



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Metropolis-style random sampling of quantum gates for the estimation of low-energy observables

Judah F. Unmuth-Yockey

Phys. Rev. D **105**, 034515 — Published 25 February 2022

DOI: [10.1103/PhysRevD.105.034515](https://doi.org/10.1103/PhysRevD.105.034515)

Metropolis-style random sampling of quantum gates for the estimation of low-energy observables

Judah F. Unmuth-Yockey*

Theory division, Fermi national accelerator laboratory, Batavia, IL 60510-5011, USA

(Dated: February 11, 2022)

We propose a quantum algorithm to compute low-energy expectation values of a quantum Hamiltonian by sampling a partition function associated with the average energy of that Hamiltonian. For any given quantum circuit-Hamiltonian pair, there is an associated average energy. The sampling is done through an accept/reject Metropolis-style algorithm on the quantum gates of the circuit itself. Observables calculated under the canonical ensemble from these samples of circuits are extrapolated from higher-energies to the ground state.

I. INTRODUCTION

In many areas of physics, we are often interested in calculating observables when the physical system is in, or is around, its lowest energy state. One particularly difficult problem of this nature is the calculation of expectation values of observables in the ground state of a quantum system. Finding the eigenvector with the smallest eigenvalue for a many-body quantum Hamiltonian using a classical computer rapidly becomes expensive as the number of bodies becomes large [1]. Quantum computing potentially alleviates some of this expense. Algorithms on quantum computers to compute the ground state of a Hamiltonian have provided a vast speed up—a quadratic improvement—over classical algorithms [2, 3]. Often, however, we are not necessarily interested in the actual eigenvector itself; it is rather a means to an end: to calculate an average of some observable.

A powerful tool to calculate gross properties—like averages—nonperturbatively comes from Markov chain Monte Carlo sampling. In the case of a classical statistical mechanics system, one can sample the configurations of a partition function to reconstruct ensemble averages under the canonical ensemble. This is what is done in lattice gauge theory, for example, where a system of fundamental particles are simulated on a Euclidean lattice such that gauge invariance is maintained exactly [4]. Average properties of the system can be computed and, when possible, Wick-rotated back into Minkowski space-time. In this context, it is important to take the continuum (and infinite-volume) limit. This is done by taking the gauge-coupling, g , to zero, or $\beta_{\text{gauge}} \sim 1/g^2 \rightarrow \infty$.

However, in practice one never performs simulations with β_{gauge} infinitely large, and instead moderate values of β_{gauge} are chosen, and the infinite- β_{gauge} limit is done by extrapolation. In this way, continuum-limit results are found from finite-lattice-spacing calculations. Here, we propose to perform a similar extrapolation when what is desired are expectation values of observables in the ground state. We propose to perform systematic calculations of observables at larger energies, and extrapolate

to their value at low energy by fitting. This is not done by a variational minimization routine [5–10], where the ground state is sought out by minimizing a cost function of expectation values; though, a procedure similar in spirit to what is proposed here is described in Ref. [11]. The ground state is not sought out using quantum imaginary time evolution [7, 12, 13] which involves a classical minimization routine based on state tomography to construct unitary circuits to implement non-unitary operators. We do not use a strictly “gate-based” method, where the ground state of a quantum Hamiltonian is calculated with some finite probability [14–18]. The method proposed here is also not a quantum Metropolis algorithm, where a thermal quantum state is sampled [19–21], or a state is adiabatically annealed [22–25] to its ground state. Instead, by using samples of averages of observables calculated away from the ground state, we extrapolate those observables to their lowest-energy values with some error, avoiding the ground state entirely.

The remainder of the article is as follows: In Sec. II we describe the algorithm in detail, and discuss its computational cost. In Sec. III we use the algorithm in an example, and compute the ground-state energy eigenvalue and magnetization for the transverse-field quantum Ising model for a small system using a classical computer. Finally in Sec. IV we provide possible future directions and work, as well as conclusions.

II. RANDOM CIRCUIT SAMPLING

The goal of the algorithm is to construct an ensemble of *circuits* whose average energies fluctuate about some equilibrium value controlled by an external parameter set by the user. This is done non-perturbatively by the random sampling of quantum gates. By controlling the target equilibrium average-energy value, observables of interest can be calculated from the averages of the circuits, and extrapolated to the low-energy sector. This is possible because of the functional form of the Boltzmann weight and the partition function, where smaller “temperatures” favor lower energy configurations. In addition, because of the variational principle, energy expectation values are upper-bounds on the ground state

* jfunmuthyockey@gmail.com

energy, so that when extrapolating to lower energies, this bound is satisfied. In this way, any state preparation is unnecessary, and instead one extrapolates the physics of interest from higher energy expectation values, to their low-energy values. We now describe the algorithm.

We begin with an initial state $|\psi_0\rangle$ for N_q qubits that is easy to prepare, such as the all-‘plus’ state, or the all-zeros state. We assume the circuit—represented by the unitary U —that acts on this state can be expressed in N_t layers, that is, $U = U^{(N_t)}U^{(N_t-1)} \dots U^{(2)}U^{(1)}$. Each layer has three sub-layers. There is a sub-layer of one-qubit gates, and two sub-layers of two-qubit gates, where each two-qubit sub-layer acts on distinct pairs of qubits. Figure 1 shows an illustration of the layer structure. The

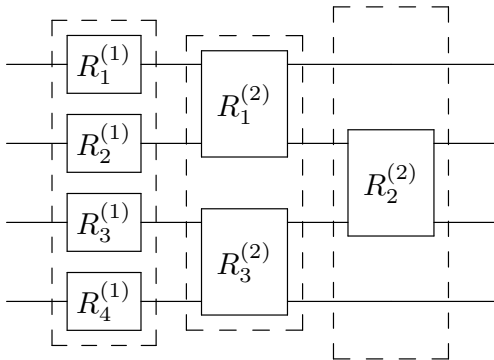


FIG. 1. A single layer with four qubits, comprised of three sub-layers which are boxed with dashed lines. One sub-layer only has single-qubit rotations indicated by $R^{(1)}$, while the next two sub-layers have two-qubit rotations, indicated by $R^{(2)}$.

one- and two-qubit gates are initialized randomly, or *e.g.* as identity matrices, from the groups $U(2)$ and $U(4)$, respectively. For the description here we also assume that each of the N_t layers is identical. *I.e.* if we know the gates in a single layer, we know them in all layers. This is simply the statement that the global unitary which is acting on $|\psi_0\rangle$ can be decomposed via Trotterization, and does not have any ‘‘time dependence.’’ At this point, a model must be chosen whose physics is the physics of interest. We specify this model with a Hamiltonian, \mathcal{H} , which is Hermitian.

We now proceed using a Metropolis-style update [26] method for the quantum gates in the circuit. The Metropolis algorithm consists of proposing a new circuit configuration, and accepting that new circuit configuration based on a probability given by the Boltzmann weight associated with the change in energy between the proposed, and current configurations [27]. With the initial state, and the initial circuit, we evolve the initial state using the circuit many times to calculate the average energy associated with that circuit, $\mathcal{E}_0 \equiv \langle \psi_0 | U^\dagger \mathcal{H} U | \psi_0 \rangle = \langle U^\dagger \mathcal{H} U \rangle$ [28]. We then choose a gate at random in a layer (it can be a fixed layer, *e.g.* the first layer, since we enforce that all layers be the same), and for that

gate we propose a new unitary matrix sampled uniformly from the Haar measure [29, 30] [31]. It may be preferable to bias how one chooses the gates to update. This would need to be investigated based on other factors *e.g.* auto-correlations, acceptance, and circuit ansatz. We then evolve the initial state with the proposed modified circuit—expressed as U' —many times to compute the new expectation value, $\mathcal{E}_f \equiv \langle U'^\dagger \mathcal{H} U' \rangle$. Using the current and proposed average energies we accept the new circuit with a probability, p , given by

$$p = \begin{cases} 1 & \text{if } \mathcal{E}_f \leq \mathcal{E}_0 \\ e^{-\beta(\mathcal{E}_f - \mathcal{E}_0)} & \text{otherwise.} \end{cases} \quad (1)$$

In Eq. (1) the parameter β is chosen by the user, and controls the equilibrium average energy value. Larger values of β push the target average energy lower. If the new circuit is chosen, the circuit is kept and the process repeats; otherwise, the proposed circuit is ignored, the altered gate is reverted back, a new gate is chosen at random and the process repeats. Sampling performed in this manner attempts to approximately compute averages with respect to the partition function,

$$Z = \int dV e^{-\beta \langle V^\dagger \mathcal{H} V \rangle} = \int dV e^{-\beta \mathcal{E}(V)} \quad (2)$$

where dV is the Haar measure for the group $U(2^{N_q})$. Because of the invariance of the Haar measure the partition function is independent of $|\psi_0\rangle$.

For a fixed value of β , at equilibrium, the circuit composition fluctuates such that the average energy is distributed about some equilibrium value, $\bar{\mathcal{E}}$. Using this equilibrium value as a reference, one can choose increasingly larger values of β , and extrapolate observables to their low-energy values.

The cost of the algorithm is as follows: Since each layer is identical, we need only remember the contents of a single layer, and a proposed gate, which is trivial on a classical computer; however, for N_q qubits, and N_t layers, there are $N_q N_t$ arbitrary one-qubit gates in the circuit, and $(N_q - 1)N_t$ arbitrary two-qubit gates in the circuit. If we define a single ‘‘sweep’’ as $N_q + (N_q - 1)$ proposals, and the number of sweeps, N_s is chosen by the user, the cost of a sweep—the execution time—scales linearly in the number of qubits.

Aside from the circuit cost associated with executing arbitrary single- and two-qubit gates, the major overhead comes from the number of measurements. The error, ϵ , on the expectation value of the energy associated with a circuit converges like $\sim 1/\sqrt{N_{\text{meas}}}$ where N_{meas} is the number of uncorrelated measurements done to compute an expectation value. Let \mathcal{N} be the number of uncorrelated expectation values recorded. Then the error, λ , on the average of expectation values converges similarly like $\sim 1/\sqrt{\mathcal{N}}$. The final error on the average expectation value scales like $\sim \sqrt{\epsilon^2 + \lambda^2}$.

III. THE ISING MODEL AS AN EXAMPLE

In this section we attempt to elucidate the procedure described in the previous section with an example: the Ising model. The one-dimensional, quantum, transverse-field Ising model Hamiltonian with open boundary conditions is given by

$$\mathcal{H}_{\text{Ising}} = - \sum_{i=1}^{N_q-1} \sigma_i^z \sigma_{i+1}^z - h_x \sum_{i=1}^{N_q} \sigma_i^x, \quad (3)$$

where h_x is the magnitude of the transverse field. This Hamiltonian will serve as the \mathcal{H} used in the algorithm. To test the basic principles behind the algorithm we attempt to classically compute the ground-state energy eigenvalue, and the expectation value of the magnetization, $M \equiv \sum_{i=1}^{N_q} \sigma_i^x$, in the ground state for a given value of h_x . We compare calculations done with the algorithm with exact results calculated using exact diagonalization.

We consider the case of $N_q = 4$ and $N_t = 6$ (although other layer depths were tried and little dependence was found for this case), and calculate observables at values of the transverse field $h_x = 1/4$, and $h_x = 3/2$, *i.e.* in the ordered, and disordered phases, respectively. The circuit, represented by a unitary U in the previous section, is comprised of six layers, and each layer consists of the three sub-layers mentioned before. We store a single layer as a list with three lists inside of it. Each list contains the current one-qubit, or two-qubit gates. A single gate is randomly chosen, and a new gate is proposed. Then the proposed circuit is made by multiplying the list with itself N_t times, the initial state is evolved with that circuit, and the energy expectation value is computed.

For the classical simulation done here, this energy expectation value is computed exactly. In practice one would need to evaluate the same circuit many times to compute an accurate value for the energy expectation value. While this must be done when using a quantum computer, here for convenience we simply calculate the expectation value exactly from the current circuit acting on the initial state. This new energy expectation value is compared with the current energy expectation value, and the accept/reject procedure is carried out. The process then repeats. Notice that in contrast to the classical Ising-spin Metropolis, where a spin-flip is the most dramatic change that can take place; here, the new unitary gate may result in a very different overall unitary evolution.

We typically perform approximately 1×10^6 sweeps and record measurements every 10 sweeps to help with auto-correlation. An example of time series data for the expectation value of the energy can be seen in Fig. 2. After removing early measurements during equilibration, we bin the data and calculate the error using single-elimination jackknife resampling [27, 32] to assess the auto-correlation. We choose a bin-size when the error saturates, and perform the remainder of the analysis under the jackknife.




FIG. 2. Example time-series data for the expectation value of the energy for the first 4000 measurements.




FIG. 3. The average energy as a function of β . The ground state energy value for the Ising model with $h_x = 3/2$ is shown as a black star. A fit is plotted with its error-band along with the asymptotic value of the fit with error bar. The largest contribution to the error on the asymptote is statistical.

From the measurements of \mathcal{E} , and $\langle M \rangle$ we compute their average values as a function of β . Measurements of the energy—really the average of the expectation values of the energy—for the case of $h_x = 3/2$ can be seen in Fig. 3. For sufficiently large β , we find the data approximately obeys a form

$$\bar{\mathcal{E}} = \frac{A}{\beta} + B \quad (4)$$

where A and B are free parameters. We also allow a floating power-law,

$$\bar{\mathcal{E}} = \frac{A}{\beta^C} + B \quad (5)$$

with C an additional free parameter, and use the B from this fit to include a systematic error from the fit choice with Eq. (4). We find good agreement between the data with both Eqs. (4) and (5), as can be seen in




FIG. 4. The average of the expectation value of the magnetization for $h_x = 3/2$. Here a fit line is drawn with its error-band, as well as the exact expectation value for the magnetization in the ground state represented as a black star. The final asymptote to the fit is also drawn with an error bar. The largest contribution to the error on the asymptote is statistical.

Fig. 3. The asymptotic value of the fit from Eq. (4) gives $\bar{\mathcal{E}}_\infty = -6.45(10)$ with a $\chi^2/\text{d.o.f.} = 0.47$, where the error comes from the statistical error, as well as a small systematic error from the choice of fit and the fit range, which are added in quadrature. This is to be compared with the exact value of $E_0 = -6.50$. For the magnetization with $h_x = 3/2$ we find slower convergence, but consistent with a power-law in β . Figure 4 shows the average of the expectation value of the magnetization. For this observable we found a fit including quadratic corrections of $1/\beta$,

$$\overline{\langle M \rangle} = \frac{D}{\beta} + \frac{E}{\beta^2} + F, \quad (6)$$

to be most suitable, where D , E , and F are fit parameters; however, we also tried a floating power-law—as in Eq. (5)—to assess a systematic error. Both fits were good, and Fig. 4 shows the Eq. (6) fit, along with the asymptotic value and an error bar. The choice of fit and fit range contributed a small systematic error to the total asymptote error. For the asymptote we find $\overline{\langle M \rangle}_\infty = 3.69(8)$, which is to be compared with the exact answer, $\langle M \rangle_{\text{exact}} = 3.66$. The fit has a $\chi^2/\text{d.o.f.} = 0.17$.

For $h_x = 1/4$, the average energy can be seen in Fig. 5. We see slight curvature, and so include corrections like $1/\beta^2$ into the fit ansatz, in addition to allowing for a single floating power. The fit with quadratic corrections in $1/\beta$, as well as the exact answer, can be seen in Fig. 5. We find for the asymptotic value $\bar{\mathcal{E}}_\infty = -3.17(15)$, which is to be compared with $\mathcal{E}_{\text{exact}} = -3.10$. Here the largest contribution to the error is systematic, with a relatively larger possible range for the asymptote depending on the fit and fit-range used. The $\chi^2/\text{d.o.f.} = 0.29$. The magnetization for $h_x = 1/4$ can be seen in figure 6. Due again to the slower convergence, we use the same fits




FIG. 5. The average energy as a function of $1/\beta$. The ground state energy value for the Ising model with $h_x = 1/4$ is shown as a black star. A fit is plotted with its error-band along with the asymptotic value of the fit with error bar. The largest contribution to the error of the asymptote are the systematic errors associated with the fit ranges and fit functions.




FIG. 6. The average of the expectation value of the magnetization for $h_x = 1/4$. Here a fit line is drawn with its error-band, as well as the exact expectation value for the magnetization in the ground state represented as a black star. The final asymptote to the fit is also drawn with an error bar. The largest contribution to the error on the asymptote is statistical.

as in the $h_x = 3/2$ case, and find we have not really reached the linear region of $1/\beta$ yet. Nevertheless, the quadratic, and floating-power fits provide a reasonable answer for the asymptote. The fit and the exact value can be seen in Fig. 6. We find for the asymptotic value $\overline{\langle M \rangle}_\infty = 0.82(2)$, which is to be compared with the exact value $\langle M \rangle_{\text{exact}} = 0.814$. The largest contribution to the error in this case is statistical. The $\chi^2/\text{d.o.f.} = 0.42$. We find for this simple system, with only a handful of β values, a reasonable extrapolation to the low-energy sector is possible.

IV. CONCLUSIONS AND FUTURE WORK

We have proposed a quantum algorithm to compute expectation values in the low-energy regime of a Hamiltonian. Rather than use a variational method with a cost-function which is minimized, or attempt to construct an exact circuit which evolves the system to its ground-state directly and use it to calculate observables, we propose to randomly sample circuits, which fluctuate about some target average energy, and extrapolate measurements to their low-energy values. The ability to sample around a target average value is a useful feature of this algorithm made possible by fixing the value of β . This allows a systematic approach towards lower energies. The cost for this is degree of control is the introduction of systematic errors in the fitting and extrapolation process, along with run-time overhead due to the nature of probabilistic sampling.

There are a few additional items to note. The first is how the accuracy of estimates scales. The most expensive aspect of the algorithm is the slow convergence of averages with the number of samples, which go like one over the square-root of the number of measurements. Ideally, one could pick a desired error and sample to that extent; however, the situation is more complicated. Since progressively larger β values are needed in the extrapolation, this will lead to a decrease in acceptance during the gate proposal, leading to a necessary increased run-time. There are many improvements on classical computers for spin and gauge models for the Metropolis algorithm which overcome this issue at “low-temperature,” *e.g.* event-driven simulations [27, 33]. Perhaps similar algorithmic advances could be made here.

Second, as one approaches criticality in a statistical quantum theory, the correlation length diverges. This typically causes large auto-correlations when one uses the naive Metropolis method, again leading to long run-times. As before, several algorithms attempt to eliminate this “critical slowing down,” *e.g.* cluster flips [34, 35], worm-like (or world-line) algorithms [36], both of which are non-local, and hybrid Monte Carlo [37]. One could imagine such advances are possible in this case as well.

In addition to the longer run-time, for the ansatz considered here, criticality could push the number of layers to become larger, since the circuit only uses local gates to create long-range correlations.

Finally, the circuit ansatz chosen here is not definitive, and other options could be investigated. The form proposed here was only done so because of its generality, such that it necessarily can reproduce the full Haar measure sum in the appropriate limit. However, other gate ansatzes may be more efficient. This direction deserves more investigation.

The circuit ansatz used here has a physical interpretation in terms of unitary time evolution by an effective Hamiltonian with nearest-neighbor two-spin interactions, and single spins coupled to external fields, since a Hermitian Hamiltonian for the gates can be given in terms of pairs, or single, Pauli operators. In this sense, during the proposal steps one adjusts the strengths of the nearest neighbor—or on-site—couplings and accepts or rejects them based on the new energy expectation value.

The possible difficulties associated with the simple accept/reject method could be cured if efficient sampling methods can be devised for the layered-gate ansatz used here. This ansatz is also a feature which could possibly be improved. The fact that efficient sampling algorithms have already been devised in the past for spin and gauge systems is promising. We leave the development of such algorithms as future work.

ACKNOWLEDGMENTS

I would like to thank Hank Lamm, Michael Wagman, Bharath Sambasivam, Jay Hubisz, Erik Gustafson, and Michael Hite for stimulating discussions. I would also like to thank Hank Lamm, Jack Laiho, Yannick Meurice, and M. Sohaib Alam for reading through drafts of the manuscript. This manuscript has been authored by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.

-
- [1] This is due to the exponential growth of the Hilbert space with system size. Symmetries and conservation laws, along with sparse storage, can alleviate this cost.
 - [2] L. K. Grover, A fast quantum mechanical algorithm for database search, in *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '96 (Association for Computing Machinery, New York, NY, USA, 1996) p. 212219.
 - [3] D. Poulin and P. Wocjan, Preparing ground states of quantum many-body systems on a quantum computer, *Phys. Rev. Lett.* **102**, 130503 (2009).
 - [4] I. Montvay and G. Münster, *Quantum Fields on a Lattice*, Cambridge Monographs on Mathematical Physics (Cambridge University Press, 1994).
 - [5] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O’Brien, A variational eigenvalue solver on a photonic quantum processor, *Nature Communications* **5**, 4213 (2014).
 - [6] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, Variational quantum algorithms, *Nature Reviews Physics* **3**, 625 (2021).
 - [7] K. Yeter-Aydeniz, B. T. Gard, J. Jakowski, S. Majumder, G. S. Barron, G. Siopsis, T. S. Humble, and R. C. Pooser, Benchmarking quantum chemistry computations with variational, imaginary time evolution, and krylov

- space solver algorithms, *Advanced Quantum Technologies* **4**, 2100012 (2021).
- [8] J. Tilly, H. Chen, S. Cao, D. Picozzi, K. Setia, Y. Li, E. Grant, L. Wossnig, I. Rungger, G. H. Booth, and J. Tennyson, The variational quantum eigensolver: a review of methods and best practices (2021), arXiv:2111.05176 [quant-ph].
- [9] J. Zhang, R. Ferguson, S. Khn, J. F. Haase, C. M. Wilson, K. Jansen, and C. A. Muschik, Simulating gauge theories with variational quantum eigensolvers in superconducting microwave cavities (2021), arXiv:2108.08248 [quant-ph].
- [10] A. C. Y. Li, M. S. Alam, T. Iadecola, A. Jahin, D. M. Kurcuoglu, R. Li, P. P. Orth, A. B. Zgler, G. N. Perdue, and N. M. Tubman, Benchmarking variational quantum eigensolvers for the square-octagon-lattice kitaev model (2021), arXiv:2108.13375 [quant-ph].
- [11] C. Cao, Y. Yu, Z. Wu, N. Shannon, B. Zeng, and R. Joynt, Energy extrapolation in quantum optimization algorithms (2021), arXiv:2109.08132 [quant-ph].
- [12] M. Motta, C. Sun, A. T. K. Tan, M. J. O'Rourke, E. Ye, A. J. Minnich, F. G. S. L. Brandão, and G. K.-L. Chan, Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution, *Nature Physics* **16**, 205 (2020).
- [13] K. Yeter-Aydeniz, G. Siopsis, and R. C. Pooser, Scattering in the ising model with the quantum lanczos algorithm, *New Journal of Physics* **23**, 043033 (2021).
- [14] Y. Ge, J. Tura, and J. I. Cirac, Faster ground state preparation and high-precision ground energy estimation with fewer qubits, *Journal of Mathematical Physics* **60**, 022202 (2019), <https://doi.org/10.1063/1.5027484>.
- [15] L. Lin and Y. Tong, Near-optimal ground state preparation, *Quantum* **4**, 372 (2020).
- [16] H. Wang, Quantum algorithm for preparing the ground state of a system via resonance transition, *Scientific Reports* **7**, 16342 (2017).
- [17] D. Lee, J. Bonitati, G. Given, C. Hicks, N. Li, B.-N. Lu, A. Rai, A. Sarkar, and J. Watkins, Projected cooling algorithm for quantum computation, *Physics Letters B* **807**, 135536 (2020).
- [18] E. J. Gustafson, Projective cooling for the transverse ising model, *Phys. Rev. D* **101**, 071504(R) (2020).
- [19] K. Temme, T. J. Osborne, K. G. Vollbrecht, D. Poulin, and F. Verstraete, Quantum metropolis sampling, *Nature* **471**, 87 (2011).
- [20] M.-H. Yung and A. Aspuru-Guzik, A quantum-quantum metropolis algorithm, *Proceedings of the National Academy of Sciences* **109**, 754 (2012), <https://www.pnas.org/content/109/3/754.full.pdf>.
- [21] J. E. Moussa, Measurement-based quantum metropolis algorithm (2019), arXiv:1903.01451 [quant-ph].
- [22] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, Optimization by simulated annealing, *Science* **220**, 671 (1983).
- [23] T. Kadowaki and H. Nishimori, Quantum annealing in the transverse ising model, *Phys. Rev. E* **58**, 5355 (1998).
- [24] E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, and D. Preda, A quantum adiabatic evolution algorithm applied to random instances of an np-complete problem, *Science* **292**, 472 (2001).
- [25] P. Hauke, H. G. Katzgraber, W. Lechner, H. Nishimori, and W. D. Oliver, Perspectives of quantum annealing: methods and implementations, *Reports on Progress in Physics* **83**, 054401 (2020).
- [26] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, Equation of state calculations by fast computing machines, *The Journal of Chemical Physics* **21**, 1087 (1953), <https://doi.org/10.1063/1.1699114>.
- [27] B. A. Berg, *Markov chain Monte Carlo simulations and their statistical analysis* (World scientific publishing co., 2004).
- [28] The circuit must be run as many times as is necessary to achieve the desired degree of accuracy on the energy expectation value. The scaling of this this accuracy is described at the end of the section.
- [29] A. Hurwitz, ber die erzeugung der invarianten durch integration, *Nachrichten von der Gesellschaft der Wissenschaften zu Gttingen, Mathematisch-Physikalische Klasse* **1897**, 71 (1897).
- [30] A. Haar, Der massbegriff in der theorie der kontinuierlichen gruppen, *Annals of Mathematics* **34**, 147 (1933).
- [31] One could also propose a new unitary somewhat 'close' to the current unitary to improve acceptance.
- [32] M. H. QUENOUILLE, NOTES ON BIAS IN ESTIMATION, *Biometrika* **43**, 353 (1956), <https://academic.oup.com/biomet/article-pdf/43/3-4/353/987603/43-3-4-353.pdf>.
- [33] A. Bortz, M. Kalos, and J. Lebowitz, A new algorithm for monte carlo simulation of ising spin systems, *Journal of Computational Physics* **17**, 10 (1975).
- [34] R. H. Swendsen and J.-S. Wang, Nonuniversal critical dynamics in monte carlo simulations, *Phys. Rev. Lett.* **58**, 86 (1987).
- [35] U. Wolff, Collective monte carlo updating for spin systems, *Phys. Rev. Lett.* **62**, 361 (1989).
- [36] N. Prokof'ev and B. Svistunov, Worm algorithms for classical statistical models, *Phys. Rev. Lett.* **87**, 160601 (2001).
- [37] S. Duane, A. Kennedy, B. J. Pendleton, and D. Roweth, Hybrid monte carlo, *Physics Letters B* **195**, 216 (1987).