

Optimal Hamiltonian Simulation by Quantum Signal Processing

Guang Hao Low and Isaac L. Chuang

Department of Physics, Center for Ultracold Atoms, and Research Laboratory of Electronics Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA

(Received 8 June 2016; revised manuscript received 13 November 2016; published 5 January 2017)

The physics of quantum mechanics is the inspiration for, and underlies, quantum computation. As such, one expects physical intuition to be highly influential in the understanding and design of many quantum algorithms, particularly simulation of physical systems. Surprisingly, this has been challenging, with current Hamiltonian simulation algorithms remaining abstract and often the result of sophisticated but unintuitive constructions. We contend that physical intuition can lead to optimal simulation methods by showing that a focus on simple single-qubit rotations elegantly furnishes an optimal algorithm for Hamiltonian simulation, a universal problem that encapsulates all the power of quantum computation. Specifically, we show that the query complexity of implementing time evolution by a d -sparse Hamiltonian \hat{H} for time-interval t with error ϵ is $\mathcal{O}[td\|\hat{H}\|_{\max} + \log(1/\epsilon)/\log\log(1/\epsilon)]$, which matches lower bounds in all parameters. This connection is made through general three-step “quantum signal processing” methodology, comprised of (i) transducing eigenvalues of \hat{H} into a single ancilla qubit, (ii) transforming these eigenvalues through an optimal-length sequence of single-qubit rotations, and (iii) projecting this ancilla with near unity success probability.

DOI: 10.1103/PhysRevLett.118.010501

Introduction.—According to Feynman, “if you want to make a simulation of nature, you’d better make it quantum mechanical, and by golly it’s a wonderful problem, because it doesn’t look so easy.” [1]. Quantum computers are based on the physics of quantum mechanics, a fundamental tenant of nature as we know it. Thus, it seems natural to expect that the design and interpretation of quantum algorithms be heavily driven by physical intuition. The adiabatic algorithm [2,3] inspired by adiabaticity, and quantum walks [4,5] inspired by locality, are prominent examples. However, many quantum algorithms, most surprisingly those for the simulation of physical systems [6], are not as similarly insightful, and successive improvements in their complexity and analysis trend towards increasing abstraction and mathematical sophistication.

Analogous to physical theories, good quantum algorithms for physics simulations should, beyond being correct, also ideally be simple. In seeking simplicity, not only is their implementation on physical machines eased, but so too could their performance and understanding be enhanced. As the essence of coherent quantum computation is the design of unitary operations with desired properties, this motivates consideration of its closest analogue in experiments: physical quantum control [7], which has a similar goal of designing quantum *response functions* [8].

This hints at a deep connection between the design of optimal quantum algorithms and the synthesis of optimal quantum control policies. While robust time-optimal control [9,10] is certainly an established tool in quantum computing, its role is often secondary to the ends: the synthesis of computing primitives, such as Clifford

gates or even the quantum Fourier transform. It would be more desirable if physical dynamics were directly applicable to generic quantum algorithms without this intermediary. Indeed, the fact that intuition of the simplest quantum control—discrete single-qubit rotations $\hat{R}_\phi(\theta) = e^{-i(\theta/2)(\hat{\sigma}_x \cos \phi + \hat{\sigma}_y \sin \phi)}$ —can extend to algorithms such as Grover search supports this notion.

This relationship is made concrete by interpreting discrete sequences of physical operations as *programs* that compute functions. In the simplest setting, chaining N identical rotations generates $\langle 1 | \hat{R}_{\pi/2}^N(\theta) | 0 \rangle = \sin(N\theta/2)$. With θ as the input, this computes the function $f(\theta) = \sin(N\theta/2)$, which may be estimated through measurement. As Pauli matrices $\hat{\sigma}_{x,y,z}$ form a complete basis for 2-by-2 matrices, generic sequences of the form

$$\begin{aligned} \hat{V}(\theta) &= \hat{R}_{\phi_N}(\theta) \hat{R}_{\phi_{N-1}}(\theta) \cdots \hat{R}_{\phi_1}(\theta), \quad \vec{\phi} \in \mathbb{R}^N, \\ &= A(\theta) \hat{1} + iB(\theta) \hat{\sigma}_z + iC(\theta) \hat{\sigma}_x + iD(\theta) \hat{\sigma}_y, \end{aligned} \quad (1)$$

which we fully characterized in [8], then compute more general functions of θ in the real A, B, C, D , through a program specified by some choice of phases $\vec{\phi}$. In fact, the isomorphism of single-qubit rotations to those on a sphere (up to a double covering), furnishes an intuitive *classical* interpretation for this simple model of quantum computation. Moreover, the quantum control in Eq. (1), being piecewise, is naturally compatible with the inherently discrete nature of fault-tolerant architectures.

Though physically appealing, the computational merit of directly exploiting the structure of single-qubit rotations, or any physical system, ultimately rests on two criteria:

(i) usefulness in solving important problems, and (ii) optimality in space and time resources. It is also particularly challenging to see how this approach could apply generally to the complex multiqubit dynamics arising in the simulation of quantum systems.

The simulation problem maps one set of physical dynamics of interest—described by Hamiltonian \hat{H} —to another physical system that can be precisely controlled. Thus one expects the role of physics to be preeminent and obvious. Following seminal work by Lloyd [11] for Hamiltonians with local interactions, and Aharonov and Ta-Shma [12] for more general sparse Hamiltonians, many celebrated results have been obtained over the years [13–19] for approximating the time evolution operator $e^{-i\hat{H}t}$ for time-interval t with error ϵ . Encouragingly, intuitive quantum walks are already part of the state of art. More, however, could be hoped from their other components.

The complexity of such quantum algorithms is usually judged by the number of queries made to a unitary quantum oracle \hat{O} that provides a description of \hat{H} . Many interesting physical systems are described by the especially well-studied model of d -sparse \hat{H} with at most d nonzero elements in every row, and the best known algorithms [18] are based on the Childs quantum walk [14], which builds upon the Szegedy walk [20], that simulates time evolution by $\arcsin(\hat{H})$, which must be linearized. The difficulty lies in finding a quantum circuit that does this with the fewest queries to \hat{O} and the fewest number of additional primitive quantum gates.

Lower bounds on the query cost are well known. The “no-fast-forwarding” theorem [13,18] demands at least $\Omega(\tau)$ queries independent of ϵ , where $\tau = td\|\hat{H}\|_{\max}$ and $\|\hat{H}\|_{\max}$ is the largest element of \hat{H} in absolute value, and impressive recent work [17,18] proved an exact error scaling of $\Theta[\log(1/\epsilon)/\log\log(1/\epsilon)]$ for $\tau = \mathcal{O}(1)$. Though this suggests a naive additive lower bound $\Omega[\tau + \log(1/\epsilon)/\log\log(1/\epsilon)]$ [18], the best algorithms to date approach these factors multiplicatively with either linear scaling in time $\mathcal{O}(\tau/\sqrt{\epsilon})$ [14] or sublogarithmic scaling in error $\mathcal{O}[\tau \log(\tau/\epsilon)/\log\log(\tau/\epsilon)]$ [17,21]. Long unanswered is the existence of an algorithm that is *additively* optimal, with implications for the relation between continuous and discrete-time models of physics, and of interest in problems [22] where τ, ϵ scale together.

We achieve precisely this with a simple algorithm that matches the additive lower bound. In fact, it also realizes the optimal trade-off between time and error; thus, no further improvement in query complexity for this formulation of Hamiltonian simulation is possible. Compared to prior art [14,17], this represents up to a square-root improvement. Moreover, the space overhead in ancilla qubits, beyond those required for the quantum walk, is reduced from scaling with some function of τ/ϵ to just 1.

Most remarkably, this is achieved by finding a class of computational problems addressed by the optimal control of the single-qubit in Eq. (1) in a very natural way. Given a

unitary \hat{W} with eigenstates $\hat{W}|u_\lambda\rangle = e^{i\theta_\lambda}|u_\lambda\rangle$, we consider the general problem of constructing a quantum circuit \hat{V}_{ideal} with transformed eigenphases

$$\hat{W} \mapsto \hat{V}_{\text{ideal}} = \sum_\lambda e^{ih(\theta_\lambda)}|u_\lambda\rangle\langle u_\lambda|, \quad (2)$$

using the fewest queries to controlled \hat{W} for any real function $h(\theta)$. We call our solution to this “quantum signal processing” (Fig. 1), and its application to d -sparse Hamiltonian simulation leads to tremendous simplification and the claimed improvements. Our success here elevates optimal discrete quantum control in general as a tool that can be rigorous and essential in the design of optimal quantum algorithms, thus providing a medium through which physical intuition may flow.

Two key properties distinguish quantum signal processing from routines that can effect similar transformations, such as quantum phase estimation [23] or linear-combination-of-unitaries [16,18,19] which require a large number of ancilla. First, is its intuitive use of just a single ancilla qubit. Second, the query complexity of the methodology is exactly the degree N of optimal trigonometric polynomial approximations to $e^{ih(\theta)}$ with error ϵ [24–27], without the decaying success probability of prior art. Analogous to digital filter design techniques in discrete-time signal processing [28], this also elegantly bridges the design of a number of quantum algorithms to the vast field of function approximation [26].

In the following, we describe the reduction of quantum signal processing to optimal quantum control, and show how to efficiently choose the phases $\vec{\phi}$ Eq. (1) such that any unitary transformation Eq. (2) is approximated with error ϵ and success probability $1 - \mathcal{O}(\epsilon)$. The essential features of the quantum walk are then reviewed to show how quantum




FIG. 1. Quantum circuits mapping (a) a sequence of single-qubit rotations $\hat{V}(\theta)$ to (d) quantum signal processing \hat{V} . Each single-qubit rotation $\hat{R}_\phi(\theta)$ is replaced by (b) \hat{U}_ϕ , built from Hadamard gates and controlled- W with eigenstates $\hat{W}|u_\lambda\rangle = e^{i\theta_\lambda}|u_\lambda\rangle$. Thus (c), \hat{U}_ϕ on input $|u_\lambda\rangle$ reduces to a single-qubit rotation $\hat{R}_\phi(\theta_\lambda)$. By linearity, \hat{V} on an arbitrary input $|\psi\rangle$ may be understood as rotations $\hat{V}(\theta_\lambda)$ controlled by a superposition of $|u_\lambda\rangle$. By some choice of single-qubit input state and measurement basis, coefficients of the $|u_\lambda\rangle$ are then rescaled by the components of the function $\hat{V}(\theta_\lambda)$ programmed by $\vec{\phi}$.

signal processing for the special case of $h(\theta) = -\tau \sin(\theta)$ solves the sparse Hamiltonian simulation problem. That this achieves lower bounds follows by analyzing the scaling between N, τ, ϵ of quantum signal processing for this $h(\theta)$.

Quantum signal processing.—All quantum algorithms require a rigorous analysis of their resource costs in space and time. Thus any form of quantum control repurposed to such ends must have a similarly rigorous characterization. Previously [8], we studied the optimal control of arbitrary sequences of single-qubit rotations in Eq. (1), provided an intuitive characterization of the functions A, B, C, D achievable by some choice of $\vec{\phi}$, and provided efficient algorithms for synthesizing all these functions and the required $\vec{\phi}$ from some partial specification. The results relevant here are

Theorem 1.—[Achievable (A, C)] \forall even $N > 0$, a choice of real functions $A(\theta), C(\theta)$ can be implemented by some $\vec{\phi} \in \mathbb{R}^N$, if and only if, all these are true:

- (i) $\forall \theta \in \mathbb{R}, A^2(\theta) + C^2(\theta) \leq 1$.
- (ii) $A(0) = 1$.
- (iii) $A(\theta) = \sum_{k=0}^{N/2} a_k \cos(k\theta), \{a_k\} \in \mathbb{R}^{N/2+1}$.
- (iv) $C(\theta) = \sum_{k=1}^{N/2} c_k \sin(k\theta), \{c_k\} \in \mathbb{R}^{N/2}$.

Moreover, $\vec{\phi}$ can be efficiently computed from $A(\theta), C(\theta)$.

Note that in Ref. [8], A, C are expressed as trigonometric polynomials, but can be rewritten as a Fourier series (iii),(iv) using Chebyshev polynomials of the first and second kind $T_k[\cos(\theta)] = \cos(k\theta)$ and $U_k(\cos \theta) = \sin[(k+1)\theta]/\sin \theta$.

We now map these results, in three steps, to quantum signal processing, which transforms an arbitrary input unitary $\hat{W} = \sum_{\lambda} e^{i\theta_{\lambda}} |u_{\lambda}\rangle\langle u_{\lambda}|$ into one with modified eigenphases $\hat{V}_{\text{ideal}} = \sum_{\lambda} e^{ih(\theta_{\lambda})} |u_{\lambda}\rangle\langle u_{\lambda}|$:

- (a) *Signal transduction* of \hat{W} into a signal unitary classically controlled by $\phi \in \mathbb{R}$:

$$\hat{U}_{\phi} = \sum_{\lambda} \hat{R}_{\phi}(\theta_{\lambda}) \otimes |u_{\lambda}\rangle\langle u_{\lambda}|. \quad (3)$$

This is implemented in Fig. 1(b) with one controlled \hat{W} , which is always possible on a quantum computer in the worst case by replacing all of its gates with a controlled version, and $\mathcal{O}(1)$ single-qubit rotations:

$$\begin{aligned} \hat{U}_{\phi} &= (e^{-i\phi\hat{\sigma}_z/2} \otimes \hat{1}) \hat{U}_0 (e^{i\phi\hat{\sigma}_z/2} \otimes \hat{1}), \\ \hat{U}_0 &= |+\rangle\langle +| \otimes \hat{1} + |-\rangle\langle -| \otimes \hat{W} \\ &= \sum_{\lambda} e^{i\theta_{\lambda}/2} \hat{R}_0(\theta_{\lambda}) \otimes |u_{\lambda}\rangle\langle u_{\lambda}|, \end{aligned} \quad (4)$$

where $|\pm\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$. As \hat{U}_{ϕ} acting on $|u_{\lambda}\rangle$ selects the rotation $\hat{R}_{\phi}(\theta_{\lambda}) = \langle u_{\lambda} | \hat{U}_{\phi} | u_{\lambda} \rangle$ as seen in Fig. 1(c), these are precisely the single-qubit ancilla rotations in Eq. (1) with rotation angle θ_{λ} controlled by the λ index, but with an additional global phase $e^{i\theta_{\lambda}/2}$.

- (b) *Signal transformation* by computing unitary functions $\hat{V}(\theta_{\lambda})$ over a superposition of θ_{λ} on the

single-qubit ancilla through the simple circuit of Fig. 1(d):

$$\hat{V} = \hat{U}_{\phi_N} \hat{U}_{\phi_{N-1}} \cdots \hat{U}_{\phi_1}, \quad \vec{\phi} \in \mathbb{R}^N, \quad N \text{ even}. \quad (5)$$

As this invokes \hat{W} a number N times, its query cost is $\mathcal{O}(N)$. Note that the unwanted phase $e^{i\theta_{\lambda}/2}$ can be uncomputed by alternating between \hat{U}_{ϕ} and $\hat{U}_{\phi+\pi}^{\dagger}$, since $\hat{R}_{\phi}(\theta) = \hat{R}_{\phi+\pi}^{\dagger}(\theta)$ and N is even.

- (c) *Signal projection* of the ancilla onto some basis, to select desired components of $\hat{V}(\theta_{\lambda})$ in Eq. (1). As the desired phase transformation can be implemented through $A(\theta), C(\theta)$, consider the input state $|+\rangle|u_{\lambda}\rangle$, and postselect on measuring $\langle + |$. Other choices are, of course, possible. This applies onto state $|u_{\lambda}\rangle$ the coefficient

$$\begin{aligned} \langle + | \hat{V} | + \rangle | u_{\lambda} \rangle &= [A(\theta_{\lambda}) + iC(\theta_{\lambda})] | u_{\lambda} \rangle, \\ p &= \min_{\theta \in \mathbb{R}} |\langle + | \hat{V}(\theta) | + \rangle|^2 \\ &= \min_{\theta \in \mathbb{R}} |A(\theta) + iC(\theta)|^2, \end{aligned} \quad (6)$$

with worst-case success probability p . Thus, (a)–(c) provide a reduction from finding quantum algorithms for approximating \hat{V}_{ideal} to finding Fourier approximations of $A(\theta) + iC(\theta)$ to $e^{ih(\theta)}$.

By applying Theorem 1 to these three steps of quantum signal processing, we now prove following theorem which furnishes the complexity of implementing \hat{V}_{ideal} given this Fourier approximation:

Theorem 2.—(Quantum signal processing) \forall real odd periodic functions $h: (-\pi, \pi) \rightarrow (-\pi, \pi)$ and even $N > 0$, let $[A(\theta), C(\theta)]$ be real Fourier series in $[\cos(k\theta), \sin(k\theta)], k = 0, \dots, N/2$, that approximate

$$\max_{\theta \in \mathbb{R}} |A(\theta) + iC(\theta) - e^{ih(\theta)}| \leq \epsilon. \quad (7)$$

Given $A(\theta), C(\theta)$, one can efficiently compute the $\vec{\phi}$ such that $\langle + | \hat{V} | + \rangle$ in Eq. (5) applies \hat{U}_{ϕ} a number N times to approximate \hat{V}_{ideal} in Eq. (2) with success probability $p \geq 1 - 16\epsilon$ and trace distance

$$\epsilon_{\text{Tr}} = \max_{|\psi\rangle} \|(\langle + | \hat{V} | + \rangle - \hat{V}_{\text{ideal}}) |\psi\rangle\| \leq 8\epsilon. \quad (8)$$

Note that the restricted symmetry of $h(\theta)$ is for consistency with the parity and periodicity of $A(\theta), C(\theta)$.

Given $A(\theta), C(\theta)$ that satisfy Eq. (7), conditions (i), (ii) of Theorem 1 will not generally be satisfied. (i) is violated as $\max_{\theta} (|A(\theta)|, |C(\theta)|) \leq 1 + \epsilon$. Thus, we rescale

$$\begin{aligned} A_1(\theta) &= A(\theta)/(1 + \epsilon), \quad C_1(\theta) = C(\theta)/(1 + \epsilon), \\ |[A_1(\theta) + iC_1(\theta)] - e^{ih(\theta)}| &\leq \epsilon/(1 + \epsilon) + \epsilon \leq 2\epsilon, \end{aligned} \quad (9)$$

at the cost of a slightly larger error 2ϵ . Note that $A_1^2 + C_1^2 \geq [(1 - \epsilon)/(1 + \epsilon)]^2$. (ii) is violated as $A_1(0) = \cos \delta \geq (1 - \epsilon)/(1 + \epsilon)$ for some $\delta \in \mathbb{R}$. Fixing this is

more involved. As $\hat{V}(\theta)$ is unitary, $A_1^2 + B^2 + C_1^2 + D^2 = 1$. We can apply the prescription in Ref. [8] using polynomial sum of squares to compute the unspecified B, D from A_1, C_1 , such that B, D are of the form (iii) and (iv), respectively. Thus, $A_1^2(0) + B^2(0) = 1$, and $|B(0)| = |\sin \delta|$. Define

$$A_2(\theta) = A_1(\theta) \cos \delta + B(\theta) \sin \delta,$$

$$|A_2(\theta) - A_1(\theta)| \leq \frac{2\epsilon}{1+\epsilon} + |B(\theta)| \frac{2\sqrt{\epsilon}}{1+\epsilon} \leq 6\epsilon. \quad (10)$$

This introduces an additional error by using the triangle inequality and $B^2 \leq 1 - A_1^2 - C_1^2 \leq 1 - [(1-\epsilon)/(1+\epsilon)]^2 = 4\epsilon/(1+\epsilon)^2$. By construction, $A_2(0) = 1$. The functions $A_2(\theta), C_1(\theta)$ thus satisfy Theorem 1. By adding the errors in Eqs. (9), (10), the distance of $\langle +|\hat{V}|+ \rangle$ from \hat{V}_{ideal} in Eq. (2) and the worst-case success probability in Eq. (6) are

$$\epsilon_{\text{Tr}} \leq \max_{\theta \in \mathbb{R}} |A_2(\theta) + iC_1(\theta) - e^{ih(\theta)}| \leq 8\epsilon,$$

$$p \geq (1 - 8\epsilon)^2 \geq 1 - 16\epsilon. \quad (11)$$

The optimality of this remarkably simple procedure for eigenphase transformation follows from its role in obtaining the best possible quantum algorithm for d -sparse Hamiltonian simulation. We now highlight essential features of a quantum walk \hat{W} constructed from the oracles \hat{O} describing sparse \hat{H} .

Childs' quantum walk [14].— \hat{W} can be constructed from oracles that specify a d -sparse Hamiltonian \hat{H} with n -qubit eigenstates $\hat{H}|\lambda\rangle = \lambda|\lambda\rangle$. Access to two oracles \hat{O}_H, \hat{O}_F is commonly assumed: \hat{O}_H accepts the input $(j, k) \in [2^n] \times [2^n]$ on $2n$ -qubit registers and returns $H_{j,k} = \langle j|\hat{H}|k\rangle$ in another m -qubit register. \hat{O}_F accepts the input $(j, l) \in [2^n] \times [d]$ on the same $2n$ -qubit registers and computes in place the column index $f(j, l) \in [2^n]$ of the l th nonzero element in the j th row of \hat{H} .

It is well known [29] that with 1 query to $\hat{O}_F, \hat{O}_H, \hat{O}_H^\dagger$ each and $\mathcal{O}[n + m \text{poly}(\log m)]$ primitive gates, one can implement an isometry \hat{T} that maps every state $|\lambda\rangle|0\rangle^{\otimes n+m+2}$ onto two eigenstates $|\lambda_\pm\rangle$ of \hat{W} :

$$\hat{T}|\lambda\rangle = (|\lambda_+\rangle + |\lambda_-\rangle)/\sqrt{2}. \quad (12)$$

Moreover, \hat{T} is constructed such that the walk $\hat{W} = i\hat{S}(2\hat{T}\hat{T}^\dagger - \hat{1})$ has eigenvalues $\hat{W}|\lambda_\pm\rangle = e^{i\theta_{\lambda_\pm}}|\lambda_\pm\rangle$,

$$\theta_{\lambda_\pm} = \pm \arcsin(\lambda/\|\hat{H}\|_{\max} d) + (1 \mp 1)\pi/2, \quad (13)$$

that depends on the \hat{H} eigenvalues λ . As \hat{W} corresponds to reflection about $\hat{T}\hat{T}^\dagger$ followed by swapping $(2n+2)$ -qubit registers with \hat{S} , its query and gate complexities are identical to \hat{T} up to constant factors.

Hamiltonian simulation is achieved by creatively applying \hat{W} some number of times to implement $|\lambda_\pm\rangle \mapsto e^{-i\lambda t}|\lambda_\pm\rangle$, independent of the \pm index. Uncomputing with \hat{T}^\dagger then maps $|\lambda_\pm\rangle$ back onto $|\lambda\rangle|0\rangle^{\otimes n+m+2}$ with the desired phase evolution. However, some difficulties arise. First, the

applied phase θ_{λ_\pm} is nonlinear in λ . Second, each eigenstate $|\lambda_\pm\rangle$ evolves under \hat{W} with phases in opposite directions. Thus uncomputing with \hat{T}^\dagger does not map $\hat{W}\hat{T}|\lambda\rangle|0\rangle^{\otimes n+m+2}$ back onto the basis $|\lambda\rangle|0\rangle^{\otimes n+m+2}$. In Ref. [18], these are overcome by approximating the unitary transformation in Eq. (2) with target function

$$h(\theta) = -\tau \sin(\theta) \Rightarrow h(\theta_{\lambda_\pm}) = -\lambda t, \quad (14)$$

resulting in the desired phase, but implemented using a technique combining a linear combination of N -controlled $\hat{W}^{1, \dots, N}$ such that the success probability decays with N . Our quantum signal processing methodology does not experience such a decay and its direct application furnishes an optimal Hamiltonian simulation algorithm.

Hamiltonian simulation.—Applying quantum signal processing in Theorem 2 to Hamiltonian simulation requires a good Fourier approximation to

$$A(\theta) + iC(\theta) \approx e^{ih(\theta)} = e^{-i\tau \sin(\theta)}, \quad (15)$$

which is provided by the Jacobi-Anger expansion [30]

$$\cos[\tau \sin(\theta)] = J_0(\tau) + 2 \sum_{k \text{ even} > 0} J_k(\tau) \cos(k\theta),$$

$$\sin[\tau \sin(\theta)] = 2 \sum_{k \text{ odd} > 0} J_k(\tau) \sin(k\theta), \quad (16)$$

where $J_k(\tau)$ are Bessel functions of the first kind. Note that these Fourier series are already in the form required by conditions (iii), (iv) of Theorem 1. As $|J_k(\tau)| \leq |\tau/2|^{|k|}/|k|!$ [30] decays rapidly with k , good approximations are obtained truncating Eq. (16) at $k > N/2$. This approximates $e^{-i\tau \sin(\theta)}$ with error shown in Ref. [18] for $\tau \leq N/2 = q - 1$ to be

$$\epsilon \leq \sum_{k=q}^{\infty} 2|J_k(\tau)| \leq \frac{4\tau^q}{2^q q!} = \mathcal{O}\left[\left(\frac{e\tau}{2q}\right)^q\right]. \quad (17)$$

Inserting into Theorem 2, the query complexity of Hamiltonian simulation follows by solving Eq. (17) for N , using the implementation of \hat{U}_ϕ in Eq. (4) with $\mathcal{O}(1)$ queries, and that \hat{V} in Eq. (5) contains N applications of \hat{U}_ϕ .

The optimality of this result for all input parameters follows from known lower bounds. Specifically, Eq. (17) is matched with a corresponding lower bound $N = \Omega(q)$ [18,31] for any q satisfying

$$\epsilon < \frac{1}{2} \left| \sin\left(\frac{\tau}{q}\right) \right|^q = \mathcal{O}\left[\left(\frac{\tau}{q}\right)^q\right]. \quad (18)$$

Note that Eqs. (17), (18) are solved by the Lambert W function [32], which captures the detailed trade-off between τ and ϵ . Its asymptotic behavior may be understood by substituting $q = (e/2)(\tau + \gamma)$, where $\tau, \gamma \geq 0$. When $\tau = \mathcal{O}(\gamma)$, one finds $\gamma = \mathcal{O}[\log(1/\epsilon)/\log \log(1/\epsilon)]$. Thus, we express the complexity of Hamiltonian simulation as

Theorem 3.—(Optimal sparse Hamiltonian simulation). A d -sparse Hamiltonian \hat{H} on n qubits with matrix elements

specified to m bits of precision can be simulated for time-interval t , error ϵ , and success probability at least $1 - 2\epsilon$ with $\mathcal{O}[td\|\hat{H}\|_{\max} + \log(1/\epsilon)/\log\log(1/\epsilon)]$ queries and a factor $\mathcal{O}[n + m\text{polylog}(m)]$ additional quantum gates.

This is valid for $\tau = \mathcal{O}[\log(1/\epsilon)/\log\log(1/\epsilon)]$ and stronger than prior art [17,18], which assumes $\tau = \mathcal{O}(1)$. Unlike most Hamiltonian simulation algorithms, the query cost is additive in the simulation length τ and the target error ϵ . As such, the τ term matches the lower bound $\Omega(\tau)$ [13,18] with no multiplicative dependence on error.

Conclusion.—We have shown that optimal quantum algorithms for Hamiltonian simulation can be remarkably simple and physically motivated. Here, physical intuition flows into the process by directly using the dynamics of discrete single-qubit rotations as a computational module, which proves to be exceptionally useful when translated into quantum signal processing. Indeed, we have focused on choosing target functions (A, C) for even N in Theorem 1, but many other choices described in Ref. [8] are possible. For example, fixed-point amplitude amplification [33] and Heisenberg-limited quantum imaging [34] are special cases for the choice $\{A(\theta) \propto T_N[\beta \cos(\theta/2)], B(\theta) = 0\}, \beta > 1$.

Directly exploiting the structured dynamics in other physical systems could lead to powerful tools for a similarly intuitive approach to rigorous and optimal quantum algorithms. The question of what other important quantum algorithms can be designed or improved in this manner is an exciting natural extension to this work.

G. H. Low and I. L. Chuang thank Cedric Yen-Yu Lin, Robin Kothari, and Matthew Hastings for insightful discussions, and acknowledge funding by the ARO Quantum Algorithms Program, the NSF CUA, and NSF RQCC Project No. 1111337.

-
- [1] R. P. Feynman, *Int. J. Theor. Phys.* **21**, 467 (1982).
 - [2] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, [arXiv:quant-ph/0001106](https://arxiv.org/abs/quant-ph/0001106).
 - [3] A. Mizel, D. A. Lidar, and M. Mitchell, *Phys. Rev. Lett.* **99**, 070502 (2007).
 - [4] A. M. Childs, *Phys. Rev. Lett.* **102**, 180501 (2009).
 - [5] A. M. Childs, D. Gosset, and Z. Webb, *Science* **339**, 791 (2013).
 - [6] I. M. Georgescu, S. Ashhab, and F. Nori, *Rev. Mod. Phys.* **86**, 153 (2014).
 - [7] K. Khodjasteh, D. A. Lidar, and L. Viola, *Phys. Rev. Lett.* **104**, 090501 (2010).
 - [8] G. H. Low, T. J. Yoder, and I. L. Chuang, [arXiv:1603.03996](https://arxiv.org/abs/1603.03996) [*Phys. Rev. X* (to be published)].
 - [9] T. Caneva, M. Murphy, T. Calarco, R. Fazio, S. Montangero, V. Giovannetti, and G. E. Santoro, *Phys. Rev. Lett.* **103**, 240501 (2009).

- [10] K. Khodjasteh and L. Viola, *Phys. Rev. Lett.* **102**, 080501 (2009).
- [11] S. Lloyd, *Science* **273**, 1073 (1996).
- [12] D. Aharonov and A. Ta-Shma, in *Proceedings of the Thirty-fifth Annual ACM Symposium on Theory of Computing*, STOC '03 (ACM, New York, 2003), p. 20.
- [13] D. W. Berry, G. Ahokas, R. Cleve, and B. C. Sanders, *Commun. Math. Phys.* **270**, 359 (2007).
- [14] A. M. Childs, *Commun. Math. Phys.* **294**, 581 (2010).
- [15] A. M. Childs and R. Kothari, *Theory of Quantum Computation, Communication, and Cryptography* (Springer, Berlin, Heidelberg, 2011), p. 94.
- [16] A. M. Childs and N. Wiebe, *Quantum Inf. Comput.* **12**, 901 (2012).
- [17] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, in *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, STOC '14 (ACM, New York, NY, USA, 2014), p. 283.
- [18] D. W. Berry, A. M. Childs, and R. Kothari, in *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on* (2015), p. 792.
- [19] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, *Phys. Rev. Lett.* **114**, 090502 (2015).
- [20] M. Szegedy, [arXiv:quant-ph/0401053](https://arxiv.org/abs/quant-ph/0401053).
- [21] R. Cleve, D. Gottesman, M. Mosca, R. D. Somma, and D. Yonge-Mallo, in *Proceedings of the Forty-first Annual ACM Symposium on Theory of Computing*, STOC '09 (ACM, New York, NY, USA, 2009), p. 409.
- [22] G. J. Sussman and J. Wisdom, *Science* **241**, 433 (1988).
- [23] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, 1st ed. (Cambridge University Press, Cambridge, England, 2004).
- [24] J. McClellan, T. Parks, and L. Rabiner, *IEEE Trans. Audio Electroacoust.* **21**, 506 (1973).
- [25] R. Pachón and L. N. Trefethen, *Bit Numer Math* **49**, 721 (2009).
- [26] M. J. D. Powell, *Approximation Theory and Methods* (Cambridge University Press, Cambridge, England, 1981).
- [27] L. N. Trefethen, *Approximation Theory and Approximation Practice* (Siam, Philadelphia, 2013).
- [28] A. Oppenheim and R. Schaffer, *Discrete-Time Signal Processing, 3rd Ed.*, Prentice-Hall Signal Processing Series (Prentice Hall, Englewood Cliffs, NJ, 2010).
- [29] D. W. Berry and A. M. Childs, *Quantum Inf. Comput.* **12**, 29 (2012).
- [30] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*, National Bureau of Standards Applied Mathematics Series 55 (U.S. Government Printing Office, Washington, DC, 1964).
- [31] R. Kothari (private communication).
- [32] R. M. Corless, G. H. Gonnet, D. E. G. Hare, D. J. Jeffrey, and D. E. Knuth, *Adv. Comput. Math.* **5**, 329 (1996).
- [33] T. J. Yoder, G. H. Low, and I. L. Chuang, *Phys. Rev. Lett.* **113**, 210501 (2014).
- [34] G. H. Low, T. J. Yoder, and I. L. Chuang, *Phys. Rev. Lett.* **114**, 100801 (2015).