


Evaluating Energy Differences on a Quantum Computer with Robust Phase Estimation

A. E. Russo¹, K. M. Rudinger¹, B. C. A. Morrison^{1,2} and A. D. Baczewski^{1,2,*}

¹Center for Computing Research, Sandia National Laboratories, Albuquerque, New Mexico 87185, USA

²Center for Quantum Information and Control (CQuIC), Department of Physics and Astronomy, University of New Mexico, Albuquerque, New Mexico 87131, USA

 (Received 20 July 2020; revised 17 March 2021; accepted 22 March 2021; published 24 May 2021)

We adapt the robust phase estimation algorithm to the evaluation of energy differences between two eigenstates using a quantum computer. This approach does not require controlled unitaries between auxiliary and system registers or even a single auxiliary qubit. As a proof of concept, we calculate the energies of the ground state and low-lying electronic excitations of a hydrogen molecule in a minimal basis on a cloud quantum computer. The denominative robustness of our approach is then quantified in terms of a high tolerance to coherent errors in the state preparation and measurement. Conceptually, we note that all quantum phase estimation algorithms ultimately evaluate eigenvalue differences.

DOI: [10.1103/PhysRevLett.126.210501](https://doi.org/10.1103/PhysRevLett.126.210501)

Introduction.—Assessing energy differences, rather than total energies, is ubiquitous in physics. Whether there is a gap between the ground and first excited state of a particular Hamiltonian is related to outstanding problems in condensed matter [1] and high energy physics [2], and it is at the heart of deep connections between many-body physics and theoretical computer science [3]. Myriad spectroscopic techniques ultimately compare the energies of two or more eigenstates of a single Hamiltonian as one among many identifying features of a particular piece of matter. This Letter is concerned with using a quantum computer for this purpose. We indicate the Hamiltonian of interest as \mathcal{H} with $N = 2^n = \dim \mathcal{H}$. The ground state of \mathcal{H} is labeled by its eigenvalue $|E_0\rangle$, and the a th eigenstate above it is $|E_a\rangle$.

By repeatedly preparing particular superpositions of two energy eigenstates, allowing them to undergo a unitary evolution $\mathcal{W}(\mathcal{H})$ [4–7], undoing the preparation, and measuring in the computational basis [see Fig. 1(b)], we can infer the difference in energy between the two eigenstates without the need for auxiliary qubits [8] or controlled unitary operations. This differs from other approaches to quantum phase estimation (QPE) [11] that use one or more auxiliary qubits to provide a ground reference for the phase accumulated on the register encoding the physical system [12–21]. Our procedure is inspired by the robust phase estimation (RPE) algorithm that was introduced to characterize and calibrate the phase (i.e., rotation angle) of a single-qubit gate [22].

A common form for $\mathcal{W}(\mathcal{H})$ is an approximation to the exponential map governing Hamiltonian evolution for a fixed time [23,24], though it may take other forms for which the phase is a known function of the eigenvalues [5,25]. While phase estimation is broadly applicable to the calculation of eigenvalues on quantum computers, the

physical significance of $\mathcal{W}(\mathcal{H})$ is a consequence of encoding the degrees of freedom of a system of interest in the Hilbert space of n qubits. While we consider the specific encoding of interacting electrons in a molecular system [26,27], we note that our results can be extended to others, including those relevant to nuclear matter [28], quantum field theories [29], and spin systems [30].

In fact, all forms of phase estimation, with or without auxiliary qubits, are not simply eigenvalue estimation but eigenvalue *difference* estimation. The operations $\mathcal{W}(\mathcal{H})$ and $\mathcal{W}(\mathcal{H} + \alpha I)$ are identical up to an undetectable global phase, $\exp[i\chi(\alpha)]$, where the form of χ depends on \mathcal{W} [31]. In order to actually estimate the phase of an eigenstate of \mathcal{W} , one must have access to a known *reference* energy level. $\Lambda(\mathcal{W})$, a singly controlled version of \mathcal{W} , is generated by a Hamiltonian of the form $0_N \oplus \mathcal{H}$, where 0_N is the $N \times N$ zero matrix. The N -fold degenerate zero-energy subspace created by 0_N allows for the estimation of the phase of *any* of the eigenstates of \mathcal{H} relative to these reference eigenstates [see Fig. 1(a)]. This is the structure of *most* QPE implementations, which we henceforth generically refer to as QPE algorithms with auxiliary qubits [32]. Part of what distinguishes RPE is that, instead of relying on the auxiliary register to relativize the phase of the unitary evolution, the relative phase is accumulated between two energy eigenstates in a uniform superposition. This allows us to avoid using an auxiliary register and controlled unitaries at the cost of requiring more complicated state preparation.

Three strengths of QPE with auxiliary qubits are (i) the relativization of the phase accumulated on the $|1\rangle$ branch of the auxiliary register to the $|0\rangle$ branch, (ii) the projection of the system register onto an energy eigenstate after a single round, and (iii) the ability to reuse that eigenstate in subsequent rounds without having to prepare it again. Point (i) is a critical advantage if one needs absolute




FIG. 1. Comparison of QPE and RPE circuits. (a) In QPE the system register is prepared in the a th eigenstate of \mathcal{H} ($U_{p,a}$) while Hadamard gates are applied to each of g auxiliary qubits. For $j \in [0, g-1]$, \mathcal{W} is applied 2^j times with the j th auxiliary qubit as a control. The inverse quantum Fourier transform (QFT^\dagger) is applied to the auxiliary register prior to measurement, which yields g bits of E_a . Each \mathcal{W} acts on a $2N$ -dimensional Hilbert space in which the control qubit provides an N -fold degenerate zero-energy subspace relative to which a phase difference accumulates between any of the N energy eigenvalues of \mathcal{H} . (b) In RPE there is no auxiliary register. First, a superposition of the a th and b th eigenstates of \mathcal{H} is prepared ($U_{p,\varphi_{ab}=0}$). Then \mathcal{W} is applied k_g times. The superposition is unprepared ($U_{p,\varphi_{ab}=\beta}^\dagger$) and all qubits are measured, yielding a sample from P_c for $\beta = 0$ and a sample from P_s for $\beta = \pi/2$. P_c and P_s encode $E_b - E_a$ in a phase θ_{ab} , defined in Eq. (2). \mathcal{W} acts on an N -dimensional Hilbert space and only energy differences between eigenstates of \mathcal{H} can be extracted from this phase.

energies, but not essential if energy differences will suffice. Further, if one knows the trace of the Hamiltonian over an M -dimensional subspace it is possible to reconstruct absolute energies from $M - 1$ -independent pairwise energy differences measured within that subspace. This is evident in the experimental results in Fig. 2. Points (ii) and (iii) are critical advantages if state preparation dominates the Hamiltonian evolution resource requirements, noting that the depth of the Hamiltonian evolution unitaries for RPE will be reduced by merit of their not needing to be controlled unitaries.

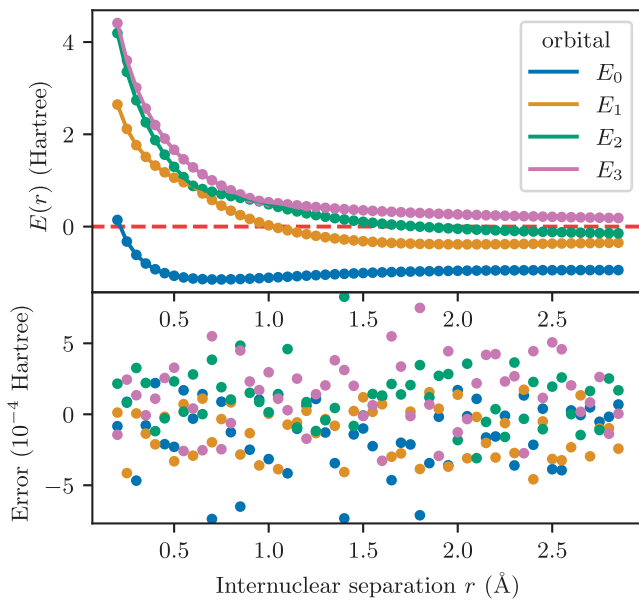


FIG. 2. Verification of RPE for evaluating energy differences in a molecule. Top: the first four energy levels of H_2 in a minimal basis, as calculated using RPE on IBM Vigo (dots) and diagonalization on a classical computer (lines). Bottom: the error in the first four energy levels relative to the result evaluated on a classical computer.

Not only does RPE offer overall circuit depth improvement, but it reduces the total number of 2-qubit operations required. These are a bottleneck in current hardware given their low fidelities relative to single-qubit gates [33,34]. Given access to a gate-level description of a circuit \mathcal{S} that implements $\mathcal{W}(\mathcal{H})$ using only arbitrary local gates and controlled NOTs (CNOTs), the most straightforward way to implement $\mathcal{W}(0_N \oplus \mathcal{H})$ is to simply turn every gate \mathcal{G} in \mathcal{S} into its singly controlled version $\Lambda(\mathcal{G})$. Though clever compilation schemes [35–40] may offer nontrivial improvements, if \mathcal{G} contains s single-qubit gates and t CNOTs it can be shown [41,42] that the overall CNOT cost of implementing $\Lambda(\mathcal{G})$ may be as bad as $6t + 2s$. We expect these benefits to be substantial for hardware with little or no error correction and restricted connectivity.

One might ask whether the need to repeat the potentially erroneous state preparation and measurement (SPAM) due to a lack of projection onto an energy eigenstate after a single round of auxiliary-qubit-free RPE is a limiting factor. A central result of this Letter is the observation that RPE’s robustness manifests as a high tolerance to SPAM errors. This suggests conditions for which this approach might be advantageously employed for quantum simulation. In particular, an advantage might be realized in the intermediate term where adiabatic [26,43] or filtering-based [44–46] state preparation can be replaced by precompiled state preparation circuits that exploit classical tractability and do not appreciably contribute to the total circuit depth.

Methods.—RPE may be thought of as a combination of Ramsey and Rabi experiments with logarithmic spacing in the number of gate repetitions [47]. This allows the phase of the gate to be learned with Heisenberg-like scaling in accuracy, without requiring any entanglement or auxiliary qubits. Additionally, RPE will still produce accurate phase estimates even when there is a significant amount of error in any of the constituent circuits’ state preparations, measurements, or gates. Accordingly, RPE has been demonstrated in experimental systems to yield highly accurate phase

estimates [48] while being robust against various noise channels [49].

While RPE concerns itself with estimating a single-qubit gate's phase [e.g., the angle θ in the gate $R_x(\theta) = \exp(-i\theta\sigma_x/2)$], this phase is actually the difference between the two eigenvalues of the Hamiltonian that generates the unitary rotation [50]. This principle can be generalized to unitary maps of dimension greater than 2, allowing for the difference between two eigenvalues of an arbitrary Hamiltonian to be estimated using RPE.

To adapt RPE to higher dimensions, one simply needs implementations of (i) $\mathcal{W}(\mathcal{H})$ and (ii) a state preparation unitary

$$U_{p,\varphi_{ab}=\beta}|0\rangle = \frac{1}{\sqrt{2}}(|E_a\rangle + e^{i\beta}|E_b\rangle) = |\varphi_{ab} = \beta\rangle, \quad (1)$$

where we specifically require $U_{p,\varphi_{ab}}$ for two values of φ_{ab} that are separated by $\pi/2$ rad. The energy difference between eigenstates a and b is related to a relative phase, $\theta_{ab} \bmod 2\pi$, accumulated while evolving with $\mathcal{W}(\mathcal{H})$ for a particular time interval that is absorbed into the units. This relative phase is encoded in the probability distributions

$$P_c(k_g\theta_{ab}) = |\langle 0|U_{p,\varphi_{ab}=0}^\dagger \mathcal{W}^{k_g}(\mathcal{H})U_{p,\varphi_{ab}=0}|0\rangle|^2 \quad (2a)$$

$$= \frac{1}{2}[1 + \cos(k_g\theta_{ab})] \quad \text{and} \quad (2b)$$

$$P_s(k_g\theta_{ab}) = |\langle 0|U_{p,\varphi_{ab}=\pi/2}^\dagger \mathcal{W}^{k_g}(\mathcal{H})U_{p,\varphi_{ab}=0}|0\rangle|^2 \quad (2c)$$

$$= \frac{1}{2}[1 + \sin(k_g\theta_{ab})], \quad (2d)$$

where the circuits that sample from these distributions are evident from Eqs. (2a) and (2c) and the functional forms of the distributions are given in Eqs. (2b) and (2d). Here k_g is the number of applications of $\mathcal{W}(\mathcal{H})$ during the g th generation. k_g is chosen with logarithmic spacing, i.e., $k_g = 2^g$, and experiments proceed by refining the estimate of θ_{ab} across generations consisting of increasing numbers of repetitions of $\mathcal{W}(\mathcal{H})$ [51].

For a fixed value of k_g , the circuits represented by Eqs. (2a) and (2c) are repeated sufficiently many times to estimate P_c and P_s from the relative frequencies of 0 and 1 outcomes. Equations (2b) and (2d), then, unambiguously specify θ_{ab} on a segment of $2\pi/k_g$ rad,

$$k_g\theta_{ab} = \text{atan2}(2P_c - 1, 2P_s - 1) \bmod 2\pi, \quad (3)$$

where atan2 accounts for the branch cuts of \arctan by tracking the signs of the x and y components. RPE uses estimates of θ_{ab} from experiments with $k_{g'}$ for $g' < g$ to select a particular segment. At each successive generation, if the right branch is chosen, the error in θ_{ab} will exhibit Heisenberg-like scaling.

A key feature of RPE is its tolerance to additive errors in P_c and P_s . In the Supplemental Material [52], we study the impact of coherent errors on state preparation ($\langle 0|U_{p,\varphi_{ab}}$) and “unpreparation” ($U_{p,\varphi_{ab}}^\dagger|0\rangle$). The parameters of the error channel under consideration are related to the deviation of the state prepared (or unprepared) relative to the target state $|\varphi_{ab}\rangle$. These include errors that generate support with erroneous amplitude (\mathcal{E}_c) and phase (\mathcal{E}_p) in the “target subspace,” i.e., $\text{span}\{|E_a\rangle, |E_b\rangle\}$, but orthogonal to $|\varphi_{ab}\rangle$. It also includes leakage errors (\mathcal{E}_l) that generate support outside of that subspace. We indicate equivalent errors occurring during unpreparation (including readout) with primed variables (e.g., \mathcal{E}'_c).

We derived worst-case bounds on the associated additive contributions to P_c and P_s and translated them into worst-case bounds on additive error in the estimate of $k_g\theta_{ab}$ [see Eq. (3)]. This additive error is indicated as δ_λ [52]. Combined with bounds on additive errors under which RPE can succeed [22,59], we identified conditions on coherent SPAM errors that permit estimation of energy differences with Heisenberg-like scaling [60]. Our results indicate a high tolerance to these errors. Prominently, there are conditions for which RPE will still succeed if as much as $\sim 9\%$ of the probability in the prepared (unprepared) state leaks outside of the target subspace.

Results.—To verify our RPE protocol for evaluating energy differences in physical simulation, we conducted a proof-of-concept experiment through the cloud-based IBM Quantum Experience [61,62]. We computed three of the independent pairwise energy differences between the four eigenstates of molecular hydrogen (H_2) in a minimal basis along its dissociation curve. Combined with a knowledge of the trace of the Hamiltonian over this subspace, we reconstructed the energy eigenvalues themselves. The results are illustrated in Fig. 2, in which it is evident that RPE succeeds in accurately computing these eigenvalues from pairwise differences. All Hamiltonian (and k_g) dependence was precompiled into two- or three-CNOT circuits for this 2-qubit demonstration, leading to k_g -independent depth circuits of at most 11 CNOTs [52]. We remark that this precompilation approach cannot provide a quantum advantage as it relies on the Hamiltonian being classically diagonalizable. A demonstration without precompilation is likely to require non-trivial quantum hardware improvements.

However, by compiling into constant-depth circuits we are able to verify that our protocol achieves the ideal scaling with k_g . This is illustrated in Fig. 3, in which we also compare the experimentally observed scaling to that predicted by circuit simulations with and without noise. Our noisy simulations are based upon calibration data furnished by IBM at the time of the experiment. The noiseless simulations provide a benchmark for the optimal performance of our circuits, with the noisy simulations

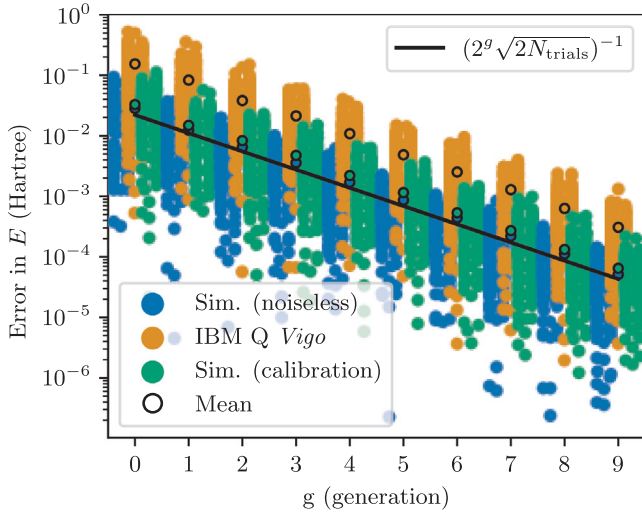


FIG. 3. Simulated and experimental distributions of errors in the H_2 energy calculation with 1024 repetitions per circuit. A swarm plot with errors from all internuclear separations and energy differences, for each generation of RPE. Experimental results on IBM *Vigo* are compared to results from circuit simulations without noise and using the calibration-based noise model supplied by IBM. Scaling of the error with $1/2^g$ is observed, consistent with Heisenberg-like scaling and indicating that the correct branch is predominantly chosen in these sequences.

suggesting that experiment will realize a relatively small deviation from this. The fact that the experiment realizes a mean error that scales with $1/2^g$ indicates that we are choosing the correct branch between successive generations, even using noisy hardware. However, the fact that the noisy simulations predict errors that are almost an order of magnitude smaller than those that are experimentally observed suggests that the furnished noise model is insufficient to predict actual hardware behavior, highlighting both the utility of more expressive noise models [63] and the relatively loose relationship between average gate infidelities and worst-case error rates [34,71]. Nevertheless, that the procedure still works in the presence of “hidden” error processes also highlights RPE’s resilience to such hidden errors.

Finally, we illustrate the denominative robustness of RPE to SPAM errors. Figure 4 presents a particular two-dimensional slice of our error model in which $\mathcal{E}_c = \mathcal{E}'_c$ and $\mathcal{E}_l = \mathcal{E}'_l$ vary. All other parameters of the model are optimized to produce a worst-case bound on the additive error in Fig. (3). This worst-case additive error is then compared to the upper bound for which the success of RPE is guaranteed. We find that, for $\mathcal{E}_c = \mathcal{E}'_c = 0$, RPE can tolerate a probability of leakage out of the target subspace in each of the preparation and measurement circuits up to $\sim 9\%$. The sensitivity to coherent state preparation errors within the target subspace is apparently higher, only tolerating individual coherent error probabilities of just




FIG. 4. Robustness against SPAM errors. The maximum additive error $|\delta_\lambda|$ in the measured angle λ used in the RPE protocol is plotted as a scaled function of the strength of state preparation errors within (\mathcal{E}_c) and outside (\mathcal{E}_l) the target subspace (see Supplemental Material for derivation [52]). We have set $\mathcal{E}_c = \mathcal{E}'_c$ and $\mathcal{E}_l = \mathcal{E}'_l$ to get a two-dimensional slice of the error bound in the four-dimensional parameter space. The upper limit for RPE protocol success $|\delta_\lambda| < (\pi/3)$ is plotted as a white line [59]. The region below this line corresponds to conditions for which RPE will succeed in spite of coherent SPAM errors. Values of $|\delta_\lambda| \geq (\pi/2)$ are plotted in black.

$\sim 4\%$, partially due to the selection of worst-case phase error within that subspace (\mathcal{E}_p).

Conclusion.—We have adapted RPE from its original application in efficiently estimating the phase of a single-qubit gate to efficiently estimating energy differences in quantum simulation. This approach to phase estimation does not require any auxiliary qubits or controlled implementations of $\mathcal{W}(\mathcal{H})$. While approaches using auxiliary qubits benefit from projection into an energy eigenstate after each round, we have shown that RPE is tolerant to errors in SPAM. We expect the long-term utility of such a protocol to be eclipsed by auxiliary-qubit-based approaches in future fault-tolerant quantum computers. However, we expect this approach to be impactful in the intermediate term, specifically, for verifying and validating quantum simulation algorithms in the era between the noisy, intermediate-scale quantum present and the fault-tolerant quantum-error-corrected future.

The in-between epoch in which we expect RPE to be most useful is one in which the capabilities of quantum computers will be typified by a number of features. A few error-corrected logical qubits might be available, but with logical error rates and connectivities that are sufficiently limited that the implementation of one-to-many controlled $\mathcal{W}(\mathcal{H})$ is not possible for the desired precision. There might also be sufficiently few logical qubits that it is possible to

classically diagonalize the Hamiltonian over a particular subspace, in which case straightforward compilation of the state preparation unitaries will be possible. Finally, RPE might be useful in diagnosing adiabatic state preparation algorithms that rely on finding a pathway between a noninteracting and interacting Hamiltonian in which the gap between the ground and first excited state remains as large as possible. As RPE allows us to efficiently evaluate this gap with limited resources, we see this as a promising application.

We gratefully acknowledge useful conversations with Shelby Kimmel, Will Kirby, Andrew Landahl, Guang Hao Low, Ojas Parekh, Nicholas Rubin, Mohan Sarovar, Rolando Somma, James Whitfield, and Nathan Wiebe. This work was supported in part by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Quantum Algorithms Team and Quantum Computing Applications Team programs. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for DOE's National Nuclear Security Administration under Award No. DE-NA0003525. Any subjective views or opinions that might be expressed in the Letter do not necessarily represent the views of the U.S. Department of Energy or the U.S. Government. A. D. B. and K. M. R. conceived of this research project. K. M. R. directed the technical work and A. D. B. directed the writing of the manuscript. A. E. R. and K. M. R. implemented the necessary routines in the PYTHON package `pyGSTi` [72,73] and ran the cloud-based experiments. B. C. A. M. and A. D. B. provided a critical assessment of the derivations. All authors contributed to the theory, analysis, and writing of this Letter.

Note added.—Recently, the authors became aware of two contemporaneous results involving phase estimation without auxiliary qubits [74,75].

*Corresponding author.
adbacze@sandia.gov

- [1] F. D. M. Haldane, *Phys. Rev. Lett.* **50**, 1153 (1983).
- [2] A. Jaffe and E. Witten, *The Millennium Prize Problems* (American Mathematical Society, 2006).
- [3] T. S. Cubitt, D. Perez-Garcia, and M. M. Wolf, *Nature (London)* **528**, 207 (2015).
- [4] S. Lloyd, *Science* **273**, 1073 (1996).
- [5] A. M. Childs, *Commun. Math. Phys.* **294**, 581 (2010).
- [6] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, *Phys. Rev. Lett.* **114**, 090502 (2015).
- [7] R. Babbush, C. Gidney, D. W. Berry, N. Wiebe, J. McClean, A. Paler, A. Fowler, and H. Neven, *Phys. Rev. X* **8**, 041015 (2018).
- [8] Throughout, we use the phrase “auxiliary qubit” in place of the phrase “ancilla qubit” following the etymological concerns raised in [9] and an alternative proposed in [10].
- [9] K. Wiesner, [arXiv:1705.06768](https://arxiv.org/abs/1705.06768).
- [10] D. Puzzioli, Ph.D. thesis, University of Waterloo, 2018.
- [11] A. Y. Kitaev, [arXiv:quant-ph/9511026](https://arxiv.org/abs/quant-ph/9511026).
- [12] R. B. Griffiths and C.-S. Niu, *Phys. Rev. Lett.* **76**, 3228 (1996).
- [13] R. Somma, G. Ortiz, J. E. Gubernatis, E. Knill, and R. Laflamme, *Phys. Rev. A* **65**, 042323 (2002).
- [14] B. L. Higgins, D. W. Berry, S. D. Bartlett, H. M. Wiseman, and G. J. Pryde, *Nature (London)* **450**, 393 (2007).
- [15] E. Knill, G. Ortiz, and R. D. Somma, *Phys. Rev. A* **75**, 012328 (2007).
- [16] N. Wiebe and C. Granade, *Phys. Rev. Lett.* **117**, 010503 (2016).
- [17] K. M. Svore, M. Hastings, and M. Freedman, *Quantum Inf. Comput.* **14**, 306 (2013).
- [18] G. H. Low, N. P. Bauman, C. E. Granade, B. Peng, N. Wiebe, E. J. Bylaska, D. Wecker, S. Krishnamoorthy, M. Roetteler, K. Kowalski *et al.*, [arXiv:1904.01131](https://arxiv.org/abs/1904.01131).
- [19] T. E. O’Brien, B. Tarasinski, and B. Terhal, *New J. Phys.* **21**, 023022 (2019).
- [20] R. D. Somma, *New J. Phys.* **21**, 123025 (2019).
- [21] L. Lin and Y. Tong, [arXiv:2102.11340](https://arxiv.org/abs/2102.11340).
- [22] S. Kimmel, G. H. Low, and T. J. Yoder, *Phys. Rev. A* **92**, 062315 (2015).
- [23] H. F. Trotter, *Proc. Natl. Acad. Sci. U.S.A.* **10**, 545 (1959).
- [24] M. Suzuki, *J. Math. Phys. (N.Y.)* **32**, 400 (1991).
- [25] D. W. Berry and A. M. Childs, *Quantum Inf. Comput.* **12**, 29 (2012).
- [26] A. Aspuru-Guzik, A. D. Dutoi, P. J. Love, and M. Head-Gordon, *Science* **309**, 1704 (2005).
- [27] Y. Cao, J. Romero, J. P. Olson, M. Degroote, P. D. Johnson, M. Kieferová, I. D. Kivlichan, T. Menke, B. Peropadre, N. P. D. Sawaya, S. Sim, L. Veis, and A. Aspuru-Guzik, *Chem. Rev.* **119**, 10856 (2019).
- [28] E. F. Dumitrescu, A. J. McCaskey, G. Hagen, G. R. Jansen, T. D. Morris, T. Papenbrock, R. C. Pooser, D. J. Dean, and P. Lougovski, *Phys. Rev. Lett.* **120**, 210501 (2018).
- [29] S. P. Jordan, K. S. Lee, and J. Preskill, *Science* **336**, 1130 (2012).
- [30] A. M. Childs, D. Maslov, Y. Nam, N. J. Ross, and Y. Su, *Proc. Natl. Acad. Sci. U.S.A.* **115**, 9456 (2018).
- [31] For example, in approaches to simulation based on Trotterization $\mathcal{W}(\mathcal{H}) \approx \exp(i\mathcal{H}t)$ and thus $\chi(\alpha) = \alpha t$.
- [32] One might also consider QPE implementations with a single auxiliary qubit and an adaptive choice of the premeasurement unitary applied to the auxiliary qubit across multiple rounds, but this approach does not inherit all of the properties of the approach in Fig. 1(a).
- [33] C. D. Bruzewicz, J. Chiaverini, R. McConnell, and J. M. Sage, *Appl. Phys. Rev.* **6**, 021314 (2019).
- [34] M. Kjaergaard, M. E. Schwartz, J. Braumüller, P. Krantz, J. I.-J. Wang, S. Gustavsson, and W. D. Oliver, *Annu. Rev. Condens. Matter Phys.* **11**, 369 (2020).
- [35] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Woerner, *Phys. Rev. A* **52**, 3457 (1995).
- [36] S. Khatri, R. LaRose, A. Poremba, L. Cincio, A. T. Sornborger, and P. J. Coles, *Quantum* **3**, 140 (2019).

- [37] D. Venturelli, M. Do, B. O’Gorman, J. Frank, E. Rieffel, K. E. Booth, T. Nguyen, P. Narayan, and S. Nanda, in *12th International Scheduling and Planning Application Workshop (SPARK)* (2019), <https://icaps19.icaps-conference.org/index.html>.
- [38] D. Maslov, *New J. Phys.* **19**, 023035 (2017).
- [39] A. Botea, A. Kishimoto, and R. Marinescu, in *Eleventh Annual Symposium on Combinatorial Search* (2018), <https://sites.google.com/view/socs18/>.
- [40] K. E. Booth, M. Do, J. C. Beck, E. Rieffel, D. Venturelli, and J. Frank, in *Twenty-Eighth International Conference on Automated Planning and Scheduling* (AAAI Press, Palo Alto, 2018).
- [41] G. Song and A. Klappenecker, [arXiv:quant-ph/0207157](https://arxiv.org/abs/quant-ph/0207157).
- [42] V. V. Shende and I. L. Markov, *Quantum Inf. Comput.* **9**, 461 (2009).
- [43] E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, and D. Preda, *Science* **292**, 472 (2001).
- [44] D. Poulin and P. Wocjan, *Phys. Rev. Lett.* **102**, 130503 (2009).
- [45] Y. Ge, J. Tura, and J. I. Cirac, *J. Math. Phys. (N.Y.)* **60**, 022202 (2019).
- [46] L. Lin and Y. Tong, *Quantum* **4**, 372 (2020).
- [47] By “logarithmic spacing” we mean that the circuit depths are, e.g., 1, 2, 4, 8, ..., as opposed to, e.g., 1, 2, 3, 4, ..., i.e., the circuit depths are spaced uniformly on a logarithmic scale.
- [48] K. Rudinger, S. Kimmel, D. Lobser, and P. Maunz, *Phys. Rev. Lett.* **118**, 190502 (2017).
- [49] A. M. Meier, K. A. Burkhardt, B. J. McMahon, and C. D. Herold, *Phys. Rev. A* **100**, 052106 (2019).
- [50] For the example of $R_x(\theta) = \exp(-i\theta\sigma_x/2)$, the Hamiltonian that generates R_x is $\theta\sigma_x/2$; by inspection, the Hamiltonian eigenvalue difference is θ .
- [51] Other spacings, i.e., $k_j \neq 2^j$ may potentially be utilized, but we do not consider those cases here.
- [52] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevLett.126.210501> for more details, which includes Refs. [53–58].
- [53] S. B. Bravyi and A. Y. Kitaev, *Ann. Phys. (Amsterdam)* **298**, 210 (2002).
- [54] J. T. Seeley, M. J. Richard, and P. J. Love, *J. Chem. Phys.* **137**, 224109 (2012).
- [55] 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.*, *Phys. Rev. X* **6**, 031007 (2016).
- [56] B. Drury and P. Love, *J. Phys. A* **41**, 395305 (2008).
- [57] B. Higgins, D. Berry, S. Bartlett, M. Mitchell, H. Wiseman, and G. Pryde, *New J. Phys.* **11**, 073023 (2009).
- [58] F. Belliard and V. Giovannetti, *Phys. Rev. A* **102**, 042613 (2020).
- [59] A. E. Russo, W. M. Kirby, K. M. Rudinger, A. D. Baczewski, and S. Kimmel, *Phys. Rev. A* **103**, 042609 (2021).
- [60] We emphasize here that the error considered here is in addition to statistical error, which is managed by taking sufficiently many samples.
- [61] IBM Quantum, <https://quantum-computing.ibm.com/>, 2021
- [62] H. Abraham *et al.*, QISKIT: An open-source framework for quantum computing (2019), <http://dx.doi.org/10.5281/zenodo.2562110>.
- [63] Such noise models can range from, e.g., one- and two-local completely positive trace-preserving maps [64,65] to “non-local” models that include cross talk errors [66,67] to time-dependent error models [68], as opposed to the uniform depolarizing error model implicit in the provided calibration data [69,70].
- [64] S. T. Merkel, J. M. Gambetta, J. A. Smolin, S. Poletto, A. D. Córcoles, B. R. Johnson, C. A. Ryan, and M. Steffen, *Phys. Rev. A* **87**, 062119 (2013).
- [65] R. Blume-Kohout, J. K. Gamble, E. Nielsen, K. Rudinger, J. Mizrahi, K. Fortier, and P. Maunz *Nat. Commun.* **8**, 14485 (2017).
- [66] K. Rudinger, T. Proctor, D. Langharst, M. Sarovar, K. Young, and R. Blume-Kohout, *Phys. Rev. X* **9**, 021045 (2019).
- [67] M. Sarovar, T. Proctor, K. Rudinger, K. Young, E. Nielsen, and R. Blume-Kohout, *Quantum* **4**, 321 (2020).
- [68] T. Proctor, M. Reville, E. Nielsen, K. Rudinger, D. Lobser, P. Maunz, R. Blume-Kohout, and K. Young, *Nat. Commun.* **11**, 5396 (2020).
- [69] E. Magesan, J. M. Gambetta, and J. Emerson, *Phys. Rev. A* **85**, 042311 (2012).
- [70] Jay M. Gambetta, A. D. Córcoles, S. T. Merkel, B. R. Johnson, John A. Smolin, Jerry M. Chow, Colm A. Ryan, Chad Rigetti, S. Poletto, Thomas A. Ohki, Mark B. Ketchen, and M. Steffen, *Phys. Rev. Lett.* **109**, 240504 (2012).
- [71] Y. R. Sanders, J. J. Wallman, and B. C. Sanders, *New J. Phys.* **18**, 012002 (2015).
- [72] E. Nielsen, R. J. Blume-Kohout, K. M. Rudinger, T. J. Proctor, L. Saldyt *et al.*, PYTHON GST implementation (pyGSTi) v. 0.9, Technical Report, Sandia National Lab. (SNL-NM), Albuquerque, NM, 2019, <https://www.osti.gov/biblio/1543289>.
- [73] E. Nielsen, K. Rudinger, T. Proctor, A. Russo, K. Young, and R. Blume-Kohout, *Quantum Sci. Technol.* **5**, 044002 (2020).
- [74] S. Lu, M. C. Bañuls, and J. I. Cirac, [arXiv:2006.03032](https://arxiv.org/abs/2006.03032).
- [75] T. E. O’Brien, S. Polla, N. C. Rubin, W. J. Huggins, S. McArdle, S. Boixo, J. R. McClean, and R. Babbush, [arXiv:2010.02538](https://arxiv.org/abs/2010.02538).