Quantum Simulations of Classical Annealing Processes

R. D. Somma,[1,*] S. Boixo,[2,3] H. Barnum,[2] and E. Knill[4]

[1]*Perimeter Institute for Theoretical Physics, Waterloo, ON N2L 2Y5, Canada*
[2]*CCS-3 (Information Sciences), Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*
[3]*Department of Physics and Astronomy, University of New Mexico, Albuquerque, New Mexico 87131, USA*
[4]*National Institute of Standards and Technology, Boulder, Colorado 80305, USA*

(Received 8 April 2008; published 26 September 2008)

We describe a quantum algorithm that solves combinatorial optimization problems by quantum simulation of a classical simulated annealing process. Our algorithm exploits quantum walks and the quantum Zeno effect induced by evolution randomization. It requires order $1/\sqrt{\delta}$ steps to find an optimal solution with bounded error probability, where δ is the minimum spectral gap of the stochastic matrices used in the classical annealing process. This is a quadratic improvement over the order $1/\delta$ steps required by the latter.

DOI: 10.1103/PhysRevLett.101.130504 PACS numbers: 03.67.Ac, 87.10.Rt, 87.55.de, 89.70.Eg

Combinatorial optimization problems (COPs) are important in almost every branch of science, from computer science to statistical physics and computational biology [1]. Each instance of a COP requires that we minimize some objective function over a search space consisting of d *configurations*. The search space may have additional structure, such as that provided by a graph, to give a notion of locality. Because d is typically exponential in the size of the problem instance, finding a solution by exhaustive search is hard in general. One can exploit the notion of locality to find solutions more quickly, but the presence of many nonoptimal local minima often prevents efficient convergence to a solution. Therefore, more efficient optimization strategies are desirable.

A well-known and often used general strategy for solving COPs is simulated annealing (SA) [2]. SA imitates the process undergone by a metal that is heated to a high temperature and then cooled slowly enough for thermal excitations to prevent it from getting stuck in local minima so that it ends up in one of its lowest-energy configurations. In SA, the objective function E of the COP plays the role of the energy, so the lowest-energy configuration is the optimum. The annealing process can be simulated with a variety of techniques. Here, we focus on discrete Markov chain Monte Carlo (MCMC) simulations as used, for example, in statistical physics [3]. MCMC simulations generate a stochastic sequence of configurations via a Markov process that, in the case of SA, converges to the Gibbs distribution at a low final temperature. More specifically, the annealing process is determined by a choice of an *annealing schedule* consisting of a finite increasing sequence of *inverse temperatures* $\beta_1 < \beta_2 < \ldots < \beta_P$, and by an associated sequence of *transition rules* $\{M_1, \ldots, M_P\}$ consisting of stochastic matrices acting on configurations. When the structure of the problem can be exploited by a good choice of transition rules, the MCMC algorithm can outperform exhaustive search.

One way to characterize the implementation complexity of SA based on MCMC simulations is to count the number of times that the transition rules must be applied before converging to the desired final distribution within an acceptable error. For simplicity, we consider regular annealing schedules with $\beta_k = (k-1)\Delta\beta$ and choose $\Delta\beta = \mathcal{O}(\delta/E_M)$, where δ is the minimum spectral gap of the matrices M_k and $E_M = \max_\sigma |E[\sigma]|$. We assume that E has been shifted so that $E \geq 0$. Let γ be the spectral gap of E, defined as the difference between the two smallest values in the range of E. By adapting arguments from Ref. [4] to the discrete-time setting, it can be shown that, for a success probability greater than $1-\epsilon$, the implementation complexity of SA is given by $\mathcal{N}_{\mathrm{SA}} = P = \mathcal{O}[\frac{E_M}{\gamma} \times \log(d/\epsilon^2)/\delta]$.

Ideally, $\mathcal{N}_{\mathrm{SA}}$ is small compared to the size d of the configuration space. Since problem instance sizes are typically polylogarithmic in d, $\mathcal{N}_{\mathrm{SA}} = \mathcal{O}[\mathrm{polylog}(d)]$ is considered efficient. Efficient $\mathcal{N}_{\mathrm{SA}}$ is obtained, for example, when computing physical properties of the N-spin ferromagnetic Ising model in an homogeneous external field [5]. However, inefficient $\mathcal{N}_{\mathrm{SA}}$ is obtained if the external field is random [6], making the problem intractable due to gaps δ that are exponentially small in N. The dependence of the complexity of MCMC on δ^{-1} is characteristic of Markov processes and may be unavoidable [7]. Thus, new methods with better scaling in δ are desirable.

Quantum mechanics provides new resources with which to attack optimization problems [8,9]. Quantum computers (QCs) can theoretically solve some problems, including integer number factorization and unstructured search, more efficiently than classical computers [10]. Still, whether a QC could solve all COPs more efficiently than is possible with classical computers is an open question. In this Letter, we show that QCs can speed up the simulation of classical annealing processes. We present a method for transforming instances of MCMC-based SA into a quantum simulated

 © 2008 The American Physical Society

annealing (QSA) algorithm for which the number of times, $\mathcal{N}_{\mathrm{QSA}}$, that the transition rules are used is $O((E_M/\gamma)^2 \log^2(d/\epsilon) \log d/(\epsilon\sqrt{\delta}))$, a quadratic improvement as a function of δ^{-1}. This improvement is most significant for hard instances where $\delta \ll 1$. The dependence on $1/\epsilon$ can be improved to polylog$(1/\epsilon)$. QSA is based on ideas and techniques from quantum walks [11] and the quantum Zeno effect, where the latter can be implemented by phase estimation or by randomization of an evolution period.

This Letter is organized as follows. First, we describe a "quantization" of a reversible, ergodic Markov chain in terms of a bipartite quantum walk. This is a similarity-transformed version of the quantum walk used in Refs. [11] to obtain quantum speedups in search problems. Second, we describe how to transform an instance of SA by adapting the annealing schedule and applying the Markov chain quantization. Finally, we analyze the complexity of QSA to determine the speedup over SA.

Quantum Walks and Markov Chains.—Discrete-time quantum walks were introduced as the quantum analogues of classical random walks [12]. We focus on the bipartite quantum walks defined in Refs. [11].

Consider a d-configuration classical system $\mathcal{S}$ with energies $E[\sigma]$ for configurations σ. Consider an ergodic, reversible Markov process on $\mathcal{S}$ with transition probabilities $p(\sigma'|\sigma) = m_{\sigma\sigma'}$ and stationary distribution π^σ. Reversibility is equivalent to the detailed balance condition $\pi^\sigma m_{\sigma\sigma'} = \pi^{\sigma'} m_{\sigma'\sigma}$. Let $\mathcal{H}$ be the quantum state space spanned by orthonormal states $|\sigma\rangle$ for configurations σ of $\mathcal{S}$. In SA, $\pi^\sigma = e^{-\beta E[\sigma]}/\mathcal{Z}$, with $\mathcal{Z} = \sum_\sigma e^{-\beta E[\sigma]}$, is the Gibbs distribution at some inverse temperature β. We assume not only a classical algorithm to efficiently sample from the distribution $m_{\sigma\sigma'}$ given σ, but an efficient quantum algorithm that computes the transformation defined by $|\sigma\rangle|\mathfrak{o}\rangle \mapsto |\sigma\rangle \sum_{\sigma'} \sqrt{m_{\sigma\sigma'}}\,|\sigma'\rangle$, with $|\mathfrak{o}\rangle$ an efficiently preparable state of $\mathcal{H}$ (e.g., a computational basis state). This transformation, which runs efficiently on arbitrary superpositions, can be constructed from an efficient classical algorithm that on input σ computes the list of nonzero $m_{\sigma\sigma'}$, to high precision, where the length of the list is polynomially bounded. This is usually available for MCMC algorithms. (The error from finite-precision approximation of $m_{\sigma\sigma'}$ and $\sqrt{m_{\sigma\sigma'}}$ is insignificant compared to other sources of error in both classical MCMC and our quantum SA algorithm.)

The bipartite quantum walk is defined on the tensor product $\mathcal{H}_A \otimes \mathcal{H}_B$ of two copies of $\mathcal{H}$. Following [11], we define isometries X and Y that map states of $\mathcal{H}$ to states of $\mathcal{H}_A \otimes \mathcal{H}_B$ by

$$X|\sigma\rangle = \sum_{\sigma'} \sqrt{m_{\sigma\sigma'}}\,|\sigma\sigma'\rangle, \quad Y|\sigma'\rangle = \sum_{\sigma} \sqrt{m_{\sigma'\sigma}}\,|\sigma\sigma'\rangle. \quad (1)$$

Let D_π be the diagonal matrix with entries π^σ on the diagonal. Let M be the matrix with entries $M_{\sigma'\sigma} = m_{\sigma\sigma'}$. From the detailed balance condition, $X^\dagger Y = D_\pi^{1/2} M D_\pi^{-1/2}$ is symmetric. It follows that $X^\dagger Y$ and M have the same eigenvalues $\lambda_0 = 1 > \lambda_1 \geq \ldots \geq \lambda_{d-1} \geq 0$. Let $|\phi_j\rangle$ be the λ_j eigenstate of $X^\dagger Y$. Then, $|\phi_0\rangle = \sum_\sigma \sqrt{\pi^\sigma}\,|\sigma\rangle$, which upon measurement in the basis $|\sigma\rangle$ has the same probability distribution as the stationary distribution of the Markov process.

Define unitary operators U_X and U_Y by

$$U_X|\sigma\rangle|\mathfrak{o}\rangle \equiv X|\sigma\rangle, \qquad U_Y|\mathfrak{o}\rangle|\sigma\rangle \equiv Y|\sigma\rangle, \quad (2)$$

with arbitrary action on other states. Let P_1 and P_2 be the projectors onto the subspaces spanned by $\{|\sigma\rangle|\mathfrak{o}\rangle\}_\sigma$ and $\{U_X^\dagger U_Y|\mathfrak{o}\rangle|\sigma\rangle\}_\sigma$, respectively. The reflection operators through P_i are defined by $R_i = 2P_i - \mathbb{1}$. A step of the bipartite quantum walk W based on M is given by $W = R_2 R_1$. This walk is related to the one used in Ref. [11] by a unitary, but π^σ-dependent, similarity transformation, which helps avoid amplitude leakage when W changes in QSA.

We now relate the spectra of the quantum walk W and the Markov chain M [11]. Define phases $\varphi_j = \arccos\lambda_j$ so that $X^\dagger Y|\phi_j\rangle = \cos\varphi_j|\phi_j\rangle$. The spectral gap of M is $\delta = 1 - \lambda_1 \leq (\varphi_1)^2/2$. From Eq. (2),

$$P_1 U_X^\dagger U_Y|\mathfrak{o}\rangle|\phi_j\rangle = \cos\varphi_j|\phi_j\rangle|\mathfrak{o}\rangle \quad (3)$$

$$P_2|\phi_j\rangle|\mathfrak{o}\rangle = \cos\varphi_j U_X^\dagger U_Y|\mathfrak{o}\rangle|\phi_j\rangle, \quad (4)$$

so W preserves the (at most) two-dimensional subspace spanned by $\{|\phi_j\rangle|\mathfrak{o}\rangle, U_X^\dagger U_Y|\mathfrak{o}\rangle|\phi_j\rangle\}$. On the Bloch sphere defined by states in this subspace, for $j \geq 1$, W acts as a $4\varphi_j$ rotation around an axis perpendicular to the defining states [13]. Thus, the eigenphases of W in this subspace are $\pm 2\varphi_j$. The eigenphase-0 states are either the *quantum stationary state* $|\psi_0\rangle = |\phi_0\rangle|\mathfrak{o}\rangle$ or orthogonal to both P_i. The goal is to prepare $|\psi_0\rangle$ so that we can sample from the stationary distribution of M by measuring the first system [14].

Note that $R_2 = U_X^\dagger U_Y S_{AB} R_1 S_{AB} U_Y^\dagger U_X$, with S_{AB} the swap operation. Therefore, W can be implemented as a sequence of quantum steps whose complexity is related to that of the MCMC steps using M, given our assumptions. For asymptotic comparison, it therefore suffices to compare the number of uses of W in the quantum algorithm to the number of uses of M in the classical algorithm.

Quantum simulated annealing.—We assume that for any $\beta \geq 0$, there is a transition matrix M_β satisfying the assumptions of the previous section and with stationary distribution $\pi_\beta^\sigma = e^{-\beta E[\sigma]}/\mathcal{Z}$. Like SA, QSA is based on an annealing schedule that we choose to consist of equally spaced inverse temperatures $\beta_k = (k-1)\Delta\beta$ for $k = 1, \ldots, Q$. Let W_k be the quantum walk step operator for M_{β_k}, $|\psi_0^k\rangle$ its quantum stationary state (*quantum Gibbs state* for β_k), and $\varphi_{1,k}$ its phase gap. The goal of QSA is to sequentially prepare $|\psi_0^{k+1}\rangle$ from $|\psi_0^k\rangle$ by means of an approximate projective measurement onto $|\psi_0^{k+1}\rangle$ [15]

realized by a simulated measurement onto the eigenbasis of W_{k+1}. The uniform superposition $|\psi_0^1\rangle$ can be prepared efficiently (e.g., for n qubits it is done by applying n Hadamard gates). If the states $|\psi_0^k\rangle$ change slowly enough, the state $|\psi_0^Q\rangle$ can be obtained with high probability of success, due to a version of the quantum Zeno effect. If β_Q is sufficiently large, $|\psi_0^Q\rangle$ is a good approximation of a uniform superposition of the ground configurations of $\mathcal{S}$ so that we can obtain such a ground configuration with high probability by measurement. The complexity of QSA is dominated by the complexity of the simulated measurements, for which we give two strategies, one based on the phase estimation algorithm (PEA) and the other on randomized applications of W_k. Both strategies' complexities are dominated by $1/\varphi_{1,k}$. The quadratic quantum speedup is due to the quadratic increase of $\varphi_{1,k}$ over the eigenvalue gap of M_k.

In the following, we describe the implementation of the projections for the Zeno effect in QSA. The use of PEA is depicted in Fig. 1(a). QSA does not need to use the result of the phase estimation, though the result could be used to terminate and restart the procedure if the measurement outcome is not $|0\rangle^{\otimes p}$. The decoherence it induces in the eigenbasis of W_{k+1} suffices to achieve the required Zeno effect. Thus, the effect of the PEA on $\mathcal{H}_A \otimes \mathcal{H}_B$ is equivalent to the one obtained by the action of r W_{k+1}'s, with r chosen uniformly at random from 0 to $2^p - 1$ [Fig. 1 (b)]. To exponentially reduce the error due to remaining coherences between $|\psi_0^{k+1}\rangle$ and orthogonal states, we repeat the random process s times, resulting in a total action

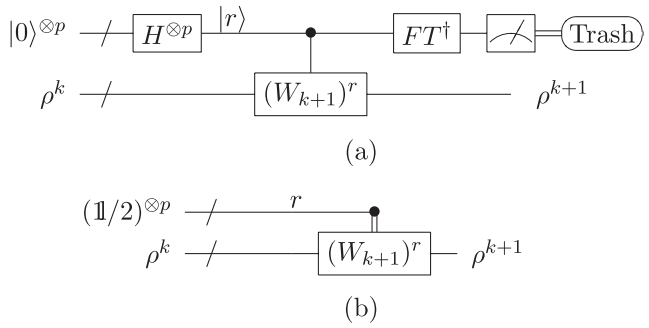

FIG. 1. (a) Phase estimation algorithm. The top p-qubit register encodes a p-bit approximation to an eigenphase of W_{k+1} on readout, and is initialized with Hadamard gates to an equal superposition state. The second register's states are in $\mathcal{H}_A \otimes \mathcal{H}_B$. A sequence of $2^p - 1$ controlled W_{k+1} operations is applied, and the first register is measured after an inverse quantum Fourier transform. If the measurement outcome is $|0\rangle^{\otimes p}$, the second register is approximately projected onto a 0-phase eigenstate of W_{k+1}. (b) Randomization procedure. If the PEA's outcome is ignored, the overall effect on $\mathcal{H}_A \otimes \mathcal{H}_B$ is equivalent to the one induced by initializing a set of p bits (first register) in a random state r, with $r \in \{0, \ldots, 2^p - 1\}$, and by acting on $\mathcal{H}_A \otimes \mathcal{H}_B$ with $(W_{k+1})^r$. Here, double vertical lines indicate classical control.

of $W_{k+1}^{\sum_{q=1}^{s} r_q}$ with $0 \le r_q \le 2^p - 1$ independently random. To prevent excessive amplitude leakage into undesirable 0-eigenphase eigenstates of W_k, we decohere the second register after each randomization step. That is, we measure $\mathcal{H}_B$ in the computational basis and discard the result. The total complexity of QSA is given by $\mathcal{O}(Q2^p s)$ walk steps.

We now choose Q, p, and s to ensure sufficiently high probability of success. Let ρ^k denote the state after the kth randomization and decoherence step. We have $\rho^1 = |\psi_0^1\rangle\langle\psi_0^1|$. Assume that $|\langle\psi_0^{k+1}|\psi_0^k\rangle|^2 \ge 1 - \mu^2$ for all k. By expanding to lowest order in $\Delta\beta$, one can verify that $\mu = \mathcal{O}(\Delta\beta E_M)$. We show by induction that for $2^p > 2^3\pi/\sqrt{2\delta}$ and $s \ge 1 + \log_2(2k)/2 = \mathcal{O}[\log(k)]$, $\langle\psi_0^k|\rho^k|\psi_0^k\rangle \ge 1 - 2k\mu^2$. Thus, if $\mu^2 < \epsilon/(4Q)$, ρ^Q is the quantum Gibbs state for $\beta = Q\Delta\beta$ with error probability at most $\epsilon/2$. We can write $\rho^k = (1-\chi)|\psi_0^k\rangle\langle\psi_0^k| + \nu_k(|\psi_0^k\rangle\langle\psi_\perp^k| + \text{H.c.}) + \chi\rho_\perp$, where $|\psi_\perp^k\rangle$ is a unit state orthogonal to $|\psi_0^k\rangle$, $\rho_\perp$ is a density matrix with support orthogonal to $|\psi_0^k\rangle$, and $\chi \le 2k\mu^2$. To make the induction argument possible, we add the induction hypothesis $\nu_k < \mu/2$. The induction hypotheses apply to ρ^1 by definition. Note that $\langle\psi_0^{k+1}|\rho^{k+1}|\psi_0^{k+1}\rangle = \langle\psi_0^{k+1}|\rho^k|\psi_0^{k+1}\rangle$. We can estimate $\langle\psi_0^{k+1}|\rho^k|\psi_0^{k+1}\rangle \ge (1-\chi)|\langle\psi_0^{k+1}|\psi_0^k\rangle|^2 - 2\nu_k|\langle\psi_0^{k+1}|\psi_0^k\rangle||\langle\psi_0^{k+1}|\psi_\perp^k\rangle| + \langle\psi_0^{k+1}|\rho_\perp|\psi^{k+1}\rangle \ge (1 - 2k\mu^2)(1-\mu^2) - 2\nu_k\mu \ge 1 - 2(k+1)\mu^2$. This establishes the main induction hypothesis for $k+1$. Before the randomization step, the density matrix's transition between $|\psi_0^{k+1}\rangle$ and the orthogonal subspace can be written in the form $\nu'(|\psi_0^{k+1}\rangle\langle\phi_\perp| + \text{H.c.})$ with unit state $|\phi_\perp\rangle$ orthogonal to $|\psi_0^{k+1}\rangle$ and the other 0-eigenphase eigenstates of W^{k+1}, because the decoherence step ensures that the support of P_1 is preserved by the operator ρ^k. The estimate on $\langle\psi_0^{k+1}|\rho^k|\psi_0^{k+1}\rangle$ implies that $\nu' \le \sqrt{2(k+1)}\mu$ by positivity of ρ^k [16]. Because $|\psi_0^{k+1}\rangle$ is stabilized by W_{k+1}, the transition is transformed by randomization to $\nu''(\langle\psi_0^{k+1}|\phi_\perp'\rangle + \text{H.c.})$ with $\nu''|\phi_\perp'\rangle = \nu'(\frac{1}{2^p}\sum_{r=0}^{2^p-1} W_{k+1}^r)^s|\phi_\perp\rangle$. In the eigenbasis of W_{k+1}, the entries of $|\phi_\perp\rangle$ are multiplied by terms with absolute values $(\frac{1}{2^p}|\sum_{r=0}^{2^p-1} e^{ir2\varphi}|)^s \le (\frac{1}{2^{p-1}|1-e^{i2\varphi}|})^s < (\frac{\pi}{2^{p-3}|\varphi|})^s < 2^{-s}$, since the relevant eigenphases 2φ satisfy $\pi/2 \ge |\varphi| \ge \sqrt{2\delta}$. Thus, the choice $s = 1 + \log_2[2(k+1)]/2$ ensures that $\nu'' < \mu/2$. Because the decoherence step preserves $|\psi_0^{k+1}\rangle$, we have $\nu_{k+1} \le \nu'' < \mu/2$. This completes the induction step of the proof. We summarize the QSA in Fig. 2.

To determine the order of the number of quantum steps $\mathcal{N}_{\text{QSA}}$ required by QSA, let β_f be the desired final inverse temperature so that $\Delta\beta = \beta_f/Q$. Choose Q to be a sufficiently large multiple of $\beta_f^2 E_m^2/\epsilon$. For optimization, we let $\beta_f = \ln[d/(2\epsilon)]/\gamma = \mathcal{O}[\log(d/\epsilon)/\gamma]$. According to the bounds at the beginning of the previous paragraph, this ensures that after measuring the final state, the probability of finding a nonoptimal configuration is at most ϵ, with a contribution of $\epsilon/2$ from the probability of being orthogonal to $|\psi_0^Q\rangle$ and $\epsilon/2$ from the Gibbs distribution's proba-

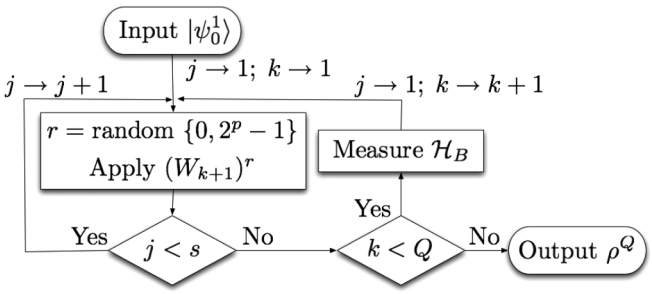

FIG. 2. Flow diagram for the QSA. The parameters r, s, and Q are chosen so that the final state ρ^Q is sufficiently close to the desired state $|\psi_0^Q\rangle$ (see text).

bility of not being optimal. Because $2^p = \mathcal{O}(1/\sqrt{\delta})$ and $s = \mathcal{O}[\log(Q)]$, we find that $\mathcal{N}_{\mathrm{QSA}} = \mathcal{O}[Q\log(Q)/\sqrt{\delta}]$ with $Q = \mathcal{O}(\beta_f^2 E_m^2/\epsilon)$ and $\beta_f = \mathcal{O}[\log(d/\epsilon)/\gamma]$. If we anticipate that $Q > d$, we can just search every configuration classically to find the optima, so we can bound $\log(Q) \leq \log(d)$ to simplify

$$\mathcal{N}_{\mathrm{QSA}} = \mathcal{O}\left[\left(\frac{E_M}{\gamma}\right)^2 \frac{\log^2(d/\epsilon)\log d}{\epsilon\sqrt{\delta}}\right]. \tag{5}$$

The dependence of $\mathcal{N}_{\mathrm{QSA}}$ on $1/\epsilon$ can be improved to polylog$(1/\epsilon)$ by repetition of QSA with an initial target error $\epsilon = 1/2$ in Eq. (5). For optimization, it suffices to repeat QSA $\mathcal{O}[\log(\epsilon)]$ many times. A different approach to prepare the desired stationary state with high probability of success is to apply a high-confidence version of the PEA [13] at the end of QSA to project onto $|\psi_0^Q\rangle$. If the projection fails, the algorithm is repeated.

Although the dependence of $\mathcal{N}_{\mathrm{QSA}}$ on E_M/γ is worse than the one appearing in classical SA, it is worth noting that unlike the inverse spectral gap $1/\delta$, in many important applications, this parameter is bounded by a constant or a polynomial in instance size.

Conclusions.—We presented a quantum algorithm based on a "quantization" of simulated annealing algorithms implemented with MCMC methods. This quantum simulated annealing (QSA) algorithm forces the state to closely follow a superposition with amplitudes derived from finite-temperature Gibbs distributions. This is accomplished by either an explicit measurement using phase estimation with quantum walk operators, or by decoherence using random applications of these operators. QSA can be used both for combinatorial optimization and for sampling from a Gibbs distribution for statistical physics applications. In contrast to SA, which scales with $\mathcal{O}(1/\delta)$, where δ is the minimal spectral gap of the transition matrices, QSA scales with $\mathcal{O}(1/\sqrt{\delta})$. Although in general the QSA does not yield a polynomial-resource algorithm, it reduces required resources by an asymptotic exponential factor for the ubiquitous hard cases, where the gap becomes exponentially small in the problem size.

We thank S. Jordan for discussions, and B. Eastin and D. Hume for their careful reading of the manuscript. This research was supported by Perimeter Institute for Theoretical Physics, by the Government of Canada through Industry Canada, and by the Province of Ontario through the Ministry of Research and Innovation. It was also carried out under the auspices of the NNSA of the US DOE at LANL under Contract No. DE-AC52-06NA25396, and with support from NSF Grant No. PHY-0653596. Contributions to this work by NIST, an agency of the US government, are not subject to copyright laws.

*rsomma@perimeterinstitute.ca

[1] W. J. Cook, W. H. Cunningham, W. R. Pulleyblank, and A. Schrijver, *Combinatorial Optimization* (J. Wiley and Sons, New York, 1998).
[2] S. Kirkpatrick, C. D. Gelett, and M. P. Vecchi, Science **220**, 671 (1983).
[3] M. E. J. Newman and G. T. Barkema, *Monte Carlo Methods in Statistical Physics* (Oxford University Press, Oxford, UK, 1999).
[4] D. W. Stroock, *An Introduction to Markov Processes* (Springer-Verlag, Berlin, 2005).
[5] M. Jerrum and A. Sinclair, SIAM J. Comput. **22**, 1087 (1993).
[6] F. Barahona, J. Phys. A **15**, 3241 (1982).
[7] D. J. Aldous, J. Lond. Math. Soc. **s2-25**, 564 (1982).
[8] B. Apolloni, N. Cesa-Bianchi, and D. de Falco, *Stochastic Processes, Physics and Geometry* (World Scientific, Singapore, 1990); T. Kadowaki and H. Nishimori, Phys. Rev. E **58**, 5355 (1998); E. Farhi *et al.*, Science **292**, 472 (2001); G. E. Santoro *et al.*, Science **295**, 2427 (2002); G. E. Santoro and E. Tosatti, Nature Phys. **3**, 593 (2007).
[9] R. D. Somma, C. D. Batista, and G. Ortiz, Phys. Rev. Lett. **99**, 030603 (2007).
[10] P. Shor, *Proceedings of the 35th Annual Symp. Found. Comp. Science* (1994) p. 116; L. K. Grover, *Proceedings of the 28th Annual ACM Symp. on the Th. Comp.* (1996), p. 212.
[11] A. Ambainis, *Proceedings of the 45th Symposium on Foundations of Computer Science* (2004), p. 22; M. Szegedy, *Proceedings of the 45th IEEE Symposium on Foundations of Computer Science* (2004), p. 32; F. Magniez, A. Nayak, J. Roland, and M. Santha, *Proceedings of the 39th Annual ACM Symposium on Theory of Computing* (2007), p. 575.
[12] A. Ambainis *et al.*, *Proceedings of the 33th Annual ACM Symposium on Theory of Computing* (2001), p. 37; see also J. Kempe, Contemp. Phys. **44**, 307 (2003), for a review and references therein.
[13] E. Knill, G. Ortiz, and R. Somma, Phys. Rev. A **75**, 012328 (2007).
[14] D. Aharonov and A. Ta-Shma, *Proceedings of the 35th Annual Symposium on the Theory of Computation (STOC)* (2003), p. 20–29; SIAM J. Comput. **37**, 47 (2007).
[15] A. M. Childs *et al.*, Phys. Rev. A **66**, 032314 (2002).
[16] For a density matrix ρ, $|\rho_{lm}| \leq \sqrt{\rho_{ll}\rho_{mm}}$, where $\rho_{ll'} := \langle l|\rho|l'\rangle$.