

# Sampling from the Thermal Quantum Gibbs State and Evaluating Partition Functions with a Quantum Computer

David Poulin<sup>1</sup> and Pawel Wocjan<sup>2</sup>

<sup>1</sup>*Département de Physique, Université de Sherbrooke, Québec, Canada, J1K 2R1*

<sup>2</sup>*School of Electrical Engineering and Computer Science, University of Central Florida, Florida 32816-2362, USA*

(Received 15 June 2009; published 24 November 2009)

We present a quantum algorithm to prepare the thermal Gibbs state of interacting quantum systems. This algorithm sets a universal upper bound  $D^\alpha$  on the thermalization time of a quantum system, where  $D$  is the system's Hilbert space dimension and  $\alpha \leq \frac{1}{2}$  is proportional to the Helmholtz free energy density. We also derive an algorithm to evaluate the partition function of a quantum system in a time proportional to the system's thermalization time and inversely proportional to the targeted accuracy squared.

DOI: 10.1103/PhysRevLett.103.220502

PACS numbers: 03.67.Ac, 24.10.Cn, 64.70.qd

The partition function at inverse temperature  $\beta$  is of central interest in statistical physics because all thermodynamic variables such as energy, heat capacity, entropy, and correlation functions can be derived from it. For interacting systems, the partition function involves a sum over an exponential number of terms, so it cannot be computed exactly except in few special cases. Thus, approximation schemes are required to estimate these physical quantities numerically.

Monte Carlo simulations are often the method of choice to evaluate thermodynamic quantities. When the system is classical, system configurations  $x$  can be generated randomly according to the thermal distribution  $\Pr(x) = \frac{e^{-\beta H(x)}}{Z(\beta)}$  by a suitable Markov chain. Sampling from this distribution provides an estimate of the partition function and other quantities of interest. The running time of the algorithm is determined by the rate of convergence of the Markov chain and the number of samples needed to lower the statistical fluctuations below a desired accuracy. The method of Ref. [1] uses a quantum computer to prepare a purification of the thermal distribution in time that scales like the inverse of the square root of the spectral gap. Building on this work, one of us has recently shown that a quantum computer also achieves a quadratic speed-up with respect to relative accuracy for evaluating partition functions [2].

Monte Carlo simulations can also be employed for quantum systems using standard mapping between  $N$ -dimensional quantum systems to  $(N + 1)$ -dimensional classical systems. However, frustrated and fermionic systems are affected by the sign problem [3] which leads to an exponential blowup of the statistical fluctuations of the sample. For this reason, Monte Carlo simulations are unreliable for these systems.

This Letter presents two complementary results. First, we demonstrate a quantum algorithm to prepare the thermal Gibbs state  $\rho(\beta) = \frac{e^{-\beta H}}{Z(\beta)}$  of any locally interacting quantum system, where  $Z(\beta) = \text{Tr}(e^{-\beta H})$  is the partition function. The problem of thermalizing a quantum system with a quantum computer was first raised in [4] where heu-

ristic methods were proposed, but no rigorous upper bounds have been derived. Our results set a universal upper bound  $D^\alpha$  for the thermalization time of a quantum system, where  $D$  is the system's Hilbert space dimension and  $\alpha \leq \frac{1}{2}$  is proportional to the Helmholtz free energy density.

Second, we present a quantum algorithm to evaluate the partition function of quantum systems within *relative* accuracy  $\epsilon$ . The complexity of the algorithm scales like the inverse of  $\epsilon^2$  ( $\epsilon$  when the system is classical) and proportionally to the thermalization time of the system. This algorithm is not affected by the sign problem and works for any locally interacting quantum system. It is a full quantum generalization of the algorithm presented in [1,2] for the evaluation of partitions functions of classical systems. The combination of our two algorithms yields a universal upper bound for the time needed to evaluate a partition function with a quantum computer.

We consider a system composed of  $n$  interacting  $d$ -level particles with a  $k$ -local Hamiltonian  $H = \sum_j h_j$ , where each term  $h_j$  has a bounded norm and acts on at most  $k$  particles ( $k$  is a constant). It is one of quantum information science's most celebrated result that  $k$ -local Hamiltonians can be simulated efficiently on a quantum computer. More precisely, for those Hamiltonians, the time evolution operator  $U(t) = e^{-iHt}$  can be approximately expressed as the product of a sequence of discrete one- and two-qubit gates, and the number of gates scales essentially linearly with the duration  $t$  of the simulated process and polynomially with the system size  $n$  [5–8].

Given the system's Hamiltonian and a temperature  $1/\beta$ , we call a *thermalization process* any (possibly time-dependent) local Hamiltonian on the system and a bath (of size scaling polynomially with the number of particles in the system) that generates, after time  $T$ , a unitary transformation  $V$  such that

$$\text{Tr}_{\text{bath}} \left\{ V \left( \frac{I}{D} \otimes |0\rangle\langle 0|_{\text{bath}} \right) V^\dagger \right\} = \rho(\beta), \quad (1)$$

or more generally that this equation holds within some accuracy  $\delta$ . The system's thermalization time  $T_{\text{th}}(\beta)$  is

the duration of the shortest thermalization process. Note that in our definition, the temperature dependence is encoded in the details of  $V$  instead of in the initial state of the bath as one would expect physically. This is to rule out trivial thermalization processes where the bath is prepared in the desired state  $\rho(\beta)$  and  $V$  simply swaps the system and the bath. Nevertheless, the temperature dependence of the thermalization process we will present could equally well be encoded in the initial state of the bath with a simple modification to our algorithm.

Thermalization processes that occur in nature constitute of local interactions (all forces of nature are two-body), so can be efficiently simulated by a quantum computer, i.e., in a time proportional to the duration of the thermalization process [5–8]. In what follows, we present a universal thermalization process and characterize its complexity. Then, we show how to use this—or any other—thermalization process to evaluate partition functions with a quantum computer.

For sake of analysis, we shift the system’s Hamiltonian by a constant  $E^*$  to ensure that  $H$  is positive. This changes the partition function by a factor  $e^{-\beta E^*}$  that we keep track of implicitly: we henceforth assume  $0 < H < E_{\max}$  where  $E_{\max}$  is a known (polynomial) upper bound of the Hamiltonian. We denote the eigenvalues and eigenvectors of the Hamiltonian  $H|a\rangle = E_a|a\rangle$ .

We will make extensive use of the system’s time evolution operator with  $t = \pi/(4E_{\max})$ , that we simply denote  $U$ . As mentioned above,  $U$  cannot be simulated exactly with a quantum computer. Instead, we can produce  $\tilde{U}$ , which is a good approximation to it. This will unavoidably limit the accuracy of our thermalization process and of the estimated partition function. Indeed, defining the effective Hamiltonian  $\tilde{H} = -\frac{i}{t} \log \tilde{U}$  (by definition of  $t$ , there is no multivalued problem), our algorithms will prepare the Gibbs state  $\tilde{\rho}(\beta)$  and evaluate the partition function  $\tilde{Z}(\beta)$  associated with  $\tilde{H}$ , not  $H$ . However, we show in Appendix A of the supplementary material [9] that an accurate simulation  $\|U - \tilde{U}\| \leq \epsilon$  leads to an accurate effective Hamiltonian  $\|H - \tilde{H}\| \leq KE_{\max}\epsilon$  (where  $K$  is a small constant). This, in turn, implies an accurate estimate of the partition function  $|\tilde{Z}(\beta) - Z(\beta)| \leq 2K\beta E_{\max}\epsilon Z(\beta)$  (see Appendix B in the supplementary material [9]) and a high fidelity Gibbs state  $F[\rho(\beta), \tilde{\rho}(\beta)] \geq 1 - K'\beta E_{\max}\epsilon$  (see Appendix C in the supplementary material [9]).

Because the exact time required to simulate  $U$  within accuracy  $\epsilon$  depends on the details of the system, we measure the complexity of our algorithms in terms of the number of times they need to implement  $U$ . Hence, the true running time of our algorithms have an additional arbitrarily weak dependence [6] on the targeted accuracy  $\epsilon$  of the partition function and Gibbs state.

Our algorithm requires a cooling schedule  $0 = \beta_0 < \beta_1 < \dots < \beta_\ell = \beta$  such that  $F_k := \frac{Z(\beta_{k+1})}{Z(\beta_k)} \geq 1/2$  for all  $k$ . Writing  $F_k = \sum_a \frac{e^{-\beta_k E_a}}{Z(\beta_k)} e^{-\Delta\beta_k E_a} = \langle e^{-\Delta\beta_k H} \rangle_{\beta_k}$  where  $\Delta\beta_k = \beta_{k+1} - \beta_k$ , we see that a polynomial-length

cooling schedule can always be constructed by choosing  $\Delta\beta_k = \log(2)/E_{\max}$ , which implies a cooling schedule of length  $\ell = E_{\max}\beta/\log 2$ , but shorter schedules can be used in most cases. Notice how the choice of a lower bound on the Hamiltonian  $E^*$  affects these ratios. For this reason, it is desirable that  $-E^*$  be as close as possible to the true ground state energy of the system to decrease the length of the cooling schedule.

The partition function is expressed as a product

$$Z(\beta_k) = Z(\beta_0)F_0F_1\dots F_{k-1} \quad (2)$$

and will be computed by evaluating each fraction and using the fact that  $Z(\beta_0) = D$ , the dimension of the system’s Hilbert space. If each ratio  $F_k$  is evaluated within accuracy  $\frac{\epsilon}{\ell}$ , the resulting relative error on  $Z(\beta)$  will be  $\mathcal{O}(\epsilon)$ . Our universal thermalization process requires evaluating the  $F_k$  sequentially: the fraction  $F_k$  is evaluated from the Gibbs state  $\rho(\beta_k)$ , and the value of  $Z(\beta_k)$  is needed to prepare  $\rho(\beta_{k+1})$ .

Our method makes use of quantum phase estimation (QPE) [10,11]. This is a transformation that operates on the “system” register and an  $m$ -qubit “energy” register initialized to the state  $|0_m\rangle$  as follows:

$$|a\rangle \otimes |0_m\rangle \xrightarrow{W} |a\rangle \otimes \sum_E f(E, E_a)|E\rangle$$

Here,  $W$  denotes the Hadamard transform  $W|0_m\rangle = 2^{-(m/2)} \sum_{r=0}^{2^m-1} |r\rangle$  and FT denotes the Fourier transform. The central gate applies  $r$  repetitions of the time evolution operator  $U$  to the system register, where  $r$  is dictated by the state of the lower register. Since  $r$  can be as large as  $2^m - 1$ , the running time of this procedure is  $2^m$ . The label  $E$  runs over the discretization of the interval  $[0, 8E_{\max})$ , i.e.  $E \in \frac{8E_{\max}}{2^m} \times \{0, 1, \dots, 2^m - 1\}$ . The function  $f(E, E_a)$  is the discrete Fourier transform of  $g(r) = e^{-itE_a r}$ , so it is highly peaked around the value  $E = E_a$ , i.e.,

$$|f(E, E_a)|^2 = \frac{\sin\left(\frac{E-E_a}{E_{\max}} \frac{\pi}{4} 2^m\right)}{\sin\left(\frac{E-E_a}{E_{\max}} \frac{\pi}{4}\right)}.$$

Thus, QPE essentially measures the energy of the system and writes it down on the energy register, up to some fluctuations associated to the width  $f(E, E_a)$ .

Let us first describe our algorithm to prepare Gibbs states by ignoring these fluctuations and return to them later. In fact, our algorithm will prepare *purified marked Gibbs states*

$$|\beta_k\rangle = \sum_a \sqrt{\frac{e^{-\beta_k E_a}}{Z(\beta_k)}} |a\rangle \otimes |\phi_a\rangle \otimes |E_a\rangle \quad (3)$$

that are defined on one system register, one scratchpad register, and one energy register. The term “purified”

refers to the fact that the state  $|\beta_k\rangle$  is pure and has a reduced system density matrix equal to the Gibbs state at inverse temperature  $\beta_k$ . The term “marked” means that the energy register encodes the energy of each system eigenvalue.

To obtain the state  $|\beta_0\rangle$ , we must first prepare the system and scratchpad in any maximally entangled state  $D^{-(1/2)}\sum_a|a\rangle\otimes|\phi_a\rangle$  and then apply QPE to the system register (ignoring fluctuations of  $f$ ). The finite temperature quantum Gibbs state  $|\beta_k\rangle$  is obtained from  $|\beta_0\rangle$  and an ancillary qubit initially in the state  $|0\rangle$  that we rotated by an angle  $\theta(E, \beta_k) = \arcsin(e^{-(\beta_k E/2)})$  conditioned on the energy register, resulting in

$$|\Psi_k\rangle = D^{-(1/2)}\sum_a|a\rangle\otimes|\phi_a\rangle\otimes|E_a\rangle\otimes|\theta(E_a, \beta_k)\rangle \quad (4)$$

$$= \sqrt{\frac{Z(\beta_k)}{D}}|\beta_k\rangle|0\rangle + \dots \quad (5)$$

where  $|\theta\rangle = \sin\theta|0\rangle + \cos\theta|1\rangle$  and the ellipsis represents terms in which the ancilla qubit is in the state  $|1\rangle$ .

We can use Grover’s algorithm [12] to amplify the overlap of  $|\Psi_k\rangle$  with the projector  $\Pi_0 = I \otimes |0\rangle\langle 0|$  associated to the subspace where the ancillary qubit is in the state  $|0\rangle$ . Indeed, applying  $q = \lfloor \sqrt{\frac{D}{Z(\beta_k)}} \rfloor$  times the sequence of two reflections  $(I - 2|\Psi_k\rangle\langle\Psi_k|)(I - 2\Pi_0)$  to the state  $|\Psi_k\rangle$  will boost the amplitude of the term  $|\beta_k\rangle\otimes|0\rangle$  from its initial value  $\sqrt{\frac{Z(\beta_k)}{D}}$  to nearly 1. However,  $q$  cannot be computed exactly because at this stage of the algorithm,  $Z(\beta_k)$  is known only within relative accuracy  $\frac{k\epsilon}{\ell}$ . This is nonetheless sufficient to amplify the amplitude of the term  $|\beta_k\rangle\otimes|0\rangle$  to  $\mathcal{O}(1)$ . An amplitude  $1 - \mathcal{O}(\frac{\epsilon}{\ell})$  can be achieved using fixed-point search [13,14] at an additional multiplicative cost of  $\log\frac{\ell}{\epsilon}$ .

The fluctuations of  $f(E, E_a)$  will in general invalidate the procedure we have described. If we perform QPE with a slightly larger number of qubits—i.e., if we use an  $(m+4)$ -qubit energy register—the probability that the estimated energy  $E$  deviates from its true value  $E_a$  by more than  $2^{-m}/t$  will be less than  $1/16$  [11]. To further suppress these fluctuations, we can perform  $\eta$  independent QPE procedures and compute the median  $M$  of the  $\eta$  results. The probability that this median value deviates from the true energy by more than  $2^{-m}/t$  is less than  $2^{-\eta}$  [15].

The evaluation of the median can be done by a coherent quantum process. There is an efficient quantum circuit on  $\eta$  energy registers and one  $m$ -qubit “median” register that maps  $|E^{(1)}\rangle\otimes|E^{(2)}\rangle\otimes\dots\otimes|E^{(\eta)}\rangle\otimes|0_m\rangle$  to  $|E^{(1)}\rangle\otimes|E^{(2)}\rangle\otimes\dots\otimes|E^{(\eta)}\rangle\otimes|M(E^{(1)}, E^{(2)}, \dots)\rangle$ . From the result stated above, we know that on input eigenstate  $|a\rangle$ , only two median register states  $|E_a^-\rangle$  and  $|E_a^+\rangle$  encoding the closest  $m$ -bit estimates of  $E_a$  can have an amplitude of magnitude larger than  $\sqrt{2^{-\eta}}$ . This can be used to prepare a “good enough” version of the infinite temperature purified marked Gibbs state

$$|\tilde{\beta}_0\rangle = D^{-(1/2)}\sum_a|a\rangle\otimes(\alpha_a|\phi_a^-\rangle|E_a^-\rangle + \beta_a|\phi_a^+\rangle|E_a^+\rangle) + |\text{bad}\rangle, \quad (6)$$

where  $|\alpha_a|^2 + |\beta_a|^2 \geq 1 - 2^{-\eta}$ ,  $|E_a^\pm - E_a| \leq 2^{-m}/t$ , and  $|\text{bad}\rangle$  represent the components where the median deviates from the true energy by more than  $2^{-m}/t$ , so  $\langle\text{bad}|\text{bad}\rangle \leq 2^{-\eta}$ . The scratchpad states  $|\phi_a^\pm\rangle$  contain the original states  $|\phi_a\rangle$  used to purify the infinite temperature Gibbs state and the  $\eta$  energy registers used for the computation of the median, whose value is encoded in the third register.

With  $m = \log_2(\beta\ell/t\epsilon)$  the error caused on the Gibbs state by the roundoff of the energy estimate will be  $\mathcal{O}(\frac{\epsilon}{\ell})$  (see Appendix C in the supplementary material [9]). Similarly, this roundoff error will lead to a  $\mathcal{O}(\frac{\epsilon}{\ell})$  relative error on our estimate of  $F_k$  (see Appendix B in the supplementary material [9]). The error due to the  $|\text{bad}\rangle$  component of Eq. (6) requires a bit more attention. The reason is that the relative weight of the good and the bad component [i.e., first and second term of Eq. (6)] is in general not preserved by Grover’s amplification. Although the norm of  $|\text{bad}\rangle$  starts out small, it can increase by a factor  $e^{\beta_k E_{\max}/2}$  during the amplification process. Setting  $\eta = \lceil \ln(\frac{\ell}{\epsilon}) + \beta_k E_{\max} \rceil / \ln(2)$  insures that the norm of  $|\text{bad}\rangle$  will remain bounded by  $\frac{\epsilon}{\ell}$  after the amplification process, so can safely be ignored. This completes our presentation of the universal thermalization process, and demonstrates that the Gibbs state can be prepared with accuracy  $\delta = \frac{\epsilon}{\ell}$  by a local process in time  $T_{\text{th}}(\beta) \in \mathcal{O}(\sqrt{\frac{D}{Z(\beta)}} \frac{\beta E_{\max}}{\delta} \log\frac{1}{\delta} [\log\frac{1}{\delta} + \beta E_{\max}])$ .

We now describe how the state  $|\beta_k\rangle$  can be used to estimate the ratio  $F_k$ . Similarly to the procedure that led to Eq. (5), we append to the quantum Gibbs state  $|\beta_k\rangle$  yet another ancillary qubit in the state  $|0\rangle$  and rotate it by an angle  $\theta(E, \Delta\beta_{k+1})$  conditioned on the energy register, resulting in

$$|\Phi_k\rangle = \sum_a \sqrt{\frac{e^{-\beta_k E_a}}{Z(\beta_k)}} |a\rangle\otimes|\phi_a\rangle\otimes|E_a\rangle\otimes|\theta(E_a, \Delta\beta_{k+1})\rangle$$

$$= \sqrt{F_k}|\beta_{k+1}\rangle|0\rangle + \dots \quad (7)$$

where the ellipsis represent terms in which the ancilla qubit is in the state  $|1\rangle$ . Clearly, the squared norm of  $\Pi_0|\Phi_k\rangle$  is equal to the quantity  $F_k$  we are trying to estimate. Thus,  $F_k$  can be estimated by quantum counting [16], i.e., using phase estimation to estimate the eigenvalue of the unitary matrix  $(I - 2|\Phi_k\rangle\langle\Phi_k|)(I - 2\Pi_0)$  associated to the Jordan block that supports  $|\Phi_k\rangle$ . An accuracy  $\frac{\epsilon}{\ell}$  requires  $\frac{\ell}{\epsilon}$  uses of the circuit that prepares  $|\Phi_k\rangle$ , so the total time required to evaluate the partition function with accuracy  $\epsilon$  with constant success probability is  $\mathcal{O}(\sqrt{\frac{D}{Z(\beta)}} \frac{\beta^5 E_{\max}^5}{\epsilon^2})$ .

More generally, any thermalization process can be used to evaluate  $F_k$ . Indeed, a purified Gibbs state can be prepared by simulating the thermalization process with a quantum computer, substituting the system’s maximally

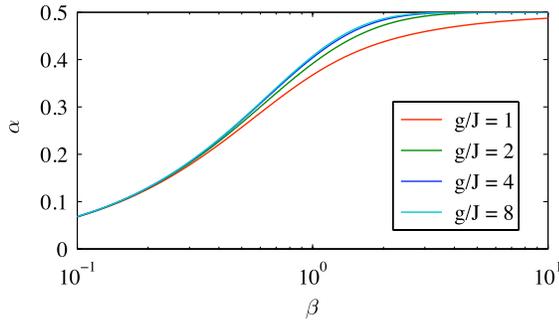


FIG. 1 (color online). Scaling exponent as a function of the inverse temperature for the one dimensional Ising model defined by the Hamiltonian  $H(J, g) = \sum_i g \sigma_x^{(i)} + J \sigma_z^{(i)} \sigma_z^{(i+1)}$  for different values of  $g/J$  ( $g/J = 1$  is critical). The absolute values of  $g$  and  $J$  are fixed by  $\|H(g, J)\| = \|H(0, 1)\|$ . Like the partition function,  $\alpha$  is symmetric under  $J \leftrightarrow g$ .

mixed state  $\frac{1}{D}$  in Eq. (1) by a maximally entangled state  $\frac{1}{\sqrt{D}} \sum_a |a\rangle \otimes |\phi_a\rangle$ . Hence, this requires a time proportional to the thermalization time of the system. This state can be marked using multiple QPEs and evaluating their median as described above. An error at most  $\epsilon$  is achieved with constant probability in time  $\mathcal{O}(T_{\text{th}}(\beta) \frac{\beta^5 E_{\text{max}}^5}{\epsilon^2})$ .

The complexity of our algorithm scales with the system size as  $\sqrt{\frac{D}{Z(\beta)}} = D^\alpha$  where  $D = d^n$  is the Hilbert space dimension and we can express the scaling exponent  $\alpha = \frac{1}{2} h(\beta) / \log(d)$  as a function of the Helmholtz free energy density  $h(\beta)$  (energy is defined relative  $E^*$ ). Figure 1 shows this exponent for the Ising model as a function of inverse temperature. For temperatures of order unity in natural units, this scaling parameter approaches  $\frac{1}{2}$ , which naturally coincides with the running time found in [17] to compute the ground state energy of the system. At higher temperatures however, the algorithm can be significantly faster. Surprisingly, the running time is shortest at the critical field  $g/J = 1$ , where the correlations in the system are the strongest.

For a classical Hamiltonian, the energy of a given configuration can be computed exactly. In that case, our thermalization process produces a Gibbs state with accuracy  $\delta$  in time  $\sqrt{\frac{D}{Z(\beta_k)}} \log \frac{1}{\delta}$ . The only approximation to the partition function involved in that case comes from quantum counting, so  $Z(\beta)$  can be evaluated with accuracy  $\epsilon$  in time  $\mathcal{O}(\sum_k \frac{\ell}{\epsilon} \sqrt{\frac{D}{Z(\beta_k)}})$ . This should be compared to the time  $\mathcal{O}(\sum_k \frac{\ell}{\epsilon} \frac{1}{\sqrt{\delta_k}})$  required by the quantum simulated annealing algorithm [1,2], where  $\delta_k$  is the spectral gap of a Markov chain whose fixed point is the Gibbs state of the system at inverse temperature  $\beta_k$ . Thus, our method offers a speed-up for the computation of a classical partition function whenever no such Markov chain can be found with  $\delta \geq Z(\beta)/D$ .

**Conclusion.**—We have presented a quantum algorithm to prepare the Gibbs state of a quantum system in a time that grows like the  $\alpha$ th power of the system's Hilbert space

dimension, where  $\alpha \leq \frac{1}{2}$  is proportional to the Helmholtz free energy density  $h(\beta)$  at inverse temperature  $\beta$ . This sets a universal upper bound on the thermalization time of quantum systems. However, our universal thermalization process fails to recognize special properties of the system—e.g., large energy gap, absence of long range correlations—that could potentially speed-up the computation. We are actively investigating this problem.

We have demonstrated how the ability to thermalize a quantum system on a quantum computer leads to an algorithm to estimate its partition function with accuracy  $\epsilon$  in a time proportional to the system's thermalization time and  $\epsilon^{-2}$ . Our method provides a complete quantum generalization of the work reported in [1,2] for partitions functions of classical systems and may provide an additional speed-up when there exist no rapidly mixing Markov chain to prepare the system's Gibbs state.

- 
- [1] R. Somma, S. Boixo, H. Barnum, and E. Knill, Phys. Rev. Lett. **101**, 130504 (2008).
  - [2] P. Wocjan, C.-F. Chiang, A. Abeyesinghe, and D. Nagaj, Phys. Rev. A **80**, 022340 (2009).
  - [3] S. Miyashita, M. Takasu, and M. Suzuki, in *Quantum Monte Carlo Methods in Equilibrium and Nonequilibrium Systems*, edited by M. Suzuki (Springer-Verlag, Berlin, 1987).
  - [4] B.M. Terhal and D.P. DiVincenzo, Phys. Rev. A **61**, 022301 (2000).
  - [5] D. Aharonov and A. Ta-Shma, *Proceedings of the 35th Annual ACM Symposium on the Theory of Computation* (ACM, San Diego, 2003), p. 20.
  - [6] D.W. Berry, G. Ahokas, R. Cleve, and B.C. Sanders, Commun. Math. Phys., **270**, 359 (2007).
  - [7] S. Lloyd, Science **273**, 1073 (1996).
  - [8] C. Zalka, Fortschr. Phys. **46**, 877 (1998).
  - [9] See EPAPS Document No. E-PRLTAO-103-016949 for supplementary material. For more information on EPAPS, see <http://www.aip.org/pubservs/epaps.html>.
  - [10] A. Kitaev, arXiv:quant-ph/9511026.
  - [11] M.A. Nielsen and I.L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, England, 2000).
  - [12] L. Grover, in *Proceedings of the 28th Annual ACM Symposium on the Theory of Computation* (ACM Press, New York, 1996), p. 212.
  - [13] L. Grover, arXiv:quant-ph/0503205.
  - [14] P. Høyer, in *The 16th Annual International Symposium on Algorithms and Computation, Sanya, China*, edited by X. Deng and D. Du (2005), Vol. 308.
  - [15] D. Nagaj, P. Wocjan, and Y. Zhang, Quantum Inf. Comput. **9**, 1053, (2009).
  - [16] G. Brassard, P. Hoyer, and A. Tapp, in *Automata, Languages and Programming, Proceedings of ICALP'98*, edited by K.G. Larsen, S. Skyum, and G. Winskel, Lect. Notes Comput. Sci. Vol. 1443 (Springer-Verlag, Berlin, Germany, 1998), p. 820.
  - [17] D. Poulin and P. Wocjan, Phys. Rev. Lett. **102**, 130503 (2009).