


Random Compiler for Fast Hamiltonian Simulation

Earl Campbell

Department of Physics and Astronomy, University of Sheffield, Sheffield S10 2TN, United Kingdom

 (Received 21 January 2019; revised manuscript received 25 June 2019; published 14 August 2019)

The dynamics of a quantum system can be simulated using a quantum computer by breaking down the unitary into a quantum circuit of one and two qubit gates. The most established methods are the Trotter-Suzuki decompositions, for which rigorous bounds on the circuit size depend on the number of terms L in the system Hamiltonian and the size of the largest term in the Hamiltonian Λ . Consequently, the Trotter-Suzuki method is only practical for sparse Hamiltonians. Trotter-Suzuki is a deterministic compiler but it was recently shown that randomized compiling offers lower overheads. Here we present and analyze a randomized compiler for Hamiltonian simulation where gate probabilities are proportional to the strength of a corresponding term in the Hamiltonian. This approach requires a circuit size independent of L and Λ , but instead depending on λ the absolute sum of Hamiltonian strengths (the ℓ_1 norm). Therefore, it is especially suited to electronic structure Hamiltonians relevant to quantum chemistry. Considering propane, carbon dioxide, and ethane, we observe speed-ups compared to standard Trotter-Suzuki of between $306\times$ and $1591\times$ for physically significant simulation times at precision 10^{-3} . Performing phase estimation at chemical accuracy, we report that the savings are similar.

DOI: [10.1103/PhysRevLett.123.070503](https://doi.org/10.1103/PhysRevLett.123.070503)

Quantum computers could be used to mimic the dynamics of other quantum systems, providing a computational method to understand physical systems beyond the reach of classical supercomputers. A quantum computation is broken down into a discrete sequence of elementary one and two qubit gates. To simulate the continuous unitary evolution of the Schrödinger equation, an approximation must be made into a finite sequence of discrete gates. The precision of this approximation can be improved by using more gates. The standard approaches are the Trotter and higher order Suzuki decompositions [1–3]. In addition to simulating dynamics, we are often interested in learning the energy spectra of Hamiltonians. Assuming a good ansatz for the ground state, we can combine quantum simulation with phase estimation to find the energy of the ground state [4] and excited states [5–7]. For a molecule with unknown electronic configuration, this is called the electronic structure problem [8,9] and it is crucially important in chemistry and material science. However, electronic structure Hamiltonians contain a very large number of terms and unfortunately the gate count of Trotter-Suzuki increases with the number of terms. While the scaling is formally efficient, the required number of gates is impractically large. An alternative to the Trotter-Suzuki method without this scaling problem would therefore have significant applications.

A recurrent theme in the literature is that stochastic noise can be less harmful than coherent noise [10,11], which hints that randomization might be useful for washing out coherent errors in circuit design. Poulin *et al.* [12] showed that randomness is especially useful in the

simulation of time-dependent Hamiltonians as it allows us to average out rapid Hamiltonian fluctuations. Campbell [13] and Hastings [14] have shown that random compiling can actually help reduce errors below what is feasible with a deterministic compiler. Since the optimization of Hamiltonian simulation circuits is a special case of compilation, one expects random compilers to be helpful in this setting. Following this line of reasoning, Childs, Ostrander, and Su [15] showed that it is useful to randomly permute the order of terms in Trotter-Suzuki decompositions. However, randomly permuted Trotter-Suzuki decompositions still suffer the same scaling problem that plagues deterministic Trotter-Suzuki; that is, the gate count depends on the number of Hamiltonian terms.

Here we propose a simple and elegant approach to Hamiltonian simulation that uses randomization to cure this scaling problem. Our proposal is similar to Trotter-Suzuki in that we implement a sequence of small rotations, without any use of ancillary qubits or complex circuit gadgets. Our key idea is to weight the probability of gates by the corresponding interaction strength in the Hamiltonian. Our simulation scheme can be seen as a Markovian process, which is inherently random but biased in such a way that we stochastically drift toward the correct unitary with high precision. For this reason, we call it the quantum stochastic drift protocol, or simply QDRIFT. Unlike any Trotter-Suzuki method, the gate count of QDRIFT is completely independent of the number of terms in the Hamiltonian. Consequently, we find that our approach can speed up quantum simulations of electronic structure Hamiltonians by several orders of magnitude

within regimes of practical interest. For the example of 60 qubit ethane, we find a speed-up of over a factor of 1000 when the approximation error is 0.001 and the simulation time is $t = 6000$ (the same simulation time often used in phase estimation [16]). In quantum chemistry, phase estimation is performed using controlled e^{iH} unitaries and here our techniques can lead to even larger resource savings.

Our analysis is limited in scope in two ways. First, we only compare against other Trotter-Suzuki decompositions. However, there are numerous approaches outside the Trotter-Suzuki family that make use of ancillary qubits and complex gadgets to obtain better asymptotic performance [17–22], such as the LCU (linear combinations of unitary) technique. Second, we only compare performance of rigorous bounds on gate counts, even though numerical studies of small systems show that far fewer gates are needed than are suggested by rigorous bounds [23–25]. Note that for the special case of local Hamiltonians, tighter analysis is possible because error propagation is localized and obeys Lieb-Robinson bounds [26,27], but, unfortunately, the electronic structure Hamiltonians are highly nonlocal.

The Hamiltonian simulation problem.—We begin by restating the problem more formally. Consider a Hamiltonian

$$H = \sum_{j=1}^L h_j H_j \quad (1)$$

decomposed into a sum of H_j each of which is Hermitian and normalized (such that the largest singular value of H_j is 1). We can always choose H_j so that the weighting h_j are positive real numbers. Herein we denote $\lambda = \sum_j h_j$ and remark that this upper bounds the largest singular value of H . The decomposition of the Hamiltonian should be such that for each H_j the unitary $e^{i\tau H_j}$ can be implemented on our quantum hardware for any τ . Our goal is then to find an approximation of e^{iH} into a sequence of $e^{i\tau H_j}$ gates up to some desired precision. We use the number of $e^{i\tau H_j}$ unitaries to quantify the cost of the quantum computation, and we aim to minimize the number of such unitaries used. In the simplest Trotter formulae, one divides $U = e^{iH}$ into r segments so that $U = U_r^r$ with $U_r = e^{iH/r}$ and uses that

$$V_r = \prod_{j=1}^L e^{i h_j H_j / r} \quad (2)$$

approaches U_r in the large r limit. Furthermore, r repetitions of V_r will approach U in the large r limit, so $V_r^r \rightarrow U$. The gate count in this sequence will be $N = Lr$, so we would like to know the smallest r that suffices to achieve a desired precision ϵ . Analytic work on this problem (we use the analysis of Refs. [15,25]) shows that the Trotter error is no more than

$$\epsilon = \frac{L^2 \Lambda^2 t^2}{2r} e^{\Lambda t / r}, \quad (3)$$

where $\Lambda := \max_j h_j$ is the magnitude of the strongest term in the Hamiltonian. Solving for r we find approximately $r \sim L^2 \Lambda^2 t^2 / 2\epsilon$ segments are needed, each segments contains L unitaries, leading to a total gate count of $N = Lr \sim L^3 (\Lambda t)^2 / 2\epsilon$. Table I compares this against other approaches including more sophisticated higher-order Suzuki decompositions. As we increase the order of the decomposition, the scaling approaches $O(L^2 \Lambda t)$, although the constant factors become rapidly worse for higher orders, so that in practice the optimal choice is usually second or fourth order. Childs, Ostrander, and Su showed that randomly permuted Trotter decompositions can further improve the gate count (see Table I).

Having reviewed the prior art of product formulae, we notice the L dependence never improved below quadratic. Therefore, Trotter decompositions are limited to simulations of quantum systems with sparse interactions, so that L must scale polynomially with the system size n . Furthermore, in chemistry problems $L = O(n^4)$ and while technically efficient, the resulting $O(n^8)$ scaling is prohibitively large. Next we turn to our protocol that eliminates this dependence.

The QDRIFT protocol.—Our full algorithm is given as pseudocode in Fig. 1. Each unitary in the sequence is selected independently from an identical distribution (IID sampling). The strength τ_j of each unitary is fixed to a constant $\tau_j = \tau := t\lambda/N$, which is independent of h_j , so we implement gates of the form $e^{i\tau H_j}$. The probability of choosing unitary $e^{i\tau H_j}$ is weighted by the interaction strength h_j , with normalization of the distribution entailing that $p_j = h_j/\lambda$. Therefore, the full circuit implemented is labeled by an ordered list of j values $\mathbf{j} = \{j_1, j_2, \dots, j_N\}$ that corresponds to unitary

$$V_{\mathbf{j}} = \prod_{k=1}^N e^{i\tau H_{j_k}}, \quad (4)$$

which is selected from the product distribution $P_{\mathbf{j}} = \lambda^{-N} \prod_{k=1}^N h_{j_k}$. While this quantum process is random,

TABLE I. Resource scaling for different product formulae (see the Supplemental Material [28] for details and caveats).

Protocol	Gate count (upper bound)
1st order Trotter DET	$O[L^3(\Lambda t)^2/\epsilon]$
2nd order Trotter DET	$O[L^{5/2}(\Lambda t)^{3/2}/\epsilon^{1/2}]$
(2k)th order Trotter DET	$O[L^{2+(1/2k)}(\Lambda t)^{1+(1/2k)}/\epsilon^{1/2k}]$
(2k)th order Trotter RANDOM	$O[L^2(\Lambda t)^{1+(1/2k)}/\epsilon^{1/2k}]$
QDRIFT (general result)	$O[(\lambda t)^2/\epsilon]$
QDRIFT (when $\lambda = \Lambda L$)	$O[L^2(\Lambda t)^2/\epsilon]$
QDRIFT (when $\lambda = \Lambda\sqrt{L}$)	$O[L(\Lambda t)^2/\epsilon]$

Input: A list of Hamiltonian terms $H = \sum_j h_j H_j$, a classical oracle function `SAMPLE()` that returns an value j from the probability distribution $p_j = h_j / (\sum_j h_j)$ and a target precision ϵ .
Output: An ordered list V_{list} of unitary gates of the form $\exp(i\tau H_j)$.

1. $\lambda \leftarrow \sum_j h_j$
2. $N \leftarrow \lceil 2\lambda^2 t^2 / \epsilon \rceil$ (or solve exact expression in appendix)
3. $i \leftarrow 0$
4. $V_{\text{list}} = \{\}$ (set gate list empty)
5. While $i < N$
 - (a) $i \leftarrow i + 1$
 - (b) $j \leftarrow \text{SAMPLE}()$
 - (c) Append $e^{i\lambda t H_j / N}$ to ordered list V_{list}
6. Return V_{list} .

FIG. 1. Pseudocode for the QDRIFT protocol.

we build into the probabilities a bias so that with many repetitions the evolution stochastically drifts towards the target unitary. Since each unitary is sampled independently, the process is entirely Markovian and we can consider the evolution resulting from a single random operation. The evolution is mathematically represented by a quantum channel that mixes unitaries as follows

$$\mathcal{E}(\rho) = \sum_j p_j e^{i\tau H_j} \rho e^{-i\tau H_j} \quad (5)$$

$$= \sum_j \frac{h_j}{\lambda} e^{i\tau H_j} \rho e^{-i\tau H_j}. \quad (6)$$

Using Taylor series expansions of the exponentials, we have that to leading order in τ ,

$$\mathcal{E}(\rho) = \rho + i \sum_j \frac{h_j \tau}{\lambda} (H_j \rho - \rho H_j) + O(\tau^2). \quad (7)$$

We compare this with the channel \mathcal{U}_N that is one N th of the full dynamics we wish to simulate, so that

$$\begin{aligned} \mathcal{U}_N(\rho) &= e^{itH/N} \rho e^{-itH/N} \\ &= \rho + i \frac{t}{N} (H\rho - \rho H) + O\left(\frac{t^2}{N^2}\right), \end{aligned} \quad (8)$$

where we have expanded out to leading order in t/N . Using that $H = \sum_j h_j H_j$, we have

$$\mathcal{U}_N(\rho) = \rho + i \sum_j \frac{th_j}{N} (H_j \rho - \rho H_j) + O\left(\frac{t^2}{N^2}\right). \quad (9)$$

Comparing \mathcal{E} and \mathcal{U}_N , we see that the zeroth and first order terms match whenever $\tau = t\lambda/N$. The higher order terms will not typically match and more careful analysis (see the Supplemental Material [28]) shows that the channels \mathcal{E} and \mathcal{U}_N differ by an amount bounded by

$$\delta \leq \frac{2\lambda^2 t^2}{N^2} e^{2\lambda t/N} \approx \frac{2\lambda^2 t^2}{N^2}, \quad (10)$$

where the first inequality is rigorous and the approximation on the right is very accurate even for modest N .

Since δ is the approximation error on a single random operation \mathcal{E} , the error of N repetitions \mathcal{E}^N relative to the target unitary U is then

$$\epsilon = N\delta \lesssim \frac{2\lambda^2 t^2}{N}. \quad (11)$$

We see the total error decreases as we increase N . Setting N to $N_{\text{QD}} = 2\lambda^2 t^2 / \epsilon$ (rounding up to nearest integer) suffices to ensure that $N\delta$ is less than the required precision ϵ . The exact value of N is easily calculated, but again the aforementioned approximation is very good.

Asymptotics comparison.—The QDRIFT approach needs approximately $2\lambda^2 t^2 / \epsilon$ gates and we include this in Table I to compare against prior methods. Since it does not explicitly depend on L , there are no sparsity constraints and this is the only known product formulae to beat the $O(L^2)$ barrier. Though one may argue that L dependence is hidden in $\lambda = \sum_j h_j$. The bounds for other Trotter-Suzuki formulae are given in terms of $\Lambda = \max_j h_j$, and these quantities are related by $\lambda \leq \Lambda L$. The worst case for QDRIFT is therefore $\lambda = \Lambda L$, which occurs for systems like the 1D nearest neighbor Heisenberg chain [15,25,32]. In this regime, QDRIFT is significantly better than first-order Trotter but the asymptotics suggest that it will be outperformed by higher-order Trotter. However, many real world systems have long range interactions that lead to $\lambda \ll \Lambda L$. For instance, if we had $\lambda \sim \Lambda \sqrt{L}$ then the QDRIFT scaling would be $O(L)$, which is comfortably better than the $O(L^2)$ that was the best prior art. While QDRIFT has significantly better L dependence, it does depend quadratically on Λt whereas higher-order Trotter approaches linear scaling in Λt . Therefore, for a fixed Hamiltonian, QDRIFT may excel for short times, but there will always be a critical t value above which it performs worse.

Numerics.—We have generated electronic structure Hamiltonians for propane, carbon-dioxide and ethane by using the OPENFERMION library [33], which naturally satisfy $\lambda \ll \Lambda L$ and so QDRIFT should perform favorably. We present our results in Fig. 2 using target precision $\epsilon = 10^{-3}$. Observe that QDRIFT offers a significant advantage at low t , which is often several orders of magnitude better than any prior Trotter-Suzuki decomposition. We remarked in our introduction that $t = 6000$ has been




FIG. 2. The number of gates used to implement $U = \exp(iHt)$ for various t and $\epsilon = 10^{-3}$ and three different Hamiltonians (energies in Hartree) corresponding to the electronic structure Hamiltonians of propane (in STO-3G basis), carbon dioxide (in 6-31g basis), and ethane (in 6-31g basis). Since the Hamiltonian contains some very small terms, one can argue that conventional Trotter-Suzuki methods would fare better if they truncate the Hamiltonian by eliminating negligible terms. For this reason, whenever simulating to precision ϵ we also remove from the Hamiltonian the smallest terms with weight summing to ϵ . This makes a fairer comparison, though in practice we found it made no significant difference to performance. For the Suzuki decompositions we choose the best from the first four orders, which is sufficient to find the optimal.

identified as relevant for phase estimation in quantum chemistry problems [16] and here we see speed-ups of $591\times$, $306\times$, and $1006\times$ for propane, carbon dioxide, and ethane (respectively). However, since QDRIFT scales worse with t than higher-order Trotter, for longer time simulations our advantage decreases and we eventually observe a crossover at times around $t = 10^7$ – 10^8 where prior methods perform better. But this crossover does not occur until the simulation time is so long that 10^{23} – 10^{25} gates are required. This is an extremely high gate count. Quantum error correction would certainly be needed and it is well known that to implement this many non-Clifford gates would require many billions of physical qubits even with generous hardware assumptions [34–37]. For these molecules, any foreseeable device performing Hamiltonian simulation would significantly benefit from using QDRIFT over standard Trotter-Suzuki.

Phase estimation.—When using phase estimation to find ground state energies, one performs many controlled- $\exp(iHt)$ rotations. Estimating energies to precision δ_E —chemical precision means $\delta_E \sim 10^{-4}$ —the largest time used is at least $t \sim \pi/\delta_E$, with slightly longer times needed to boost the inherent success probability of phase estimation. Note that the Trotter error ϵ is not directly connected to δ_E but instead contributes to the failure probability. Running phase estimation several times allows us to handle modest failure probabilities, so in practice ϵ can be much larger than δ_E . Therefore, the relevant ϵ and t regime for phase estimation matches the regime where QDRIFT performs well in simulation tasks. We provide a detailed analysis of phase estimation in the Supplemental Material [28], which shows that QDRIFT offers 2–3 orders of magnitude improvement when the failure probability of a single run is 5%.

Diamond norm distance.—An important technicality is that for a random circuit the appropriate measure of error ϵ is the diamond norm distance [38]. If we instead consider a specific instance of a randomly chosen unitary V_j in

Eq. (4), then the error will typically (on average) be much larger than ϵ , with standard statistical arguments (see e.g., Ref. [12]) suggesting it would be closer to $\sqrt{\epsilon}$. It is counterintuitive that the random circuit error is considerably less than the error of any particular unitary, so let us elaborate. If we initialize the quantum computer in state $|\psi\rangle$, then QDRIFT leads to state $|\Psi_j\rangle = V_j|\psi\rangle$ with probability P_j . If our experimental setup forgets (erases from memory) which unitary was implemented, then it prepares the mixed state

$$\rho = \mathcal{E}^N(|\psi\rangle\langle\psi|) = \sum_j P_j V_j |\psi\rangle\langle\psi| V_j^\dagger = \sum_j P_j |\Psi_j\rangle\langle\Psi_j|. \quad (12)$$

Since this channel is ϵ close in diamond distance to the ideal channel \mathcal{U} , it follows that ρ is ϵ close in trace norm distance to the target state $\mathcal{U}(|\psi\rangle\langle\psi|) = U|\psi\rangle\langle\psi|U^\dagger$. Trace norm distance is the relevant quantity because it ensures that if we perform a measurement, then the probabilities of the outcomes (on state ρ) do not differ by more than 2ϵ from the ideal probability given by $U|\psi\rangle$. Provided we estimate expectation values over several runs, each using a new and independent, randomly generated unitary, the precision of our estimate will be governed by ϵ rather than the looser $\sqrt{\epsilon}$ bound obtained without use of the diamond norm.

Discussion.—A common setting is where H_j are taken as tensor products of Pauli spin operators, then e^{itH_j} can be realized using Clifford gates and a single-qubit Pauli Z rotation [39]. When performing quantum error correction, the resource overhead of Clifford gates is negligible [34,35] whereas the single-qubit Pauli Z rotation must be decomposed into a large number of single-qubit T and Clifford gates. One further advantage of QDRIFT is that it consumes many Pauli rotations of exactly the same angle, allowing the use of adder-circuit catalysis that significantly reduces T -counts [40,41]. This is especially true when the Pauli rotations belong to the Clifford hierarchy [42], since

one then has the option of directly distilling magic states providing the rotation without further compilation [43–46]. Interestingly, Duclos-Cianci and Poulin [44] give a short discussion of how their magic state distillation protocol could be used in a Hamiltonian simulation scheme using a modified-Trotter decomposition where the gates all have the same τ value. While they allude to such a Hamiltonian simulation protocol, they do not provide any details or error analysis and nor did they suggest that randomization would be part of the protocol.

This work was supported by the EPSRC (Grant No. EP/M024261/1). We thank Simon Benjamin, Xiao Yuan, and Sam McArdle for discussions on the electronic structure problem and providing molecular Hamiltonians taken from OPENFERMION. For regular discussions on Hamiltonian simulation we thank John Clark, David White, Ben Jones, and George O’Brien. We thank Yuan Su for sharing details regarding Ref. [15]. For comments on the manuscript, we thank Dominic Berry.

-
- [1] M. Suzuki, *Phys. Lett.* **146A**, 319 (1990).
 [2] M. Suzuki, *J. Math. Phys. (N.Y.)* **32**, 400 (1991).
 [3] D. W. Berry, G. Ahokas, R. Cleve, and B. C. Sanders, *Commun. Math. Phys.* **270**, 359 (2007).
 [4] D. S. Abrams and S. Lloyd, *Phys. Rev. Lett.* **83**, 5162 (1999).
 [5] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O’Brien, *Nat. Commun.* **5**, 4213 (2014).
 [6] O. Higgott, D. Wang, and S. Brierley, *Quantum* **3**, 156 (2019).
 [7] T. O’Brien, B. Tarasinski, and B. Terhal, *New J. Phys.* **21**, 023022 (2019).
 [8] A. Aspuru-Guzik, A. D. Dutoi, P. J. Love, and M. Head-Gordon, *Science* **309**, 1704 (2005).
 [9] S. McArdle, S. Endo, A. Aspuru-Guzik, S. Benjamin, and X. Yuan, [arXiv:1808.10402](https://arxiv.org/abs/1808.10402).
 [10] J. J. Wallman and J. Emerson, *Phys. Rev. A* **94**, 052325 (2016).
 [11] G. C. Knee and W. J. Munro, *Phys. Rev. A* **91**, 052327 (2015).
 [12] D. Poulin, A. Qarry, R. Somma, and F. Verstraete, *Phys. Rev. Lett.* **106**, 170501 (2011).
 [13] E. Campbell, *Phys. Rev. A* **95**, 042306 (2017).
 [14] M. B. Hastings, *Quantum Inf. Comput.* **17**, 488 (2017).
 [15] A. M. Childs, A. Ostrander, and Y. Su, [arXiv:1805.08385](https://arxiv.org/abs/1805.08385).
 [16] D. Wecker, B. Bauer, B. K. Clark, M. B. Hastings, and M. Troyer, *Phys. Rev. A* **90**, 022305 (2014).
 [17] D. W. Berry and A. M. Childs, *Quantum Inf. Comput.* **12**, 29 (2012).
 [18] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, in *Forum of Mathematics, Sigma* (Cambridge University Press, Cambridge, England, 2017), Vol. 5.
 [19] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, *Phys. Rev. Lett.* **114**, 090502 (2015).
 [20] D. W. Berry, A. M. Childs, and R. Kothari, in *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on* (IEEE, New York, 2015), pp. 792–809.
 [21] G. H. Low and I. L. Chuang, [arXiv:1610.06546](https://arxiv.org/abs/1610.06546).
 [22] R. Babbush, C. Gidney, D. W. Berry, N. Wiebe, J. McClean, A. Paler, A. Fowler, and H. Neven, *Phys. Rev. X* **8**, 041015 (2018).
 [23] D. Poulin, M. B. Hastings, D. Wecker, N. Wiebe, A. C. Doherty, and M. Troyer, [arXiv:1406.4920](https://arxiv.org/abs/1406.4920).
 [24] R. Babbush, J. McClean, D. Wecker, A. Aspuru-Guzik, and N. Wiebe, *Phys. Rev. A* **91**, 022311 (2015).
 [25] A. M. Childs, D. Maslov, Y. Nam, N. J. Ross, and Y. Su, *Proc. Natl. Acad. Sci. U.S.A.* **115**, 9456 (2018).
 [26] J. Haah, M. B. Hastings, R. Kothari, and G. H. Low, [arXiv:1801.03922](https://arxiv.org/abs/1801.03922).
 [27] A. M. Childs and Y. Su, [arXiv:1901.00564](https://arxiv.org/abs/1901.00564).
 [28] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevLett.123.070503> for detailed error bounds and analysis of phase estimation protocol, which includes Refs. [29,30,31].
 [29] R. Cleve, A. Ekert, C. Macchiavello, and M. Mosca, *Proc. R. Soc. Ser. A* **454**, 339 (1998).
 [30] B. L. Higgins, D. W. Berry, S. D. Bartlett, H. M. Wiseman, and G. J. Pryde, *Nature (London)* **450**, 393 (2007).
 [31] 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:1809.05523](https://arxiv.org/abs/1809.05523).
 [32] Y. Nam and D. Maslov, [arXiv:1805.04645](https://arxiv.org/abs/1805.04645).
 [33] J. R. McClean, I. D. Kivlichan, D. S. Steiger, Y. Cao, E. S. Fried, C. Gidney, T. Häner, V. Havlíček, Z. Jiang, M. Neeley *et al.*, [arXiv:1710.07629](https://arxiv.org/abs/1710.07629).
 [34] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, *Phys. Rev. A* **86**, 032324 (2012).
 [35] J. O’Gorman and E. T. Campbell, *Phys. Rev. A* **95**, 032338 (2017).
 [36] M. Reiher, N. Wiebe, K. M. Svore, D. Wecker, and M. Troyer, *Proc. Natl. Acad. Sci. U.S.A.* **114**, 7555 (2017).
 [37] E. Campbell, A. Khurana, and A. Montanaro, [arXiv:1810.05582](https://arxiv.org/abs/1810.05582).
 [38] J. Watrous, *The Theory of Quantum Information* (Cambridge University Press, Cambridge, England, 2018).
 [39] N. J. Ross and P. Selinger, *Quantum Inf. Comput.* **16**, 901 (2016).
 [40] C. Gidney, *Quantum* **2**, 74 (2018).
 [41] M. Beverland, E. Campbell, M. Howard, and V. Kliuchnikov, [arXiv:1904.01124](https://arxiv.org/abs/1904.01124).
 [42] D. Gottesman and I. L. Chuang, *Nature (London)* **402**, 390 (1999).
 [43] A. J. Landahl and C. Cesare, [arXiv:1302.3240](https://arxiv.org/abs/1302.3240).
 [44] G. Duclos-Cianci and D. Poulin, *Phys. Rev. A* **91**, 042315 (2015).
 [45] E. T. Campbell and J. O’Gorman, *Quantum Sci. Technol.* **1**, 015007 (2016).
 [46] E. T. Campbell and M. Howard, *Quantum* **2**, 56 (2018).