

Speedup via quantum sampling

Pawel Wocjan* and Anura Abeyesinghe

School of Electrical Engineering and Computer Science, University of Central Florida, Orlando, Florida 32816, USA

(Received 8 September 2008; published 31 October 2008)

The Markov-chain Monte Carlo method is at the heart of efficient approximation schemes for a wide range of problems in combinatorial enumeration and statistical physics. It is therefore very natural and important to determine whether quantum computers can speed up classical mixing processes based on Markov chains. To this end, we present a quantum algorithm, making it possible to prepare a quantum sample—i.e., a coherent version of the stationary distribution of a reversible Markov chain. Our algorithm has a significantly better running time than that of a previous algorithm based on adiabatic-state generation. We also show that our methods provide a greater speedup over a recently proposed method for obtaining the ground states of (classical) Hamiltonians.

DOI: 10.1103/PhysRevA.78.042336

PACS number(s): 03.67.Ac, 87.10.Rt

I. INTRODUCTION

Randomization plays a crucial role in the design of efficient algorithms for computing approximate solutions to problems in combinatorial enumeration and statistical physics that are known to be #P-complete. Important examples are randomized polynomial-time approximation schemes for evaluating the permanent of a non-negative matrix [1], the volume of a convex polytope [2], and the partition functions of the monomer-dimer and ferromagnetic Ising systems [3,4]. The centerpiece of all these algorithms is the Markov-chain Monte Carlo (MCMC) method, making it possible to approximately sample from a particular probability distribution π over a large set Ω .

In the MCMC method, one constructs a sparse, ergodic Markov chain (stochastic matrix) P on the state space Ω such that its stationary distribution π is the desired probability distribution and then, starting from some initial state x , repeatedly applies P so that the resulting probability distribution over Ω after τ steps is sufficiently close to π . The required number of step τ is referred to as the mixing time. Bounding the mixing time of the Markov chain is often the major technical hurdle in proving the running time of the overall algorithm. This problem can be reduced to estimating the spectral gap δ of P as described in the following.

We refer the reader to [5,6] for more details on Markov chains. Let P be an ergodic (i.e., irreducible and aperiodic) reversible Markov chain with finite state space Ω and stationary distribution $\pi=(\pi(x))_{x\in\Omega}$. Let $P^{(t)}(x,y)$ denote the t -step transition probability from x to y and $P^{(t)}(x,\cdot)$ the probability distribution if we start in x and apply the Markov chain t times. To determine how fast the stationary distribution is approached from some initial probability distribution, we have to look at the spectral properties of the transition matrix P .

It follows by the Perron-Frobenius theorem that the stationary distribution π is the unique (left) eigenvector of P —i.e., $\pi^T P = \pi^T$ —with associated eigenvalue $\lambda_0=1$. Let $\{\lambda_j: 1 \leq j \leq N-1\}$, $\lambda_j \in \mathbb{R}$, denote the remaining eigenvalues

(not necessary distinct), where $N=|\Omega|$. It also follows that these eigenvalues satisfy $|\lambda_j| < 1$ for $1 \leq j \leq N-1$. Let us order the eigenvalues such that

$$1 = \lambda_0 > |\lambda_1| \geq |\lambda_2| \geq \dots \geq |\lambda_{N-1}| \geq 0.$$

Let δ denote the spectral gap of P —i.e., $\delta=1-|\lambda_1|$. The rate of convergence to the stationary distribution is governed by the spectral gap ([5], p. 61).

The variation distance from initial state x is

$$d^{(t)}(x) = \frac{1}{2} \sum_{y \in \Omega} |P^{(t)}(x,y) - \pi(y)|.$$

For $\epsilon \in (0,1)$, let

$$\tau_\epsilon(x) = \min\{t: d^{(t)}(x) \leq \epsilon \text{ for all } t' \geq t\},$$

$$\tau_\epsilon = \max_{x \in \Omega} \tau_\epsilon(x).$$

These quantities satisfy the following inequalities:

$$\frac{1}{2\delta} \log(2\epsilon)^{-1} \leq \tau_\epsilon$$

$$\tau_\epsilon(x) \leq \frac{1}{\delta} \{\log[\pi(x)]^{-1} + \log \epsilon^{-1}\}.$$

Given the fact the MCMC method is at the heart of so many efficient classical algorithms, it is very natural and important to determine to what extent quantum computers can speed up classical mixing. It is possible to prepare the quantum sample $|\pi\rangle$ —i.e., a coherent version of the stationary distribution where the amplitudes are given by the square roots of the probabilities of the states by quantized Markov chains [7,8]. The running time of this algorithm is $O(1/\sqrt{\delta\pi(x)})$. Basically, this algorithm is a Grover search where marking of the target state is the reflection around the quantum sample $|\pi\rangle$. This reflection can be realized with cost $O(1/\sqrt{\delta})$.

Unfortunately, this running time is too high for applications where the state space Ω is exponentially large and $\pi(x)$ can be exponentially small. The question whether a quantum

*wocjan@eecs.ucf.edu

speed up to $O(1/\sqrt{\delta} \log[1/\pi(x)])$ is possible has been examined in [9]. The author proposes a method based on quantum walks that decohere under repeated randomized measurements to attack this problem. He shows that this speedup is indeed achievable for the decoherent quantum walk on a periodic lattice \mathbb{Z}_n^d . However, the question whether this speedup is achievable for arbitrary Markov chains remains an important open problem.

In this paper, we propose a different method for speeding up classical mixing processes. We show how to efficiently prepare the quantum sample $|\pi\rangle$ provided that we have a sequence of slowly varying Markov chains in the following sense: (i) there are Markov chains $P_0, P_1, \dots, P_r = P$ with stationary distributions $\pi_0, \pi_1, \dots, \pi_r$ such that distributions of adjacent Markov chains are sufficiently close and (ii) the quantum sample $|\pi_0\rangle$ can be prepared efficiently. The idea of quantum-state generation based on slowly varying Markov chains was proposed in [10], where adiabatic techniques were used to create the quantum samples $|\pi_i\rangle$ sequentially. The resulting running time is $\Omega(1/\delta)$; i.e., it does not provide the speedup that would reduce δ to $\sqrt{\delta}$.

We improve the running time by (a) using Szegedy's quantum walk operators instead of Hamiltonians derived from the Markov chains and (b) preparing the intermediate quantum samples by amplitude amplification. The resulting running time is worse than Richter's conjectured running time of $O(1/\sqrt{\delta} \log[1/\pi(x)])$. However, in some situations, our running time is better than the classical lower bound on the mixing time.

We then apply our method to the special case of simulated annealing. We obtain a better running time than that of a recently proposed method for this purpose [11,12].

Their method makes use of Szegedy's quantum walk operators and the quantum Zeno effect. It yields a quadratic speedup with respect to the spectral gap and provides an overall speedup with respect to the classical algorithm for simulated annealing. However, if one applied their method to the general case considered in [10], one would also obtain the quadratic speedup with respect to the spectral gap, but not the *overall* speedup. The problem is that the quantum Zeno effect would result in an exponential slowdown in the general case.¹ This is avoided by amplitude amplification in our approach, which can lead to an overall speedup in the general case.

The paper is organized as follows. In Sec. II we present a modified version of Szegedy's quantum analog of an ergodic reversible Markov chain whose unique eigenvector with eigenvalue 1 (on the relevant computational subspace) is the quantum sample of the stationary distribution. We also describe its spectral properties in detail. In Sec. III we present a primitive for preparing quantum states based on amplitude amplification and also a primitive for implementing approxi-

mately phase gates that are needed for amplitude amplification. In Sec. IV we use the primitives to obtain a quantum method for preparing the quantum sample of an arbitrary reversible Markov chain provided that we have a sequence of slowly varying Markov chains. In Sec. V we show that our approach makes it possible to prepare quantum samples of Boltzmann-Gibbs distributions more efficiently.

II. QUANTUM ANALOG OF CLASSICAL ERGODIC REVERSIBLE MARKOV CHAINS

We refer the reader to [7,8] for more details on the quantization. Let $\mathcal{H} = \mathbb{C}^N \otimes \mathbb{C}^N$. The basis states of \mathcal{H} are denoted by $|xy\rangle$ for $x, y \in \Omega$. For $x \in \Omega$, define the normalized vectors

$$|p_x\rangle = \sum_{y \in \Omega} \sqrt{p_{xy}} |y\rangle,$$

where p_{xy} denotes the transition probability from x to y . A *quantum update* is any unitary U that satisfies

$$U|x\rangle|0\rangle = |x\rangle|p_x\rangle$$

for some fixed state $0 \in \Omega$ and all $x \in \Omega$. We refer to the cost to realize U and its inverse U^\dagger as the quantum update cost.

To construct the quantum walk, we define the subspaces

$$\mathcal{A} = \text{span}\{|x\rangle|0\rangle : x \in \Omega\},$$

$$\mathcal{B} = U^\dagger S U \mathcal{A},$$

where S denotes the unitary operator swapping the two tensor components of \mathcal{H} . For $\mathcal{K} = \mathcal{A}, \mathcal{B}$, denote by $\Pi_{\mathcal{K}}$ the orthogonal projection onto \mathcal{K} and by

$$R_{\mathcal{K}} = 2\Pi_{\mathcal{K}} - I$$

the reflection around \mathcal{K} .

Definition 1 (quantum walk). The quantum walk $W(P)$ based on the classical reversible Markov chain P is defined to be the unitary operation (rotation)

$$W(P) = R_{\mathcal{B}} \cdot R_{\mathcal{A}}. \quad (1)$$

The quantum walk $W(P)$ can be realized by applying both U and U^\dagger twice:

$$W(P) = U^\dagger \cdot S \cdot U \cdot R_{\mathcal{A}} \cdot U^\dagger \cdot S \cdot U \cdot R_{\mathcal{A}}.$$

Our definition of $W(P)$ is equal to that in [11]. This is different from the definition used in [7,8]:

$$\tilde{W}(P) = R_{\tilde{\mathcal{B}}} \cdot R_{\tilde{\mathcal{A}}},$$

where

$$\tilde{\mathcal{A}} = \text{span}\{|x\rangle|p_x\rangle : x \in \Omega\} = U \mathcal{A},$$

$$\tilde{\mathcal{B}} = \text{span}\{|p_x\rangle|x\rangle : x \in \Omega\} = U \mathcal{B}.$$

Since $W(P)$ and $\tilde{W}(P)$ are equal up to conjugation by U , we can apply the spectral analysis from [7] to determine the spectrum of $W(P)$. We refer to the subspace $\mathcal{A} + \mathcal{B}$ as the

¹The method based on the quantum Zeno effect is only efficient if the stationary distributions of adjacent Markov chains are very close to each other. In contrast, amplitude amplification works as long as the distances of adjacent stationary distributions are bounded from below by the inverse of some polynomial in the input length. The latter condition is referred to as slowly varying.

busy subspace and to its orthogonal complement—i.e., $\mathcal{A}^\perp \cap \mathcal{B}^\perp$ —as the idle subspace. Clearly, the operator $W(P)$ acts as identity on the idle subspace. On the busy subspace, the spectrum of $W(P)$ is as follows.

Theorem 1. Let P be a time-reversible Markov chain. Let $\theta_1, \dots, \theta_M \in (0, \frac{\pi}{2})$ be such that $|\lambda_1| = \cos(\theta_1), \dots, |\lambda_M| = \cos \theta_M$, where $M \leq N-1$ and the remaining eigenvalues are equal to 0—i.e., $\lambda_{M+1}, \dots, \lambda_{N-1} = 0$.

(i) On $\mathcal{A} \cap \mathcal{B}$ the operator $W(P)$ acts as the identity I . This subspace is one dimensional and is spanned by the eigenvector $|\pi\rangle|0\rangle$ where

$$|\pi\rangle = \sum_x \sqrt{\pi_x} |x\rangle$$

is the quantum sample of the stationary distribution π of P .

(ii) On $\mathcal{A} \cap \mathcal{B}^\perp$ and $\mathcal{A}^\perp \cap \mathcal{B}$ the operator $W(P)$ acts as $-I$. The dimensions of $\mathcal{A} \cap \mathcal{B}^\perp$ and $\mathcal{A}^\perp \cap \mathcal{B}$ are equal to $N-1-M$ —i.e., the dimension of the kernel of P .

(iii) On $\mathcal{A} + \mathcal{B}$ those eigenvalues of $W(P)$ that have a non-zero imaginary part are exactly $e^{\pm 2i\theta_1}, \dots, e^{\pm 2i\theta_M}$ with the same multiplicity.

(iv) $W(P)$ has no other eigenvalues on $\mathcal{A} + \mathcal{B}$.

Proof. This follows from [8] (Theorem 4) and the results in [7] (Sec. 12).

In abuse of notation, we often use $|\pi\rangle$ instead of $|\pi\rangle|0\rangle$. In the following, we always stay in the busy subspace $\mathcal{A} + \mathcal{B}$. This is important because we want to obtain the unique eigenvector $|\pi\rangle$ and not any other eigenvector with eigenvalue 1 contained in the idle subspace.

The phase gap of $\Delta(P)$ of $W(P)$ is defined to be $2\theta_1$. This is motivated by the above theorem since the angular distance of 1 from any other eigenvalue (corresponding to an eigenvector in the busy subspace) is at least $\Delta(P)$. The phase gap satisfies

$$\Delta(P) \geq |1 - e^{2i\theta_1}| = 2\sqrt{1 - \lambda_1^2} \geq 2\sqrt{\delta(P)}.$$

This inequality is at the heart of the quadratic speedup due to quantum walks.

III. PRIMITIVES

A. Preparation via amplitude amplification

We use Grover’s $\frac{\pi}{3}$ -amplitude amplification (fixed-point search), making it possible to drive a source state to the desired target state by applying a sequence of phase gates [13]. It is a special case of the general approach to amplitude amplification based on the phase matrix that was introduced in [14].

Lemma 1. Let $|t_i\rangle$ and $|t_{i+1}\rangle$ be two arbitrary quantum states in \mathbb{C}^d with $|\langle t_i | t_{i+1} \rangle|^2 \geq p$ for some p with $0 < p \leq 1$. Denote by Π_i the projection on the subspace spanned by $|t_i\rangle$ and by Π_i^\perp the projection onto the orthogonal subspace. Let $\omega = e^{(\pi/3)i}$. Define the unitaries

$$R_i = \omega \Pi_i + \Pi_i^\perp,$$

$$R_{i+1} = \omega \Pi_{i+1} + \Pi_{i+1}^\perp.$$

Define the unitaries $U_{i,m}$ recursively as follows:

$$U_{i,0} = I,$$

$$U_{i,m+1} = U_{i,m} \cdot R_i \cdot U_{i,m}^\dagger \cdot R_{i+1} \cdot U_{i,m}.$$

Then, at the m th level of recursion we have

$$|\langle t_{i+1} | U_{i,m} | t_i \rangle|^2 \geq 1 - (1-p)^{3^m}.$$

The unitaries in $\{R_i, R_i^\dagger, R_{i+1}, R_{i+1}^\dagger\}$ are used at most 3^m times.

Note that the running time of the algorithm is worse than that of Grover’s algorithm (the latter is also a special case of amplitude amplification based on the phase matrix). The reason why we cannot use Grover’s algorithm is as follows. Since we only have a lower bound on the overlap between $|t_i\rangle$ and $|t_{i+1}\rangle$, we do not know how many Grover iterations we have to apply without overshooting. Therefore, we have to employ the version of Grover’s algorithm considered in [15], making it possible to obtain the target state even if the overlap is not known. However, the problem is that for this algorithm we have to prepare the initial state *several* times. This prevents us from using Grover’s algorithm because it is absolutely necessary for our primitive discussed below that the initial state is prepared only once.

Corollary 1. Let $|t_0\rangle, \dots, |t_r\rangle$ be arbitrary quantum states in \mathbb{C}^d with $|\langle t_i | t_{i+1} \rangle|^2 \geq p$ for $i=0, \dots, r-1$. Given the state $|t_0\rangle$, we can prepare a state $|\tilde{t}_r\rangle$ such that

$$\| |\tilde{t}_r\rangle - |t_r\rangle \| \leq \epsilon_1,$$

for any $\epsilon_1 > 0$, by invoking the unitaries from $\{R_i, R_i^\dagger : i=0, \dots, r\}$ no more than

$$L = \frac{12r \log(2r/\epsilon_1)}{\log[1/(1-p)]}$$

times.

Proof. Set $q=1-p$ and $M=3^m$. For $i=0, \dots, r-1$, define $|t'_{i+1}\rangle = U_{i,m} |t_i\rangle$. It follows from Lemma 1 that $|t'_{i+1}\rangle$ can be expressed as

$$|t'_{i+1}\rangle = \alpha |t_{i+1}\rangle + \beta |t_{i+1}^\perp\rangle,$$

where α and β are two probability amplitudes with $|\alpha| \geq \sqrt{1-q^M}$ and $|\beta| \leq \sqrt{q^M}$, and $|t_{i+1}^\perp\rangle$ is some state with $\langle t_{i+1}^\perp | t_{i+1} \rangle = 0$. Consequently, we have

$$\| |t'_{i+1}\rangle - |t_{i+1}\rangle \| \leq 1 - \sqrt{1-q^M} + \sqrt{q^M} \leq 2\sqrt{q^M}.$$

For $i=0, \dots, r-1$, define

$$|\tilde{t}_{i+1}\rangle = \prod_{j=0}^i U_{j,m} |t_0\rangle.$$

The task is now to show how to choose m so that

$$\| |t_r\rangle - |\tilde{t}_r\rangle \| \leq \epsilon_1.$$

To do this, we use induction. The base step is

$$\| |\tilde{t}_1\rangle - |t_1\rangle \| \leq 2\sqrt{q^M}.$$

The inductive step is

$$\begin{aligned} \|\tilde{t}_r\rangle - |t_r\rangle\| &= \|\tilde{t}_r\rangle - U_{r-1;m}|t_{r-1}\rangle + U_{r-1;m}|t_{r-1}\rangle - |t_r\rangle\| \\ &\leq \|U_{r-1;m}\|\|\tilde{t}_{r-1}\rangle - |t_{r-1}\rangle\| + \|\tilde{t}'_r\rangle - |t_r\rangle\| \\ &\leq \|\tilde{t}_{r-1}\rangle - |t_{r-1}\rangle\| + 2\sqrt{q^M}. \end{aligned}$$

We obtain

$$\|\tilde{t}_r\rangle - |t_r\rangle\| \leq 2r\sqrt{q^M}.$$

To make the norm distance less or equal to ϵ_1 , it always suffices to choose M to be the smallest power of 3 satisfying

$$M \geq \frac{2 \log(2r/\epsilon_1)}{\log[1/(1-p)]}.$$

It follows that the unitaries from the set $\{R_i, R_i^\dagger : i=0, \dots, r+1\}$ are used at most $2rM$ times. This number is bounded from above by

$$L = \frac{12r \log(2r/\epsilon_1)}{\log[1/(1-p)]}.$$

B. Approximate phase gates

In this section we consider the case where the states $|t_0\rangle, |t_1\rangle, \dots, |t_r\rangle$ in Corollary 1 are quantum samples of stationary distributions. We show how to approximately implement the required phase transformation using quantum walks and a variant of the phase estimation algorithm.

Lemma 2. Let W be a unitary acting on \mathbb{C}^d with unique eigenvector $|\psi_0\rangle$ with eigenvalue $\lambda_0=1$. Denote the remaining eigenvectors and eigenvalues of W by $|\psi_j\rangle$ and $\lambda_j = e^{2\pi i \phi_j}$ for $j=1, \dots, d-1$, respectively. Let

$$\Delta = \min_{j=1, \dots, d-1} |\phi_j|$$

be the phase gap of W . Let

$$a = \lceil \log(1/\Delta) \rceil,$$

$$c = \lceil \log(1/\sqrt{\epsilon_2}) \rceil,$$


for some $\epsilon_2 > 0$. Then, there is a quantum circuit V acting on $\mathbb{C}^d \otimes (\mathbb{C}^2)^{\otimes ac}$ that invokes the controlled- W gate at most $2^a c$ times and has the properties

$$V|\psi_0\rangle|0\rangle^{\otimes ac} = |\psi_0\rangle|0\rangle^{\otimes ac},$$

$$V|\psi_j\rangle|0\rangle^{\otimes ac} = \sqrt{1-\epsilon_2}|\psi_j\rangle|\chi_j\rangle + \sqrt{\epsilon_2}|\psi_j\rangle|0\rangle^{\otimes ac},$$

where $|\chi_j\rangle$ are some unit vectors in $(\mathbb{C}^2)^{\otimes ac}$ with $\langle 0 \dots 0 | \chi_j \rangle = 0$ for $j=1, \dots, d-1$.

Proof. First, we apply the phase estimation circuit U with a ancilla qubits as depicted below. The circuit U invokes the controlled- W gate $2^a - 1$ times:



We have

$$\begin{aligned} U|\psi_j\rangle|0 \dots 0\rangle^{\otimes a} &= |\psi_j\rangle \otimes \text{DFT}^\dagger \left(\frac{1}{\sqrt{2^a}} \sum_{m=0}^{2^a-1} e^{2\pi i m \phi_j} |m\rangle \right) \\ &= |\psi_j\rangle \otimes \frac{1}{2^a} \sum_{m, m'=0}^{2^a-1} e^{2\pi i m \phi_j} e^{-2\pi i m m' / 2^a} |m'\rangle. \end{aligned}$$

The amplitude $\alpha_{m'}$ of the state $|m'\rangle$ is

$$\frac{1}{2^a} \sum_m e^{2\pi i (\phi_j - m' / 2^a) m} = \frac{1}{2^a} \frac{1 - e^{2\pi i (2^a \phi_j - m')}}{1 - e^{2\pi i (\phi_j - m' / 2^a)}}.$$

Observe that for $j=0$

$$\alpha_{m'} = \begin{cases} 1 & \text{if } m' = 0, \\ 0 & \text{if } m' \neq 0, \end{cases}$$

and so

$$U|\psi_0\rangle|0\rangle^{\otimes a} = |\psi_0\rangle|0\rangle^{\otimes a}.$$

Now consider the case $j \neq 0$. To bound $|\alpha_0|$, we use the inequality $|1 - e^{ix}| \geq 2|x|/\pi$ whenever $-\pi \geq x \geq \pi$. We obtain

$$\begin{aligned} |\alpha_0| &= \frac{1}{2^a} \left| \frac{1 - e^{2\pi i 2^a \phi_j}}{1 - e^{2\pi i \phi_j}} \right| \leq \frac{1}{2^{a-1}} \left| \frac{1}{1 - e^{2\pi i \phi_j}} \right| \\ &\leq \frac{1}{2^{a-1}} \frac{\pi}{2 \times 2\pi |\phi_j|} = \frac{1}{2^{a+1} |\phi_j|} \leq \frac{1}{2^{a+1} \Delta} \leq \frac{1}{2}. \end{aligned}$$

We conclude that for $j \neq 0$ we have

$$U|\psi_j\rangle|0\rangle^{\otimes a} = \alpha|\psi_j\rangle|\chi_j\rangle + \alpha_0|\psi_j\rangle|0\rangle^{\otimes a},$$

where $\langle 0 \dots 0 | \chi_j \rangle = 0$, $|\alpha| > \sqrt{3}/2$, and $|\alpha_0| \leq 1/2$.

Since we are only interested in the amplitude of the state $|0\rangle^{\otimes a}$ on the ancilla qubits, we can replace the inverse discrete Fourier transform DFT^\dagger in the phase estimation circuit by the Walsh-Hadamard transform $H^{\otimes a}$. This is seen as follows. Observe that both $H^{\otimes a}$ and DFT create the uniformly weighted superposition of all computational basis states when applied to $|0\rangle^{\otimes a}$, implying that

$$\langle 0 \dots 0 | \text{DFT}^\dagger | \phi \rangle = \langle 0 \dots 0 | H^{\otimes a} | \phi \rangle$$

for an arbitrary state $|\phi\rangle$ on the ancilla register.

Second, we reduce the “error amplitude” to $\sqrt{\epsilon_2}$ by applying the circuit U c times, using a new block of a ancillas each time. Let V be the resulting circuit. V invokes the controlled- W gate $(2^a - 1)c$ times. For $j \neq 0$, we have

$$V|\psi_j\rangle|0\rangle^{\otimes ac} = \sqrt{1-\epsilon_2}|\psi_j\rangle|\chi_j\rangle + \sqrt{\epsilon_2}|\psi_j\rangle|0\rangle^{\otimes ac},$$

as desired. For $j=0$, we have

$$V|\psi_0\rangle|0\rangle^{\otimes ac} = |\psi_0\rangle|0\rangle^{\otimes ac}.$$

This completes the proof.

Corollary 2. Let W be a unitary acting on \mathbb{C}^d with unique eigenvector $|\psi_0\rangle$ with eigenvalue $\lambda_0=1$. Denote the remaining eigenvectors and eigenvalues of W by $|\psi_j\rangle$ and $\lambda_j = e^{2\pi i\varphi_j}$ for $j=1, \dots, d-1$, respectively. Let

$$\Delta = \min_{j=1, \dots, d-1} |\varphi_j|$$

be the phase gap of W . Let Π be the projector onto the space spanned by $|\psi_0\rangle$ and Π^\perp the projector onto the orthogonal complement. Let R be the unitary that acts on \mathbb{C}^d as follows

$$R = \omega\Pi + \Pi^\perp.$$

Let

$$a = \lceil \log(1/\Delta) \rceil,$$

$$c = \lceil \log(1/\sqrt{\epsilon_2}) \rceil,$$

for some $\epsilon_2 > 0$. Then, there is a quantum circuit \tilde{R} acting on $\mathbb{C}^d \otimes (\mathbb{C}^2)^{\otimes ac}$ that invokes the controlled- W gate $2^{a+1}c$ times and has the following properties: for $j=0$,

$$\tilde{R}|\psi_0\rangle|0\rangle^{\otimes ac} = (R|\psi_0\rangle)|0\rangle^{\otimes ac},$$

and for $j \neq 0$,

$$\tilde{R}|\psi_j\rangle|0\rangle^{\otimes ac} = (R|\psi_j\rangle)|0\rangle^{\otimes ac} + |\xi\rangle,$$

where $|\xi\rangle$ is some error vector in $\mathbb{C}^d \otimes (\mathbb{C}^2)^{\otimes ac}$ with

$$\|\xi\| \leq 2\sqrt{\epsilon_2}.$$

Proof. Let

$$\tilde{R} = V^\dagger \cdot (I_d \otimes Q) \cdot V,$$

where Q is the phase gate

$$Q = \omega|0\rangle\langle 0|^{\otimes ac} + (I - |0\rangle\langle 0|^{\otimes ac})$$

that acts on the ancilla register.

For $j=0$, it clear that

$$\tilde{R}|\psi_0\rangle|0\rangle^{\otimes ac} = (R|\psi_0\rangle)|0\rangle^{\otimes ac} = \omega|\psi_0\rangle|0\rangle^{\otimes ac}.$$

Let us now analyze the action of \tilde{R} for $j \neq 0$. The state after the application of V is

$$V|\psi_j\rangle|0\rangle^{\otimes ac} = \sqrt{1-\epsilon_2}|\psi_j\rangle|\chi_j\rangle + \sqrt{\epsilon_2}|\psi_j\rangle|0\rangle^{\otimes ac},$$

where $\langle \chi_j | 0 \rangle^{\otimes ac} = 0$. The state after the application of $I_d \otimes Q$ is

$$\begin{aligned} & \sqrt{1-\epsilon_2}|\psi_j\rangle|\chi_j\rangle + \sqrt{\epsilon_2}\omega|\psi_j\rangle|0\rangle^{\otimes ac} \\ &= \sqrt{1-\epsilon_2}|\psi_j\rangle|\chi_j\rangle + \sqrt{\epsilon_2}\omega|\psi_j\rangle|0\rangle^{\otimes ac} \\ & \quad + \sqrt{\epsilon_2}|\psi_j\rangle|0\rangle^{\otimes ac} - \sqrt{\epsilon_2}|\psi_j\rangle|0\rangle^{\otimes ac} \\ &= |\psi_j\rangle \otimes (\sqrt{1-\epsilon_2}|\chi_j\rangle + \sqrt{\epsilon_2}|0\rangle^{\otimes ac}) + |\xi'\rangle, \end{aligned}$$

where

$$|\xi'\rangle = \sqrt{\epsilon_2}(\omega-1)|\psi_j\rangle|0\rangle^{\otimes ac}.$$

We have $\|\xi'\| \leq 2\sqrt{\epsilon_2}$. In the final step, the application of V^\dagger leads to the state

$$|\psi_j\rangle|0\rangle^{\otimes ac} + V^\dagger|\xi'\rangle = |\psi_j\rangle|0\rangle^{\otimes ac} + |\xi\rangle,$$

with $\|\xi\| \leq 2\sqrt{\epsilon_2}$. We conclude

$$\tilde{R}|\psi_j\rangle|0\rangle^{\otimes ac} = |\psi_j\rangle|0\rangle^{\otimes ac} + |\xi\rangle = (R|\psi_j\rangle)|0\rangle^{\otimes ac} + |\xi\rangle.$$

It follows from Lemma 2 that V and V^\dagger invoke the controlled- W gate $(2^a-1)c$ times. So \tilde{R} invokes the controlled- W gate $2(2^a-1)c < 2^{a+1}c$ times.

IV. QUANTUM SAMPLING

We now use Corollaries 1 and 2 to prove the following theorem.

Theorem 2. Let P_0, P_1, \dots, P_r be classical Markov chains with stationary distributions $\pi_0, \pi_1, \dots, \pi_r$ and spectral gaps $\delta_0, \delta_1, \dots, \delta_r$, respectively. Assume the stationary distributions of adjacent Markov chains are close to each other in the sense that their stationary distributions π_i and π_{i+1} are close with respect to fidelity—i.e.,

$$\left(\sum_{x \in \Omega} \sqrt{\pi_i(x)} \sqrt{\pi_{i+1}(x)} \right)^2 = |\langle \pi_i | \pi_{i+1} \rangle|^2 \geq p$$

for $i=0, \dots, r-1$,

$$\min\{\delta_i : i=0, \dots, r\} \geq \delta,$$

and we can prepare the quantum sample $|\pi_0\rangle$.

Then, for any $\epsilon > 0$, there is a quantum sampling algorithm, making it possible to sample according to a probability distribution $\tilde{\pi}_r$ that is close to π_r with respect to the total variation distance—i.e., $D(\tilde{\pi}_r, \pi_r) \leq \epsilon$.

The algorithm invokes the controlled- W_i operators at most $2^{a+1}cL$ times where

$$L = \frac{12r \log(8r/\epsilon)}{\log[1/(1-p)]}$$

$$a = \lceil \log(1/\Delta) \rceil$$

$$c = \left\lceil \log \left(\frac{96r \log(8r/\epsilon)}{\epsilon \log[1/(1-p)]} \right) \right\rceil.$$

Proof. Corollary 1 shows that, given the initial state $|\pi_0\rangle$, we can prepare a state $|\tilde{\pi}_r\rangle$ with $\|\pi_r - |\tilde{\pi}_r\rangle\| \leq \epsilon_1$ by invoking the unitaries from the set $\{R_i, R_i^\dagger : i=0, \dots, r-1\}$ no more than

$$L = \frac{12r \log(2r/\epsilon_1)}{\log[1/(1-p)]}$$

times.

In Corollary 1 we assumed that we can implement these exactly. However, in reality we can only implement the operators \tilde{R}_i and their inverses \tilde{R}_i^\dagger as described in Corollary 2. This approximation adds an error vector $|\xi\rangle$ every time an operator \tilde{R}_i or \tilde{R}_i^\dagger is applied, where $\|\xi\| \leq 2\sqrt{\epsilon_2}$ for some $\epsilon_2 > 0$.

Let $|\tilde{\psi}\rangle$ be the state obtained by implementing Corollary 2 using \tilde{R}_i to approximate R_i . Then, since these operators or their inverses are invoked no more than L times, we have

$$|\tilde{\psi}\rangle = |\tilde{\pi}_r\rangle_{\mathcal{H}}|0\rangle_{\mathcal{A}}^{\otimes ac'} + |\xi\rangle,$$

where $c' = \lceil \log(1/\sqrt{\epsilon_2}) \rceil$ and $|\xi\rangle$ is some vector with

$$\|\xi\| \leq 2L\sqrt{\epsilon_2}.$$

\mathcal{H} is the Hilbert space that our quantum samples reside in, and \mathcal{A} is the Hilbert space of the ancilla qubits that are required to implement the approximate phase gates \tilde{R}_i and their inverses.

Let

$$|\psi\rangle = |\pi_r\rangle|0\rangle^{\otimes ac}$$

be the ideal state. We choose $\epsilon_1 = \epsilon/4$ and $\epsilon_2 = \epsilon^2/(64L^2)$ so that

$$\begin{aligned} \|\psi\rangle - |\tilde{\psi}\rangle &\leq \|\psi\rangle - |\tilde{\pi}_r\rangle|0\rangle^{\otimes ac}\| + \|\tilde{\pi}_r\rangle|0\rangle^{\otimes ac} - |\tilde{\psi}\rangle\| \\ &\leq \epsilon_1 + 2L\sqrt{\epsilon_2} = \epsilon/4 + \epsilon/4 = \epsilon/2. \end{aligned}$$

For each $x \in \Omega$, we define the projector

$$\Lambda_x = |x\rangle\langle x| \otimes |0\rangle\langle 0|^{\otimes ac}$$

acting on $\mathcal{H} \otimes \mathcal{A}$. Let

$$\Lambda_0 = I_{\mathcal{H}, \mathcal{A}} - I_{\mathcal{H}} \otimes |0\rangle\langle 0|^{\otimes ac}.$$

Let $\Omega' = \Omega \cup \{0\}$. Observe that the desired distribution π is equal to the probability distribution given by

$$\pi(x) = \|\Lambda_x|\psi\rangle\|^2.$$

Our protocol yields the probability distribution

$$\tilde{\pi}(x) = \|\Lambda_x|\tilde{\psi}\rangle\|^2.$$

We now bound the total variation distance between π and $\tilde{\pi}$ from above. For a subset $S \subseteq \Omega'$, let

$$\Lambda_S = \sum_{x \in S} \Lambda_x.$$

We have

$$\begin{aligned} D(\pi, \tilde{\pi}) &= \max_{S \subseteq \Omega'} |\pi(S) - \tilde{\pi}(S)| \\ &= \max_{S \subseteq \Omega'} \left| \|\Lambda_S|\psi\rangle\|^2 - \|\Lambda_S|\tilde{\psi}\rangle\|^2 \right| \\ &\leq 2 \max_{S \subseteq \Omega'} \left| \|\Lambda_S|\psi\rangle\| - \|\Lambda_S|\tilde{\psi}\rangle\| \right| \leq 2\|\psi\rangle - |\tilde{\psi}\rangle\| \leq \epsilon. \end{aligned}$$

It follows from Corollaries 1 and 2 that we invoke the controlled- W_i operators or their inverses at most $2^{a+1}cL$ times.

V. QUANTUM SIMULATED ANNEALING

The Metropolis algorithm refers to a general construction that transforms any irreducible Markov chain on state space

Ω to a time-reversible Markov chain with a required stationary distribution. We consider the case where the desired stationary distribution is equal to the Boltzmann-Gibbs distribution π_β of some (classical) Hamiltonian H at (inverse) temperature β . We show how to prepare the corresponding quantum sample $|\pi_\beta\rangle$ using Theorem 2. The resulting algorithm has a better running time than that of the recently proposed algorithm based on the quantum Zeno effect [11,12].

For completeness, we give a short description of the Metropolis algorithm. The presentation is based on [16] (Lemma 10.8). For a finite state space Ω and neighborhood structure $\{N(x) : x \in \Omega\}$, let $N = \max_{x \in \Omega} |N(x)|$. Let M be any number such that $M \geq N$. For all $x \in \Omega$, let $E(x) \geq 0$ be the energy of the state x . The desired stationary distribution is $\pi_\beta = (\pi_\beta(x) : x \in \Omega)$, where

$$\pi_\beta(x) = \frac{\exp(-E(x)\beta)}{Z_\beta}$$

is the probability of state x and

$$Z_\beta = \sum_{x \in \Omega} \exp[-\beta E(x)]$$

denotes the partition function at temperature β .

Consider the Markov chain P_β whose transition probabilities p_{xy} are

$$\min\{1, \exp[(E(y) - E(x))\beta]\}/M \quad \text{if } x \neq y, y \in N(x),$$

$$0 \quad \text{if } x \neq y, y \in N(x),$$

$$1 - \sum_{z \neq x} p_{xz} \quad \text{if } x = y.$$

Then, if this chain is irreducible and aperiodic, P_β is time reversible and its stationary distribution is given by the Boltzmann-Gibbs probabilities $\pi_\beta(x)$. Let δ_β denote its spectral gap.

Let H be the Hamiltonian defined by

$$H = \sum_{x \in \Omega} E(x)|x\rangle\langle x|.$$

We need the following simple lemma that characterizes how the quantum sample changes when the temperature is increased.

Lemma 3. The quantum samples of the stationary distributions of the above Metropolis process at temperatures β and $\beta + \Delta\beta$ satisfy

$$|\langle \pi_\beta | \pi_{\beta + \Delta\beta} \rangle|^2 \geq \exp(-\|H\|\Delta\beta)$$

for all β and all $\Delta\beta$.

Proof. We have

$$\begin{aligned} \langle \pi_\beta | \pi_{\beta + \Delta\beta} \rangle &= \sum_{x \in \Omega} \frac{\exp[-\beta E(x)/2]}{\sqrt{Z_\beta}} \frac{\exp[-(\beta + \Delta\beta)E(x)/2]}{\sqrt{Z_{\beta + \Delta\beta}}} \\ &\geq \sum_{x \in \Omega} \frac{\exp[-\beta E(x)]}{Z_\beta} \exp[-E(x)\Delta\beta/2] \\ &\geq \exp(-\|H\|\Delta\beta/2). \end{aligned}$$

The first inequality follows from the fact that $Z_\beta > Z_{\beta+\Delta\beta}$. By taking the square, we obtain the desired result.

The following corollary follows directly from Theorem 2 by observing that $p=1/e$ if we set $\Delta\beta=1/\|H\|$ in Lemma 3.

Corollary 3. Let $r=\beta\|H\|$, $\beta_i=i/\|H\|$ for $i=0, \dots, r$,

$$\delta \leq \min\{\delta_{\beta_i}; i=1, \dots, r\}$$

be a lower bound on the smallest spectral gap and Δ the phase gap corresponding δ . Then, for any $\epsilon > 0$, there is a quantum algorithm that outputs the states x according to a probability distribution $\tilde{\pi}_\beta$, with

$$D(\tilde{\pi}_\beta, \pi_\beta) \leq \epsilon.$$

The algorithm invokes the operators from $\{W_{\beta_i}\}$ at most

$$2^{a+1}cL$$

times where

$$L = \frac{12\beta\|H\|\log(8\beta\|H\|/\epsilon)}{\log[e/(e-1)]}$$

$$a = \lceil \log(1/\Delta) \rceil$$

$$c = \left\lceil \log \left(\frac{96\beta\|H\|\log(8\beta\|H\|/\epsilon)}{\epsilon \log[e/(e-1)]} \right) \right\rceil.$$

Lemma 4. Let H be a Hamiltonian acting on a state space of cardinality d with spectral gap γ . Let Π be the projector onto the eigensubspace corresponding to the minimal eigenvalue. Then

$$|\langle \pi_\beta | \Pi | \pi_\beta \rangle|^2 \geq 1 - \epsilon_3$$

provided that the inverse temperature satisfies

$$\beta \geq \frac{1}{\gamma} \log \left(\frac{(1 - \epsilon_3)d}{\epsilon_3} \right). \quad (2)$$

Proof. The worst case occurs when the ground state $|g\rangle$ is unique and all other states have energy $\gamma + E(g)$ where $E(g)$ is the ground-state energy. In this case the probability of obtaining the ground state when measuring $|\pi_\beta\rangle$ in the computational basis is

$$p = \frac{1}{(d-1)e^{-\gamma\beta} + 1}.$$

To make this probability greater or equal to $1 - \epsilon_3$, it suffices to choose the inverse temperature β as in Eq. (2).

By choosing $\epsilon=1/4$ in Corollary 3 and $\epsilon_3=1/4$ in Lemma 4 we obtain the following corollary.

Corollary 4. There is a quantum algorithm that outputs a ground state of H with probability greater than $1/2$. It invokes the operators from $\{W_{\beta_i}; i=1, \dots, r\}$ at most

$$\frac{1}{\Delta} \frac{\|H\|}{\gamma} \log d \log \left(\frac{\|H\|}{\gamma} \log d \right) \quad (3)$$

times.

Let us explain how the above algorithm differs from the algorithm based on the quantum Zeno effect [12]. That algorithm has the running time

$$\frac{1}{\Delta'} \left(\frac{\|H\|}{\gamma} \right)^2 \log^3 d, \quad (4)$$

where Δ' is the phase gap corresponding to the minimal spectral gap in the sequence of Markov chains.

Both algorithms make use of Szegedy's quantum walk operators to obtain a speedup over the classical case due to the quadratic relation between phase gaps and spectral gaps. The reduction from $(\|H\|/\gamma)^2$ in (4) to $\|H\|/\gamma$ in (3) is due to the advantage of amplitude amplification over the quantum Zeno effect. Note also that $\Delta' \leq \Delta$ because the change in temperature between adjacent Markov chains is $\Delta\beta' = O(\gamma/(\|H\|^2 \log d))$ in [12] and $\Delta\beta=1/\|H\|$ in our algorithm. Roughly speaking, amplitude amplification makes it possible to make "bigger" jumps (i.e., bigger changes in temperature) than the quantum Zeno effect, without decreasing the success probability.

VI. CONCLUSIONS AND DISCUSSIONS

We have presented a simple quantum algorithm, making it possible to prepare quantum samples of stationary distributions of arbitrary slowly varying Markov chains. It significantly improves upon a previous algorithm for that purpose based on adiabatic generation [10]. It also provides a speedup over a recently proposed method for preparing quantum samples of Boltzmann-Gibbs distributions of classical Hamiltonians [12].

ACKNOWLEDGMENTS

The authors would like to thank Chen-Fu Chiang for helpful comments. P. W. and A. A. gratefully acknowledge the support by NSF Grant Nos. CCF-0726771 and CCF-0746600.

[1] M. Jerrum, A. Sinclair, and E. Vigoda, *J. ACM* **51**, 671 (2004).
 [2] M. Dyer, A. Frieze, and R. Kannan, *J. ACM* **38**, 1 (1991).
 [3] M. Jerrum and A. Sinclair, *SIAM J. Comput.* **18**, 1149 (1989).
 [4] M. Jerrum and A. Sinclair, *SIAM J. Comput.* **22**, 1087 (1993).

[5] A. Sinclair, *Algorithms for Generation and Counting—A Markov Chain Approach* (Birkhauser, Boston, 1992).
 [6] M. Jerrum, *Counting, Sampling and Integrating: Algorithms and Complexity*, Lectures in Mathematics (Birkhauser, Boston, 2003).
 [7] M. Szegedy, in *Proceedings of the 45th Annual IEEE*

- Symposium on Foundations of Computer Science* (IEEE, New York, 2004), pp. 32–41.
- [8] F. Magniez, A. Nayak, J. Roland, and M. Santha, in *Proceedings of the 39th Annual ACM Symposium on Theory of Computing* (ACM, Washington, DC, 2007), pp. 575–584.
- [9] P. C. Richter, *Phys. Rev. A* **76**, 042306 (2007).
- [10] D. Aharonov and A. Ta-Shma (2003), *Proc. of the 35th Annual ACM Symposium on Theory*, pp. 20–29.
- [11] R. Somma, S. Boixo, and H. Barnum, e-print arXiv:0712.1008.
- [12] R. Somma, S. Boixo, H. Barnum, and E. Knill, e-print arXiv:0804.1571.
- [13] L. Grover, e-print arXiv:quant-ph/0503205.
- [14] P. Hoyer, *Lect. Notes Comput. Sci.* **3827**, 308 (2005).
- [15] M. Boyer, G. Brassard, P. Hoyer, and A. Tapp, *Fortschr. Phys.* **46**, 493 (1998).
- [16] M. Mitzenmacher and E. Upfal, *Probability and Computing—Randomized Algorithms and Probabilistic Analysis* (Cambridge University Press, Cambridge, England, 2005).