Hamiltonian simulation with nearly optimal dependence on all parameters

Dominic W. Berry* Andrew M. Childs† Robin Kothari‡

We present an algorithm for sparse Hamiltonian simulation that has optimal dependence on all parameters of interest (up to log factors). Previous algorithms had optimal or near-optimal scaling in some parameters at the cost of poor scaling in others. Hamiltonian simulation via quantum walk has optimal dependence on the sparsity d at the expense of poor scaling in the allowed error ϵ. In contrast, an approach based on fractional-query simulation provides optimal scaling in ϵ at the expense of poor scaling in d. Here we combine the two approaches, achieving the best features of both. By implementing a linear combination of quantum walk steps with coefficients given by Bessel functions, our algorithm achieves near-linear scaling in $\tau := d\|H\|_{\max}t$ and sublogarithmic scaling in $1/\epsilon$. Our dependence on ϵ is optimal, and we prove a new lower bound ruling out sublinear dependence on τ.

The problem of simulating the dynamics of quantum systems was the original motivation for quantum computers [14] and remains one of their major potential applications. The first explicit quantum simulation algorithm [16] gave a method for simulating Hamiltonians that are sums of local interaction terms. Aharonov and Ta-Shma gave an efficient simulation algorithm for the more general class of sparse Hamiltonians [2], and much subsequent work has given improved simulations [4, 5, 6, 7, 8, 12, 17, 18]. Sparse Hamiltonians include most physically realistic Hamiltonians as a special case. In addition, sparse Hamiltonian simulation can also be used to design other quantum algorithms [9, 10, 15]. For example, it was used to convert the algorithm for evaluating a balanced binary NAND tree with n leaves [13] to the discrete-query model [10]. A later algorithm obtained slightly improved scaling with a different approach [3].

In the Hamiltonian simulation problem, we are given an n-qubit Hamiltonian H (a Hermitian matrix), an evolution time t, and a precision $\epsilon > 0$, and are asked to implement the unitary operation e^{-iHt} up to error at most ϵ (as quantified by the diamond norm distance). We say that H is d-sparse if it has at most d nonzero entries in any row. In the sparse Hamiltonian simulation problem, H is specified by a black box that takes input $(j,\ell) \in [2^n] \times [d]$ (where $[d] := \{1,\dots,d\}$) and outputs the location and value of the ℓth nonzero entry in the jth row of H. Specifically, as in [5], we assume access to an oracle O_H acting as $O_H|j,k,z\rangle = |j,k,z \oplus H_{jk}\rangle$ for $j,k \in [2^n]$ and bit strings z representing entries of H, and another oracle O_F acting as $O_F|j,\ell\rangle = |j,f(j,\ell)\rangle$, where $f(j,\ell)\colon [2^n] \times [d] \to [2^n]$ is a function giving the row index of the ℓth nonzero element in column j. Note that the form of O_F assumes that the locations of the nonzero entries of H can be computed in place. This is possible if we can efficiently compute both $(j,\ell) \mapsto f(j,\ell)$ and the reverse map $(j,f(j,\ell)) \mapsto \ell$, which holds in typical applications of sparse Hamiltonian simulation. Alternatively, if f provides the nonzero elements in order, we can compute the reverse map with only a $\log d$ overhead by binary search.

At present, the two best algorithms for sparse Hamiltonian simulation are one based on a Szegedy quantum walk [5, 8] and another based on simulating an unconventional model of query complexity called the fractional-query model [6]. The quantum walk approach has query complexity $O(d\|H\|_{\max}t/\sqrt{\epsilon})$, which is linear in both the sparsity d and the evolution time t. (Here $\|H\|_{\max}$ denotes the largest entry of H in absolute value.) However, this approach has poor dependence on the allowed error ϵ. In contrast, the fractional-query approach has query complexity $O\big(\tilde{\tau}\frac{\log(\tilde{\tau}/\epsilon)}{\log\log(\tilde{\tau}/\epsilon)}\big)$, where $\tilde{\tau} := d^2\|H\|_{\max}t$. This approach gives exponentially

*Department of Physics and Astronomy, Macquarie University

†Department of Combinatorics & Optimization and Institute for Quantum Computing, University of Waterloo; Department of Computer Science, Institute for Advanced Computer Studies, and Joint Center for Quantum Information and Computer Science, University of Maryland

‡David R. Cheriton School of Computer Science and Institute for Quantum Computing, University of Waterloo; Center for Theoretical Physics, Massachusetts Institute of Technology

better dependence on the error at the expense of quadratically worse dependence on the sparsity and logarithmically worse dependence on the evolution time. Considering the fundamental importance of quantum simulation, it is desirable to achieve the best features of both approaches.

In this work, we combine the two approaches, giving the following.

Theorem 1. *A d-sparse Hamiltonian H acting on n qubits can be simulated for time t within error ϵ with*

$$O\left(\tau \frac{\log(\tau/\epsilon)}{\log\log(\tau/\epsilon)}\right) \tag{1}$$

queries and

$$O\left(\tau[n+\log(\tau/\epsilon)] \frac{\log(\tau/\epsilon)}{\log\log(\tau/\epsilon)}\right) \tag{2}$$

additional 2-qubit gates, where $\tau := d\|H\|_{\max} t$.

This result provides a strict improvement over the method of [6], removing a factor of d in τ, and thus providing near-linear instead of superquadratic dependence on d.

We also prove a lower bound showing that any algorithm must use $\Omega(\tau)$ queries. While a lower bound of $\Omega(t)$ was known previously [4], our new lower bound shows that the complexity must be at least linear in the product of the sparsity and the evolution time. Our proof is similar to a previous limitation on the ability of quantum computers to simulate non-sparse Hamiltonians [11]: by replacing each edge in the graph of the Hamiltonian by a complete bipartite graph $K_{d,d}$, we effectively boost the strength of the Hamiltonian by a factor of d at the cost of increasing the sparsity by a factor of d. Combining this result with the error-dependent lower bound of Ref. [6], we find a lower bound as follows.

Theorem 2. *For any $\epsilon, t > 0$, integer $d \geq 2$, and fixed value of $\|H\|_{\max}$, there exists a d-sparse Hamiltonian H such that simulating H for time t with precision ϵ requires $\Omega\big(\tau + \frac{\log(1/\epsilon)}{\log\log(1/\epsilon)}\big)$ queries.*

Thus our result is near-optimal in either τ or ϵ on its own. However, our upper bound (1) has a product, whereas the lower bound (2) has a sum. It remains unclear how to close the gap between these bounds. Intriguingly, a slight modification of our technique gives another algorithm with the following complexity.

Theorem 3. *For any $\alpha \in (0, 1]$, a d-sparse Hamiltonian H acting on n qubits can be simulated for time t within error ϵ with $O\big(\tau^{1+\alpha/2} + \tau^{1-\alpha/2}\log(1/\epsilon)\big)$ queries.*

This result provides a nontrivial tradeoff between the parameters t, d, and ϵ, and suggests that further improvements to such tradeoffs may be possible.

Overview of algorithms. Our algorithm uses a Szegedy quantum walk as in Refs. [5, 8], but with a linear combination of different numbers of steps. Such an operation can be implemented using techniques that were developed to simulate the fractional-query model [6]. This allows us to introduce a desired phase more accurately than with the phase estimation approach of [5, 8]. As in [6], we first implement the approximated evolution for some time interval with some amplitude and then use oblivious amplitude amplification to make the implementation deterministic, facilitating simulations for longer times.

References [5, 8] define a quantum walk step U that depends on the Hamiltonian H to be simulated. This operation can be implemented using a state preparation procedure that only requires one call to the sparse Hamiltonian oracle, avoiding the need to decompose H into terms as in product-formula approaches. The walk takes place in an enlarged state space that includes two copies of the space acted on by H. Each eigenstate of H is mapped onto a superposition of two eigenstates $|\mu_\pm\rangle$ of the quantum walk step U.

The relevant eigenvalues $\mu_\pm$ of U are related to the eigenvalues λ of H via $\mu_\pm = \pm e^{\pm i \arcsin(\lambda/Xd)}$, where $X \geq \|H\|_{\max}$ is a parameter that can be increased to create a lazy quantum walk. For small λ/Xd, the steps of the quantum walk for either μ_+ or μ_- yield a phase factor that is nearly proportional to that for the Hamiltonian evolution. However, the phase deviates from the desired value since the function $\arcsin(x)$ is not precisely linear about $x = 0$.

An obvious way to increase the accuracy is to increase X above its minimum value of $\|H\|_{\max}$. However, the number of steps of the quantum walk is $O(tXd)$, so increasing X results in a less efficient simulation. Another approach is to use phase estimation to correct the phase factor [5, 8], but this approach still gives polynomial dependence on $1/\epsilon$.

Instead, we propose using a superposition of steps of the quantum walk to effectively linearize the arcsin function. Specifically, rather than applying U, we apply $V_k := \sum_{m=-k}^{k} a_m U^m$ for some coefficients $a_{-k}, \ldots, a_k$. We show that the coefficients can be chosen by considering the generating function for the Bessel function [1, 9.1.41],

$$\sum_{m=-\infty}^{\infty} J_m(z)\mu_\pm^m = \exp\left[\frac{z}{2}\left(\mu_\pm - \frac{1}{\mu_\pm}\right)\right] = e^{i\lambda z/Xd}, \tag{3}$$

where the second equality follows from the relationship between λ and $\mu_\pm$ (note in particular that $\mu_+ - \frac{1}{\mu_+} = \mu_- - \frac{1}{\mu_-}$). Thus the ability to perform the operation $\sum_{m=-\infty}^{\infty} J_m(-x)U^m$ would allow us to exactly implement the evolution under H for time x/Xd. Remarkably, the coefficients required in the expansion are the same for both μ_+ and μ_-, so there is no need to distinguish the spaces. By truncating the sum to some finite range $\{-k, \ldots, k\}$, we obtain an expression in which each term can be performed using at most k queries. Because the Bessel function falls off exponentially for large $|m|$, we can obtain error at most ϵ with a cutoff k that is only logarithmic in $1/\epsilon$.

To perform V_k, we prepare ancilla qubits in a superposition state $|\chi_k\rangle$, then perform a controlled operation $\sum_{m=-k}^{k} |m\rangle\langle m| \otimes U^m$. Following a projective measurement that gives result $|\chi_k\rangle$, the operation is V_k with success probability approximately $1/a^2$, where $a := \sum_{m=-k}^{k} |a_m|$. Rather than using V_k to implement evolution over the entire time, we break the time up into "segments" and use V_k to achieve the time evolution for each segment. The probability of success may be boosted to 1 using the oblivious amplitude amplification procedure introduced in Ref. [6]. Such an exact implementation enables many segments to be performed consecutively with complexity linear in the number of segments.

The complexity of our algorithm is the number of segments (tXd/x) times the complexity for each segment (k) times the number of steps needed for oblivious amplitude amplification (a). We have some freedom in choosing x, which controls the amount of evolution time simulated by each segment. To obtain near-linear dependence on the evolution time t, we choose $x = O(1)$. Then amplitude amplification requires $O(1)$ steps, and the number of segments needed is $O(\tau)$, giving the linear factor in (1). The value of k needed to achieve overall error at most ϵ is logarithmic in τ/ϵ, yielding the logarithmic factor in (1).

An alternative approach is to use a larger segment that scales with τ. Choosing $x = \tau^\alpha$ for $\alpha \in (0, 1]$, we need $k = O(\tau^\alpha + \log(1/\epsilon))$. Then we require $O(\tau^{1-\alpha})$ segments and $O(\tau^{\alpha/2})$ steps of amplitude amplification, giving the scaling presented in Theorem 3.

Discussion. Our technique for Hamiltonian simulation combines ideas from quantum walks and fractional-query simulation to provide improved performance over both previous techniques. As a result, it provides near-optimal scaling in all parameters of interest. In particular, the scaling is only slightly superlinear in $\tau = d\|H\|_{\max}t$, whereas we have proven that linear scaling is optimal. Furthermore, the method has sublogarithmic scaling in the allowed error, which was proven to be optimal in Ref. [6].

Nevertheless, there is still a gap between the complexity of our algorithm and the lower bound in (2), as they involve different tradeoffs between τ and ϵ. It remains open whether the performance can be further improved, perhaps to give performance similar to (2), although we can rule out scaling strictly as in (2).

Our technique can potentially be used for the more general task of operation conversion, in which we use one quantum operation to implement another. In our work, we convert a step of a quantum walk to Hamiltonian evolution, whereas in Ref. [15] the task is to convert Hamiltonian evolution to an inverse. One approach to operation conversion is to use phase estimation. Here we have shown that a superposition of operations can provide far better performance.

References.

[1] Milton Abramowitz and Irene A. Stegun, *Handbook of mathematical functions*, National Bureau of Standards, 1964.

[2] Dorit Aharonov and Amnon Ta-Shma, *Adiabatic quantum state generation and statistical zero knowledge*, Proceedings of the 35th ACM Symposium on Theory of Computing, pp. 20–29, 2003, arXiv:quant-ph/0301023.

[3] Andris Ambainis, Andrew M. Childs, Ben W. Reichardt, Robert Špalek, and Shengyu Zhang, *Any AND-OR formula of size N can be evaluated in time $N^{1/2+o(1)}$ on a quantum computer*, SIAM Journal on Computing **39** (2010), no. 6, 2513–2530, arXiv:quant-ph/0703015 and arXiv:0704.3628.

[4] Dominic W. Berry, Graeme Ahokas, Richard Cleve, and Barry C. Sanders, *Efficient quantum algorithms for simulating sparse Hamiltonians*, Communications in Mathematical Physics **270** (2007), no. 2, 359–371, arXiv:quant-ph/0508139.

[5] Dominic W. Berry and Andrew M. Childs, *Black-box Hamiltonian simulation and unitary implementation*, Quantum Information and Computation **12** (2012), no. 1–2, 29–62, arXiv:0910.4157.

[6] Dominic W. Berry, Andrew M. Childs, Richard Cleve, Robin Kothari, and Rolando D. Somma, *Exponential improvement in precision for simulating sparse Hamiltonians*, Proceedings of the 46th ACM Symposium on Theory of Computing, pp. 283–292, 2014, arXiv:1312.1414.

[7] Andrew M. Childs, *Quantum information processing in continuous time*, Ph.D. thesis, Massachusetts Institute of Technology, 2004.

[8] _____, *On the relationship between continuous- and discrete-time quantum walk*, Communications in Mathematical Physics **294** (2010), no. 2, 581–603, arXiv:0810.0312.

[9] Andrew M. Childs, Richard Cleve, Enrico Deotto, Edward Farhi, Sam Gutmann, and Daniel A. Spielman, *Exponential algorithmic speedup by quantum walk*, Proceedings of the 35th ACM Symposium on Theory of Computing, pp. 59–68, 2003, arXiv:quant-ph/0209131.

[10] Andrew M. Childs, Richard Cleve, Stephen P. Jordan, and David Yonge-Mallo, *Discrete-query quantum algorithm for NAND trees*, Theory of Computing **5** (2009), no. 5, 119–123, arXiv:quant-ph/0702160.

[11] Andrew M. Childs and Robin Kothari, *Limitations on the simulation of non-sparse Hamiltonians*, Quantum Information and Computation **10** (2010), no. 7-8, 669–684, arXiv:0908.4398.

[12] Andrew M. Childs and Nathan Wiebe, *Hamiltonian simulation using linear combinations of unitary operations*, Quantum Information and Computation **12** (2012), no. 11-12, 901–924, arXiv:1202.5822.

[13] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann, *A quantum algorithm for the Hamiltonian NAND tree*, Theory of Computing **4** (2008), no. 8, 169–190, arXiv:quant-ph/0702144.

[14] Richard P. Feynman, *Simulating physics with computers*, International Journal of Theoretical Physics **21** (1982), no. 6–7, 467–488.

[15] Aram W. Harrow, Avinatan Hassidim, and Seth Lloyd, *Quantum algorithm for linear systems of equations*, Physical Review Letters **103** (2009), no. 15, 150502, arXiv:0811.3171.

[16] Seth Lloyd, *Universal quantum simulators*, Science **273** (1996), no. 5278, 1073–1078.

[17] David Poulin, Angie Qarry, Rolando D. Somma, and Frank Verstraete, *Quantum simulation of time-dependent Hamiltonians and the convenient illusion of Hilbert space*, Physical Review Letters **106** (2011), no. 17, 170501, arXiv:1102.1360.

[18] Nathan Wiebe, Dominic W. Berry, Peter Høyer, and Barry C. Sanders, *Simulating quantum dynamics on a quantum computer*, Journal of Physics A **44** (2011), no. 44, 445308, arXiv:1011.3489.