

Preparing thermal states of quantum systems by dimension reduction

E. Bilgin* and S. Boixo†

* *Caltech, Institute for Quantum Information, Pasadena CA 91125, USA*

† *Harvard University, Cambridge, MA 02138, USA*

See arXiv:1008.4162v2

Abstract

We present an algorithm that prepares thermal Gibbs states of one dimensional quantum systems on a quantum computer without any memory overhead, and in a time significantly shorter than other known alternatives. Specifically, the time complexity is dominated by the quantity $N^{\|h\|/T}$, where N is the size of the system, $\|h\|$ is a bound on the operator norm of the local terms of the Hamiltonian (coupling energy), and T is the temperature. Given other results on the complexity of thermalization, this overall scaling is likely optimal. For higher dimensions, our algorithm lowers the known scaling of the time complexity with the dimension of the system by one.

Many open problems in condensed matter physics concern strongly correlated quantum many-body systems. These are typically not solvable analytically, and we have to resort to numerical simulations. Unfortunately, numerical methods tend to fail for general hamiltonians on these systems, due to the exponential scaling of the dimension of the corresponding Hilbert space. This problem is one of the main motivations for the quest of quantum computers. Indeed, quantum computers can efficiently simulate unitary evolutions of quantum many-body systems with local interactions [5, 11], because they can inherently deal with exponentially large Hilbert spaces.

Nevertheless, the preparation of the desired initial state is still a difficult problem in general [9, 8, 12, 1, 15]. There have been several proposals to tackle this problem [17, 16, 4, 14, 18]. Some significant alternatives have worse complexity scaling than ours [17, 14], while others apply to a restricted set of systems [4, 18]. The quantum metropolis algorithm [16], in particular, might often be faster, but lacks complexity bounds. The classical algorithm proposed in [6] can be used to prepare 1D quantum thermal states with only a polynomial time complexity overhead with respect to our method, but its (classical) memory requirements scale exponentially with inverse temperature, and it does not extend to higher dimensional systems.

The time complexity of our method for one dimensional systems is dominated by the quantity $N^{\|h\|/T}$, where N is the number of subsystems, T is the temperature, and $\|h\|$ is a bound on the operator norm of the local terms of the Hamiltonian, the interaction strength. Note that this scaling is polynomial in N . The memory of the quantum computer scales simply with N , an exponential improvement over general classical algorithms. Our algorithm can also be massively parallelized, and when run in a cellular automaton architecture the memory scales as $N^{\|h\|/T}$, but the total time would be linear in N (the total number of steps would still be the same). In two and higher dimensions, our method lowers the number of effective dimensions by one. This results in an exponential speedup, but the exponential scaling with N remains.

The overall scaling appears to be optimal: the known complexity of thermalizing 1D quantum systems makes a guaranteed polynomial scaling with temperature highly unlikely [1, 15]. We also expect the grouping of $\|h\|/T$ in the exponent by dimensional analysis. In other words, the relevant temperature scale is set by the Hamiltonian.

In our method, we first thermalize small regions, which we merge recursively until we have thermalized the whole system. The merging fails with some probability, but only when the failures are close to the end of this recursive procedure do we need to rebuild big sections. A careful error analysis shows that each of these merging operations can be implemented with a cost independent of the system size

and the quantum correlation length, resulting in a running time that is only polynomial in the system size and independent of the correlation length. This method trivially generalizes to higher dimensions and reduces the scaling of the cost with the system dimension by one compared to a direct projection.




Figure 1: The procedure to thermalize an 8-qubit chain. After thermalizing individual qubits at level $k = 0$, we pair them up and merge them by adding the Hamiltonian that connects the two qubits. This procedure is then repeated recursively as we merge two already thermalized regions of size 2^k at level k to obtain a thermalized chain of size 2^{k+1} at level $k + 1$.

We implement the merging perturbatively. Assume that we are given access to copies of $\rho \propto e^{-\beta H}$ (from previous steps). The Hamiltonian H corresponds to the halves to be merged, but the procedure is more general. We show how to generate (with high probability) the state $\rho^{(1)} \propto e^{-\beta(H+h)}$, where h corresponds to the link between the two halves. We then repeat the same process to produce the sequence

$$\rho = \rho^{(0)} \rightarrow \rho^{(\epsilon)} \rightarrow \rho^{(2\epsilon)} \rightarrow \dots \rightarrow \rho^{(1)}. \quad (1)$$

For every transformation in the sequence the success probability is

$$p \geq 1 - \epsilon\beta\|h\|, \quad (2)$$

with trace norm error $\mathcal{O}(\epsilon^2\beta^2\|h\|^2)$ and cost (evolution time)

$$\mathcal{O}(\log(1/(\epsilon\beta\|h\|)))/(\epsilon^2\beta\|h\|^2). \quad (3)$$

The average number of steps until we generate a complete sequence (see Eq. (1)) without failures is $\langle m \rangle \in \mathcal{O}(e^{\beta\|h\|})$ ¹. Each time that we fail we need to produce two new thermal regions to be merged. The average number of failures is $\langle \alpha \rangle \in \mathcal{O}(e^{\beta\|h\|})$. We analyze the average number of steps $\langle \tau(k) \rangle$ required to prepare a thermalized chain of length 2^k at level k (see Fig. 1). We get $\langle \tau(\log N) \rangle \in \mathcal{O}(\exp((\beta\|h\| + \log 2)\log N))$. Similarly, the total error is $\mathcal{O}(N\epsilon\beta^2\|h\|^2)$. If we choose $\epsilon = \bar{\epsilon}/(N\beta^2\|h\|^2)$, we get a total error of $\mathcal{O}(\bar{\epsilon})$ in trace-norm. Finally, using Eq. (3) for the evolution time of each step, we obtain the dominant contribution to the total evolution time $\beta N^{\beta\|h\|}/\bar{\epsilon}^2$.

¹This is also known from the theory of success runs. We give the average cost, but the tail has an exponential decay rate, so the worst case cost is similar (see, for instance [2]).

We have presented an algorithm that prepares a thermal state of a 1D quantum system in time polynomial in the system size and exponential in the inverse temperature (as required by the existence of QMA-complete ground state problems in 1D). This algorithm can be generalized into D dimensions. At level k of the recursion, we have built squares (for 2D) or cubes that are now merged. We do not get polynomial scaling with system size for $D > 1$ because the intersection of two neighboring regions goes like N^{D-1} . Note that this is to be expected because there exist 2D ground states with constant gap that encode the solution to NP-complete problems. A careful analysis confirms that the time complexity is dominated by the operations at the top level, and the dominating factor is $\beta e^{2\beta \|h\| DN^{D-1}}/\bar{\epsilon}^2$. This is an exponential speedup from the known $\exp(O(N^D))$. The memory requirements still scale with the number of sites of the model, N^D .

There are also several possible improvements to the scaling of this algorithm. If one is interested in thermalizing a classical system with a small quantum perturbation one can first solve for the classical part of the Hamiltonian. Then, one would only need to use projections for the quantum perturbation. Also, if one is interested in thermalizing a quantum system with short-ranged quantum correlations, one can also use belief propagation [7, 10, 13, 3] to reduce the storage requirements from $\mathcal{O}(N)$ qubits to $\mathcal{O}(l \log(N))$, where l is a constant related to the quantum correlation length. This can be done by tracing out parts of the blocks that do not share any entanglement with the boundary to be merged. The time complexity of the algorithm remains the same as before. Note that the cost (both in memory and time) of the classical algorithm of quantum belief propagation for 1D quantum systems is exponential in $l \sim 1/T$. This bound is only heuristic, and similar to [6].

References

- [1] Dorit Aharonov, Daniel Gottesman, Sandy Irani, and Julia Kempe. The power of quantum systems on a line. *Commun. Math. Phys.*, 287(1):41–65, 2009.
- [2] N. Balakrishnan and Markos V. Koutras. *Runs and Scans with Applications*. Wiley-Interscience, 1 edition, 2001.
- [3] Ersen Bilgin and David Poulin. Coarse-grained belief propagation for simulation of interacting quantum systems at all temperatures. *Phys. Rev. B*, 81(5):054106, 2010.
- [4] M. Cramer and J. Eisert. A quantum central limit theorem for non-equilibrium systems: exact local relaxation of correlated states. *New Journal of Physics*, 12(5):055020–+, 2010.
- [5] R. P. Feynman. Simulating physics with computers. *Int. J. Theor. Phys.*, 21(6):467–488, 1982.
- [6] M. B. Hastings. Solving gapped hamiltonians locally. *Phys. Rev. B*, 73(8):085115–13, 2006.
- [7] M. B. Hastings. Quantum belief propagation: An algorithm for thermal quantum systems. *Phys. Rev. B*, 76(20):201102–4, 2007.
- [8] J. Kempe, A. Kitaev, and O. Regev. The complexity of the local hamiltonian problem. *SIAM J. Comput.*, 35(5):1070, 2006.
- [9] A. Yu. Kitaev, A. H. Shen, and M. N. Vyalyi. *Classical and Quantum Computation*. Amer. Math. Soc., 2002.
- [10] M.S. Leifer and D. Poulin. Quantum graphical models and belief propagation. *Ann. Phys.*, 323(8):1899–1946, 2008.
- [11] Seth Lloyd. Universal quantum simulators. *Science*, 273(5278):1073–1078, 1996.

- [12] Roberto Oliveira and Barbara M Terhal. The complexity of quantum spin systems on a two-dimensional square lattice. *Quant. Inf, Comp. Vol. 8, No. 10, pp. 0900-0924 (2008)*, 2005.
- [13] David Poulin and Ersen Bilgin. Belief propagation algorithm for computing correlation functions in finite-temperature quantum many-body systems on loopy graphs. *Phys. Rev. A*, 77(5):052318–5, 2008.
- [14] David Poulin and Pawel Wocjan. Sampling from the thermal quantum gibbs state and evaluating partition functions with a quantum computer. *Phys. Rev. Lett.*, 103(22):220502, 2009.
- [15] Norbert Schuch, Ignacio Cirac, and Frank Verstraete. Computational difficulty of finding matrix product ground states. *Phys. Rev. Lett.*, 100(25):250501–4, 2008.
- [16] K. Temme and *et al.* Quantum Metropolis Sampling. 2009.
- [17] Barbara M. Terhal and David P. DiVincenzo. Problem of equilibration and the computation of correlation functions on a quantum computer. *Phys. Rev. A*, 61(2):022301, 2000.
- [18] M.-H. Yung and *et al.* Simulation of Classical Thermal States on a Quantum Computer: A Transfer Matrix Approach. 2010.