

A quantum–quantum Metropolis algorithm

Man-Hong Yung (翁文康) and Alán Aspuru-Guzik¹

Department of Chemistry and Chemical Biology, Harvard University, Cambridge MA, 02138

Edited by Frank A. Wilczek, Center for Theoretical Physics, Cambridge, MA, and approved November 16, 2011 (received for review July 20, 2011)

The classical Metropolis sampling method is a cornerstone of many statistical modeling applications that range from physics, chemistry, and biology to economics. This method is particularly suitable for sampling the thermal distributions of classical systems. The challenge of extending this method to the simulation of arbitrary quantum systems is that, in general, eigenstates of quantum Hamiltonians cannot be obtained efficiently with a classical computer. However, this challenge can be overcome by quantum computers. Here, we present a quantum algorithm which fully generalizes the classical Metropolis algorithm to the quantum domain. The meaning of quantum generalization is twofold: The proposed algorithm is not only applicable to both classical and quantum systems, but also offers a quantum speedup relative to the classical counterpart. Furthermore, unlike the classical method of quantum Monte Carlo, this quantum algorithm does not suffer from the negative-sign problem associated with fermionic systems. Applications of this algorithm include the study of low-temperature properties of quantum systems, such as the Hubbard model, and preparing the thermal states of sizable molecules to simulate, for example, chemical reactions at an arbitrary temperature.

Metropolis method | statistical physics | quantum computing | quantum simulation | simulated annealing

Interacting many-body problems cover a wide range of applications in many areas in science and technology. These are best exemplified by the Ising spin model, which was originally developed for explaining ferromagnetism and its associated phase transitions (1). The Ising model was later found to be related to many practical applications, for example, error-correcting codes, image restoration, associative memory, and optimization problems (see, e.g., ref. 2). However, there is no efficient method for finding solutions to many-body problems in general. For example, the problem of finding the ground state of the Ising model with arbitrary couplings is known to be a nondeterministic polynomial (NP)-complete problem, meaning that, if an efficient algorithm for solving the Ising model exists, then all of the problems in the class of NP, for example the graph isomorphism problem and some variants of the traveling salesman problem, can be solved efficiently as well (see, e.g., ref. 3).

The most fundamental challenge for solving many-body problems is that they generally require an exponentially large amount of spatial and temporal computing resources to find the solutions as the system size increases (4). Although no universal method for solving general many-body problems has been found, powerful classical methods such as Markov-chain Monte Carlo (MCMC) or quantum Monte Carlo (QMC) have been invented and proven to be successful in many applications (see, e.g., ref. 5). These methods, however, have certain limitations. For example, the running time of MCMC scales as $O(1/\delta)$ (6), where δ is the gap of the transition matrix. For problems such as spin glasses (7) where δ is vanishingly small as a function of system size, MCMC becomes computationally inefficient. QMC methods, on the other hand, suffer from the negative-sign problem (8), making them inefficient for problems involving fermions, e.g., electronic problems. Despite the many efforts toward improvement that have been made (9), this limitation is still one of the biggest challenges in QMC (10).

Background of Quantum Simulation

One of the most important goals in the field of quantum computation, as proposed by Feynman (11), is to look for methods or algorithms that can solve these many-body problems more efficiently. A promising solution is quantum simulation, which aims at employing a controllable quantum system to simulate the behaviors of the target quantum system. Quantum simulation can solve the problem of the spatial requirement for solving many-body problems. Quantum simulation can be implemented either by dedicated (or analog) quantum simulators (12) or with universal (or digital) quantum computers (13). For analog quantum simulation, high-precision experimental techniques are required for faithful simulation, and therefore involves many engineering challenges. An example is the use of neutral atoms trapped in optical lattices (14) to simulate the low-temperature properties of the Hubbard or Bose–Hubbard model.

Digital Quantum Simulation. To look for algorithms based on the special properties of the laws of quantum mechanics, top-down approaches aim to improve the existing classical algorithms by combining them with elementary quantum algorithms. For the bottom-up approach, in the context of quantum simulation, the only known example that quantum algorithms can achieve exponential gain over classical methods is the simulation of time dynamics (15, 16). It was also discovered (17) that such an approach is particularly suitable for simulating single or many-particle systems where the Hamiltonians consist of two terms, namely kinetic energy and potential energy, for example, atomic or molecular systems (18). On the other hand, however, attempt to create quantum algorithms to solve the ground-state (19, 20) or thermal-state (21–23) problems of unstructured Hamiltonians failed to show an exponential gain over classical approaches. However, this lack of an exponential advantage does not mean that quantum computers fail to show advantages over classical computers in quantum simulation.

Variational Methods. In fact, many problems in physics and chemistry do exhibit certain symmetries or structures. These features allow us to take the top-down approach to generalize methods that are proven to be successful in classic computing to design quantum algorithms. For example, the variational wavefunctions, such as the Bardeen–Cooper–Schrieffer wavefunction in superconductivity and the Hartree–Fock solutions in the electronic problems of atomic and molecular systems, are good approximations to the exact ground states of the corresponding Hamiltonians to some extent. With the quantum phase estimation algorithm (24, 25), it is possible to project the exact ground states from the approximate solutions with high probabilities. Further theoretical investigations on the efficiency of this method on molecular systems have been made (26–28), and a related experimental implementation has been realized (29).

Author contributions: M.-H.Y. and A.A.-G. designed research, performed research, and wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

¹To whom correspondence should be addressed. E-mail: aspuru@chemistry.harvard.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1111758109/-DCSupplemental.

Markov-Chain-Based Methods. For problems where clever approximating solutions are unavailable, one may still be able to solve them efficiently by using Markov-chain-based Monte Carlo methods. In the context of quantum computation, Szegedy described a method to quantize classical Markov chains (30). The key result is that a quadratic speedup $O(1/\sqrt{\delta})$ in the gap δ of the transition matrix is possible. This idea was later adapted to other algorithms (31, 32) for preparing thermal states of classical Hamiltonians, based on the idea of quantum simulated annealing (QSA) (31).

In fact, Markov chains that correspond to classical Hamiltonians are relatively easy to construct, as all the eigenstates of classical Hamiltonians are in the computational basis $|x\rangle$. However, preparation of the thermal states of quantum Hamiltonians (20, 21, 33, 34), where the eigenvectors are in nontrivial superposition of the computational basis vectors, is a much more challenging problem for classical computation. Classically, one would need to solve for the full eigenvalue problem (i.e., look for all eigenvalues and eigenvectors) for the quantum Hamiltonian first, which often requires more computational resources than directly computing the corresponding partition functions. Quantum mechanically, any attempt for a straightforward generalization of the classical method by Szegedy would encounter the problem due to the inability to clone quantum states (35).

Quantum Metropolis Methods. The key strategy that has been employed so far to overcome the difficulty of obtaining information about eigenvectors of quantum Hamiltonians with a quantum computer is the use of the quantum phase estimation algorithm (24); the eigenvalues of a quantum Hamiltonian can be recorded without explicitly knowing the detailed structure of the eigenvectors. Based on this idea, Terhal and DiVincenzo (33) made an attempt to extend the Metropolis algorithm (5) to the quantum domain in order to avoid the negative-sign problem in QMC. However, one limitation of their result is that the key step involves too many energy nonlocal transitions, making it inefficient for solving practical problems. Recently, Temme et al. (34) addressed this problem by introducing a measurement-based rejection method, which makes the algorithm more realistic for simulating many-body problems. This result shows that quantum computers can solve many-body problems beyond the limits of classical computers. However, the underlying Markov chains of this algorithm are still classical in nature, which means that the scaling of the running time is still $O(1/\delta)$.

Motivation, Results, and Outline of the Quantum–Quantum Metropolis Algorithm (Q²MA)

To summarize, before this work, the existing quantum algorithms were capable of either (i) preparing thermal states of classical systems with quantum speedup (31, 32), or (ii) preparing thermal states of quantum systems without quantum speedup (33, 34); this suggests that the potential of quantum computers may not have been fully exploited. This problem motivates us to further explore the power of quantum computation by providing a quantum algorithm that allows us to prepare thermal states of quantum systems with quantum speedup.

Instead of following the path developed in the work of Temme et al. (34), we took the following strategy: First of all, a natural question to ask is whether Szegedy’s quantum speedup is restricted to employing classical Hamiltonians, where the eigenstates simply coincide with the computational basis, making cloning of the classical information possible. In this manuscript, we extend the Szegedy approach to be able to deal with quantum Hamiltonians; we introduce a quantum version of the method of Markov-chain quantization which is combined with the QSA procedure (31). Formally, we propose a quantum Metropolis algorithm that exhibits a quadratic quantum speedup as a function of the eigenvalue gap δ of the corresponding Metropolis Markov matrix M for any quantum Hamiltonian H . Note that, because

classical Hamiltonians belong to a subclass of quantum Hamiltonians, our algorithm is also applicable to classical systems.

Lastly, we note that this algorithm is still limited by the Markov-chain gap δ , which implies that there are problems, especially ground-state problems of frustrated systems, where this algorithm cannot solve efficiently. This result is expected, as otherwise, hard problems (20) in the quantum Merlin–Arthur complexity class (and even NP) would be solved by this algorithm. Whether such a quantum algorithm could exist is still one of the major open questions in quantum information science.

Summary of the Main Results. The quantum Metropolis algorithm constructs the thermal state, $\rho_{\text{th}} = e^{-\beta H} / \text{Tr}(e^{-\beta H})$, for any given quantum (or classical) Hamiltonian H , where the system to be simulated is prepared at a given temperature T ($\beta \equiv 1/k_B T$), with a quadratic quantum speedup $O(1/\sqrt{\delta})$ in the gap δ of the corresponding Markov matrix M . We call this Q²MA because it shows a quantum speedup for the Markov chain of thermal-state preparation for a quantum system being simulated. For comparison, the construction method by Temme et al. (34) is restricted by the no-cloning theorem insofar as the information of an eigenstate cannot be retrieved after accepting the proposed move in the Metropolis step. We relax this restriction by adopting a dual representation where the basis states consist of pairs of eigenstates related by time-reversal operation (see Eq. 8). For the cases where the physical system being simulated is time-reversal invariant, the pair of the eigenstates forming the basis vector become identical. This construction effectively circumvents the restriction imposed by the no-cloning theorem.

Outline of the Key Ideas of Q²MA. The goal of the proposed Q²MA is to prepare the coherent encoding of the thermal state (CETS),

$$|\alpha_0\rangle = \sum_{i=0}^{N-1} \sqrt{e^{-\beta E_i} / \mathcal{Z}} |i\rangle, \quad [1]$$

where $|i\rangle \equiv |\varphi_i\rangle|\tilde{\varphi}_i\rangle|0\rangle$, the state $|\varphi_i\rangle$ is the eigenstate of a quantum Hamiltonian H , associated with the eigenvalue E_i , and $\mathcal{Z} = \text{Tr}(e^{-\beta H})$ is the partition function. The state $|\tilde{\varphi}_i\rangle$ is the time-reversal counterpart (complex conjugate) of $|\varphi_i\rangle$, which is the eigenstate of the corresponding time-reversal Hamiltonian $\tilde{H} \equiv H^*$ with the same eigenenergy E_i . After tracing out all of the ancilla quantum bits (qubits), the CETS is equivalent to $\rho_{\text{th}} = (1/\mathcal{Z}) \sum_i e^{-\beta E_i} |\varphi_i\rangle\langle\varphi_i|$. An outline of this algorithm is shown in Fig. 1. The steps go as follows: A maximally entangled state $N^{-1/2} \sum_{x=0}^{N-1} |x\rangle|x\rangle$ is the input state of the algorithm. Alternative input states are possible as long as they have a finite overlap with the CETS, but for simplicity we will focus on only one particular type of input state here. This maximally entangled state can be mathematically transformed into a state (see Eq. 8) that spans the dual basis related to the Hamiltonian H to be simulated. Note that this state represents the infinite temperature $\beta = 0$ state. Next, a generalization of Szegedy’s operator W (see Eq. 12) will be constructed and applied to the input state. The main purpose of W is to project out the CETS at a lower temperature through the quantum phase estimation algorithm. A similar procedure is then repeated for projecting out a CETS at an even lower temperature, and so on. The whole procedure for “adiabatically” preparing from a high-temperature CETS to a lower-temperature one is known as QSA (31). When the ancilla qubits are traced out, we obtain the thermal density matrix $\rho_{\text{th}} = e^{-\beta H} / \text{Tr}(e^{-\beta H})$, which accomplishes the goal of the Q²MA.

Organization of the Paper. We first summarize the features of Markov chains, Metropolis sampling, and Szegedy’s method that are related to this work, in a way that become applicable to both

Outline of the Q² Metropolis Algorithm

- (i) Prepare the maximally entangled state

$$\left(\frac{1}{\sqrt{N}}\right) \sum_x |x\rangle |x\rangle$$

- (ii) Change to the eigenbasis of the Hamiltonian

$$\left(\frac{1}{\sqrt{N}}\right) \sum_i |\varphi_i\rangle |\tilde{\varphi}_i\rangle$$

- (iii) Set $j=0$, then phase estimation with the Szegedy operator

$$W_{\beta_j} = (2\Lambda_2 - I)(2\Lambda_1 - I)$$

- (iv) After a projective measurement, we get the CETS

$$|\alpha_0^j\rangle = \sum_{i=0}^{N-1} (e^{-\beta_j E_i} / Z_j)^{1/2} |\varphi_i\rangle |\tilde{\varphi}_i\rangle$$

- (v) Trace out the ancilla qubits

$$\rho = e^{-\beta H} / \text{Tr}(e^{-\beta H})$$

QSA: $j \rightarrow j+1$

Fig. 1. Outline of the Q²MA for preparing thermal states of quantum Hamiltonians H . (i) A maximally entangled state, $(1/\sqrt{N})\sum_x |x\rangle |x\rangle$, of n qubits is first prepared. (ii) Then, it is mathematically transformed into a maximally entangled state, $(1/\sqrt{N})\sum_i |\varphi_i\rangle |\tilde{\varphi}_i\rangle$ (see Eq. 8), associated with the Hamiltonian H . (iii) A schedule of QSA (31) is then applied to this state, lowering the simulated inverse temperature from zero to β . The QSA algorithm is based on the phase estimation algorithm using the Szegedy operator W defined in Eq. 12. (iv) The resulting state is a CETS described in Eq. 1. (v) The final state becomes the thermal state $\rho = e^{-\beta H} / \text{Tr}(e^{-\beta H})$, after tracing out all of the ancilla qubits.

classical and quantum systems. Next, we outline the generalization of the method of Markov-chain quantization by Szegedy (30) to make it applicable to quantum Hamiltonians H . Then, we analyze the structure of the Szegedy operator W in detail. Finally, we combine the results above with the method of quantum simulated annealing (31) to generate thermal states of quantum systems at any finite temperature T . We describe nonessential technical details in the *SI Text*.

Review of the Classical Metropolis Method

In the standard Metropolis scheme as employed in classical computation (5), the thermal Gibbs distribution $e^{-\beta E_i} / \mathcal{Z}$ of certain an eigenstate $|\varphi_i\rangle$ of a certain Hamiltonian H , associated with the eigenenergy E_i , at an inverse temperature $\beta = 1/k_B T$, is generated through a Markov-chain procedure, in which the matrix element m_{ij} of the corresponding Markov transition matrix M refers to the transition probability from the eigenstate $|\varphi_i\rangle$ to $|\varphi_j\rangle$. The equilibrium (stationary) distribution of the Markov-chain procedure, $\pi_i \equiv e^{-\beta E_i} / \mathcal{Z}$, satisfies the following detailed balance condition:

$$\pi_i m_{ij} = \pi_j m_{ji}. \quad [2]$$

For classical Hamiltonians, the eigenstate vector $|\varphi_i\rangle$ is simply one of the vectors in the computational basis $\{|x\rangle\}$, but for quantum Hamiltonians, $|\varphi_i\rangle$ is, in general, a nontrivial superposition of the computational basis vectors.

The Metropolis Filter. The Metropolis method solves problems by using an efficient way to construct a Markov chain so that the final distribution of certain random variables is identical to the distribution $\{\pi_i\}$. A well-known solution that can satisfy the de-

tailed balance condition [2] is given by the following choice: $m_{ji} = s_{ij} z_{ji}$, where $s_{ij} = s_{ji}$ is any symmetrical transition probability, and

$$z_{ij} = \min\{1, e^{-\beta(E_j - E_i)}\} \quad [3]$$

is sometimes called the Metropolis filter, which determines the acceptance or rejection of a proposed change of the eigenstates by comparing the energy difference between them.

In a practical implementation of the Metropolis method (e.g., when applied to the Ising model), one starts with an arbitrary initial configuration of spins and then applies a random transition (e.g., single spin-flip) from one configuration (eigenstate) to another, and compares the energy between the new eigenstate and the old one. According to the rule in Eq. 3, if the new eigenstate has a lower energy, then the move is accepted. Otherwise, the move is accepted only with a probability distribution given by the ratio of the corresponding Boltzmann factors—i.e., $e^{-\beta(E_j - E_i)}$. It is important to note that, in the whole operation, the explicit knowledge of the partition function \mathcal{Z} , which usually cannot be computed efficiently, is not required.

Performance of the Metropolis Method. The performance of the Metropolis method depends on many factors. For many cases, it depends on the properties of the underlying transition matrix M of the Markov chain, especially the eigenvalues λ_k , which are all positive and bounded by one, and the largest eigenvalue is always equal to one. For convenience, we order them as $\lambda_0 = 1 > \lambda_1 \geq \dots \geq \lambda_{N-1} > 0$. The problem is solved when the Markov-chain procedure produces the stationary distribution $\{\pi_i\}$. The convergence time of a Markov chain is limited by the eigenvalue gap $\delta \equiv 1 - \lambda_1$ of the transition matrix M as $O(1/\delta)$ (6). For problems such as spin glasses (7), this gap can be very small, and hence it will necessarily take a long time to solve. The purpose of this work is to show that the running time of quantum algorithms can be improved to $O(1/\sqrt{\delta})$, and that we can fully extend the Metropolis sampling algorithm into the quantum domain (i.e., for quantum Hamiltonians and with quantum speedup).

Summary of Szegedy's Method

The idea of Szegedy's method (30) of Markov-chain quantization is to construct a unitary operator W that contains the information of a Markov matrix, which turns out to have two important properties: (i) the CETS of Eq. 1 is one of the eigenstates of W , and (ii) the eigenvalue gap scales $O(\sqrt{\delta})$. These two properties are important for generating classical thermal states through annealing (31, 32).

The first step of Szegedy's method is to encode the information about the transition matrix elements m_{ij} into a pair of unitary operators, which acts on two sets of qubits:

$$U_X |i\rangle |0\rangle = \sum_j \sqrt{m_{ij}} |i\rangle |j\rangle, \quad [4]$$

$$U_Y |j\rangle |0\rangle = \sum_i \sqrt{m_{ji}} |i\rangle |j\rangle, \quad [5]$$

where $|i\rangle$ and $|0\rangle \equiv |000\dots\rangle$ are state vectors in the computational basis. Let us define a more compact notation $|i,0\rangle \equiv |i\rangle |0\rangle$. Using the detailed balance condition in Eq. 2, we have $\langle i,0 | U_X^\dagger U_Y |j,0\rangle = \sqrt{m_{ij} m_{ji}} = \sqrt{\pi_i m_{ij}} / \sqrt{\pi_j}$, which implies that, within the subspace $\{|i,0\rangle\}$, the matrix $U_X^\dagger U_Y$ is therefore a similar matrix to the Markov matrix M (i.e., same eigenvalue spectrum $\{\lambda_k\}$). Based on this property, the Szegedy operator is defined as

$$W \equiv (2\Lambda_2 - \mathcal{I})(2\Lambda_1 - \mathcal{I}), \quad [6]$$

where \mathcal{I} is the identity operator, $\Lambda_1 \equiv \sum_i |i, 0\rangle\langle i, 0|$, and $\Lambda_2 \equiv U_X^\dagger U_Y \Lambda_1 U_Y^\dagger U_X$. Now define a set of eigenvectors $|\phi_k\rangle$ such that $\Lambda_1 U_X^\dagger U_Y \Lambda_1 |\phi_k\rangle = \lambda_k |\phi_k\rangle$. By direct substitution, the Szegedy operator is block diagonal and can be decomposed into 2×2 matrices in the subspace $\{|\phi_k\rangle, U_X^\dagger U_Y |\phi_k\rangle\}$:

$$w_k = \begin{bmatrix} \cos(2\theta_k) & -\sin(2\theta_k) \\ \sin(2\theta_k) & \cos(2\theta_k) \end{bmatrix}, \quad [7]$$

where $\cos \theta_k \equiv \lambda_k$. For $k = 0$ (i.e., $\lambda_0 = 1$), we have $w_0 |\phi_0\rangle = |\phi_0\rangle$, where $|\phi_0\rangle = \sum_i \sqrt{\pi_i} |i, 0\rangle$ contains the stationary distribution of the Markov chain (see *SI Text*), which is point (i) above. On the other hand, the eigenvalues of w_k are $\exp(\pm 2i\theta_k)$, and the eigenvalue gap $\Delta \equiv |1 - e^{2i\theta_1}|$ satisfies the following inequality: $\Delta = 2\sqrt{1 - \lambda_1^2} \geq 2\sqrt{\delta}$, which corresponds to point (ii) and leads to the quadratic speedup $O(1/\sqrt{\delta})$ (31, 32).

Generalization of the Markov-Chain Quantization Method to Quantum Hamiltonians

The Markov-chain quantization method by Szegedy (30) discussed above is applicable to classical Hamiltonians only, because it was formulated in the computational basis. To extend Szegedy's method to make it applicable to quantum Hamiltonians, we start with the following procedure: first we prepare a set of n qubits initialized in the state $(1/\sqrt{2})^n (|0\rangle + |1\rangle)^{\otimes n}$, or equivalently $N^{-1/2} \sum_{x=0}^{N-1} |x\rangle$, where $N = 2^n$. By performing a bit-by-bit controlled-NOT gate on a set of n ancilla qubits initialized in the state $|000\dots 0\rangle$, we obtain the maximally entangled state $N^{-1/2} \sum_{x=0}^{N-1} |x\rangle|x\rangle$, which can be formally expressed as an entangled state of the form (*SI Text*)

$$|\phi_0\rangle \equiv \frac{1}{\sqrt{N}} \sum_{i=0}^{N-1} |\varphi_i\rangle |\tilde{\varphi}_i\rangle, \quad [8]$$

where $|\varphi_i\rangle = \sum_{x=0}^{N-1} \langle x|\varphi_i\rangle |x\rangle$ is an energy eigenstate of H (i.e., $H|\varphi_i\rangle = E_i|\varphi_i\rangle$), and $|\tilde{\varphi}_i\rangle \equiv \sum_{x=0}^{N-1} \langle \varphi_i|x\rangle |x\rangle$ is the time-reversal counterpart (complex conjugate) of $|\varphi_i\rangle$, which is the eigenstate of the corresponding time-reversal Hamiltonian $\tilde{H} \equiv H^*$ with the same eigenenergy E_i (i.e., $\tilde{H}|\tilde{\varphi}_i\rangle = E_i|\tilde{\varphi}_i\rangle$). As we shall see, the pair structure $|\varphi_i\rangle|\tilde{\varphi}_i\rangle$ of the eigenvectors provides a unique feature for our construction of the quantum algorithm. For example, through the phase estimation algorithm (PEA), the value of eigenvalue E_i can be obtained either from $|\varphi_i\rangle$ or $|\tilde{\varphi}_i\rangle$; this is the key property introduced in this paper, which relaxes the constraints of the previous quantum Metropolis algorithm (34) in applying it to Szegedy's method.

In the following, for the purpose of clarity, we shall assume that the PEA can be applied perfectly, in the sense that each eigenstate can be uniquely identified by a unique eigenvalue. It turns out that the problem of degeneracy of the eigenenergies of a quantum Hamiltonian alone is not necessarily a hurdle for the construction of our algorithm. We leave our discussion on the effects of degeneracy on this algorithm to the *SI Text*.

Encoding the Information of the Markov Chain. We have prepared the initial input state [8] of the Q²MA. Now, we have to define the required operations: Let us include in our Hilbert space an extra qubit initialized in the state $|0\rangle$ and define a more compact notation, $|i\rangle \equiv |\varphi_i\rangle|\tilde{\varphi}_i\rangle|0\rangle$. In the Q²MA, the information of a Markov chain (i.e., z_{ij}) is encoded in a pair of unitary operators U_X and U_Y in the following way (see Eqs. 4 and 5, and see *SI Text* for their detailed construction):

$$U_X|i\rangle = \sum_k (\sigma_{ik}|\varphi_k\rangle|\varphi_i\rangle|0\rangle + \gamma_{ik}|\varphi_i\rangle|\varphi_k\rangle|1\rangle), \quad [9]$$

$$U_Y|j\rangle = \sum_m (\sigma_{jm}|\varphi_m\rangle|\varphi_j\rangle|0\rangle + \gamma_{jm}|\varphi_m\rangle|\varphi_j\rangle|1\rangle), \quad [10]$$

where $\sigma_{ik} \equiv \alpha_{ki} \sqrt{z_{ik}}$, $\gamma_{ik} \equiv \alpha_{ki} \sqrt{1 - z_{ik}}$, and $\alpha_{ki} \equiv \langle \varphi_k|K|\tilde{\varphi}_i\rangle$. Here, K is an unitary operator, and like the spin-flip operator σ_x in the classical Metropolis method, K causes transitions from one eigenstate to another. Here, z_{ik} is the Metropolis filter defined in Eq. 3. Note that U_Y is related to U_X by a controlled-SWAP operation.

The physical meaning of U_X is that, when the system is in the eigenstate $|\varphi_i\rangle$, the transition probability for it to make a transition to another eigenstate $|\varphi_k\rangle$ is $|\sigma_{ik}|^2 \propto z_{ik}$. This operation is consistent with the rules in the classical Metropolis method. The meaning of U_Y is not as clear as U_X ; the way U_Y is constructed is mainly to make the combination $U_X^\dagger U_Y$ contain the same eigenvalue spectrum as the Markov matrix M (see Eq. 11).

Connection with the Markov Matrix. To see an explicit connection between the operators U_X and U_Y with the Metropolis Markov matrix M , we will apply the detailed balance condition [2] to the product of U_X^\dagger and U_Y , which are the key elements in constructing Szegedy's operator W (Eq. 12). From Eqs. 9 and 10, for $j \neq i$, we have $\langle j|U_X^\dagger U_Y|i\rangle = |\alpha_{ji}|^2 (z_{ij}z_{ji})^{1/2}$, where we used $\langle \varphi_i|\tilde{\varphi}_j\rangle = \langle \varphi_j|\tilde{\varphi}_i\rangle$. Note that the right-hand side is equal to $(m_{ij}m_{ji})^{1/2}$. On the other hand, $\langle i|U_X^\dagger U_Y|i\rangle = |\alpha_{ii}|^2 + \sum_k |\alpha_{ki}|^2 (1 - z_{ik})$, which, as we shall see, can be interpreted as the probability of not undergoing a transition. Now, using Eq. 2, we obtain the following decomposition:

$$\langle j|U_X^\dagger U_Y|i\rangle = \langle j|D_\pi^{1/2}|j\rangle \langle j|M|i\rangle \langle i|D_\pi^{-1/2}|i\rangle, \quad [11]$$

where $D_\pi \equiv \sum_{j=0}^{N-1} \pi_j |j\rangle\langle j|$ is a diagonal matrix, and $M \equiv \sum_{i,j} m_{ij} |j\rangle\langle i|$, with $m_{ii} \equiv \langle i|U_X^\dagger U_Y|i\rangle$ and $m_{ij} \equiv |\alpha_{ji}|^2 z_{ij}$ for $j \neq i$, is the transition matrix of the Markov chain. Within the subspace $\{|i\rangle\}$, Eq. 11 implies that $U_X^\dagger U_Y$ and M are similar matrices, which means that they have the same set of eigenvalues λ_k . The stationary distribution $\pi_i = e^{-\beta E_i} / \mathcal{Z}$ of M is the same as the Gibbs distribution for the quantum system to be simulated. This property is relevant to the construction of Szegedy's operator W , which is defined in the following section.

Construction of the Generalized Szegedy Operator

Now, we have shown in Eq. 11 that the product $U_X^\dagger U_Y$ is a similarity transform of a Markov matrix M within the subspace $\{|i\rangle\}$. We will need to take one more step to see how this property can lead us to be able to make good use of it. Following Szegedy (30), the results above allow us to construct a unitary operator (see Eq. 6)

$$W \equiv (2\Lambda_2 - \mathcal{I})(2\Lambda_1 - \mathcal{I}), \quad [12]$$

where \mathcal{I} is the identity operator, and the two projectors are defined by $\Lambda_1 \equiv \sum_{i=0}^{N-1} |i\rangle\langle i|$ and $\Lambda_2 \equiv U_X^\dagger U_Y \Lambda_1 U_Y^\dagger U_X$. Note that the structure of W is similar to those reflectors in other Grover-type algorithms. It therefore comes as no surprise that one can achieve a quadratic quantum speedup. However, the quantum speedup achieved here is quite different from that in the Grover-search problem, where the quantum speedup refers to the gain in the reduction of computational steps compared with the problem size. Here, the quantum speedup refers to the gain in the amplification of the Markov gap from δ to $\sqrt{\delta}$, which has no direct relationship with the problem size.

Table 1. Comparison of various Markov-chain-based algorithms for thermal-state preparation (see text for definitions)

Methods	Hamiltonian	Input	Output	Quantum speedup
Q Metropolis I (33)	quantum	any ρ	ρ_{th}	no
Q Metropolis II (34)	quantum	any ρ	ρ_{th}	no
Q Markov chain I (31)	classical	$ +\rangle^{\otimes n}$	CETS I	quadratic
Q Markov chain II (32)	classical	$ +\rangle^{\otimes n}$	CETS I	quadratic
Q ² MA (this work)	quantum	$ \phi_0\rangle$	CETS II	quadratic

Spectral Properties of the Generalized Szegedy Operator. The Szegedy operator W will be employed to perform the projection of the CETS (see Eq. 1) in the phase estimation algorithm. In the following, we will show that the “ground state” of W is indeed the CETS. The spectral properties of W can be seen in the following way: Define $|\alpha_k\rangle \equiv \sum_{i=0}^{N-1} a_{ki}|i\rangle$ to be the eigenvectors of $\Lambda_1 U_X^\dagger U_Y \Lambda_1$; the eigenvalue equation can be written as $\Lambda_1 U_X^\dagger U_Y |\alpha_k\rangle = \lambda_k |\alpha_k\rangle$. On the other hand, using the fact (see SI Text) that $\Lambda_1 U_X^\dagger U_Y \Lambda_1 = \Lambda_1 U_Y^\dagger U_X \Lambda_1$, we have $\Lambda_2 |\alpha_k\rangle = \lambda_k U_X^\dagger U_Y |\alpha_k\rangle$. These two equations suggest that, if we start with vectors within Λ_1 , W can be block-diagonalized into subspace of 2×2 matrices w_k spanned by the basis $\{|\alpha_k\rangle, U_X^\dagger U_Y |\alpha_k\rangle\}$. Explicitly (see Eq. 7),

$$w_k = \begin{bmatrix} \cos(2\theta_k) & -\sin(2\theta_k) \\ \sin(2\theta_k) & \cos(2\theta_k) \end{bmatrix}, \quad [13]$$

where $\cos \theta_k \equiv \lambda_k$. Note that the eigenvalues of w_k is $e^{\pm i\theta_k}$. In the case of $k = 0$, where $\lambda_0 = 1$ (or $\theta_0 = 0$), $w_0 = \mathcal{I}$ is simply an identity matrix.

From Eq. 11 and the properties of the Markov matrix (see details in SI Text), the $k = 0$ state is the CETS in Eq. 1. Recall that $|i\rangle \equiv |\varphi_i\rangle|\tilde{\varphi}_i\rangle|0\rangle$, the state $|\alpha_0\rangle$ becomes the Gibbs thermal state $\rho_{\text{th}} = e^{-\beta H} / \text{Tr}(e^{-\beta H})$ when the other qubits are traced out.

Quadratic Quantum Speedup. One of the most important features about W is that the minimum eigenvalue gap $\Delta_{\min} \equiv |2\theta_1|$ of W is larger than two times the square root of the gap $\delta \equiv 1 - \lambda_1$ of the transition matrix M (using $2\theta \geq |1 - e^{2i\theta}| = 2\sqrt{1 - \cos^2 \theta}$):

$$\Delta_{\min} \geq 2\sqrt{\delta}, \quad [14]$$

which is the mathematical origin of the quadratic speedup of Szegedy’s algorithm. Recall that, for classical Markov-chain algorithms, the performance scales as $O(1/\delta)$ (6). As we shall see next, when we apply W in the phase estimation algorithm, the performance scales as $O(1/\sqrt{\delta})$.

This section completes our discussion on the necessary tools needed for the following discussion, and we are ready to go through the method of QSA (31).

Quantum Simulated Annealing

Our quantum simulation algorithm goes as follows: Given a quantum Hamiltonian H and any finite temperature T , the goal is to obtain the corresponding CETS of the form in Eq. 1, from the initial state defined in Eq. 8, which can be readily prepared from the state $|000\dots 0\rangle$, and which can be considered as the infinite-temperature ($\beta = 0$) state. To achieve this goal, we can use the method of QSA (31). For completeness, we outline the basic strategy and summarize the related results in the following paragraph. Technical details are included in the SI Text.

Let us suppose that we aim to obtain the CETS corresponding to the temperature β . The strategy of the QSA method is to prepare a sequence, $j = 0, 1, 2, \dots, d$, of $d + 1$ of coherent thermal states, $|\alpha_0^j\rangle = \sum_{i=0}^{N-1} (e^{-\beta E_i} / \mathcal{Z}_j)^{1/2} |i\rangle$, at a time; the temperature

$\beta_j \equiv (j/d)\beta$ of the coherent thermal state is lowered in each step. The basic idea for implementing this procedure is that, for a sufficiently small change of β , $\Delta\beta \equiv \beta/d$, one can show that the CETS $|\alpha_0^j\rangle$ has a good overlap with another CETS $|\alpha_0^{j+1}\rangle$ at a lower temperature. Explicitly, the fidelity between the two states is bounded by the following expression: $|\langle \alpha_0^{j+1} | \alpha_0^j \rangle|^2 \geq 1 - \epsilon_0$, where ϵ_0 is bounded by $O(\Delta\beta^2 \|H\|^2)$ (SI Text).

The next step in the QSA is to consider a projective measurement operator $\Pi_{k+1} = \sum_{k=0}^{N-1} |\alpha_k^{j+1}\rangle \langle \alpha_k^{j+1}|$ constructed in the eigenbasis of $|\alpha_k^{j+1}\rangle$. When Π_{k+1} is applied to the CETS $|\alpha_0^j\rangle$, it will result in the lower-temperature state $|\alpha_0^{j+1}\rangle$ with a high probability. The same procedure is then repeated for the other CETS at lower temperatures. For the entire process,

$$|\alpha_0^0\rangle \xrightarrow{\Pi_1} |\alpha_0^1\rangle \xrightarrow{\Pi_2} \dots \xrightarrow{\Pi_d} |\alpha_0^d\rangle. \quad [15]$$

The total error accumulated (see details in SI Text) $\epsilon = d\epsilon_0 < O(\beta^2 \|H\|^2/d)$ can be suppressed by refining the step size, which is analogous to the quantum Zeno effect applied to a slowly varying basis.

Now, using the machinery we have developed, the operator* W_{j+1} can be used to construct such a projective measurement, through the phase estimation algorithm (32) (see also SI Text). Alternatively, one may employ the method in ref. 31, which is equivalent to some projective measurements, where W_{j+1} is applied multiple times randomly. In any case, the number of controlled W_{j+1} is at most $O(1/\sqrt{\delta})$ for each step of annealing, which is a quadratic speedup relative to classical Markov chains.

Comparison with Other Related Quantum Algorithms

To summarize, we have described a quantum Metropolis algorithm that extends Szegedy’s method of classical Markov-chain quantization to the quantum domain and provides a quadratic quantum speedup $O(1/\sqrt{\delta})$ in the gap δ of the transition matrix M . This extension is achieved by adopting a dual representation where the set of basis states consists of pairs of eigenstates $|\varphi_i\rangle|\tilde{\varphi}_i\rangle$ related by the time-reversal operation. The algorithm requires twice the amount of system qubits, and the number of ancilla qubits has a logarithmic dependence on temperature $O[\log(1/T)]$. Details for comparison are summarized in the SI Text.

A comparison of various Markov-chain-based quantum algorithms is summarized in Table 1. To facilitate reading, we explain several definitions and terminologies below: (i) the term “classical Hamiltonian” (e.g., Ising spin model) refers to the cases where the eigenstates are the same as the basis states in the computational basis; the term “Quantum Hamiltonians” refer to the cases where the eigenstates generally contain superposition states of the states in the computational basis. Therefore, the eigenstates of the classical Hamiltonians are by definition known, whereas those for quantum Hamiltonians are not. (ii) The symbol ρ refers to a density matrix, $|+\rangle \equiv (|0\rangle + |1\rangle)/\sqrt{2}$, and the quantum state $|\phi_0\rangle$ as defined by Eq. 8). (iii) The thermal density matrix is $\rho_{\text{th}} = e^{-\beta H} / \text{Tr}(e^{-\beta H})$, CETS II is defined in Eq. 1), and

*The W operator defined in [12] for $|\alpha_k^{j+1}\rangle$.

CETS I (see also ref. 23) is similar to CETS II, but the basis states $\{|i\rangle\}$ are replaced by the computational basis. (iv) By “quantum speedup”, we consider only the quantum speedup with respect to the gap δ of the transition matrix of the Markov chain. It is not known whether other types of quantum speedup exist.

Conclusion

In conclusion, we presented a quantum algorithm that completes the generalization of the classical Metropolis method to the quantum domain. Similar to ref. 34, the advantages of this quantum algorithm over classical algorithms can be exponential because there is no need to explicitly solve for the eigenvalues and eigenvectors using classical algorithms for the quantum

Hamiltonians being simulated. Finally, as the application of the Metropolis method to quantum Hamiltonians can be considered as a special case of quantum maps (operations), it may be possible that the results presented here could be generalized to allow quantum speedup for a much broader class of quantum maps.

ACKNOWLEDGMENTS. We thank S. Boixo, J. McClean, D. Poulin, R. D. Somma, K. Temme, J. D. Whitfield, and P. Wocjan for insightful discussions, and we are grateful to the following funding sources: Croucher Foundation (M.-H.Y); Defense Advanced Research Planning Agency under the Young Faculty Award N66001-09-1-2101-DOD35CAP, The Camille and Henry Dreyfus Foundation, and the Sloan Foundation (A.A.-G).

- Brush SG (1967) History of the Lenz-Ising model. *Rev Mod Phys* 39:883–893.
- Nishimori H (2001) *Statistical Physics of Spin Glasses and Information Processing: An Introduction* (Oxford Univ Press, Oxford).
- De las Cuevas G, Dür W, Van den Nest M, Martin-Delgado MA (2011) Quantum algorithms for classical lattice models. *New J Phys* 13:093021.
- Kohn W (1999) Nobel Lecture: Electronic structure of matter—wave functions and density functionals. *Rev Mod Phys* 71:1253–1266.
- Thijssen JM (1999) *Computational Physics* (Cambridge Univ Press, Cambridge, UK).
- Aldous DJ (1982) Some inequalities for reversible Markov chains. *J London Math Soc* s2-25:564–576.
- Mezard M, Parisi G, Virasoro MA (1987) *Spin Glass Theory and Beyond* (World Scientific, Teaneck, NJ).
- Troyer M, Wiese U-J (2005) Computational complexity and fundamental limitations to fermionic quantum Monte Carlo simulations. *Phys Rev Lett* 94:170201–170204.
- Kalos M, Pederiva F (2000) Exact Monte Carlo method for continuum fermion systems. *Phys Rev Lett* 85:3547–3551.
- Foulkes W, Mitas L, Needs R, Rajagopal G (2001) Quantum Monte Carlo simulations of solids. *Rev Mod Phys* 73:33–83.
- Feynman RP (1982) Simulating physics with computers. *Int J Theor Phys* 21:467–488.
- Buluta I, Nori F (2009) Quantum simulators. *Science* 326:108–111.
- Kassal I, Whitfield JD, Perdomo-Ortiz A, Yung M-H, Aspuru-Guzik A (2011) Simulating chemistry using quantum computers. *Annu Rev Phys Chem* 62:185–207.
- Greiner M, Mandel O, Esslinger T, Hänsch TV, Bloch I (2002) Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms. *Nature* 415:39–44.
- Lloyd S (1996) Universal quantum simulators. *Science* 273:1073–1078.
- Lidar D, Wang H (1999) Calculating the thermal rate constant with exponential speed-up on a quantum computer. *Phys Rev E Stat Nonlin Soft Matter Phys* 59:2429–2438.
- Zalka C (1998) Simulating quantum systems on a quantum computer. *Proc R Soc Lond A* 454:313–322.
- Kassal I, Jordan SP, Love PJ, Mohseni M, Aspuru-Guzik A (2008) Polynomial-time quantum algorithm for the simulation of chemical dynamics. *Proc Natl Acad Sci USA* 105:18681–18686.
- Ortiz G, Gubernatis J, Knill E, Laflamme R (2001) Quantum algorithms for fermionic simulations. *Phys Rev A* 64:022319–022322.
- Poulin D, Wocjan P (2009) Preparing Ground States of Quantum Many-Body Systems on a Quantum Computer. *Phys Rev Lett* 102:130503–130506.
- Poulin D, Wocjan P (2009) Sampling from the Thermal Quantum Gibbs State and Evaluating Partition Functions with a Quantum Computer. *Phys Rev Lett* 103:220502–220505.
- Bilgin E, Boixo S (2010) Preparing Thermal States of Quantum Systems by Dimension Reduction. *Phys Rev Lett* 105:170405–170408.
- Yung M-H, Nagaj D, Whitfield J, Aspuru-Guzik A (2010) Simulation of classical thermal states on a quantum computer: A transfer-matrix approach. *Phys Rev A* 82:060302(R)–060305.
- Abrams D, Lloyd S (1999) Quantum algorithm providing exponential speed increase for finding eigenvalues and eigenvectors. *Phys Rev Lett* 83:5162–5165.
- Aspuru-Guzik A, Dutoi AD, Love PJ, Head-Gordon M (2005) Simulated quantum computation of molecular energies. *Science* 309:1704–1707.
- Wang HF, Kais S, Aspuru-Guzik A, Hoffmann MR (2008) Quantum algorithm for obtaining the energy spectrum of molecular systems. *Phys Chem Chem Phys* 10:5388–5393.
- Weis L, Pittner J (2010) Quantum computing applied to calculations of molecular energies: CH2 benchmark. *J Chem Phys* 133:194106–194115.
- Whitfield J, Biamoto J, Aspuru-Guzik A (2011) Simulation of electronic structure Hamiltonians using quantum computers. *Mol Phys* 109:735–750.
- Li Z, et al. (2011) Solving Quantum ground-state problems with nuclear magnetic resonance. *Sci Rep* 1:88.
- Szegedy M (2004) Quantum speed-up of Markov chain based algorithms. *45th Annual IEEE Symposium on Foundations of Computer Science* pp 32–41.
- Somma R, Boixo S, Barnum H, Knill E (2008) Quantum simulations of classical annealing processes. *Phys Rev Lett* 101:130504–130507.
- Wocjan P, Abeyesinghe A (2008) Speedup via quantum sampling. *Phys Rev A* 78:042336–042343.
- Terhal B, DiVincenzo D (2000) Problem of equilibration and the computation of correlation functions on a quantum computer. *Phys Rev A* 61:22301–22322.
- Temme K, Osborne TJ, Vollbrecht KG, Poulin D, Verstraete F (2011) Quantum Metropolis sampling. *Nature* 471:87–90.
- Wootters WK, Zurek WH (1982) A single quantum cannot be cloned. *Nature* 299:802–803.