contract No. DE-NA0003525 with the U.S. Department of Energy (DOE). The authors own all right, title and interest in and to the article and are solely responsible for its contents. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this article or allow others to do so, for United States Government purposes. The DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan https://www.energy.gov/downloads/doe-public-access-plan.

[1] S. Lloyd, Science **273**, 1073 (1996).
[2] A. Aspuru-Guzik, A. D. Dutoi, P. J. Love, and M. Head-Gordon, Science **309**, 1704 (2005).
[3] S. P. Jordan, K. S. Lee, and J. Preskill, Science **336**, 1130 (2012).
[4] D. Poulin and P. Wocjan, Physical review letters **102**, 130503 (2009).
[5] N. M. Tubman, C. Mejuto-Zaera, J. M. Epstein, D. Hait, D. S. Levine, W. Huggins, Z. Jiang, J. R. McClean, R. Babbush, M. Head-Gordon, *et al.*, arXiv preprint arXiv:1809.05523 (2018).
[6] Y. Ge, J. Tura, and J. I. Cirac, Journal of Mathematical Physics **60**, 022202 (2019).
[7] L. Lin and Y. Tong, Quantum **4**, 372 (2020).
[8] J. Lemieux, G. Duclos-Cianci, D. Sénéchal, and D. Poulin, Physical Review A **103**, 052408 (2021).
[9] M. Reiher, N. Wiebe, K. M. Svore, D. Wecker, and M. Troyer, Proceedings of the national academy of sciences **114**, 7555 (2017).
[10] R. Babbush, C. Gidney, D. W. Berry, N. Wiebe, J. McClean, A. Paler, A. Fowler, and H. Neven, Physical Review X **8**, 041015 (2018).
[11] J. Lee, D. W. Berry, C. Gidney, W. J. Huggins, J. R. McClean, N. Wiebe, and R. Babbush, PRX Quantum **2**, 030305 (2021).
[12] Y. Su, D. W. Berry, N. Wiebe, N. Rubin, and R. Babbush, PRX Quantum **2**, 040332 (2021).
[13] Y. Su, H.-Y. Huang, and E. T. Campbell, Quantum **5**, 495 (2021).
[14] V. von Burg, G. H. Low, T. Häner, D. S. Steiger, M. Reiher, M. Roetteler, and M. Troyer, Physical Review Research **3**, 033055 (2021).
[15] A. M. Childs, D. Maslov, Y. Nam, N. J. Ross, and Y. Su, Proceedings of the National Academy of Sciences **115**, 9456 (2018).
[16] A. M. Childs and Y. Su, Physical review letters **123**, 050503 (2019).
[17] M. C. Tran, Y. Su, D. Carney, and J. M. Taylor, PRX Quantum **2**, 010323 (2021).
[18] N. Klco and M. J. Savage, Physical Review A **99**, 052335 (2019).
[19] H. Lamm, S. Lawrence, Y. Yamauchi, N. Collaboration, *et al.*, Physical Review D **100**, 034518 (2019).
[20] A. F. Shaw, P. Lougovski, J. R. Stryker, and N. Wiebe, Quantum **4**, 306 (2020).
[21] M. A. Nielsen and I. Chuang, "Quantum computation and quantum information," (2002).
[22] A. Y. Kitaev, A. Shen, and M. N. Vyalyi, *Classical and quantum computation*, 47 (American Mathematical Soc., 2002).
[23] J. Kempe, A. Kitaev, and O. Regev, SIAM journal on computing **35**, 1070 (2006).
[24] R. P. Feynman, International Journal of Theoretical Physics **21**, 467 (1982).
[25] B. O'Gorman, S. Irani, J. Whitfield, and B. Fefferman, PRX Quantum **3**, 020322 (2022).
[26] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'brien, Nature communications **5**, 1 (2014).
[27] P. J. O'Malley, R. Babbush, I. D. Kivlichan, J. Romero, J. R. McClean, R. Barends, J. Kelly, P. Roushan, A. Tranter, N. Ding, *et al.*, Physical Review X **6**, 031007 (2016).
[28] This cost isn't always negligible. As has been noted in Ref. [56], sometimes even the cost of preparing simple input states can dominate simulation costs unless certain simplifying assumptions are made.
[29] A. Y. Kitaev, arXiv preprint quant-ph/9511026 (1995),

10.48550/arXiv.quant-ph/9511026.
[30] S. Kimmel, G. H. Low, and T. J. Yoder, Physical Review A **92**, 062315 (2015).
[31] N. Wiebe and C. Granade, Physical review letters **117**, 010503 (2016).
[32] T. E. O'Brien, S. Polla, N. C. Rubin, W. J. Huggins, S. McArdle, S. Boixo, J. R. McClean, and R. Babbush, PRX Quantum **2**, 020317 (2021).
[33] A. E. Russo, K. M. Rudinger, B. C. Morrison, and A. D. Baczewski, Physical Review Letters **126**, 210501 (2021).
[34] L. Lin and Y. Tong, PRX Quantum **3**, 010318 (2022).
[35] G. Wang, D. Stilck-França, R. Zhang, S. Zhu, and P. D. Johnson, arXiv preprint arXiv:2209.06811 (2022), https://doi.org/10.48550/arXiv.2209.06811.
[36] We expect that further improvements can be made by adopting the aforementioned "more elaborate strategies" but these are likely to be constant-factor improvements rather than asymptotic ones.
[37] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, arXiv preprint quant-ph/0001106 (2000), 10.48550/arXiv.quant-ph/0001106.
[38] T. Albash and D. A. Lidar, Reviews of Modern Physics **90**, 015002 (2018).
[39] K. Wan and I. Kim, arXiv preprint arXiv:2004.04164 (2020), 10.48550/arXiv.2004.04164.
[40] A. Gilyén, Y. Su, G. H. Low, and N. Wiebe, in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing* (2019) pp. 193–204.
[41] S. Lee, J. Lee, H. Zhai, Y. Tong, A. M. Dalzell, A. Kumar, P. Helms, J. Gray, Z.-H. Cui, W. Liu, *et al.*, arXiv preprint arXiv:2208.02199 (2022).
[42] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, Physical Review A **86**, 032324 (2012).
[43] D. Litinski, Quantum **3**, 205 (2019).
[44] Y. Dong, X. Meng, K. B. Whaley, and L. Lin, Physical Review A **103**, 042419 (2021).
[45] L. K. Grover, in *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing* (1996) pp. 212–219.
[46] G. Brassard, P. Hoyer, M. Mosca, and A. Tapp, Contemporary Mathematics **305**, 53 (2002).
[47] See Supplemental Materials for more details, which includes Refs. [60-66]
[48] G. H. Low and I. L. Chuang, Physical review letters **118**, 010501 (2017).
[49] L. Kocia, F. A. Calderon-Vargas, M. D. Grace, A. B. Magann, J. B. Larsen, A. D. Baczewski, and M. Sarovar, arXiv preprint arXiv:2209.06242 (2022), 10.48550/arXiv.2209.06242.
[50] R. B. Stinchcombe, Journal of Physics C: Solid State Physics **6**, 2459 (1973).
[51] Open boundary conditions will have very similar costs.
[52] S. Sivarajah, S. Dilkes, A. Cowtan, W. Simmons, A. Edgington, and R. Duncan, Quantum Sci. Technol. **6**, 014003 (2020).
[53] Y.-F. Y. Yao and J. T. Kummer, Journal of Inorganic and Nuclear Chemistry **29**, 2453 (1967).
[54] R. Stevens and J. G. P. Binner, Journal of Materials Science **19**, 695 (1984).
[55] R. Collongues, D. Gourier, A. Kahn, J. Boilot, P. Colomban, and A. Wicker, Journal of Physics and Chemistry of Solids **45**, 981 (1984).
[56] A. Delgado, P. A. M. Casares, R. d. Reis, M. S. Zini, R. Campos, N. Cruz-Hernández, A.-C. Voigt, A. Lowe, S. Jahangiri, M. A. Martin-Delgado, J. E. Mueller, and J. M. Arrazola, arXiv (2022), 10.48550/arXiv.2204.11890, 2204.11890.
[57] R. Babbush, D. W. Berry, J. R. McClean, and H. Neven, npj Quantum Information **5**, 1 (2019).
[58] Z. Ding and L. Lin, arXiv preprint arXiv:2211.11973 (2022), 10.48550/arXiv.2211.11973.
[59] K. Gratsea, C. Sun, and P. D. Johnson, arXiv preprint arXiv:2212.09492 (2022), 10.48550/arXiv.2212.09492.
[60] L. Lin and Y. Tong, Quantum **4**, 361 (2020), 1910.14596v4.
[61] A. M. Childs and N. Wiebe, Quantum Information and Computation **12**, 0901 (2012), 1202.5822.
[62] P. Selinger, Quantum Info. Comput. **15**, 159–180 (2015).
[63] P. Niemann, R. Wille, and R. Drechsler, Quantum Information Processing **19**, 317 (2020).
[64] D. W. Berry, M. Kieferová, A. Scherer, Y. R. Sanders, G. H. Low, N. Wiebe, C. Gidney, and R. Babbush, npj Quantum Information **4**, 1 (2018).
[65] Actually, the filtering step succeeds with probability $> 1/2$, so on average the step needs to be performed twice.
[66] P. A. M. Casares, R. Campos, and M. A. Martin-Delgado, Quantum **6**, 768 (2022).