

Quantum Metropolis sampling

K. Temme¹, T. J. Osborne², K. G. Vollbrecht³, D. Poulin⁴ & F. Verstraete¹

The original motivation to build a quantum computer came from Feynman¹, who imagined a machine capable of simulating generic quantum mechanical systems—a task that is believed to be intractable for classical computers. Such a machine could have far-reaching applications in the simulation of many-body quantum physics in condensed-matter, chemical and high-energy systems. Part of Feynman’s challenge was met by Lloyd², who showed how to approximately decompose the time evolution operator of interacting quantum particles into a short sequence of elementary gates, suitable for operation on a quantum computer. However, this left open the problem of how to simulate the equilibrium and static properties of quantum systems. This requires the preparation of ground and Gibbs states on a quantum computer. For classical systems, this problem is solved by the ubiquitous Metropolis algorithm³, a method that has basically acquired a monopoly on the simulation of interacting particles. Here we demonstrate how to implement a quantum version of the Metropolis algorithm. This algorithm permits sampling directly from the eigenstates of the Hamiltonian, and thus evades the sign problem present in classical simulations. A small-scale implementation of this algorithm should be achievable with today’s technology.

Since the early days of quantum mechanics, it has been clear that there is a fundamental difficulty in studying many-body quantum systems: the configuration space, or Hilbert space, of a collection of particles grows exponentially with the number of particles. Many of the important breakthroughs in quantum physics during the twentieth century resulted from efforts to address this problem, leading to fundamental theoretical and numerical methods to approximate solutions of the many-body Schrödinger equation. However, most of these methods are limited to weakly interacting particles; unfortunately, it is precisely when the interactions are strong that the most interesting physics arises. Notable examples include high-transition-temperature superconductors, electronic structure in large molecules and quark confinement in quantum chromodynamics.

This problem with configuration space is not unique to quantum mechanics: the task of simulating interacting classical particles is challenging for the same reason. It was only with the advent of computers in the 1950s that a systematic way of simulating classical many-body systems was made possible. In their seminal paper³, Metropolis *et al.* devised a general method to calculate the properties of any substance comprising individual molecules with classical statistics. This paper is a cornerstone in the simulation of interacting systems and has had a huge influence on a wide variety of fields (see, for example, refs 4–6). The Metropolis method can also be used to simulate certain quantum systems by means of a ‘quantum-to-classical map’⁷. Unfortunately, this quantum Monte Carlo method is only scalable when the mapping conserves the positivity of the statistical weights, and fails in the case of fermionic systems as a result of the infamous sign problem⁷.

As the reality of quantum computers comes closer, it is crucial to revisit the original motivation of Feynman for building a quantum simulator and to develop a general method, suitable for quantum computing machines, to calculate the properties of any substance comprising interacting quantum molecules. Such an algorithm would

have a multitude of applications. In quantum chemistry, it could be used to compute the electronic binding energy as a function of the coordinates of the nuclei, thus solving the central problem of interest. In condensed-matter physics, it could be used to characterize the phase diagram of the Hubbard model as a function of filling factor, interaction strength and temperature. Finally, it could conceivably be used to predict the mass of elementary particles, solving a central problem in high-energy physics.

The seminal work of Lloyd² demonstrated that a quantum computer can reproduce the dynamical evolution of any quantum many-body system. It did not address, however, the crucial problem of initial conditions: how to prepare the quantum computer efficiently in a state of physical interest such as a thermal (Gibbs) or ground state. Ground states could in principle be prepared using the quantum phase estimation algorithm^{8,9}, but this method is in general not scalable, because it requires a variational state with a large overlap with the ground state. Methods are known for systems with frustration-free interactions¹⁰ and systems that are adiabatically connected to trivial Hamiltonians¹¹, but such conditions are not generically satisfied. Suggestions have been made of how a quantum computer could sample from the thermal state of a system. One¹² is related to the Metropolis rule but left open the problem of how to overcome the no-cloning result and construct local updates that can be rejected. This shortcoming immediately leads to an exponential running time of the algorithm¹². A second¹² approach to preparing thermal states is by simulating the system’s interaction with a heat bath. However, this procedure seems to produce large errors when run on a quantum computer with finite resources, and a precise framework to describe these errors seems to be out of reach. Moreover, certain systems such as polymers¹³, binary mixtures¹⁴ and critical spin chains^{15,16} experience extremely slow relaxation when put into interaction with a heat bath. The Metropolis dynamics solves this problem by allowing transformations that are not physically achievable, speeding up relaxation by many orders of magnitude and bridging the microscopic and relaxation timescales; this freedom is to a large extent responsible for the tremendous empirical success of the Metropolis method.

In this Letter, we propose a direct quantum generalization of the classical Metropolis algorithm and show how one iteration of the algorithm can be implemented in polynomial time on a quantum computer. Our quantum algorithm is not affected by the sign problem and can be used to prepare ground and thermal states of generic quantum many-body systems, bosonic and fermionic. Like the classical Metropolis algorithm, the quantum Metropolis algorithm is not expected to reach the ground state of an arbitrary Hamiltonian in polynomial time. The ability to prepare the ground state of a general Hamiltonian in polynomial time would allow the solution of quantum Merlin Arthur (QMA)-complete problems^{17,18}, which is highly unlikely. However, for realistic physical systems, the convergence rate of the classical Metropolis algorithm is often very good, and it is conceivable that the same is also true for the quantum Metropolis algorithm. It also inherits all the flexibility and versatility of the classical method, leading, for instance, to a quantum generalization of simulated annealing⁶.

¹Vienna Center for Quantum Science & Technology, Fakultät für Physik, Universität Wien, 1090 Wien, Austria. ²Institute of Theoretical Physics, Gottfried Wilhelm Leibniz Universität Hannover, 30167 Hannover, Germany. ³Max Planck Institut für Quantenoptik, 85748 Garching, Germany. ⁴Département de Physique, Université de Sherbrooke, Sherbrooke, Québec J1K 2R1, Canada.




Figure 1 | Building blocks for the quantum algorithm. **a**, The first step of the quantum circuit: the input is an arbitrary state, $|\psi\rangle$, and two r -qubit registers initialized to $|0\rangle^r$. Quantum phase estimation, Φ , is applied to the state and the second register. The energy value in this register is then copied to the first register by a sequence of Controlled NOT gates. An inverse quantum phase estimation (Φ^\dagger) is then applied to the state and the second register. **b**, The elementary step in the quantum circuit: the input is the eigenstate $|\psi_i\rangle$ with energy register $|E_i\rangle$ and two registers initialized to $|0\rangle^r$ and $|0\rangle$, respectively. The unitary transformation C is then applied, followed by a quantum phase estimation step and the coherent Metropolis gate W . The state evolves as

follows: $|\psi_i\rangle|E_i\rangle|0\rangle|0\rangle \rightarrow C|\psi_i\rangle|E_i\rangle|0\rangle|0\rangle = \sum_k x_k^i |\psi_k\rangle|E_i\rangle|0\rangle|0\rangle \rightarrow \sum_k x_k^i |\psi_k\rangle|E_i\rangle|E_k\rangle|0\rangle \rightarrow \sum_k x_k^i \sqrt{f_k^i} |\psi_k\rangle|E_i\rangle|E_k\rangle|1\rangle + \sum_k x_k^i \sqrt{1-f_k^i} |\psi_k\rangle|E_i\rangle|E_k\rangle|0\rangle$ with $f_k^i = \min(1, \exp(-\beta(E_k - E_i)))$. **c**, The binary measurement checks whether the energy of the state $|\psi\rangle$ is the same as the energy of the original one, $|\psi_i\rangle$. This is done by using an extra register containing phase estimation ancillas, a step that checks whether or not the energy is equal to E_i , and finally an undoing of the phase estimation step that preserves coherence.

To set the stage for the quantum Metropolis algorithm, let us first recall the classical version. We can assume for definiteness that the system is composed of n two-level particles, that is, Ising spins. A lattice of 100 spins has 2^{100} different configurations, so it is inconceivable to average them all. The key insight of Metropolis *et al.*³ was to set up a rapidly mixing Markov chain obeying detailed balance that samples from the configurations with the most significant probabilities. This can be achieved by randomly transforming an initial configuration to a new one (for example by flipping a randomly selected spin): if the energy of the new configuration, E_{new} , is lower than the original, E_{old} , we retain the move, but if the energy is larger we retain the move only with probability $\exp(\beta(E_{\text{old}} - E_{\text{new}}))$, where β is the inverse temperature.

The challenge we address is to set up a similar process in the quantum case, that is, to initiate an ergodic random walk on the eigenstates of a given quantum Hamiltonian with the appropriate Boltzmann weights. In analogy to a spin flip, the random walk can be realized by a random local unitary transformation, and the ‘move’ should be accepted or rejected following the Metropolis rule. There are, however, three obvious complications. First, we do not know what the eigenvectors of the Hamiltonian are (this is one of the problems that we want to solve). Second, certain operations, such as energy measurements, are fundamentally irreversible in quantum mechanics, but the Metropolis method requires rejecting, and hence undoing, certain transformations. Third, it is necessary to devise a criterion which proves that the fixed point of the quantum random walk is the Gibbs state.

To address the first obstacle, we assume for simplicity that the Hamiltonian has non-degenerate eigenvalues, E_i , and denote the corresponding eigenvectors $|\psi_i\rangle$. In the Supplementary Information, we show that those conditions are unnecessary. We can use the phase estimation algorithm^{8,19,20} to prepare a random energy eigenstate and measure the energy of a given eigenstate. Then each quantum Metropolis step (Fig. 1) takes as input an energy eigenstate $|\psi_i\rangle$ with known energy E_i and applies a random local unitary transformation C , creating the superposition $C|\psi_i\rangle = \sum_k x_k^i |\psi_k\rangle$. The transformation C could be a bit flip at a random location, as in the classical setting, or some other simple transformation. The phase estimation algorithm is then used in a coherent way, producing $\sum_k x_k^i |\psi_k\rangle|E_k\rangle$, where $|E_k\rangle$ is an extra register encoding the energy in binary format. At this point, we could measure the second register to read out the energy E_k and accept or reject the move following the Metropolis prescription. However, such an energy measurement would involve an irreversible collapse of the wave function, making it impossible to return to the original configuration in the case of a reject step.

Classically, we overcome this second obstacle by keeping a copy of the original configuration in the computer’s memory, allowing a rejected move to be easily undone. Unfortunately, this solution is ruled out in the quantum setting by the no-cloning theorem²¹. The key to the solution is to engineer a measurement that reveals as little information as possible about the new state, and therefore only slightly disturbs it. This can be achieved by a measurement that only reveals one bit of information—accept or reject the move—rather than a full energy




Figure 2 | Quantum Metropolis stochastic map. The circuit corresponds to a single application of the map \mathcal{E} . The first step, E , prepares an eigenstate of the Hamiltonian. The second step, Q , measures whether we want to accept or reject the proposed update. In the case of rejection, the complete quantum circuit comprises a sequence of measurements of the Hermitian projectors Q_i and P_i .

The recursion is aborted whenever the outcome P_1 is obtained, which indicates that we have returned to a state with the same energy as the input. Because each iteration has a constant success probability, the overall probability of obtaining the outcome P_1 approaches one exponentially as the number of iterations increases.

the gap for more complex Hamiltonians remains a challenging open problem. Also, it is well known that the choice of updates, $\{C\}$, can have a drastic impact on the convergence rate of the Markov chain in the classical setting. Finding good updates in the quantum setting is a very interesting open question, although the above example suggests that the problem might be simpler in the quantum case than in the classical case. The algorithm can be seen as a classical random walk on the eigenstates of the Hamiltonian. All samples are thus computed with respect to the actual eigenstates. This is why our method is suitable for the simulation of fermionic systems by exploiting the Jordan–Wigner transformation²⁶, as discussed in ref. 27. The fermionic sign problem is therefore not an issue for the quantum Metropolis algorithm. It is worth noting that an additional quadratic speed-up might be achievable using the methods of refs 28–30.

Received 29 September; accepted 21 December 2010.

1. Feynman, R. Simulating physics with computers. *Int. J. Theor. Phys.* **21**, 467–488 (1982).
2. Lloyd, S. Universal quantum simulators. *Science* **273**, 1073–1078 (1996).
3. Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H. & Teller, E. Equation of state calculation by fast computing machines. *J. Chem. Phys.* **21**, 1087 (1953).
4. Durr, S. *et al.* Ab initio determination of light hadron masses. *Science* **322**, 1224–1227 (2008).
5. Geman, S. & Geman, D. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Trans. Pattern Anal. Mach. Intell.* **6**, 721–741 (1984).
6. Kirkpatrick, S., Gelatt, C. D. & Vecchi, M. P. Optimization by simulated annealing. *Science* **220**, 671–680 (1983).
7. Suzuki, M. (ed.) *Quantum Monte Carlo Methods in Equilibrium and Nonequilibrium Systems* (Springer Ser. Solid-State Sci. 74, Springer, 1987).
8. Abrams, D. S. & Lloyd, S. Quantum algorithm providing exponential speed increase for finding eigenvalues and eigenvectors. *Phys. Rev. Lett.* **83**, 5162–5165 (1999).
9. Aspuru-Guzik, A., Dutoi, A. D., Love, P. J. & Head-Gordon, M. Simulated quantum computation of molecular energies. *Science* **309**, 1704–1707 (2005).
10. Verstraete, F., Wolf, M. M. & Cirac, J. I. Quantum computation and quantum-state engineering driven by dissipation. *Nature Phys.* **5**, 633–636 (2009).
11. Farhi, E. *et al.* A quantum adiabatic evolution algorithm applied to random instances of an NP-complete problem. *Science* **292**, 472–475 (2001).
12. Terhal, B. M. & DiVincenzo, D. P. Problem of equilibration and the computation of correlation functions on a quantum computer. *Phys. Rev. A* **61**, 022301 (2000).
13. Binder, K. *Monte Carlo and Molecular Dynamics Simulations in Polymer Science* (Oxford Univ. Press, 1995).
14. Liu, J. & Luijten, E. Rejection-free geometric cluster algorithm for complex fluids. *Phys. Rev. Lett.* **92**, 035504 (2004).
15. Swendsen, R. H. & Wang, J.-S. Nonuniversal critical dynamics in Monte Carlo simulations. *Phys. Rev. Lett.* **58**, 86–88 (1987).
16. Evertz, H. G. The loop algorithm. *Adv. Phys.* **52**, 1–66 (2003).
17. Kitaev, A. Y., Shen, A. H. & Vaynskiy, M. N. *Classical and Quantum Computation* (American Mathematical Society, 2002).
18. Aharonov, D. & Naveh, T. Quantum NP - a survey. Preprint at (<http://arxiv.org/abs/quant-ph/0210077>) (2002).
19. Kitaev, A. Y. Quantum computations: algorithms and error correction. *Russ. Math. Surv.* **52**, 1191–1249 (1997).
20. Cleve, R., Ekert, A., Macchiavello, C. & Mosca, M. Quantum algorithms revisited. *Proc. R. Soc. Lond. A* **454**, 339–354 (1998).
21. Wootters, W. K. & Zurek, W. H. A single quantum cannot be cloned. *Nature* **299**, 802–803 (1982).
22. Marriott, C. & Watrous, J. Quantum Arthur–Merlin games. *Comput. Complex.* **14**, 122–152 (2005).
23. Oliveira, R. & Terhal, B. M. The complexity of quantum spin systems on a two-dimensional square lattice. *Quant. Inf. Comp.* **8**, 900–924 (2008).
24. Aharonov, D., Gottesman, D., Irani, D. & Kempe, J. The power of quantum systems on a line. *Commun. Math. Phys.* **287**, 41–65 (2009).
25. Schuch, N. & Verstraete, F. Computational complexity of interacting electrons and fundamental limitations of density functional theory. *Nature Phys.* **5**, 732–735 (2009).
26. Jordan, P. & Wigner, E. Über das Paulische Äquivalenzverbot. *Zeit. Phys. A* **47**, 631–651 (1928).
27. Abrams, D. S. & Lloyd, S. Simulation of many-body Fermi systems on a universal quantum computer. *Phys. Rev. Lett.* **79**, 2586–2589 (1997).
28. Szegedy, M. In *Proc. Annu. IEEE Symp. Found. Comput. Sci.* 32–41 (IEEE, 2004).
29. Somma, R. D., Boixo, S., Barnum, H. & Knill, E. Quantum simulations of classical annealing processes. *Phys. Rev. Lett.* **101**, 130504 (2008).
30. Poulin, D. & Wocjan, P. Sampling from the thermal quantum Gibbs state and evaluating partition functions with a quantum computer. *Phys. Rev. Lett.* **103**, 220502 (2009).

Supplementary Information is linked to the online version of the paper at www.nature.com/nature.

Acknowledgements We would like to thank S. Bravyi, C. Dellago, J. Kempe and H. Verschelde for discussions. Part of this work was done during a workshop at the Erwin Schrödinger Institute for Mathematical Physics. K.T. was supported by the FWF programme CoQuS. T.J.O. was supported, in part, by EPSRC. K.G.V. is supported by DFG FG 635. D.P. is partly funded by NSERC, MITACS and FQRNT. F.V. is supported by the FWF grants FoQuS and ViCoM, by the European grant QUEVADIS and by the ERC grant QUERG.

Author Contributions This project was started by T.J.O. and F.V. five years ago and rejuvenated by an inspiring visit of K.G.V. with K.T. and F.V. in 2009. The connection to the QMA amplification scheme of Marriott and Watrous was made by D.P. at the Erwin Schrödinger Institute, and the project was finalized by K.T. and F.V. by proving quantum detailed balance.

Author Information Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Readers are welcome to comment on the online version of this article at www.nature.com/nature. Correspondence and requests for materials should be addressed to F.V. (fverstraete@gmail.com).