

Fast Quantum Algorithm for Spectral Properties

David Poulin,¹ Alexei Kitaev,² Damian S. Steiger,³ Matthew B. Hastings,^{4,5} and Matthias Troyer^{3,5}

¹*Département de Physique & Institut Quantique, Université de Sherbrooke, J1K 2R1, Canada*

²*California Institute of Technology, Pasadena, CA 91125, USA*

³*Theoretische Physik, ETH Zurich, 8093 Zurich, Switzerland*

⁴*Station Q, Microsoft Research, Santa Barbara, CA 93106-6105, USA*

⁵*Quantum Architecture and Computation Group, Microsoft Research, Redmond, WA 98052, USA*

(Dated: December 22, 2017)

We present two techniques that can greatly reduce the number of gates required for ground state preparation in quantum simulations. The first technique realizes that to prepare the ground state of some Hamiltonian, it is not necessary to implement the time-evolution operator: any unitary operator which is a function of the Hamiltonian will do. We propose one such unitary operator which can be implemented exactly, circumventing any Taylor or Trotter approximation errors. The second technique is tailored to lattice models, and is targeted at reducing the use of generic single-qubit rotations, which are very expensive to produce by distillation and synthesis fault-tolerantly. In particular, the number of generic single-qubit rotations used by our method scales with the number of parameters in the Hamiltonian, which contrasts with a growth proportional to the lattice site required by other techniques.

The simulation of quantum systems is one of the main foreseen applications of quantum computers [1, 2], it can be used to predict the properties of materials [3–6], molecules [7], quantum fields [8], etc. There are two conceptual ingredients to a quantum simulation, corresponding roughly to statics and dynamics. The static part consists in preparing an initial state of physical interest. Because we are typically interested in low-energy properties, we will often require to prepare the quantum computer in the ground state $|\phi_0\rangle$ of the simulated system's Hamiltonian H . Dynamics consists in reproducing the effect of the time-evolution operator $U(t) = e^{-iHt}$. While methods to produce the time-evolution operator have been known for decades [1, 2], the static problem is often the bottleneck of quantum simulations.

Simulating the dynamics of a quantum system on a quantum computer requires a quantum circuit which (approximately) reproduces the time-evolution operator $U(t) = e^{-iHt}$. The two standard approaches are the Trotter-Suzuki decomposition [2, 8] and Taylor expansions [9], in addition to recent advances based on quantum signal processing [10], all of which realize an approximation of $U(t)$ with controllable systematic errors. While the ideas presented in this Letter are primarily intended to simplify the static problem, the unary encoding technique we present can also be used to reduce the number of gates in dynamical simulations based on Taylor series.

The ability to reproduce the dynamics $U(t)$ of a quantum system on a quantum computer also provides a means to solve the static problem. This is because, employing quantum phase estimation [11–13], the circuit implementing $U(t)$ can be used to perform a projective energy measurement. Thus, given a good approximation $|\tilde{\phi}_0\rangle$ of the ground state $|\phi_0\rangle$, an energy measurement should collapse the wavefunction onto the ground state with reasonable probability $|\langle\phi_0|\tilde{\phi}_0\rangle|^2$.

The state $|\tilde{\phi}_0\rangle$ could be obtained from various approximations such as mean-field. When no simple good ap-

proximation $|\tilde{\phi}_0\rangle$ exists, we can express the Hamiltonian $H = H_0 + V$ as a sum of some simple term H_0 of which we can easily prepare the ground state, and an “interaction” term V . Then, the Hamiltonian $H(g) = H_0 + gV$ has a simple ground state at $g = 0$ and its ground state at $g = 1$ is the one we want to prepare. We can thus initialize the quantum computer in the ground state $|\phi_0(g = 0)\rangle$ of $H(0)$, and perform a sequence of measurements which, with high probability, will result in the state $|\phi_0(1)\rangle$. The main idea [14] is that for $|g - g'|$ sufficiently small, the probability $|\langle\phi_0(g)|\phi_0(g')\rangle|^2$ that a measurement of $H(g')$ performed on $|\phi_0(g)\rangle$ results in the ground energy should be close to 1. We thus choose a sequence of interaction strengths $0 = g_0 < g_1 < g_2 < \dots < g_N = 1$ and perform a sequence of energy measurement using phase estimation of the time-evolution operator $U(j, t) = e^{-iH(g_j)t}$. As predicted by the quantum Zeno effect, a sequence of almost identical measurements should drag the state in the measurement basis with high probability $\prod_{j=1}^N |\langle\phi_0(g_{j-1})|\phi_0(g_j)\rangle|^2$. This method can be shown to be equivalent to adiabatic evolution [15, 16], but has the advantage that only time evolution under a *time-independent* Hamiltonian is required. Other methods to prepare the ground state include quantum Metropolis sampling [17, 18], and Grover search [19].

The time-evolution operator is used jointly with quantum phase estimation to realize an energy measurement. This simply builds on the fact that H and $U(t)$ share the same eigenstates $|\phi_k\rangle$ and, for sufficiently small $t \leq \pi/\|H\|$, have a one-to-one mapping between their eigenvalues, E_k and $e^{-iE_k t}$ respectively. Clearly, this is true for many other function of the Hamiltonian $f(H)$, in particular functions that are Lipschitz continuous at E_0 . The first technique we present uses ideas from quantum walks [20], Taylor expansions [9], and the related Qubitization [21] to build a unitary transformation W which is a simple function of H . The unitary transformation W has the advantage of being implementable exactly in

an ideal quantum circuit using roughly the same number of gates as required by a single Trotter step or a single segment of a first-order Taylor approximation.

The second technique we present becomes motivated once we realize that not all gates in a quantum circuit have the same cost. In standard approaches to fault-tolerant quantum computation [22–25], Clifford gates (CNOT, Hadamard, $S = \text{diag}(1, i)$) are realized in an intrinsically fault-tolerant way. An extra gate is obtained by magic state distillation [26–30]. It can be either a Toffoli = control-control-not, or $T = \text{diag}(1, e^{i\pi/4})$, or $\sqrt{\text{SWAP}}$, or control-SWAP. These gates share the property of belonging to the third level of the Clifford hierarchy [31], so they can be distilled by standard methods and together with the Clifford gates they form a universal gate set. Thus, any single-qubit rotation can be approximated by a sequence of universal gates in a protocol called gate synthesis [32–36]. The overhead associated with distillation and synthesis scales poly-logarithmically with the inverse targeted accuracy δ , and the constant pre-factors are large, making the cost of third-level gates substantially larger than Clifford gates, and the cost of generic single-qubit rotation even more so.

All the simulations methods mentioned above make use of the fact that the Hamiltonian is sparse in some natural basis, or typically that it can be written as a sum of terms $H = \sum_{j=1}^N \alpha_j P_j$ where N grows polynomially with the system size n (the number of particles or lattice size). For instance, if there can be at most k qubits involved in each interaction term, N would scale like $n^k 3^k$. Since there are N real parameters β_j in this Hamiltonian, the quantum circuit for $U(t)$ requires a number of generic single-qubit rotations that scales like N . In lattice models however, such as encountered in condensed-matter physics or lattice field and gauge theories, many of the α_j have the same value. For instance the Hubbard model $H_{\text{Hubbard}} = -t \sum_{\langle i,j \rangle, \sigma} c_{i,\sigma}^\dagger c_{j,\sigma} + U \sum_j c_{i\uparrow}^\dagger c_{i\uparrow} c_{i\downarrow}^\dagger c_{i\downarrow}$ or the Ising model $H_{\text{Ising}} = g \sum_j X_j + J \sum_{\langle i,j \rangle} Z_i Z_j$ each contain only a single real parameter U/t and J/g respectively.

The above parameter counting argument suggest that it may in principle be possible to realize the time-evolution operator of such systems using only a constant number of generic single-qubit rotations. The second technique we present realizes this expectation by replacing these N generic single-qubit rotations by third-level gates, hence avoiding the cost of synthesis. If we denote by $C_D(\delta)$ the cost of distilling a third-level gate and $C_S(\delta)$ the cost of synthesising a generic single-qubit rotation, then the cost of our method is $mC_D(\delta)C_S(\delta) + NC_D(\delta)$ where m is the number of α_j with distinct values, while each Trotter step requires $NC_D(\delta)C_C(\delta)$. Our method can straightforwardly be incorporated into the Taylor series method to obtain similar savings.

Spectrum by Quantum Walk — To prepare the ground state of Hamiltonian H , our approach is to realize a sim-

ple quantum circuit W which does not implement a unitary time evolution $U(t)$ but some other function of the Hamiltonian. Without loss of generality, we assume that the Hamiltonian is non-negative $H \geq 0$ and that it can be expressed as $H = \sum_{j=0}^N \alpha_j P_j$ where the P_j are multi-qubit Pauli operators and $P_0 = I$. We rescale the Hamiltonian by a factor $\mathcal{N} = \sum_{j=0}^N |\alpha_j| \in \mathcal{O}(N)$ and note

$$\bar{H} = \frac{H}{\mathcal{N}} = \sum_j |\beta_j|^2 P_j, \quad (1)$$

and it follows that $\sum_j |\beta_j|^2 = 1$. Note that any sign of α_j can be absorbed in the definition of P_j . Obviously this rescaling does not affect the eigenstates, but it does change the spectral gap by a factor \mathcal{N} and this will be important when comparing this algorithms to ones based on Trotter expansions.

The unitary transformation W we construct acts on $n + \log(N+1)$ qubits, i.e., the n *system* qubits and $\log(N+1)$ *control* qubits whose basis states $|j\rangle$ are in one-to-one correspondence with the $(N+1)$ terms P_j of the Hamiltonian. There exists an invariant subspace of W on which the spectrum of W is a simple function of H . By initializing the quantum computer to that subspace, we thus obtain the desired effect.

Following [9, 21], define $|\beta\rangle$, B , S and V as follows

$$|\beta\rangle = B|0\rangle = \sum_j \beta_j |j\rangle, \quad (2)$$

$$S = B^\dagger (I - 2|0\rangle\langle 0|) B = (I - 2|\beta\rangle\langle\beta|) \otimes I, \quad \text{and} \quad (3)$$

$$V = \sum_j |j\rangle\langle j| \otimes P_j. \quad (4)$$

For any eigenstate $\bar{H}|\phi_k\rangle = E_k|\phi_k\rangle$, both S and V preserve the subspace spanned by the orthonormal states

$$|\varphi_k^0\rangle = \sum_j \beta_j |j\rangle \otimes |\phi_k\rangle \quad \text{and} \quad (5)$$

$$|\varphi_k^1\rangle = \frac{1}{\sqrt{1 - E_k^2}} (V - E_k) |\varphi_k^0\rangle. \quad (6)$$

Using the equation $S^2 = V^2 = I$, it is easy to show that in the above basis

$$S = \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix}, \quad V = \begin{pmatrix} E_k & \sqrt{1 - E_k^2} \\ \sqrt{1 - E_k^2} & -E_k \end{pmatrix}. \quad (7)$$

We define the unitary walk operator $W = SV$ which has eigenvalues $e^{\pm i\theta_k}$ and eigenstates $|\varphi_k^\pm\rangle = (|\varphi_k^0\rangle \pm i|\varphi_k^1\rangle)/\sqrt{2}$, where $\cos \theta_k = E_k$. Thus, by preparing an initial state of the form $B|0\rangle \otimes |\tilde{\phi}_0\rangle$ and performing phase estimation of W , we obtain either eigenvalue $\pm\theta_k$ with probability $|\langle\phi_k|\tilde{\phi}_0\rangle|^2$, and the post-measurement state is $|\varphi_k^\pm\rangle$. We can also use iterative phase estimation [37] which simply uses an ancillary qubit C prepare in the state $|+\rangle$, performs ${}^C W$, and measure the ancilla in the $|\pm\rangle$ basis, whose conditional probabilities

are $p_{\pm|k} = \frac{1}{2}(1 \pm E_k)$. Repeating reveals E_k and prepares one of the states $|\varphi_k^{\pm}\rangle$.

While the states $|\varphi_k^{\pm}\rangle$ are not eigenstates of H , we can nonetheless use them to compute expectation values. For instance, say we want to estimate $\langle\phi_k|\sigma|\phi_k\rangle$ for some multi-qubit Pauli operator σ . Observe that

$$\langle\varphi_k^{\pm}|\sigma|\varphi_k^{\pm}\rangle = \frac{1}{2} \left(1 + \frac{\Gamma_{\sigma} - E_k^2}{1 - E_k^2} \right) \langle\phi_k|\sigma|\phi_k\rangle \quad (8)$$

where $\Gamma_{\sigma} = \sum_j |\beta_j|^2 (-1)^{\sigma \star P_j}$ and $\sigma \star P_j = 0$ when σ and P_j commute and $\sigma \star P_j = 1$ when they anti-commute. Since E_k is known from the state preparation and Γ_{σ} can be computed efficiently, this gives us access to any static correlation functions.

Alternatively, if we insist on preparing an eigenstate of H rather than $|\varphi_k^{\pm}\rangle$, we have two solutions. First, ignoring the sign in the phase during phase estimation directly furnishes $|\varphi_k^0\rangle$ as the post-measurement state. Second, we can simply apply B^{\dagger} and measure if the control qubits are returned to the state $|0\rangle$. It is straightforward to verify that this occurs with probability $\frac{1}{2}$ and that in case of a positive outcome the resulting state of the system qubits is $|\phi_k\rangle$. In the case of a negative outcome, the quantum computer is collapsed onto the state $|\varphi_k^1\rangle$. At this point we can perform quantum phase estimation of W which will randomly project onto $|\varphi_k^{\pm}\rangle$, and iterate. The probability of a positive outcome after ℓ iterations is $1 - \frac{1}{2}^{\ell}$.

In either case, these procedures make the scheme fully compatible with the Zeno ground state preparation outlined above [14]. Had we instead chosen to perform adiabatic evolution with the operator $W(g)$ itself, we would have had to worry about the spectral gap to the states orthogonal to the space spanned by the $|\varphi_k^{\pm}\rangle$. But by completing a deterministic projection as described in this paragraph, we are guaranteed to always remain in this invariant subspace.

Fewer Single-Qubit Rotations by Unary Encoding — The walk operator W is composed of V and B . The complexity of transformation V stems from its multiple control, i.e. it applies a Pauli operator conditioned on one of N possible values. Each application of V requires $\mathcal{O}(N \log N)$ Toffoli gates. The complexity of B stems from its dependency on N real numbers β_j . So just by parameter counting, B requires $\Omega(N)$ generic single-qubit gates. But this last argument breaks down when several of the β_j are equal, as naturally occurs in condensed matter systems and lattice gauge theories.

Suppose the (rescaled) Hamiltonian takes the form

$$\bar{H} = |\beta_0|^2 I + \sum_{k=1}^K |\beta_k|^2 \sum_{j=1}^{N_k} P_j^k \quad (9)$$

such that $\sum_k N_k = N$ and $\sum_{k=0}^K |\beta_k| N_k = 1$, which is a special case of Eq. (1) where the $|\beta_j|^2$ are restricted to K distinct values. In lattice models with short-range interactions, K is a small constant and all N_k are roughly

equal to the number of particles n (deviations are caused by lattice boundaries). To reduce the number of generic single-qubit rotations used to implement W , we will increase the number of control qubits from $\log(N+1)$ to N , which we partition into K registers, with register k of size N_k . The algorithm will proceed exactly as above, except that the control register will be encoded in a unary basis, i.e. $|j\rangle = |00\dots 0100\dots 0\rangle$ where the 1 is at position j , and the state $|j=0\rangle = |00\dots 0\rangle$.

To implement B in this representation, we use a K -qubit rotation to prepare the state $\sum_k \sqrt{N_k} \beta_k |m_k\rangle$ where $m_1 = 1$ and $m_{k+1} = m_k + N_k$, and we use the unary representation. In other words, this state is a superposition of a 1 in the first position of each register and requires K generic single-qubit rotations. Then, using N $\sqrt{\text{SWAP}}$ gates, we delocalize the 1 in each register to obtain

$$|\beta\rangle = \beta_0 |0\rangle + \sum_k \beta_k \sum_{j=1}^{N_k} |m_k + j\rangle. \quad (10)$$

To achieve this, we assume that N_k is a power of 2 and we simply apply the gate $\sqrt{\text{SWAP}}$ in a binary-tree like structure, noting that $\sqrt{\text{SWAP}}(\alpha|0\rangle + \beta|1\rangle)|0\rangle = \alpha|00\rangle + \frac{\beta}{\sqrt{2}}(|10\rangle + |01\rangle)$. When N_k is not a power of 2, we can simply pad by adding terms $P_j^k = I$ for $j > N_k$.

The total number of $\sqrt{\text{SWAP}}$ gates used to implement B is thus upper bounded by $2N$. Note that to perform quantum phase estimation, the circuit W will generally need to be applied conditioned on some additional control register. To realize this, all we need to do is to control the single-qubit rotations used to prepare the state $\sum_k \sqrt{N_k} \beta_k |m_k\rangle$. The other gates do not need to be controlled because they act trivially in the absence of these initial single-qubit rotations.

The transformation V is straightforward to implement in the unary representation. It consists in applying the product of control- P_j^k where the control bit is the j th bit of the k 'th control register. Because P_j^k is a Pauli operator, control- P_j^k is a Clifford gate, so V is entirely composed of Clifford gates. To summarize, the original binary encoding uses $\log N$ control qubits, $\mathcal{O}(N)$ generic single-qubit rotations, and $\mathcal{O}(N \log N)$ third-level gates (Toffoli) while the unary encoding uses N control qubits, $\mathcal{O}(1)$ generic single-qubit rotations, and $\mathcal{O}(N)$ third-level gates ($\sqrt{\text{SWAP}}$).

Lattices with Long-Range Couplings — The unary encoding provides an advantage when there are only a few distinct coupling strengths in the model, while the binary encoding is best when most couplings are different. Lattice models with long-range interaction provide an intermediate regime, where there are many distinct couplings, but each one is repeated an extensive number of times. For instance, a long-range 1D Ising interaction $J \sum_{i < j} \sigma_z^i \sigma_z^j / (i - j)^{\alpha}$ has $K = n$ distinct coupling strengths $J_k = J/k^{\alpha}$, and each strength is repeated n times, falling somewhere between the unary and binary encoding regimes.

A good compromise is to prepare a n -qubit *coupling* state $\sum_k \sqrt{n} J_k |k\rangle$ to encode each coupling strength J_k in unary representation, and use a $\log n$ qubit register in state $\frac{1}{\sqrt{n}} \sum_i |i\rangle$ to encode the pair of sites i and $i+k$ that are coupled in binary representation. The transformation V can be realized by swapping qubits i and $i+k$ into position 0 and 1, applying a pair of Pauli operators controlled by the state of the coupling qubit, and unswapping the two qubits to their original position. Swapping a qubit into a given position can be realized with $\mathcal{O}(n)$ uses of the third-level gate control-SWAP, in a circuit of depth n (or $\log n$ using $n/2$ additional ancillary qubits). Thus, this scheme requires $\mathcal{O}(K)$ generic single-qubit rotations, $\mathcal{O}(Kn)$ third-level gates and $\mathcal{O}(K + \log n)$ ancillary qubits, which compares respectively to $\mathcal{O}(Kn)$, $\mathcal{O}(Kn \log(Kn))$, and $\mathcal{O}(\log(Kn))$ for the binary encoding and $\mathcal{O}(K)$, $\mathcal{O}(Kn)$ and $\mathcal{O}(Kn)$ for the unary encoding.

Discussion — We will now discuss the advantages and drawbacks of the methods we proposed above. The main drawback of using W instead of the time-evolution operator is the rescaling $\mathcal{N} \in \mathcal{O}(N)$ of the spectral gap. To compensate for this rescaling, we will need to apply the operator W a total of $\mathcal{N}t$ times to reach the same energy resolution that is achieved with $U(t)$. So the advantage of the proposed method depends on the number of Trotter steps that are needed to achieve a constant accuracy approximation of $U(t)$.

Consider first a lattice model with short range interactions. The analysis of Appendix B in [38] shows that the number of Trotter steps required to implement $U(t)$ to accuracy δ is $N_T \in \mathcal{O}(t\sqrt{\frac{n}{\delta}})$. To resolve the ground state energy – which is separated from the rest of the spectrum by an energy gap Δ – we need $t \gg 1/\Delta$ and $\delta \ll \Delta$, which yields $N_T \in \mathcal{O}(\frac{\sqrt{n}}{\Delta^{3/2}})$. Since each Trotter step requires $\mathcal{O}(n)$ generic single-qubit rotations, the total cost of implementing an energy measurement capable of resolving the ground state is dominated by $\mathcal{O}((\frac{n}{\Delta})^{3/2})$ single-qubit rotations. On the other hand, the combination of schemes proposed here require $\mathcal{O}(\frac{\mathcal{N}}{\Delta}) = \mathcal{O}(\frac{n}{\Delta})$ single-qubit rotations and $\mathcal{O}(n\frac{\mathcal{N}}{\Delta}) = \mathcal{O}(\frac{n^2}{\Delta})$ third-level gates. Thus, our approach could offer important advantages at intermediate values of n , when the gate synthesis cost is dominant. We note that this analysis depends on the error criteria used. The Trotter error may scale differently if one is interested not in the wave function but only in expectation values of local observables.

Second, consider a quantum chemistry problem. By the nature of Coulomb's force, once written in second quantized form $H = \sum_{pq} h_{pq} c_p^\dagger c_q + \sum_{pqrs} V_{pqrs} c_p^\dagger c_q^\dagger c_r c_s$, the Hamiltonian has $\mathcal{O}(n^4)$ real parameters. Since every gate in a quantum circuit has a (small) constant number of free parameter, we clearly need $\mathcal{O}(n^4)$ generic single-qubit rotations to implement a single Trotter step $e^{-iH\delta}$. The analysis of [38] shows that the number of Trotter steps required to implement $U(t)$ to accuracy δ is $N_T \in \mathcal{O}(t\frac{n^5}{\delta})$. The above reasoning thus leads to a

total number of generic single-qubit rotations $\mathcal{O}(\frac{n^9}{\Delta^{3/2}})$. The unary encoding is not of much help here, but using the transformation W instead of $U(t)$ to perform an energy measurement requires a total of $\mathcal{O}(n^4\frac{\mathcal{N}}{\Delta}) = \mathcal{O}(\frac{n^8}{\Delta})$ single-qubit rotations. This is a substantial improvement over Trotter approaches.

Note however that the previous paragraph was formulated in terms in upper bounds. It was found empirically that the number of Trotter steps can be much smaller in practice [39]. Indeed, for a range of small molecules, it was observed that the number of Trotter steps needs to scale as $N_T \in \mathcal{O}(t\frac{n^\gamma}{\sqrt{\delta}})$ with $\gamma \sim 2.5$ instead of $\gamma = 5$ revealed by the upper bound. This yields an overall scaling of $\mathcal{O}(\frac{n^{\gamma+4}}{\Delta^{3/2}})$ to resolve the ground state by Trotter approximation. Using the same range of small molecules as used in [39], we find that \mathcal{N} scales like n^η with $\eta \sim 1.6$ as opposed to $\eta = 4$ suggested by the upper bound. This yield an overall scaling of $\mathcal{O}(\frac{n^{\eta+4}}{\Delta})$ to resolve the ground state using our technique, leading roughly to the same savings by a factor $\frac{n}{\sqrt{\Delta}}$ as derived from the rigorous upper bounds above. It will be interesting to compare both approaches for large molecules – which will be possible once we have quantum computers with a large number of logical qubits.

Taylor expansions schemes [9] rely on breaking up the time-evolution operator $U(t) = U(t/r)^r$ into r segments, and approximating each segment by a Taylor series

$$U(t/r) \approx \sum_{k=0}^K \frac{1}{k!} (-iHt/r)^k. \quad (11)$$

To resolve the ground state, r needs to be greater than $\frac{\|H\|}{\Delta}$, and $K \sim \log(\frac{\|H\|}{\Delta^2})$. The scheme then proceeds by implementing a unitary transformation W similar to the one constructed here, but with the sum of Eq. (11) replacing the Hamiltonian. While the Hamiltonian has N terms, this sum has $\sim NK$ terms, so requires K times more gates. So in summary, our scheme has all the advantages of the Taylor series, except that it only needs a first order $K = 1$ expansion and uses a single segment, thus requires $Kr \sim \frac{\|H\|}{\Delta} \log(\frac{\|H\|}{\Delta^2})$ times fewer gates.

Finally, we see that the main cost of using W stems from the rescaling factor \mathcal{N} , which is equal to the sum of the absolute value of the terms of the Hamiltonian. To improve the scaling, we need a different function $f(H)$ with a smaller Lipschitz constant near E_0 . For instance, the Taylor series and quantum signal processing approaches [40] allow us to synthesize functions where $f(H)$ are degree- K polynomials. However, the reduction in \mathcal{N} comes with an multiplicative cost $\mathcal{O}(K)$. We leave this problem open for future research.

Acknowledgements We thank Thomas Häner for interesting discussions and Sergey Bravyi for feedback, and acknowledge support by the Swiss National Science Foundation, the NCCR QSIT, Canada's NSERC, the Canadian Institute for Advanced Research, Caltech's IQIM, and the Simons Foundation.

-
- [1] Richard P. Feynman, “Simulating physics with computers,” *International Journal of Theoretical Physics*, **21**, 467–488 (1982).
- [2] Seth Lloyd, “Universal quantum simulators,” *Science*, **273**, 1073–1077 (1996).
- [3] Daniel S. Abrams and Seth Lloyd, “Simulation of many-body Fermi systems on a universal quantum computer,” *Physical Review Letters*, **79**, 2586 (1997).
- [4] R. Somma, G. Ortiz, J. E. Gubernatis, E. Knill, and R. Laflamme, “Simulating physical phenomena by quantum networks,” *Physical Review A*, **65**, 042323 (2002).
- [5] Bela Bauer, Dave Wecker, Andrew J. Millis, Matthew B. Hastings, and Matthias Troyer, “Hybrid quantum-classical approach to correlated materials,” *Physical Review X*, **6**, 031045 (2016).
- [6] Dave Wecker, Matthew B. Hastings, Nathan Wiebe, Bryan K. Clark, Chetan Nayak, and Matthias Troyer, “Solving strongly correlated electron models on a quantum computer,” *Physical Review A*, **92**, 062318 (2015).
- [7] Alán Aspuru-Guzik, Anthony D. Dutoi, Peter J. Love, and Martin Head-Gordon, “Simulated quantum computation of molecular energies,” *Science*, **309**, 1704–1707 (2005).
- [8] Stephen P. Jordan, Keith S. M. Lee, and John Preskill, “Quantum algorithms for quantum field theories,” *Science*, **336**, 1130–1133 (2012).
- [9] Dominic W. Berry, Andrew M. Childs, Richard Cleve, Robin Kothari, and Rolando D. Somma, “Simulating hamiltonian dynamics with a truncated taylor series,” *Physical Review Letters*, **114**, 090502 (2015).
- [10] Guang Hao Low and Isaac L. Chuang, “Optimal hamiltonian simulation by quantum signal processing,” *Physical Review Letters*, **118**, 010501 (2017).
- [11] A. Yu. Kitaev, “Quantum measurements and the abelian stabilizer problem,” [arXiv preprint quant-ph/9511026 \(1995\)](#).
- [12] Richard Cleve, Artur Ekert, Chiara Macchiavello, and Michele Mosca, “Quantum algorithms revisited,” in *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, Vol. 454 (The Royal Society, 1998) pp. 339–354.
- [13] Daniel S. Abrams and Seth Lloyd, “Quantum algorithm providing exponential speed increase for finding eigenvalues and eigenvectors,” *Physical Review Letters*, **83**, 5162 (1999).
- [14] R. D. Somma, S. Boixo, Howard Barnum, and E. Knill, “Quantum simulations of classical annealing processes,” *Physical Review Letters*, **101**, 130504 (2008).
- [15] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, and Michael Sipser, “Quantum computation by adiabatic evolution,” (2000), [quant-ph/0001106](#).
- [16] D. Aharonov and A. Ta-Shma, “Adiabatic quantum state generation and statistical zero knowledge,” *Proc. 35th Annual ACM Symp. on Theo. Comp.*, 20 (2003).
- [17] Kristan Temme, Tobias J. Osborne, Karl G. Vollbrecht, David Poulin, and Frank Verstraete, “Quantum metropolis sampling,” *Nature*, **471**, 87–90 (2011).
- [18] Man-Hong Yung and Alán Aspuru-Guzik, “A quantum–quantum metropolis algorithm,” *Proceedings of the National Academy of Sciences*, **109**, 754–759 (2012).
- [19] David Poulin and Pawel Wocjan, “Preparing ground states of quantum many-body systems on a quantum computer,” *Physical Review Letters*, **102**, 130503 (2009).
- [20] Mario Szegedy, “Quantum speed-up of Markov chain based algorithms,” in *Foundations of Computer Science, 2004. Proceedings. 45th Annual IEEE Symposium on (IEEE, 2004)* pp. 32–41.
- [21] Guang Hao Low and Isaac L. Chuang, “Hamiltonian simulation by Qubitization,” [arXiv preprint arXiv:1610.06546 \(2016\)](#).
- [22] Dorit Aharonov and Michael Ben-Or, “Fault-tolerant quantum computation with constant error,” in *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing (ACM, 1997)* pp. 176–188.
- [23] Peter W. Shor, “Fault-tolerant quantum computation,” in *Foundations of Computer Science, 1996. Proceedings., 37th Annual Symposium on (IEEE, 1996)* pp. 56–65.
- [24] Andrew M. Steane, “Efficient fault-tolerant quantum computing,” *Nature*, **399**, 124–126 (1999).
- [25] Robert Raussendorf and Jim Harrington, “Fault-tolerant quantum computation with high threshold in two dimensions,” *Physical Review Letters*, **98**, 190504 (2007).
- [26] Emanuel Knill, “Fault-tolerant postselected quantum computation: Schemes,” [arXiv preprint quant-ph/0402171 \(2004\)](#).
- [27] Sergey Bravyi and Alexei Kitaev, “Universal quantum computation with ideal clifford gates and noisy ancillas,” *Physical Review A*, **71**, 022316 (2005).
- [28] Cody Jones, “Multilevel distillation of magic states for quantum computing,” *Physical Review A*, **87**, 042305 (2013).
- [29] Jeongwan Haah, Matthew B. Hastings, D. Poulin, and D. Wecker, “Magic state distillation with low space overhead and optimal asymptotic input count,” *Quantum*, **1**, 31 (2017).
- [30] Matthew B. Hastings and Jeongwan Haah, “Distillation with sublogarithmic overhead,” [arXiv preprint arXiv:1709.03543 \(2017\)](#).
- [31] Daniel Gottesman and Isaac L. Chuang, “Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations,” *Nature*, **402**, 390–393 (1999).
- [32] Christopher M. Dawson and Michael A. Nielsen, “The Solovay-Kitaev algorithm,” [Quantum Information & Computation](#), **6**, 81–95 (2006).
- [33] 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,” *Physical Review Letters*, **110**, 190502 (2013).
- [34] Neil J. Ross and Peter Selinger, “Optimal ancilla-free Clifford+T approximation of z-rotations,” [arXiv preprint arXiv:1403.2975 \(2014\)](#).
- [35] Alex Bocharov, Martin Roetteler, and Krysta M. Svore, “Efficient synthesis of universal repeat-until-success quantum circuits,” *Physical Review Letters*, **114**, 080502 (2015).
- [36] Alex Bocharov, Martin Roetteler, and Krysta M. Svore, “Efficient synthesis of probabilistic quantum circuits with fallback,” *Physical Review A*, **91**, 052317 (2015).
- [37] Alexei Yu Kitaev, Alexander Shen, and Mikhail N Vyalyi, *Classical and quantum computation*, Vol. 47 (Ameri-

- can Mathematical Society Providence, 2002).
- [38] Dave Wecker, Bela Bauer, Bryan K. Clark, Matthew B. Hastings, and Matthias Troyer, “Gate-count estimates for performing quantum chemistry on small quantum computers,” *Physical Review A*, **90**, 022305 (2014).
- [39] David Poulin, Matthew B. Hastings, Dave Wecker, Nathan Wiebe, Andrew C. Doherty, and Matthias Troyer, “The trotter step size required for accurate quantum simulation of quantum chemistry,” *Quantum Information and Computation*, **15**, 361–384 (2015).
- [40] Guang Hao Low and Isaac L. Chuang, “Hamiltonian simulation by uniform spectral amplification,” *arXiv preprint arXiv:1707.05391* (2017).