

Classical Algorithms, Correlation Decay, and Complex Zeros of Partition Functions of Quantum Many-Body Systems

Aram W. Harrow
MIT
Cambridge, USA
aram@mit.edu

Saeed Mehraban
Institute for Quantum Information
and Matter, Caltech
Pasadena, USA
mehraban@caltech.edu

Mehdi Soleimanifar
MIT
Cambridge, USA
mehdis@mit.edu

ABSTRACT

We present a quasi-polynomial time classical algorithm that estimates the partition function of quantum many-body systems at temperatures above the thermal phase transition point. It is known that in the worst case, the same problem is NP-hard below this point. Together with our work, this shows that the transition in the phase of a quantum system is also accompanied by a transition in the hardness of approximation. We also show that in a system of n particles above the phase transition point, the correlation between two observables whose distance is at least $\Omega(\log n)$ decays exponentially. We can improve the factor of $\log n$ to a constant when the Hamiltonian has commuting terms or is on a 1D chain. The key to our results is a characterization of the phase transition and the critical behavior of the system in terms of the complex zeros of the partition function. Our work extends a seminal work of Dobrushin and Shlosman on the equivalence between the decay of correlations and the analyticity of the free energy in classical spin models. On the algorithmic side, our result extends the scope of a recent approach due to Barvinok for solving classical counting problems to quantum many-body systems.

CCS CONCEPTS

• **Theory of computation** → **Quantum computation theory; Approximation algorithms analysis.**

KEYWORDS

quantum many-body systems, partition function, decay of correlations, complex zeros, thermal phase transition, approximation algorithms, Hamiltonian complexity

ACM Reference Format:

Aram W. Harrow, Saeed Mehraban, and Mehdi Soleimanifar. 2020. Classical Algorithms, Correlation Decay, and Complex Zeros of Partition Functions of Quantum Many-Body Systems. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC '20), June 22–26, 2020, Chicago, IL, USA*. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3357713.3384322>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or to publish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

STOC '20, June 22–26, 2020, Chicago, IL, USA

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-6979-4/20/06...\$15.00
<https://doi.org/10.1145/3357713.3384322>

1 INTRODUCTION

At low temperatures, the main characteristics of many-body systems in condensed matter physics or quantum chemistry are captured in the structure of the ground state of their Hamiltonian. The computational complexity of estimating the ground state energy has been extensively studied through numerous works. In particular, it has been shown that in the worst case, for many physically relevant systems including even a two-local Hamiltonian on a one-dimensional (1D) chain, estimating the ground state energy is QMA-complete [1]. On the other hand, there is a host of classical algorithms for efficiently estimating the ground state energy in certain restricted instances. For examples, the ground state of a gapped Hamiltonian on a 1D chain can be estimated reliably [2]. Also, weaker approximation algorithms are known when the Hamiltonian is on a lattice or it has a dense interaction graph [8].

While at low temperatures the system is in the vicinity of the ground space, at finite temperatures, the state of the system is a mixture of different excited states. In thermal equilibrium, a quantum system characterized by a local Hamiltonian H is in the Gibbs (or thermal) state $\rho = \exp(-\beta H)/Z_\beta(H)$, where β is the inverse of temperature and $Z_\beta(H) = \text{tr}[\exp(-\beta H)]$ is the *partition function* of the system. A natural equivalent to the ground state energy at finite temperatures is the *free energy* which is defined as $F_\beta(H) = -1/\beta \log Z_\beta(H)$. As the temperature decreases ($\beta \rightarrow \infty$), the free energy approaches the ground energy of the Hamiltonian. Many useful statistical properties of the system including the free energy and other quantities such as the entropy can be obtained from the partition function and its derivatives. However, exactly evaluating the partition function is known to be #P-hard. Hence, in order to characterize the finite-temperature behavior of the system, it is crucial to have efficient algorithms that *approximate* this quantity.

Our starting point for finding such approximation algorithms is based on the observation that the phenomenon of the thermal phase transition is an obstacle for finding *efficient* algorithms. Consider a quantum many-body system that consists of n qudits interacting according to a local Hamiltonian H . As the temperature of this system increases, meaning $\beta \rightarrow 0$, the Gibbs state ρ approaches the maximally mixed state $\mathbb{1}/d^n$. Thus, in this case, finding the partition function is trivial since $Z_{\beta=0}(H) = d^n$. On the other hand, this problem becomes significantly harder at lower temperatures. In particular, as $\beta \rightarrow \infty$, the Gibbs state approaches the ground space of the Hamiltonian H and the free energy $F_\beta(H)$ approaches the ground energy which is known to be QMA-hard to estimate. Hence, we see that the computational hardness of estimating the partition function (or equivalently the free energy) depends on the

inverse temperature β and goes through a transition from being trivial to QMA-hard as β increases.

In statistical physics, however, another transition occurs as β increases, namely, the transition in the *phase* of the system. At the thermal phase transition point, certain physical properties of the system undergo an abrupt change. An example of such a transition is when a magnetic material that consists of a network of interacting spins goes from the *ferromagnetic* to the *paramagnetic* phase. In the ferromagnetic phase (the *ordered* phase), most spins are pointing in the same direction and their net magnetic effect is non-zero, whereas in the paramagnetic phase (the *disordered* phase), the spins are distributed equally in opposite directions making their net magnetic effect zero. A distinction between these phases is the presence of *long-range correlations* in the ordered phase and the *decay of correlations* in the disordered phase. This transition does *not* happen gradually as β varies. On the contrary, the phase of the system changes suddenly at some critical inverse temperature β_c known as the phase transition point.

Does the computational hardness of estimating the partition function also undergo an abrupt change at the same transition point? This question has been studied in the *classical Ising* or *hard-core model*, and the answer is known to be affirmative. For these systems, there are efficient algorithms for estimating the partition function when $\beta < \beta_c$ [40, 45] whereas by a result of Sly and Sun [41, 42] the same problem is NP-hard under a randomized reduction for $\beta > \beta_c$.

Hence, it appears that the thermal phase transition poses a barrier to obtaining efficient algorithms, and we need a framework for characterizing this phenomenon. There are at least two methods for such purpose. One, which is the basis of our algorithm, stems from analyzing the locus of the complex zeros of the partition function. Another seemingly different method involves the decay of correlations in the Gibbs state of the system. In this work, we study the interface between the decay of correlations and locus of complex zeros and investigate their algorithmic implications. One central theme of our work is to establish an equivalence between these two concepts. From a physics perspective, this unifies two different characterizations of the phase transition phenomenon. From a computational point of view, this shows that these two distinct frameworks for developing approximation algorithms for the partition function indeed cover the same range of temperatures.

2 MAIN RESULTS

2.1 Algorithm for Quantum Partition Function

Recall that the *classical Ising* model is identified as a weighted graph whose edges have weights $J_{ij} \in \mathbb{R}$ and each vertex i takes a spin value $s_i \in \{\pm 1\}$. For this model, the partition function at inverse temperature β is the sum $Z_\beta = \sum_{s_1, \dots, s_n \in \{\pm 1\}} e^{-\beta \sum_{i \sim j} J_{ij} s_i s_j}$ over different spin configurations. In physics, such a model is identified using the Hamiltonian $H = \sum_{i \sim j} J_{ij} Z_i Z_j$ where $Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$ and Z_i is the tensor product of Z at vertex i with identity matrices elsewhere. This partition function can be equivalently formulated as $Z_\beta(H) = \text{tr}[\exp(-\beta H)]$. For a *quantum* many-body system with Hamiltonian H , the partition function is defined similarly as $Z_\beta(H) = \text{tr}[\exp(-\beta H)] = \sum_k \exp(-\beta E_k)$, where each E_k is an

eigenvalue of the Hamiltonian H . These eigenvalues are also called the energy levels of the Hamiltonian H .

If β is real, the terms $\exp(-\beta E_k)$ are all strictly positive, and hence the partition function $Z_\beta(H)$ is strictly positive itself. However, this changes when β is allowed to be complex. In that case, the terms $\exp(-\beta E_k)$ acquire complex phases that when added together might cancel each other and make the partition function zero. We call the solutions of $Z_\beta(H) = 0$ for $\beta \in \mathbb{C}$ the *complex zeros* of the partition function.

The significance of these zeros becomes more clear if one looks at the free energy $F_\beta(H)$. The zeros of $Z_\beta(H)$ are the *singularities* of $\log Z_\beta(H) = -\beta F_\beta(H)$. Since $Z_\beta(H) \neq 0$ when β is real, we see that all these singularities are located in the complex plane and the free energy is analytic near the real axis. As the number of particles n grows, the number and location of these points change. Perhaps rather surprisingly, some of these singularities approach the real axis in the limit of a large number of particles $n \rightarrow \infty$. The first point on the positive real axis where these zeros converge in the large n limit is called the critical inverse temperature and denoted by β_c (see Figure 1). This critical temperature separates different phases of matter and important quantities such as the free energy become non-analytic in the vicinity of β_c . The study of these complex zeros in connection with the phase transition phenomenon in classical Ising models was initiated by Lee and Yang [31] and later extended by Fisher [20]. This approach is one of the few rigorous methods available in the theory of phase transitions.

Complex zeros can also be studied in relation to designing approximation algorithms. For instance, one can go beyond partition functions and consider complex roots of high-degree polynomials that appear in combinatorics such as estimating the permanent of a matrix. Recently, there has been a surge of interest in studying these complex zeros in theoretical computer science due to their algorithmic applications. In particular, a new approach based on the truncated Taylor expansion introduced by Barvinok [5] directly connects the locus of the complex zeros to approximation algorithms for counting problems. In this work, we extend the scope of this method by applying it to quantum many-body systems.

We first state the condition on the location of zeros that we use in our approximation algorithm. Under this condition, it is guaranteed that the inverse temperature β at which the partition function is estimated is connected to $\beta = 0$ by a path in the complex plane that avoids the complex zeros along its way with a significant margin. Even though this algorithm works for any such path, we restrict our attention to the physically-relevant case when this zero-free region contains the real β -axis. Hence, we define:

DEFINITION 1 (VICINITY OF REAL β AXIS). *The δ -neighborhood of the interval $[0, \beta]$ for some $\beta \in \mathbb{R}^+$ is a region of the complex plane defined as $\Omega_{\delta, \beta} = \{z \in \mathbb{C} : \exists z' \in [0, \beta], |z - z'| \leq \delta\}$ (see Figure 1 for an example of such a region).*

DEFINITION 2 (ANALYTICITY CONDITION). *For a system of n particles with a local Hamiltonian H , we define:*

- (1) *A δ -neighborhood $\Omega_{\delta, \beta}$ of the interval $[0, \beta]$ (see Definition 1) is called *zero-free* if $\forall \beta' \in \Omega_{\delta, \beta}$, the partition function $Z_{\beta'}(H) \neq 0$ and moreover, $|\log Z_{\beta'}(H)| \leq O(n)$.*
- (2) *Equivalently, the free energy $F_\beta(H)$ is called δ -analytic along $[0, \beta]$ if $\Omega_{\delta, \beta}$ is a zero-free region.*

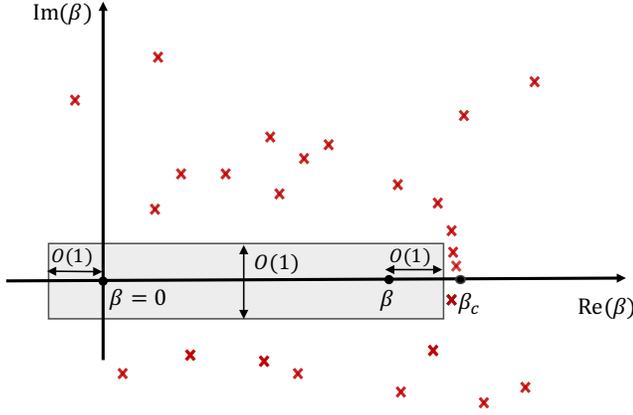


Figure 1: The location of complex zeros of the partition function, the critical point β_c , and the zero-free region near the real axis (as in Definition 2). The free energy is analytic in this region.

While the condition $|\log Z_\beta(H)| \leq O(n)$ is satisfied on the real β axis, it may not hold in the complex plane close to the zeros of $Z_\beta(H)$. If the partition function is a polynomial, which is the case for the classical Ising model, this condition follows when β is constantly far from the complex zeros. In general, though, we need to include this as an independent assumption.

DEFINITION 3. An approximation algorithm for the partition function $Z_\beta(H)$ takes as input the description of the local Hamiltonian H , the inverse temperature β , and a parameter ε and gives an estimate $\tilde{Z}_\beta(H)$ with ε -multiplicative error, i.e.

$$\left| \tilde{Z}_\beta(H) - Z_\beta(H) \right| \leq \varepsilon Z_\beta(H). \quad (1)$$

This is, up to unimportant constants, equivalent to finding an ε -additive error for $\log Z_\beta(H)$ or $F_\beta(H)$.

We now state our first result which shows that the framework of the truncated Taylor expansion [5] can be naturally extended to also estimate quantum partition functions.

THEOREM 4. There is a deterministic classical algorithm that takes a local Hamiltonian H and a number ε as inputs, runs in time $n^{O(\log(n/\varepsilon))}$, and outputs a value within ε -multiplicative error of the partition function $Z_\beta(H)$ at inverse temperature β as long as the free energy is δ -analytic along the $[0, \beta]$ line (see Definition 2) for some $\delta = O(1)$.

The result of Theorem 4 allows us to approximate the partition function, as long as $Z_\beta(H)$ is evaluated for β inside the zero-free region $\Omega_{\delta, \beta}$. Hence, the main challenge in designing an approximation algorithm is to find the zero-free region or the critical point β_c for a Hamiltonian H . This can be achieved in certain systems such as the classical Ising model by using their specific structure [7, 33]. In general, though, it is a hard problem to exactly find this region given an arbitrary Hamiltonian. One can compare this with when a 1D quantum system is assumed to have a constant spectral gap. Under this condition, there is an efficient algorithm for estimating the ground energy. However, it has been shown that validating this

condition, i.e. determining if a Hamiltonian is gapped or not, is undecidable in the worst case [16].

In our next result, we show that the partition function can be approximated using Theorem 4 at high constant temperatures. In particular, we show that for any geometrically-local Hamiltonian H , there exists a zero-free disk of radius β_0 around $\beta = 0$ for some constant $\beta_0 \leq \beta_c$ which depends only on the geometric parameters of H . Here, we say a Hamiltonian is geometrically local if the local terms in H act on neighboring qudits that are located on a D -dimensional lattice $\Lambda \subset \mathbb{Z}^D$.

THEOREM 5. There exists a real constant β_0 such that for all $\beta \in \mathbb{C}$ with $|\beta| \leq \beta_0$, the partition function $Z_\beta(\Lambda)$ of a geometrically-local Hamiltonian H does not vanish and $|\log |Z_\beta(\Lambda)|| \leq O(n)$.

As mentioned earlier, without assuming a specific structure on the family of Hamiltonians, extending the zero-free region beyond β_0 in Theorem 5 seems implausible. Hence, an alternative approach is to show the absence of zeros at a given β by assuming the validity of other conditions such as the decay of correlations. This along with other results in this direction is the subject of the next section.

2.2 Correlation Decay in the Gibbs State

Another signature of the thermal phase transition is the appearance of long-range order in the system. In the example of a magnetic system, below the phase transition in the ordered phase, distant spins are correlated and point in the same direction, whereas in the disordered phase, the correlations between disjoint parts of the system decay exponentially with their distance. More precisely, we define the exponential decay of correlations as:

DEFINITION 6 (CORRELATION DECAY CONDITION). The Gibbs state $\rho_\beta(H)$ of a geometrically-local Hamiltonian H at inverse temperature β exhibits an exponential decay of correlations if for any two disjoint observables O_1 and O_2 there exist constants ξ and c such that

$$\begin{aligned} & \left| \text{tr} [\rho_\beta(H) O_1 O_2] - \text{tr} [\rho_\beta(H) O_1] \text{tr} [\rho_\beta(H) O_2] \right| \\ & \leq c \|O_1\| \|O_2\| e^{-\text{dist}(O_1, O_2)/\xi}. \end{aligned} \quad (2)$$

What is the relation between the decay of correlations and the complex zeros of the partition function? Note that the former involves correlations in the system at a real temperature while the latter concerns the complex temperature features of the partition function. Could it be that these two apparently distinct characterizations are indeed equivalent?

Besides its physical significance, the correlation decay property is also crucially used in many approximation algorithms both in classical [45] and quantum [9] settings. Hence, we see that there are two different approaches for estimating the partition function: using the absence of complex zeros versus relying on the decay of correlations. A third approach is to use Markov chain Monte Carlo (MCMC) sampling algorithms for this purpose. An important question is whether the range of temperatures that these widely different approaches cover is the same.

The equivalence between the exponential decay of correlations (fast mixing in space) and fast convergence of MCMC algorithms (fast mixing in time) has been established before for classical systems [18, 44]. Here, we consider the equivalence between the absence of complex zeros and the exponential decay of correlations.

This question has been recently raised in [33] where one direction of this equivalence was proved for the special case of the classical Ising model [33]. Here, we study this question in the context of quantum many-body systems.

We build on a work of Dobrushin and Shlosman [17] who proved this equivalence for arbitrary translationally-invariant classical systems. We are not aware of any application of the proof techniques in this paper or [17] that appear in the recent results on approximate counting problems. Hence, the methods developed here might be of independent interest.

Before stating our result, we need to generalize the definition of the zero-free (or analytic) region in Definition 2 to a multivariable setting. Given a local Hamiltonian $H = \sum_{i=1}^m H_i$, we define complex variables z_1, \dots, z_m such that each z_i roughly equals β plus some small *complex deviation*. Hence, instead of working with functions of βH such as $\exp(-\beta H)$, we consider functions of $\sum_{i=1}^m z_i H_i$ as in $\exp(-\sum_{i=1}^m z_i H_i)$. For a fixed inverse temperature β and maximum deviation δ , we denote the set of such tuples (z_1, \dots, z_m) by $\Gamma_{\delta, \beta}$. By varying β from zero to some constant β and taking the union of the corresponding $\Gamma_{\delta, \beta}$, the set $\Omega_{\delta, \beta}$ is obtained.

As discussed earlier, the critical temperature β_c corresponds to the thermal phase transition point, where complex zeros of the partition function approach the real axis. Note that even with deviations, we do not want any of the variables z_i to exceed β_c . More precisely, we have the following definition:

DEFINITION 7 (VICINITY OF REAL β AXIS, MULTIVARIATE VERSION). Let $\Gamma_{\delta, \beta}$ be the set $\{(z_1, \dots, z_m) : \forall i \in [m], z_i \in \mathbb{C}, |z_i - \beta| \leq \delta\}$. We define $\Omega_{\delta, \beta}$ to be $\Omega_{\delta, \beta} = \bigcup_{\beta' \in \mathbb{R}^+, \beta' < \beta/(1+\delta)} \Gamma_{\delta, \beta'}$.

We also define the perturbed Gibbs state as follows.

DEFINITION 8 (COMPLEX PERTURBED GIBBS STATE). The δ -perturbed Gibbs state of a local Hamiltonian $H = \sum_{i=1}^m H_i$ at inverse temperature β is defined as

$$\rho_{\vec{z}}(H) = \frac{e^{-\sum_{i=1}^m z_i H_i}}{\text{tr}[e^{-\sum_{i=1}^m z_i H_i}]}, \quad \vec{z} = (z_1, \dots, z_m) \in \Gamma_{\delta, \beta} \quad (3)$$

where $\Gamma_{\delta, \beta}$ is defined in Definition 7.

The analyticity condition we consider for the result in this section is stronger than the one used in the approximation algorithm in Definition 3 and Theorem 4. Previously we only included systems with open boundary conditions in our analysis, but here we also need to allow for other boundary conditions. This is not restricted to the quantum case, and Dobrushin and Shlosman use similar conditions in their proof for classical systems [17]. The precise statement of our condition is the following:

DEFINITION 9 (STRONG ANALYTICITY). The free energy of a geometrically-local Hamiltonian H is strongly δ -analytic at β if for any local operator $N \geq 0$ with $\|N\| = 1$, there exists a constant c such that

$$\left| \log \left(\text{tr} \left[e^{-\sum_{i=1}^m z_i H_i} N \right] \right) \right| \leq cn, \quad \forall (z_1, \dots, z_m) \in \Gamma_{\delta, \beta}. \quad (4)$$

Assuming Definition 9, we show that the absence of complex zeros around some real β implies the exponential decay of correlations at that β .

THEOREM 10 (ANALYTICITY IMPLIES DECAY OF CORRELATIONS). Let $\rho_{\beta}(H)$ be the Gibbs state of a geometrically-local Hamiltonian at inverse temperature β in the strong analyticity region $\Gamma_{\delta, \beta}$ given in Definition 9. This state has the decay of correlation property as in Definition 6 in any of the following cases:

- (i.) The distance between the observables O_1 and O_2 is at least $\Omega(\log n)$,
- (ii.) The Hamiltonian H is the sum of mutually commuting local terms, or
- (iii.) The Hamiltonian H is defined on a 1D chain.

The class of commuting Hamiltonians includes important examples such as stabilizer Hamiltonians like the Toric code, Color code, or Levin-Wen model [32].

To see the motivation for the strong analyticity (as in Definition 9) used in this theorem, first note that when restricted to classical systems, this condition reduces to the one used in [17]. There the operator N sets the boundary conditions which fixes the value of certain spins in the system before computing the partition function, or more generally, finding the Gibbs distribution. A natural question then is how varying these boundary conditions affects the distribution. In particular, the *uniqueness of the Gibbs distribution* refers to the case that in the limit of a large number of particles, changing distant spins has a negligible effect on the distribution of spins in a finite region. Hence, a *unique* Gibbs distribution can be defined for such systems. This condition is not satisfied at all temperatures, and below the critical temperature, multiple Gibbs distributions exist. Thus, it seems natural to include the boundary conditions in the partition function when studying its complex zeros and the critical behavior of the system in general.

For quantum systems, one can think of fixing the boundary spin values by projecting them onto a specific state or more generally by post-selecting after a local measurement has been performed. Hence, $\text{tr} \left[\exp \left(-\sum_{i=1}^m z_i H_i \right) N \right]$ is the partition function of the normalized Gibbs state after conditioning on the measurement outcome associated with N . Notice that, in principle, the state of the spins after post-selection can be entangled. As we will see, this causes technical difficulties in extending the classical results to the quantum regime.

Finally, we mention that the validity of Definition 9 can be shown in the high temperature regime ($\beta \leq \beta_0$) by a slight modification of the argument in the proof of Theorem 5.

Proving the converse of Theorem 10 turns out to be more challenging. Nevertheless, we can give evidence for this direction by generalizing the result of [17] to classical systems that are not translationally invariant, and also quantizing certain steps in the proof. By fully establishing this equivalence, one could rigorously confirm the physical intuition that a quantum system enters the disordered phase at the point where the free energy becomes analytic.

THEOREM 11. Let H be a geometrically-local Hamiltonian of a classical spin system, i.e. the local terms H_i are all diagonal in the same product basis. For this system, the exponential decay of correlations given in Definition 6 implies the absence of zeros near the real axis as in Definition 2.

In these results, we consider general Hamiltonians without specializing to specific models. Tighter results have been obtained

before for specific models [13, 33, 38]. For instance, [33] proves that a variant of the decay of correlations property in the classical Ising model implies analyticity. This is based on a different framework than ours and uses the self-avoiding walk trees and a carefully designed recurrence as introduced in [46]. The analysis of [33] is fine-tuned for the Ising model and implies stronger results than those we achieve. This includes exactly matching the phase transition point even with arbitrary boundary conditions. Another approach, based on truncating a convergent cluster expansion, is used in [13] to obtain the decay of correlations for hard-core models on bipartite graphs.

2.3 Two-Local Hamiltonian and Lee-Yang Zeros

For our last result, we switch gears and focus on a specific family of 2-local Hamiltonians. We again use the idea of extrapolation, but this time, our extrapolation parameter instead of β is the strength of the external magnetic field applied to the system in the z -direction. The physical motivation is that when the system is subject to a large external field in a specific direction (the z -direction in our case), all spins align themselves in that direction, and estimating the properties of the system becomes trivial. On the other hand, as we move to smaller fields, the other interaction terms between the particles gain significance, making the problem non-trivial. Our result is an approximation algorithm for the quantum XXZ model with the following Hamiltonian:

DEFINITION 12. *The anisotropic XXZ Hamiltonian on an interaction graph $G = (V, E)$ is given by*

$$H(\mu) = - \sum_{(i,j) \in E} \left(J_{ij}(X_i X_j + Y_i Y_j) + J_{ij}^{zz} Z_i Z_j \right) - \mu \sum_{i \in V} Z_i. \quad (5)$$

We find an approximation algorithm for this model. This is stated in the following theorem. This model is outside the family of ferromagnetic systems considered previously in [11] and to the best of our knowledge, no efficient algorithm was previously known for estimating its partition function.

THEOREM 13. *There is a deterministic algorithm that runs in $n^{O(\log(n/\epsilon))}$ time and outputs an ϵ -multiplicative approximation to the partition function of the anisotropic XXZ model (see Definition 12) in the ferromagnetic regime, i.e. when $J_{ij}^{zz} \geq |J_{ij}|$, and μ is an arbitrary non-zero constant.*

3 PRELIMINARIES

Consider a D -dimensional lattice $\Lambda \subset \mathbb{Z}^D$ containing n sites with a d -dimensional particle (qudit) on each site. The Hilbert space is $\mathcal{H} = \bigotimes_{i \in \Lambda} \mathcal{H}_i$ where \mathcal{H}_i is the local Hilbert space of site i . For a region $A \subseteq \Lambda$, we denote its size by $|A|$ and its complement by \bar{A} . The diameter of A is defined to be $\text{diam}(A) = \sup\{\text{dist}(x, y) : x, y \in A\}$. The interaction of these particles is described by a local Hamiltonian H that has the following form:

$$H = \sum_{X \subseteq \Lambda} H_X. \quad (6)$$

Each term H_X is a Hermitian operator with operator norm at most h that is acting non-trivially only on the sites in X . We denote this by writing $\text{supp}(H_X) = X$. The local terms H_X do not necessarily commute with each other. Similarly, we define $H_A = \sum_{X \subseteq A} H_X$

to be the Hamiltonian restricted to a region $A \subseteq \Lambda$. We denote the number of local terms in the Hamiltonian by m and often also write $H = \sum_{i=1}^m H_i$. The 1-norm of an operator O is denoted by $\|O\|_1 = \text{tr}[\sqrt{O^\dagger O}]$ and its operator norm, i.e. the largest singular value, by $\|O\|$.

In order to impose geometric locality on the interactions between the particles, we consider the interactions that satisfy the following condition.

DEFINITION 14 (GEOMETRICALLY-LOCAL HAMILTONIANS). *A Hamiltonian $H = \sum_{X \subseteq \Lambda} H_X$ such that $\text{supp}(H_X) = \emptyset$ when $\text{diam}(X) > R$ or $|X| > \kappa$ is called a (κ, R) -local Hamiltonian. We call κ the locality and R the range of H . We use the words geometrically-local and (κ, R) -local interchangeably when κ, R are kept constant.*

This should be contrasted with the case where $\text{supp}(H_X) = \emptyset$ when $|X| > \kappa$ but there is no restriction on $\text{diam}(X)$. In order to distinguish between these two, we use the terms *geometrically-local* versus *local* throughout this paper. We also focus mostly on geometrically-local Hamiltonians with a *finite range* R , but most of our results also apply to Hamiltonians with interactions that decay fast enough, for example, with some exponential rate.

REMARK 15. *In general, the locality κ of a geometrically-local Hamiltonian on a D -dimensional lattice Λ can be bounded as $\kappa \leq O(R^D)$, which is the size of a ball of diameter R . Nevertheless, we treat both κ and R as independent parameters in this paper.*

Our approximation algorithm is based on the truncated Taylor series method which is summarized in the following proposition. The proof can be found in the full version [21], but also see the discussion in the next section.

PROPOSITION 16 (TRUNCATED TAYLOR SERIES FOR BOUNDED FUNCTIONS AND POLYNOMIALS). *We denote a disk of radius b centered at the origin in the complex plane by Δ_b , that is $\Delta_b = \{z \in \mathbb{C} : |z| \leq b\}$.*

- (1) *Let $f(z)$ be a complex function that is analytic and bounded as $|f(z)| \leq M$ when $z \in \Delta_b$ for a constant $b > 1$. Then the error of approximating $f(z)$ by a truncated Taylor series of order K for all $|z| \leq 1$ is bounded by*

$$\left| f(z) - \sum_{k=0}^K a_k z^k \right| \leq \frac{M}{b^K (b-1)}, \quad |z| \leq 1. \quad (7)$$

- (2) *Assume b is fixed and there is a deterministic algorithm that finds the coefficients a_k in time $O(N^k)$ for some parameter N . Then there exists a deterministic algorithm with running time $N^{O(\log(M/\epsilon))}$ that outputs an ϵ -additive approximation for $f(z)$.*
- (3) *[cf. [5]] Let $f(z) = \log(g(z))$ for some polynomial $g(z)$ of degree N that does not vanish when $z \in \Delta_b$. The error of approximating $f(z)$ by a truncated Taylor series of order K for $|z| \leq 1$ is bounded by $\frac{N}{K+1} \frac{1}{b^K (b-1)}$.*
- (4) *[cf. [5]] Assuming b is fixed, there exists a deterministic algorithm with running time $N^{O(\log(N/\epsilon))}$ that outputs an ϵ -additive approximation for $\log(g(z))$.*

4 SKETCH OF OUR PROOFS

Here, we give an overview of our proof techniques. More details can be found in the full version of this paper [21].

Sketch of the proof for Theorem 10. The technique used in the proof of Theorem 10 is inspired by the extrapolation idea of Theorem 4 and also the proof of the similar statement for the classical systems due to [17].

One major issue that appears in the proof of this Theorem and Theorem 11 is the handling of entangled boundary conditions. To address this, we consider the Gibbs state after a subset of spins have been measured. This means we work with partition functions of the form $\text{tr}[\exp(-\beta H)N]$ for some positive semi-definite operator N . We then define a function $f(\beta)$ that measures the correlation between disjoint observables O_1 and O_2 . This function is defined in a slightly different way than the *covariance* form in (2) and is tuned to have specific properties. In particular, we show that at $\beta = 0$, the value of this function is zero, i.e. $f(0) = 0$. This is expected intuitively since the system is in the maximally mixed state at $\beta = 0$ and particles are distributed independently at random. However, we further show that the low-order derivatives of this function up to $O(\text{dist}(O_1, O_2))$ are all zero at $\beta = 0$, i.e.

$$\left. \frac{d^k f(\beta)}{d\beta^k} \right|_{\beta=0} = 0, \quad \text{for } k = 0, 1, \dots, O(\text{dist}(O_1, O_2)). \quad (8)$$

Hence, this function looks very flat around the origin. Additionally, we prove that $f(\beta)$ is an analytic function in the zero-free region. Finally, we show that this together with the constraints on the derivatives imply that the value of $f(\beta)$, which shows how correlated O_1 and O_2 are, remains exponentially small when moving from the origin to a constant β .

This gives us an upper bound $\propto n \exp(-\text{dist}(O_1, O_2)/\xi)$ on the amount of correlation. The extra factor of n makes this bound exponentially small when $\text{dist}(O_1, O_2) = \Omega(\log n)$.

REMARK 17. *Even with the extra factor of n , our bound remains useful for algorithmic applications such as in [9]. There one needs to split the system into computationally tractable smaller pieces and solve the problem for those pieces locally. The error of this strategy can be bounded using the exponential decay of correlations. To keep this error less than $1/\text{poly}(n)$, one needs to choose the distances to be $O(\log n)$ which is the regime that our result covers.*

In classical systems, one can remove the constraint $\text{dist}(O_1, O_2) = \Omega(\log n)$ by using the Markov property of the Gibbs states. This property is known not to (exactly) hold in the quantum case. We can get around this issue in certain instances. This includes when the Hamiltonian consists of commuting terms or when it is defined on a 1D chain. In both cases, using either the commutativity of local terms or quantum belief propagation [23] (refer to the full version of the paper [21] for the precise statement), we show that by removing the interaction terms acting on particles that are far from the observables O_1 and O_2 , the correlations between O_1 and O_2 do not change by much. Hence, the system size reduces to the number of particles in the vicinity of the two observables. This number replaces the prefactor n we had before and is negligible compared to the exponential factor $\exp(-\text{dist}(O_1, O_2)/\xi)$. Thus, for these systems, the decay of correlations holds even when $\text{dist}(O_1, O_2)$ is a constant. In higher dimensions, using quantum belief propagation results in an error proportional to the size of the boundary which restricts its application for our purpose.

Sketch of the proofs for Theorem 5 and Theorem 11. We first introduce a core idea which plays a central role in the proofs of both Theorem 5 and Theorem 10. For ease of notation, we denote the partition function of a geometrically-local Hamiltonian H defined over a D -dimensional lattice $\Lambda \subset \mathbb{Z}^D$ by $Z_\beta(\Lambda)$. The particles are located on the vertices of this lattice.

In Theorem 5, our goal is to show that $Z_\beta(\Lambda) \neq 0$ inside a disk of radius β_0 , i.e. for $\beta \in \mathbb{C}$ where $|\beta| \leq \beta_0$ for some constant β_0 . We consider a series of sublattices $\emptyset = \Lambda_0 \subset \Lambda_1 \subset \Lambda_2 \subset \dots \subset \Lambda_n = \Lambda$ such that each sublattice Λ_i has one fewer vertex than Λ_{i+1} . By convention, we let $Z_\beta(\emptyset) = 1$. As long as the sublattice Λ_i has only a constant number of particles, we can always ensure $Z_\beta(\Lambda_i) \neq 0$ by choosing β to be a sufficiently small constant. One might worry that by adding more particles, the partition function vanishes.

Our main contribution is to prove this does not happen. We do so by showing that the partition function after involving new particles does not become smaller than a constant fraction of the partition function before adding the particles. In other words, we show there exists a constant $c > 1$ such that

$$|Z_\beta(\Lambda_{i+1})| \geq c^{-1} |Z_\beta(\Lambda_i)|, \quad i \in \{1, 2, \dots, n-1\}. \quad (9)$$

By repeatedly applying this bound, we obtain the following exponentially small (yet sufficiently large for our purposes) lower bound on the partition function of the whole system:

$$|Z_\beta(\Lambda)| \geq c^{-n} |Z_\beta(\Lambda_1)|. \quad (10)$$

This leads to the bound given in Theorem 5. This lower bound is obtained using a method known as the *cluster expansion*. Since the local terms in the Hamiltonian of a quantum system do not necessarily commute, applying this method becomes quite technical. The cluster expansion that we use is due to Hastings [22, 28], which represents the operator $\exp(H)$ as sum of products of local terms H_i . This allows us to express $Z_\beta(\Lambda_{i+1})$ in terms of $Z_\beta(\Lambda_i)$ plus some small correction terms that account for the interaction terms acting on the added particle. This by itself does not lead to a bound on the partition function. Our technical contribution is to use an inductive proof to connect such a decomposition to the lower bound (9).

A similar strategy is used in the proof of Theorem 11 which is based on the result of [17] for translationally-invariant classical systems. We essentially show a similar bound to (9) on how much the partition function can shrink after adding new particles. Here, instead of cluster expansions, we use the exponential decay of correlations to show such a lower bound. However, notice that the decay of correlations is a property of the system at a *real* β , whereas we want to bound the absolute value of the partition function at some *complex* β . There are multiple steps in the proof before we can get around this issue.

One crucial step is to reduce the proof of the analyticity of the free energy to a condition that roughly speaking states that changing the value of a spin in the system only causes a small relative change in the partition function of the system even for complex β . We prove this by isolating the effect of this spin flip from the rest of the system using the decay of correlations. This requires removing the imaginary part of β for all the interactions in the vicinity of the flipped spin and bounding the resulting error.

This overall approach involves a subtle use of the boundary conditions in the spin system. In the quantum case, this means

applying local projectors (or more generally a positive operator) to the Gibbs state before evaluating the partition function. These projectors can in general be entangled which makes using this proof technique more challenging for quantum systems.

Sketch of the proof for Theorem 4. The basis of our algorithm in Theorem 4 is the following observation. It is computationally easy to find the partition function and its derivatives at $\beta = 0$. Note that in a system of n qudits, $Z_{\beta=0}(H) = d^n$ and its derivatives are

$$\left. \frac{d^k Z_{\beta}(H)}{d\beta^k} \right|_{\beta=0} = (-1)^k \text{tr}[H^k]. \quad (11)$$

Since the local Hamiltonian H equals $\sum_{i=1}^m H_i$ for some $m = \text{poly}(n)$, its k th power H^k is also the sum of $n^{O(k)}$ many local terms, i.e.

$$H^k = \sum_{j=1}^{n^{O(k)}} H_j^{(k)}, \quad (12)$$

where $H_j^{(k)}$ is a product of k local terms H_i . Each of the new terms $H_j^{(k)}$ acts on a region that is at most k times larger than the support of the original terms H_i which is still some constant. We can find $\text{tr}[H^k]$ by adding $n^{O(k)}$ many terms like $\text{tr}[H_j^{(k)}]$. Furthermore, since the support of each $H_j^{(k)}$ is $\leq k$, each term $\text{tr}[H_j^{(k)}]$ can be computed in time $2^{O(k)}$. Hence, the derivatives of (11) can be computed in time $n^{O(k)}$.

How can the solution at $\beta = 0$ be used to estimate the one at some non-zero β ? We use a technique due to Barvinok [4, 6] that has been applied to similar counting problems. The idea is to *extrapolate* this solution at $\beta = 0$ to find $Z_{\beta}(H)$ at some non-zero β where the problem is non-trivial. The extrapolation is done simply by using a *truncated Taylor expansion* of $\log Z_{\beta}(H)$ at $\beta = 0$. Since our goal is to find the partition function with some ε -multiplicative error, it is sufficient to estimate $\log Z_{\beta}(H)$ within ε -additive error.

The main barrier to the reliability of this algorithm is establishing the fast convergence of the Taylor expansion. Such a Taylor expansion is only valid when $\log Z_{\beta}(H)$ remains a complex-analytic function, meaning the extrapolation is done along a path contained in the zero-free region. This is precisely the condition stated in Definition 2. Under this assumption, the Taylor theorem along with the bound $|\log Z_{\beta}(H)| \leq O(n)$ that we get from being in the zero-free region give

$$\left| \log Z_{\beta}(H) - \sum_{k=0}^{K-1} \frac{1}{k!} \left. \frac{d^k \log Z_{\beta}(H)}{d\beta^k} \right|_{\beta=0} \right| \leq c_1 n e^{-c_2 K} \quad (13)$$

for some constants c_1, c_2 . The running time of computing the terms in this expansion is dominated by that of finding the derivatives which, as mentioned earlier, takes time $n^{O(K)}$. To get an additive error of ε for $\log Z_{\beta}(H)$, it suffices to choose $K = O(\log(n/\varepsilon))$, resulting in a quasi-polynomial time algorithm.

The running time of this algorithm depends exponentially on the distance between the zeros and the extrapolation path. This allows us to clearly see why our algorithm fails beyond the phase transition point. If we try to extrapolate to $\beta \geq \beta_c$, we need to find a zero-free region that avoids the “armor” of zeros that are

concentrated around the real axis at β_c . This results in a zero-free region with a vanishing width. Hence, the running time blows up, which matches our expectation from the NP hardness result above β_c [42].

Sketch of the proof for Theorem 13. Thus far we have only considered complex zeros of the partition function as a function of β . These are often called Fisher zeros [20]. One can, however, fix β and consider the partition function as a function of other parameters in the Hamiltonian. When that parameter is the strength of the external magnetic field denoted by μ , these zeros are called Lee-Yang zeros [31]. In a pioneering result, Lee and Yang showed that for ferromagnetic systems, the locus of these zeros can be exactly determined and they are all on the imaginary axis in the complex μ -plane.

A generalization of this theorem has been proved for a class of 2-local quantum systems including the anisotropic Heisenberg model [43]. The result follows by mapping the quantum system to a classical spin system and applying a Lee-Yang type argument to the classical model.

Knowing the location of the complex zeros, we use the extrapolation algorithm to estimate the solution at a constant μ by finding the low-order derivatives of the partition function at $\mu = 0$. We can apply this to the quantum XXZ model given in (5).

5 PREVIOUS WORK

Classical Statistical Physics and Combinatorial Counting: The Gibbs distribution and partition function appear naturally in combinatorial optimization, statistical physics, and machine learning. In particular, the classical *Ising model* has been studied extensively within these areas. These studies have resulted in various probabilistic and deterministic approximation algorithms for this model and its variants. In the following, we summarize some of these results.

Most notable and the first rigorously proven efficient algorithm for the Ising model is the result of Jerrum and Sinclair [24] that uses a Markov chain Monte Carlo (MCMC) sampling algorithm to estimate the partition function in the ferromagnetic regime on arbitrary graphs. More generally, it has been shown that one can set up Markov chains for sampling from the Gibbs distribution that mix rapidly *if and only if* the correlations decay exponentially. This is known as the equivalence of mixing in time and mixing in space [18, 44].

Another approach uses the decay of correlations in the Gibbs distribution. This property essentially allows one to decompose the interaction graph of the system into smaller computationally tractable pieces, and then combine the results of the computation on those pieces to find the overall partition function. In contrast to the MCMC approach, algorithms based on the decay of correlations can be deterministic. This approach, for instance, has led to efficient deterministic algorithms for the hard-core model up to the hardness threshold [45] and the antiferromagnetic Ising model [40].

There is a recent conceptually different approach to estimating the partition function, which is the basis of this work. This approach views the partition function as a high-dimensional polynomial and uses the truncated Taylor expansion to extend the solution at a computationally easy point to a non-trivial regime of parameters.

Since its introduction [5], this method has been used to obtain deterministic algorithms for various interesting problems such as the ferromagnetic and antiferromagnetic Ising models [34, 37] on bounded graphs.

The question of the relation between the analyticity of the free energy and the decay of correlations was recently considered in [33] where the authors show that the correlation decay implies the absence of Fisher zeros near the real axis. A more general statement has been proved by Dobrushin and Shlosman [17] for translationally-invariant classical systems.

Quantum Many-Body Systems: The problem of estimating the partition function and correlation decay in quantum systems has also been studied in the past. We review some of these results here.

There are various results (e.g., [14, 39]) that estimate the partition function by sampling from the Gibbs state using a quantum computer (also known as quantum Gibbs sampling). The best known bound on the running time of these algorithms is exponential in the number of particles. This running time can be reduced if we assume other conditions. For example, [25] shows that a strong form of the decay of correlations implies an efficient quantum Gibbs sampler for commuting Hamiltonians. If in addition to the decay of correlations we add the decay of quantum conditional mutual information, then this result can be extended to non-commuting Hamiltonians [9]. Turning these quantum algorithms into classical ones results in an $n^{\text{polylog}(n)}$ running time. Although we cannot directly compare these results with our algorithm due to different conditions that are imposed, the $n^{O(\log n)}$ running time that we achieve outperforms that of these algorithms.

Considering the success of approximation schemes for the classical statistical problems, it is desirable to import those results to evaluate the thermal properties of interacting quantum many-body systems. This indeed can be done for some models like the quantum transverse field Ising model [10, 15] or the quantum XY model [11] in the ferromagnetic regime using what is called the *quantum-to-classical mapping*. However, this approach only works for a restricted set of Hamiltonians known as stoquastic Hamiltonians in which all off-diagonal matrix elements are real and non-positive. This set is known to be restricted in many ways. For example, estimating the ground state energy of a general quantum Hamiltonian is QMA-complete but the problem is in AM when we restrict to stoquastic Hamiltonians.

Establishing the decay of correlations in the Gibbs state has also been studied in quantum settings. In particular, it has been shown that the Gibbs state has this property in the 1D translationally invariant case [3] or above some constant temperature in higher dimensions [28]. Thus, in these regimes, there exist efficient representations for the state of the system using a tensor network ansatz like matrix product states or projected entanglement pair states [22, 28, 36]. However, this does not necessarily imply an efficient algorithm that finds and faithfully manipulates these tensor networks.

The decay of conditional mutual information is another property of the Gibbs state that has been rigorously proved for 1D systems [26] and conjectured for higher dimensions. This result has been used to find algorithmic schemes for preparing the Gibbs state on a quantum computer [9] or estimating the free energy in 1D

[27, 30]. A recent result of [29] uses cluster expansions along with a technique very similar to the one we use in Theorem 10 (i.e. showing the low-order derivatives of the correlation function are zero) to establish the decay of conditional mutual information above some constant temperature.

6 DISCUSSION AND OPEN QUESTIONS

Our work raises many questions that we leave for future work. Here we mention some of them.

(1) Perhaps the most immediate problem is to fully establish (or refute) the connection between the decay of correlations and the absence of zeros. There are at least two directions to pursue. First, it would be interesting to prove the exponential decay of correlations in the zero-free region of non-commuting Hamiltonians in higher dimensions. Currently we can only show this when the distance of the observables is $\Omega(\log n)$. It seems for this to work, the region of applicability of certain tools such as quantum belief propagation needs to be extended to the complex regime.

Additionally, establishing the absence of zeros in quantum systems when the correlations decay exponentially is also open. A first step might be to prove this for commuting Hamiltonians or 1D chains. In this work, we have already extended some parts of the proof of this statement for the classical systems to commuting Hamiltonians, but it seems to complete the proof, a more careful analysis of the entangled boundary conditions is required.

(2) While we focus on the covariance form of the correlations (2), one can also consider quantum conditional mutual information (qCMI) as a measure of correlations. Using the absence of zeros to prove the decay of qCMI is another interesting question. This would extend the result of [29] to lower temperatures down to the phase transition point. Since the approach of [29] resembles some of the techniques we use, this looks like a promising direction.

(3) Is there some range of temperatures or Hamiltonian parameters that a quantum computer cannot efficiently sample from the Gibbs state but the extrapolation technique still works? At least, when the parameter of interest is temperature, this depends on the fate of the previous questions we mentioned, i.e. showing that the decay of correlations and qCMI are necessary for the absence of zeros. The result of [9] implies an efficient quantum sampler under the same conditions. Are there other parameters besides temperature for which one can show a separation between these notions?

(4) Is it possible to improve the lower bound we obtained for the critical point β_c in Theorem 5 without using other conditions such as the decay of correlations? In general, what is the computational hardness of determining the thermal phase transition point β_c ?

(5) Can the running time of our algorithm be improved for specific systems to polynomial time? For this to be done, we need an algorithm that given a local Hamiltonian $H = \sum_{\ell=1}^{O(n)} H_{\ell}$, computes $\text{tr}[H^k]$ in time $O(n \cdot 2^{O(k)})$ instead of the current $n^{O(k)}$ running time. This has been achieved for the classical Ising model [34, 37] by relating the derivatives of the partition function to combinatorial objects that can be efficiently counted. Another approach for an improved running time is recently introduced in [29], where the authors apply a multivariate version of the cluster expansion that we use to compute the derivatives of $\log Z_{\beta}(H)$ efficiently.

(6) Can we use the extrapolation idea to avoid the *sign problem*? The easy regime, which includes the starting point of the extrapolation, could be a regime of parameters where the Hamiltonian is sign-free and MCMC algorithms yield a good estimate, whereas the end point is where the sign problem exists. A candidate parameter for extrapolation is the chemical potential. There are important physical systems such as lattice gauge theories for which at zero chemical potential the partition function is sign-free while there is a severe sign problem for non-zero chemical potentials.

(7) Barvinok's approach has been used to obtain approximation algorithms for other problems related to quantum computing [12, 19, 35]. Are there other relevant applications for this method?

ACKNOWLEDGEMENTS

We thank Fernando Brandão, Kohtaro Kato, Tomotaka Kuwahara, Zeph Landau, Milad Marvian, and John Wright for helpful discussions. This work was funded by NSF grants CCF-1452616, CCF-1729369, PHY-1818914; ARO contract W911NF-17-1-0433; and a Samsung Advanced Institute of Technology Global Research Partnership. The Institute for Quantum Information and Matter is an NSF Physics Frontiers Center.

REFERENCES

- [1] AHARONOV, D., GOTTESMAN, D., IRANI, S., AND KEMPE, J. The power of quantum systems on a line. *Communications in Mathematical Physics* 287, 1 (2009), 41–65.
- [2] ARAD, I., LANDAU, Z., VAZIRANI, U., AND VIDICK, T. Rigorous RG algorithms and area laws for low energy eigenstates in 1D. *Comm. Math. Phys.* 356, 1 (2017), 65–105.
- [3] ARAKI, H. Gibbs states of a one dimensional quantum lattice. *Communications in Mathematical Physics* 14, 2 (1969), 120–157.
- [4] BARVINOK, A. Computing the partition function for cliques in a graph. *Theory of Computing. An Open Access Journal* 11 (2015), 339–355.
- [5] BARVINOK, A. *Combinatorics and complexity of partition functions*, vol. 30 of *Algorithms and Combinatorics*. Springer, Cham, 2016.
- [6] BARVINOK, A. Computing the permanent of (some) complex matrices. *Foundations of Computational Mathematics* 16, 2 (2016), 329–342.
- [7] BARVINOK, A., AND SOBERÓN, P. Computing the partition function for graph homomorphisms. *Combinatorica* 37, 4 (2017), 633–650.
- [8] BRANDÃO, F., AND HARROW, A. Product-state approximations to quantum ground states. In *Proceedings of the 2013 ACM Symposium on Theory of Computing—STOC 2013* (2013), ACM, New York, pp. 871–880.
- [9] BRANDÃO, F., AND KASTORYANO, M. J. Finite correlation length implies efficient preparation of quantum thermal states. *Communications in Mathematical Physics* (2016), 1–16.
- [10] BRAVYI, S. Monte Carlo simulation of stoquastic Hamiltonians. *Quantum Information & Computation* 15, 13–14 (2015), 1122–1140.
- [11] BRAVYI, S., AND GOSSET, D. Polynomial-time classical simulation of quantum ferromagnets. *Physical Review Letters* 119, 10 (2017), 100503.
- [12] BRAVYI, S., GOSSET, D., AND MOVASSAGH, R. Classical algorithms for quantum mean values. *arXiv preprint arXiv:1909.11485* (2019).
- [13] CANNON, S., AND PERKINS, W. Counting independent sets in unbalanced bipartite graphs. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms* (2020), SIAM, pp. 1456–1466.
- [14] CHOWDHURY, A. N., AND SOMMA, R. Quantum algorithms for Gibbs sampling and hitting-time estimation. *Quantum Information & Computation* 17, 1–2 (2017), 41–64.
- [15] CROSSON, E., AND HARROW, A. W. Rapid mixing of path integral Monte Carlo for 1d stoquastic Hamiltonians. *arXiv preprint arXiv:1812.02144* (2018).
- [16] CUBITT, T. S., PEREZ-GARCIA, D., AND WOLF, M. M. Undecidability of the spectral gap. *Nature* 528, 7581 (2015), 207.
- [17] DOBRUSHIN, R., AND SHLOSMAN, S. Completely analytical interactions: constructive description. *Journal of Statistical Physics* 46, 5–6 (1987), 983–1014.
- [18] DYER, M., SINCLAIR, A., VIGODA, E., AND WEITZ, D. Mixing in time and space for lattice spin systems: a combinatorial view. *Random Structures Algorithms* 24, 4 (2004), 461–479.
- [19] ELДАР, L., AND MEHRABAN, S. Approximating the permanent of a random matrix with vanishing mean. In *59th Annual IEEE Symposium on Foundations of Computer Science—FOCS 2018*. IEEE Computer Soc., Los Alamitos, CA, 2018, pp. 23–34.
- [20] FISHER, M. The nature of critical points. *Lecture notes in Theoretical Physics* 7c (1965), 1–159.
- [21] HARROW, A., MEHRABAN, S., AND SOLEIMANIFAR, M. Classical algorithms, correlation decay, and complex zeros of partition functions of quantum many-body systems. *arXiv preprint arXiv:1910.09071* (2019).
- [22] HASTINGS, M. Solving gapped hamiltonians locally. *Physical Review B* 73, 8 (2006), 085115.
- [23] HASTINGS, M. Quantum belief propagation: An algorithm for thermal quantum systems. *Physical Review B* 76, 20 (2007), 201102.
- [24] JERRUM, M., AND SINCLAIR, A. Polynomial-time approximation algorithms for the Ising model. *SIAM Journal on Computing* 22, 5 (1993), 1087–1116.
- [25] KASTORYANO, M., AND BRANDÃO, F. Quantum Gibbs samplers: the commuting case. *Communications in Mathematical Physics* 344, 3 (2016), 915–957.
- [26] KATO, K., AND BRANDÃO, F. Quantum approximate Markov chains are thermal. *Communications in Mathematical Physics* 370, 1 (2019), 117–149.
- [27] KIM, I. Markovian matrix product density operators: Efficient computation of global entropy. *arXiv preprint arXiv:1709.07828* (2017).
- [28] KLIESCH, M., GOGOLIN, C., KASTORYANO, M., RIERA, A., AND EISERT, J. Locality of temperature. *Physical Review X* 4, 3 (2014), 031019.
- [29] KUWAHARA, T., KATO, K., AND BRANDÃO, F. G. Clustering of conditional mutual information for quantum gibbs states above a threshold temperature. *arXiv preprint arXiv:1910.09425* (2019).
- [30] KUWAHARA, T., AND SAITO, K. Polynomial-time classical simulation for one-dimensional quantum gibbs states. *arXiv preprint arXiv:1807.08424* (2018).
- [31] LEE, T.-D., AND YANG, C.-N. Statistical theory of equations of state and phase transitions. ii. lattice gas and Ising model. *Physical Review* 87, 3 (1952), 410.
- [32] LEVIN, M., AND WEN, X.-G. String-net condensation: A physical mechanism for topological phases. *Physical Review B* 71, 4 (2005), 045110.
- [33] LIU, J., SINCLAIR, A., AND SRIVASTAVA, P. Fisher zeros and correlation decay in the Ising model. In *10th Innovations in Theoretical Computer Science—ITCS*, 2019.
- [34] LIU, J., SINCLAIR, A., AND SRIVASTAVA, P. The Ising partition function: zeros and deterministic approximation. *Journal of Statistical Physics* 174, 2 (2019), 287–315.
- [35] MANN, R., AND BREMNER, M. Approximation algorithms for complex-valued Ising models on bounded degree graphs. *arXiv preprint arXiv:1806.11282* (2018).
- [36] MOLNAR, A., SCHUCH, N., VERSTRAETE, F., AND CIRAC, I. Approximating gibbs states of local Hamiltonians efficiently with projected entangled pair states. *Physical Review B* 91, 4 (2015), 045138.
- [37] PETERS, H., AND REGTS, G. Location of zeros for the partition function of the Ising model on bounded degree graphs. *arXiv preprint arXiv:1810.01699* (2018).
- [38] PETERS, H., AND REGTS, G. On a conjecture of Sokal concerning roots of the independence polynomial. *Michigan Math. J.* 68, 1 (2019), 33–55.
- [39] POULIN, D., AND WOJCIAN, P. Sampling from the thermal quantum gibbs state and evaluating partition functions with a quantum computer. *Physical Review Letters* 103, 22 (2009), 220502.
- [40] SINCLAIR, A., SRIVASTAVA, P., AND THURLEY, M. Approximation algorithms for two-state anti-ferromagnetic spin systems on bounded degree graphs. *Journal of Statistical Physics* 155, 4 (2014), 666–686.
- [41] SLY, A. Computational transition at the uniqueness threshold. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science—FOCS 2010*. IEEE Computer Soc., Los Alamitos, CA, 2010, pp. 287–296.
- [42] SLY, A., AND SUN, N. The computational hardness of counting in two-spin models on d -regular graphs. In *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science—FOCS 2012*. IEEE Computer Soc., Los Alamitos, CA, 2012, pp. 361–369.
- [43] SUZUKI, M., AND FISHER, M. Zeros of the partition function for the Heisenberg, ferroelectric, and general Ising models. *Journal of Mathematical Physics* 12, 2 (1971), 235–246.
- [44] WEITZ, D. *Mixing in time and space for discrete spin systems*. ProQuest LLC, Ann Arbor, MI, 2004. Thesis (Ph.D.)—University of California, Berkeley.
- [45] WEITZ, D. Counting independent sets up to the tree threshold. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing—STOC 2006* (2006), ACM, New York, pp. 140–149.
- [46] WEITZ, D. Counting independent sets up to the tree threshold. In *STOC'06: Proceedings of the 38th Annual ACM Symposium on Theory of Computing* (2006), ACM, New York, pp. 140–149.