

On Sampling from Ising Models with Spectral Constraints*

Andreas Galanis ✉

University of Oxford, UK

Alkis Kalavasis ✉

Yale University, New Haven, CT, USA

Anthimos Vardis Kandiros ✉

Massachusetts Institute of Technology, Cambridge, MA, USA

Abstract

We consider the problem of sampling from the Ising model when the underlying interaction matrix has eigenvalues lying within an interval of length γ . Recent work in this setting has shown various algorithmic results that apply roughly when $\gamma < 1$, notably with nearly-linear running times based on the classical Glauber dynamics. However, the optimality of the range of γ was not clear since previous inapproximability results developed for the antiferromagnetic case (where the matrix has entries ≤ 0) apply only for $\gamma > 2$.

To this end, Kunisky (SODA'24) recently provided evidence that the problem becomes hard already when $\gamma > 1$ based on the low-degree hardness for an inference problem on random matrices. Based on this, he conjectured that sampling from the Ising model in the same range of γ is NP-hard.

Here we confirm this conjecture, complementing in particular the known algorithmic results by showing NP-hardness results for approximately counting and sampling when $\gamma > 1$, with strong inapproximability guarantees; we also obtain a more refined hardness result for matrices where only a constant number of entries per row are allowed to be non-zero. The main observation in our reductions is that, for $\gamma > 1$, Glauber dynamics mixes slowly when the interactions are all positive (ferromagnetic) for the complete and random regular graphs, due to a bimodality in the underlying distribution. While ferromagnetic interactions typically preclude NP-hardness results, here we work around this by introducing in an appropriate way mild antiferromagnetism, keeping the spectrum roughly within the same range. This allows us to exploit the bimodality of the aforementioned graphs and show the target NP-hardness by adapting suitably previous inapproximability techniques developed for antiferromagnetic systems.

2012 ACM Subject Classification Theory of computation \rightarrow Problems, reductions and completeness; Theory of computation \rightarrow Random walks and Markov chains; Theory of computation \rightarrow Generating random combinatorial structures

Keywords and phrases Ising model, spectral constraints, Glauber dynamics, mean-field Ising, random regular graphs

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2024.70

Category RANDOM

Related Version *Full Version:* <https://arxiv.org/abs/2407.07645>

Funding Vardis Kandiros was supported by a Fellowship of the Eric and Wendy Schmidt Center at the Broad Institute of MIT and Harvard and by the Onassis Foundation-Scholarship ID: F ZP 016-1/2019-2020.

* For the purpose of Open Access, the author has applied a CC BY public copyright licence to any Author Accepted Manuscript (AAM) version arising from this submission.



© Andreas Galanis, Alkis Kalavasis, and Anthimos Vardis Kandiros;
licensed under Creative Commons License CC-BY 4.0

Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2024).

Editors: Amit Kumar and Noga Ron-Zewi; Article No. 70; pp. 70:1–70:14



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

The Ising model with a symmetric interaction matrix $J \in \mathbb{R}^{N \times N}$ is a probability distribution μ_J over $\{-1, 1\}^N$ with

$$\mu_J(\sigma) = \frac{1}{Z_J} \exp\left(\frac{1}{2}\sigma^\top J\sigma\right) \text{ for all vectors } \sigma \in \{-1, 1\}^N,$$

where the normalizing constant $Z_J = \sum_{\sigma \in \{-1, 1\}^N} \exp\left(\frac{1}{2}\sigma^\top J\sigma\right)$ is the *partition function* of the model. The most well-studied setting for the Ising model is when the underlying matrix J corresponds to the adjacency matrix of a graph G , scaled by a real parameter β which corresponds to the (inverse) temperature;¹ for $\beta > 0$ the model is called ferromagnetic, and antiferromagnetic otherwise. The more general setting with non-uniform weights in the entries of J arises frequently in statistical learning settings, see, e.g., [9, 30, 25, 12].

The Ising model is the most fundamental example of a spin system, capturing how local interactions affect the global macroscopic behaviour, see [48, 41, 40, 33] for applications in various areas. From a computer science perspective, sampling from the Ising model plays a key role in various learning and inference problems. Understanding the limits of efficient sampling has therefore been a major focus in the literature, yielding new algorithmic techniques as well as exploring the power of classical algorithms (such as Glauber dynamics) and their connections to phase transitions in statistical mechanics; we briefly review some of the relevant literature below.

The prototypical setting where the problem of sampling for the Ising model has been studied is lattices (such as \mathbb{Z}^2), where the landscape for Markov-chain algorithms has been well-understood [38, 39, 37]. Random graph models have also been considered more recently such as sparse random graphs [43, 13, 6, 16, 36] or the Sherrington-Kirkpatrick model [18, 17, 26]. More closely related to the setting considered in this paper is the case of general graphs. In the ferromagnetic case, where the entries of J are all nonnegative, the classical algorithm by Jerrum and Sinclair [28] gives a poly-time sampler (albeit with a relatively large running-time polynomial), see also [24, 19]. In the antiferromagnetic case, the problem is more interesting for bounded-degree graphs, where in the case of uniform weights the existence of polynomial-time algorithms is connected to the uniqueness threshold, see [46, 47, 21, 45, 35].

Recently, the development of spectral independence [4, 1] has given tight results on the performance of Glauber dynamics. This has led to nearly linear-time algorithms in various settings, see e.g., [11, 10, 18, 3, 31] and has made it possible to connect the performance of Glauber dynamics with the eigenvalues of the underlying matrix J . In this direction, [18, 3] show that Glauber dynamics is fast mixing when $\lambda_{\max}(J) - \lambda_{\min}(J) < 1$ which significantly improves upon the standard Dobrushin's uniqueness condition (the latter only applies when $\sum_j |J_{ij}| < 1$ for all $i \in [N]$).

On the other side, the optimality of these algorithmic results in terms of the spectrum is less clear. It is known [34, 14] that Glauber dynamics mixes slowly in the complete graph for temperatures $\beta > 1$, which corresponds precisely to the condition $\lambda_{\max}(J) - \lambda_{\min}(J) > 1$ by taking J to be the adjacency matrix of the N -vertex complete graph, scaled by β/N . This does not however translate in a straightforward way to hardness results and does not preclude the possibility that various alternative methods could potentially go beyond the 1-gap, see, e.g., [44, 31, 27] for some recent approaches using variational methods. To this

¹ Note that in this parametrization $\mu_J(\sigma) \propto \exp(\frac{1}{2}\beta\sigma^\top A\sigma)$, where A is the adjacency matrix of the graph.

end, Kunisky [32] gave further evidence that $\lambda_{\max}(J) - \lambda_{\min}(J) > 1$ is hard for sampling via a reduction to hypothesis testing in a Wishart negatively-spiked matrix model that involves random matrices (which is known to resist low-degree algorithms [5]). Kunisky also posed the conjecture that in fact NP-hardness for sampling under spectral constraints should hold when $\lambda_{\max}(J) - \lambda_{\min}(J) > 1$. To add a bit to the mystery, it is noteworthy that the inapproximability results for the antiferromagnetic case (mentioned earlier) only apply roughly when $\lambda_{\max}(J) - \lambda_{\min}(J) > 2$, see also below for a more detailed discussion.

1.1 Our results

Our aim in this work is to address Kunisky's conjecture and close the gap between algorithmic and NP-hardness results. In particular, we answer in the affirmative the conjecture in [32], obtaining NP-hardness results that complement the algorithmic results of [18, 3]. This completes the program initiated in [32], i.e., showing that Glauber is effectively optimal for "general-purpose" Ising model sampling, and clarifies the picture in terms of the computational complexity landscape under spectral constraints.

To formally state the result, we define the following computational problem.

Problem: SPECTRALISING(γ)

Input: A symmetric matrix $J \in \mathbb{R}^{N \times N}$, with $\lambda_{\max}(J) - \lambda_{\min}(J) < \gamma$.

Output: The partition function $Z_J = \sum_{\sigma \in \{-1, +1\}^N} \exp\left(\frac{1}{2} \sigma^\top J \sigma\right)$.

► **Theorem 1.** *Fix any real $\gamma > 1$. Then, it is NP-hard to approximate SPECTRALISING(γ), even within an exponential factor 2^{cN} for some constant $c = c(\gamma) > 0$.*

This confirms Conjecture 1.9 of [32] and complements the algorithm of [18, 3]. Using Theorem 1, we get the following result using the standard reduction [29] from counting to sampling (the problem is self-reducible under scaling of the matrix J). Recall, the total variation distance between probability distributions μ and ν is defined as $\text{TV}(\mu, \nu) = \frac{1}{2} \|\mu - \nu\|_1$.

► **Corollary 2.** *For every real $\gamma > 1$, the following holds. Suppose there is a poly-time sampler that, on input a symmetric matrix $J \in \mathbb{R}^{N \times N}$ with $\lambda_{\max}(J) - \lambda_{\min}(J) < \gamma$ and $\delta > 0$, returns a configuration σ whose distribution is within TV distance δ from μ_J . Then NP = RP.*

As we will explain next, it is also possible to obtain a more refined version of Theorem 1, for the restricted case where each row of the interaction matrix J has at most d non-zero entries, for some fixed integer $d \geq 4$.

Problem: BOUNDED SPECTRALISING(d, γ)

Input: A symmetric matrix $J \in \mathbb{R}^{N \times N}$, with $\leq d$ non-zero entries per row and $\lambda_{\max}(J) - \lambda_{\min}(J) < \gamma$.

Output: The partition function $Z_J = \sum_{\sigma \in \{-1, +1\}^N} \exp\left(\frac{1}{2} \sigma^\top J \sigma\right)$.

► **Theorem 3.** *Fix any integer $d \geq 4$ and real $\gamma > \frac{1}{2} \ln\left(1 + \frac{2}{d-3}\right)(d-1 + 2\sqrt{d-2})$. Then, it is NP-hard to approximate BOUNDED SPECTRALISING(d, γ), even within an exponential factor 2^{cN} for some constant $c = c(\gamma) > 0$.*

Note that when taking the limit $d \rightarrow \infty$ in the above bound, we recover the spectral condition $\gamma > 1$ of Theorem 1, so asymptotically the bound is tight; we are not aware of algorithmic results that apply specifically to the d -sparse setting under the spectral condition.

We remark further that applying the results of [47, 21] would yield hardness only in the setting where $\gamma > d \ln(1 + \frac{2}{d-2})$ (see [32, Section 1.2] for a detailed description on how to translate the results), so Theorem 3 improves on this by roughly a factor of 2 asymptotically. It should be noted however that the setting in these results is more restrictive (negative weights, which have the same value on all edges) and hence not directly comparable.

1.2 Techniques

Before giving the proofs, we explain briefly the main idea behind Theorem 1, the idea for Theorem 3 is almost identical, modulo the gadget used in the reduction.

The key ingredient in obtaining Theorem 1 is to exploit the slow mixing of Glauber dynamics on the complete graph in a suitable way. Recall that [34] showed exponential mixing time for Glauber dynamics on the N -vertex complete graph when the weights on the edges are ferromagnetic equal to β/N (entry-wise) for any $\beta > 1$ (note that the corresponding matrix J has $\lambda_{\max}(J) - \lambda_{\min}(J) = \beta$). Intuitively, the slow mixing is caused because the distribution exhibits bimodality, i.e., it is concentrated around two modes/“phases” corresponding roughly to the all-plus and all-minus configurations (see Section 2 for more details). Therefore, we would like to use the binary behaviour of the complete graph as a gadget in the reduction. The main trouble here is caused by the ferromagnetic interactions which cannot typically be related to NP-hard problems; by contrast, in the antiferromagnetic case $\beta < 0$, the max-probability configurations in the Ising distribution correspond to maximum cuts (when J encodes the adjacency matrix of a graph), and the respective gadgets in the constructions had bipartite structure.²

Hence, in order to get NP-hardness, we need to introduce some “mild” antiferromagnetism (small negative weights): mild to keep the spectrum unchanged and antiferromagnetic to allow us to reduce from an NP-hard problem (we will use MaxCut); this is quite different than the approach of [32] where the positive and negative entries in the constructed instance are more heavily mixed up (randomly). At this stage, the main observation is that the previous reductions used in the antiferromagnetic case [46, 47, 21] can accommodate this relatively easily; the only difference here is that we need to use small negative weights to connect disjoint copies of the gadgets, and amplify their effect using appropriately-sized matchings; conveniently, since the matchings (with the small weights on their edges) correspond to a low-rank perturbation, the spectrum of the underlying matrix is close to that of the complete graph.

The proof of Theorem 3 is almost identical. The main difference needed to make our construction sparse is to use a random d -regular graph as the gadget, which is known to exhibit slow mixing when $\beta > \beta_d := \frac{1}{2} \ln(1 + \frac{2}{d-2})$ [22, 13, 43, 42], with a similar bimodal behaviour to that of the complete graph for $\beta > 1$. Relative to the spectrum, the well-known result of Friedman [20] shows that the adjacency matrix A of a random d -regular graph satisfies w.h.p. $\lambda_{\max}(A) - \lambda_{\min}(A) \leq \lambda_d + \epsilon$ for any constant $\epsilon > 0$, where $\lambda_d := d + 2\sqrt{d-1}$. For technical reasons (see Remark 10 for details), we need to actually use a $(d-1)$ -regular graph as a gadget in the reduction, so the argument sketched above yields NP-hardness when $\gamma > \beta_{d-1} \lambda_{d-1}$ and $d \geq 4$.

² As a side note, we remark that the factor-2 gap from the antiferromagnetic setting (mentioned below Theorem 3) comes from the use of bipartite gadgets in these results, which have a symmetric spectrum around zero and hence effectively double the range of the eigenvalues.

1.3 Outline and Discussion

We give the details of the gadget in Section 2 and the reduction in Section 3.1. This gives a self-contained proof of Theorem 1; for Theorem 3 the argument is identical modulo the use of the (random) d -regular graph as the gadget, for which we need to import a couple of non-trivial results from the literature.

As a final remark before proceeding to the proofs, it would be interesting to explore whether the statistical hardness perspective from [32] (or some variant) perhaps applies to other counting/sampling problems where NP-hardness results are unlikely, such as approximating the number of independent sets in a bipartite graph [15], or approximating the partition function of the ferromagnetic Potts model [23]. Another related question is whether such statistical hardness results can be invoked on sparse random graph models where the spectral threshold $\lambda_{\max}(J) - \lambda_{\min}(J) = 1$ (that applies to worst-case instances) is known not to be tight (see [10, 31, 36]).

2 The Gadget of Theorem 1

Our main gadget will be a clique graph $K_n = (V, E)$ with n vertices, where $V = \{1, 2, \dots, n\}$. We will consider n to be an absolute (large) constant that we will choose later. For a small integer $t > 0$, let $S \subseteq V$ be an arbitrary subset of V with $|S| = t$. Let $r = n - t$. Intuitively, S contains the nodes that will be used to connect the gadgets with each other.

We define the phase of the configuration $\sigma \in \{-1, 1\}^n$ on $V \setminus S$ as

$$Y_\sigma = \mathbf{1}\left\{\sum_{i \in V \setminus S} \sigma_i > 0\right\} - \mathbf{1}\left\{\sum_{i \in V \setminus S} \sigma_i \leq 0\right\}.$$

Note that the phase of a configuration is defined using only the spins in $V \setminus S$. For any fixed $\beta > 0$, consider solutions to the equation

$$\ln \frac{1 - \alpha}{\alpha} + 2\beta(2\alpha - 1) = 0 \tag{1}$$

for $\alpha \in [0, 1]$. It is not hard to see that for $\beta > 1$ there are exactly three solutions $\alpha = q^-, 1/2, q^+$ which satisfy $q^+ - 1/2 = 1/2 - q^- > 0$. Using these, we define the product measure Q_S^+ (resp. Q_S^-) on configurations on S , where each spin takes the value $+1$ with probability q^+ , and -1 with probability $1 - q^+$ (resp. q^- and $1 - q^-$). Concretely, for $\tau \in \{-1, +1\}^S$, we have

$$Q_S^\pm(\tau) = (q^\pm)^{\frac{\sum_{i \in S} \tau_i + t}{2}} (1 - q^\pm)^{\frac{t - \sum_{i \in S} \tau_i}{2}} = (q^\pm(1 - q^\pm))^{t/2} \left(\frac{q^\pm}{1 - q^\pm}\right)^{\frac{\sum_{i \in S} \tau_i}{2}}. \tag{2}$$

We now state a lemma that presents the basic properties of the Ising model on our gadget graph. A similar lemma appears in the seminal results of [46, 47]. Informally, the lemma states that conditioned on the phase of the spins in $V \setminus S$, the spins in S behave almost *independently* from each other, with bias depending on the phase.

► **Lemma 4.** *Let $\beta > 1$. Then, for any real $\epsilon > 0$ and integer $t \geq 1$, for all sufficiently large integers $n = n(t, \epsilon)$ such that $n - t$ is odd, the following holds for the Ising model with interaction matrix $J \in \mathbb{R}^{n \times n}$ given by $J = \frac{\beta}{n-t} \mathbf{1}\mathbf{1}^\top$, where $\mathbf{1}$ is the n -dimensional vector with all ones.*

Let $S \subseteq [n]$ be a subset of the vertices with $|S| = t$. Then:

1. *The phases on $V \setminus S$ appear with the same probability, i.e., $\Pr_{\sigma \sim \mu_J}[Y_\sigma = +] = \Pr_{\sigma \sim \mu_J}[Y_\sigma = -] = 1/2$.*

70:6 On Sampling from Ising Models with Spectral Constraints

2. Conditioned on the phase, the joint distribution of the spins in S is approximately given by the product distribution Q_S^\pm , i.e.,

$$\text{for any } \tau \in \{-1, +1\}^S, \text{ it holds that } \Pr_{\sigma \sim \mu_J} [\sigma_S = \tau \mid Y_\sigma = \pm] = (1 \pm \epsilon) Q_S^\pm(\tau).$$

Proof. Let $r = n - t$. For r odd (as in the statement of the lemma), we have by symmetry that the phases appear with equal probability. So, we focus on proving the second item. For a vector x with entries $+1$ or -1 , we denote by $|x|$ the sum of its entries.

Let $\alpha \in [0, 1]$ be such that αr is an integer. For a configuration $\tau \in \{-1, +1\}^S$, let $Z^\alpha(\tau)$ be the contribution to the partition function of configurations σ with αr spins from $V \setminus S$ set to $+1$, $(1 - \alpha)r$ spins from $V \setminus S$ set to -1 and $\sigma_S = \tau$. Concretely,

$$Z^\alpha(\tau) = \sum_{\sigma \in \{-1, +1\}^V; \sigma_S = \tau, |\sigma_{V \setminus S}| = (2\alpha - 1)r} \exp\left(\frac{1}{2} \sigma^\top J \sigma\right).$$

The number of configurations σ with $\sigma_S = \tau$ and exactly αr of the spins in $V \setminus S$ equal to 1 is $\binom{r}{\alpha r}$. Using that $J = \frac{\beta}{r} \mathbf{1}\mathbf{1}^\top$, for each such σ , we have $\frac{1}{2} \sigma^\top J \sigma = \frac{\beta}{r} (|\sigma_{V \setminus S}| + |\tau|)^2 = \frac{\beta}{2r} ((2\alpha - 1)r + |\tau|)^2$. So,

$$Z^\alpha(\tau) = \binom{r}{\alpha r} \exp\left(\frac{\beta}{2} (2\alpha - 1)^2 r + \beta(2\alpha - 1)|\tau| + \frac{\beta}{2r} |\tau|^2\right). \quad (3)$$

We use the well-known approximation of the binomial coefficient using Stirling's approximation. This yields, for any $\alpha \in [0, 1]$, that

$$\binom{r}{\alpha r} = \exp(rH(\alpha) + o(r)), \quad (4)$$

where $H(\alpha) := -\alpha \ln \alpha - (1 - \alpha) \ln(1 - \alpha)$ is the binary entropy function. Asymptotically in r , we can also ignore the term $\exp(\frac{\beta}{2r} |\tau|^2)$, so we obtain that

$$Z^\alpha(\tau) = \exp(rf(\alpha) + o(r)) \text{ where } f(\alpha) := H(\alpha) + \frac{\beta}{2} (2\alpha - 1)^2. \quad (5)$$

The function $f(\alpha)$ plays a key role since for large r it controls the asymptotic order of $Z^\alpha(\tau)$. The important point, as we will see below, is that the global maximum of f is attained for $\alpha = q^\pm$.

Indeed, we have

$$f'(\alpha) = -\ln(\alpha) + \ln(1 - \alpha) + 2\beta(2\alpha - 1)$$

and $f''(\alpha) = -\frac{1}{\alpha(1-\alpha)} + 4\beta$. Since f'' has at most two zeros, we have that f' has at most three distinct zeros and hence f has at most three critical points. For $\beta > 1$, we have $f'(1/2) = 0$ and $f''(1/2) = -4 + 4\beta > 0$, so f has a local minimum at $\alpha = 1/2$; therefore, the maximum of f in the interval $[0, 1]$ is attained at some point $\alpha \neq 1/2$. Using the symmetry of f around $\alpha = 1/2$, there must be at least two global maxima, one in the interval $(0, 1/2)$ and $(1/2, 1)$. Since f has at most three critical points (and $1/2$ is one of them), we conclude that there are exactly two critical points/maxima other than $\alpha = 1/2$, which must therefore be the values q^+, q^- as defined in (1).

We are now ready to establish the second item of the lemma. We will argue about the $+$ phase, but the other phase is completely symmetric. Let $\tau, \tau' \in \{-1, 1\}^S$ be two configurations of spins in S . We have that

$$\frac{\Pr[\sigma_S = \tau \mid Y(\sigma_{V \setminus S}) = +]}{\Pr[\sigma_S = \tau' \mid Y(\sigma_{V \setminus S}) = +]} = \frac{\sum_{\alpha > 1/2} Z^\alpha(\tau)}{\sum_{\alpha > 1/2} Z^\alpha(\tau')}. \quad (6)$$

We will show that the sums in the numerator and denominator are dominated by α values that are close to q^+ . First, note that since q^+ is the unique global maximum of $f(\alpha)$ in the interval $[1/2, 1]$, for any arbitrarily small constant $\delta > 0$, there is $\eta > 0$ such that $f(\alpha) \leq f(q^+) - 3\eta$ for all $\alpha > 1/2$ with $\alpha \notin [q^+ - \delta, q^+ + \delta]$. We pick $\delta > 0$ sufficiently small and $r > 0$ sufficiently large so that $\exp(4\beta t\delta + \beta \frac{t^2}{r}) < \epsilon/2$. Since $|\tau| \leq t$, it follows that for r large enough it holds that

$$\sum_{\alpha > 1/2; |\alpha - q^+| > \delta} Z^\alpha(\tau) \leq \exp(r(f(q^+) - 2\eta)).$$

By the continuity of f , for $\alpha = q^+ + O(1/r)$ we have $f(\alpha) = f(q^+) + O(1/r)$ and therefore

$$\sum_{\alpha > 1/2; |\alpha - q^+| \leq \delta} Z^\alpha(\tau) \geq \exp(r(f(q^+) - \eta)).$$

It follows that

$$\frac{\sum_{\alpha > 1/2; |\alpha - q^+| > \delta} Z^\alpha(\tau)}{\sum_{\alpha > 1/2; |\alpha - q^+| \leq \delta} Z^\alpha(\tau)} \leq \exp(-\eta r) \leq \epsilon/2. \quad (7)$$

for all sufficiently large r . Thus,

$$\begin{aligned} \frac{\sum_{\alpha > 1/2} Z^\alpha(\tau)}{\sum_{\alpha > 1/2} Z^\alpha(\tau')} &\leq \frac{\sum_{\alpha > 1/2} Z^\alpha(\tau)}{\sum_{|\alpha - q^+| \leq \delta} Z^\alpha(\tau')} = \frac{\sum_{\alpha > 1/2} Z^\alpha(\tau)}{\sum_{|\alpha - q^+| \leq \delta} Z^\alpha(\tau)} \cdot \frac{\sum_{|\alpha - q^+| \leq \delta} Z^\alpha(\tau)}{\sum_{|\alpha - q^+| \leq \delta} Z^\alpha(\tau')} \\ &\leq (1 + \epsilon/2) \frac{\sum_{|\alpha - q^+| \leq \delta} Z^\alpha(\tau)}{\sum_{|\alpha - q^+| \leq \delta} Z^\alpha(\tau')}, \end{aligned} \quad (8)$$

where the last inequality follows from (7).

On the other hand, for any α with $|\alpha - q^+| \leq \delta$, using (3) we get

$$\begin{aligned} \frac{Z^\alpha(\tau)}{Z^\alpha(\tau')} &= \exp\left(\beta(2\alpha - 1)(|\tau| - |\tau'|) + \frac{\beta(|\tau|^2 - |\tau'|^2)}{2r}\right) \\ &\leq \exp(4\beta t\delta + \beta \frac{t^2}{r}) \\ &\exp(\beta(2q^+ - 1)(|\tau| - |\tau'|)) \leq (1 + \epsilon/2) \exp(\beta(2q^+ - 1)(|\tau| - |\tau'|)), \end{aligned} \quad (9)$$

where the last inequality follows from the choice of δ and r . Using the definition (2) and the fact that q^+ is a solution of (1), i.e., that $f'(q^+) = 0$, we have that

$$\exp(\beta(2q^+ - 1)(|\tau| - |\tau'|)) = \left(\frac{q^+}{1 - q^+}\right)^{\frac{|\tau| - |\tau'|}{2}} = \frac{Q_S^+(\tau)}{Q_S^+(\tau')}.$$

Hence, from (9) we obtain that $\frac{Z^\alpha(\tau)}{Z^\alpha(\tau')} \leq (1 + \epsilon/2) \frac{Q_S^+(\tau)}{Q_S^+(\tau')}$. Since this holds for all α with $|\alpha - q^+| \leq \delta$, we have

$$\frac{\sum_{|\alpha - q^+| \leq \delta} Z^\alpha(\tau)}{\sum_{|\alpha - q^+| \leq \delta} Z^\alpha(\tau')} \leq (1 + \epsilon/2) \frac{Q_S^+(\tau)}{Q_S^+(\tau')}. \quad (10)$$

Combining this with (6) and (8), we obtain that

$$\frac{\Pr[\sigma_S = \tau | Y(\sigma_{V \setminus S}) = +]}{\Pr[\sigma_S = \tau' | Y(\sigma_{V \setminus S}) = +]} \leq (1 + \epsilon) \frac{Q_S^+(\tau)}{Q_S^+(\tau')}.$$

70:8 On Sampling from Ising Models with Spectral Constraints

By interchanging the roles of τ, τ' , we also obtain the inverse inequality, so

$$(1 - \epsilon) \frac{Q_S^+(\tau)}{Q_S^+(\tau')} \leq \frac{\Pr[\sigma_S = \tau | Y(\sigma_{V \setminus S}) = +]}{\Pr[\sigma_S = \tau' | Y(\sigma_{V \setminus S}) = +]} \leq (1 + \epsilon) \frac{Q_S^+(\tau)}{Q_S^+(\tau')}. \quad (11)$$

For $\tau \in \{-1, 1\}^S$, observe that we can expand the ratio

$$\frac{\Pr[\sigma_S = \tau | Y = +]}{Q_S^+(\tau)} = \frac{\sum_{\tau'} Q_S^+(\tau') \Pr[\sigma_S = \tau | Y = +]}{\sum_{\tau'} Q_S^+(\tau) \Pr[\sigma_S = \tau' | Y = +]}$$

so using that $\min_i \frac{a_i}{b_i} \leq \frac{\sum_i a_i}{\sum_i b_i} \leq \max_i \frac{a_i}{b_i}$ for non-negative $(a_i)_i, (b_i)_i$, we obtain from (11) that

$$\left| \frac{\Pr[\sigma_S = \tau | Y = +]}{Q_S^+(\tau)} - 1 \right| \leq \max_{\tau'} \left| \frac{Q_S^+(\tau') \Pr[\sigma_S = \tau | Y = +]}{Q_S^+(\tau) \Pr[\sigma_S = \tau' | Y = +]} - 1 \right| \leq \epsilon.$$

This finishes the proof. \blacktriangleleft

3 Proofs of Main Results

3.1 Proof of Theorem 1

Let $\gamma > 1$ and $\beta = (1 + \gamma)/2 > 1$. Following the technique in [46, 47, 21], we reduce MAXCUT on 3-regular graphs to SPECTRALISING(γ).

Consider a 3-regular graph $H = (V_H, E_H)$ with $|V_H| = m$ vertices, an instance of MAXCUT. Let G be the clique graph on n vertices, with a subset S of the vertices with $|S| = t$ that will be used as terminals (cf. Lemma 4); for convenience, we assume that $t > 0$ is a multiple of 3 (with $n \gg 3t$). We construct an instance H^G of SPECTRALISING(γ) as follows:

- We replace each node $v \in V_H$ with a distinct copy of the gadget clique graph G . In particular, for any $v \in V_H$, consider a copy $G_v = (W_v, E_v)$ of the gadget G ; each edge in E_v has weight $w_+ = \beta/r$ as in Lemma 4, where recall that $\beta = (1 + \gamma)/2 > 1$. For each $v \in V_H$, let $S_v \subseteq W_v$ be a subset of the vertices in G_v of size $t = n - r$. Let \hat{H}^G be the disjoint union of the G_v 's for $v \in H$. Note that the number of vertices of \hat{H}^G is nm .
- We now describe how to encode the edges of H using connections between the gadgets (which will complete the construction of H^G). Assume that the node $u \in V_H$ has neighbors v_1, v_2, v_3 in H , i.e., $(u, v_i) \in E_H, i = 1, \dots, 3$. Then, we partition S_u into subsets S_u^i of size $t/3$ each. Each subset S_u^i corresponds to one of the three neighbors of u . Then, for each $i = 1, 2, 3$, we add a perfect matching between S_u^i and the corresponding subset $S_{v_i}^j$ of S_{v_i} that corresponds to u . The weight of each of these edges in the matching will be $w_- = (1 - \gamma)/5 < 0$, since $\gamma > 1$. This antiferromagnetic structure across different copies will be crucial in order to approximate $\maxcut(H)$ by approximating the partition function of H^G .

Let J be the adjacency matrix of the weighted graph H^G . We first show that the spectrum of J has the desired properties, i.e., that $\lambda_{\max}(J) - \lambda_{\min}(J) < \gamma$.

▷ **Claim 5 (Structure of H^G).** The symmetric matrix $J = D + E \in \mathbb{R}^{nm \times nm}$, where D is a block diagonal matrix where the matrix of each block of size $n \times n$ is $\frac{\beta}{r} \mathbf{1}\mathbf{1}^\top$ and E contains in each row exactly one non-zero element of magnitude $(1 - \gamma)/5$.

Proof. By construction since n is the number of vertices of the gadget and m is the number of vertices of the input graph. \blacktriangleleft

▷ **Claim 6 (Spectrum Preservation).** For any integer $t > 0$, there exists $n(t, \gamma) > 0$, such that for $n > n(t, \gamma)$ it holds $|\lambda_{\max}(J) - \lambda_{\min}(J)| < \gamma$.

Proof. We will use Claim 5. Using Weyl's inequality (see Chapter 3 in [7]), which controls the eigenspectrum of a matrix under small perturbations of the entries, we have that for any i , it holds that $|\lambda_i(J) - \lambda_i(D)| \leq \|E\|$, where $\|E\|$ is the spectral norm of E . By definition, E has one element in each row of absolute value $(\gamma - 1)/5$, so $\|E\| \leq \frac{\gamma-1}{5}$. It follows that

$$\begin{aligned} |\lambda_{\max}(J) - \lambda_{\min}(J)| &\leq |\lambda_{\max}(J) - \lambda_{\max}(D)| + |\lambda_{\max}(D) - \lambda_{\min}(D)| + |\lambda_{\min}(D) - \lambda_{\min}(J)| \\ &\leq \frac{2(\gamma-1)}{5} + \frac{n}{r} \frac{1+\gamma}{2}. \end{aligned} \tag{12}$$

In the above we used the well-known fact that the spectrum of D is the spectrum of each of the blocks, which, in turn, is equal to

$$\lambda_{\max}(D) - \lambda_{\min}(D) = \frac{n}{r} \frac{1+\gamma}{2},$$

since each block is a rank-1 matrix. Now, recall that $n = r + t$, so by choosing r sufficiently large we can make $\frac{n}{r} < \frac{6\gamma+4}{5\gamma+5}$, which implies that the right hand side in (12) is $< \gamma$. ◁

We next show that if we could approximate Z_J within an arbitrarily small exponential factor in poly-time, we would obtain a PTAS for $\text{maxcut}(H)$. This part of the argument is largely based on the techniques of [47]; we first state the following lemma whose proof is given for completeness in the full version.

► **Lemma 7.** *It holds that*

$$(1 - 4\epsilon)^m 2^{-m} \leq \frac{Z_{H^G} / Z_{\hat{H}^G}}{A^{3mt/2} (B/A)^{\text{maxcut}(H)t/3}} \leq (1 + 4\epsilon)^m,$$

where $B > A > 0$ are constants depending only on γ .

With these pieces at hand, we are now ready to complete the reduction for Theorem 1, which we restate here for convenience.

► **Theorem 1.** Fix any real $\gamma > 1$. Then, it is NP-hard to approximate $\text{SPECTRALISING}(\gamma)$, even within an exponential factor 2^{cN} for some constant $c = c(\gamma) > 0$.

Proof. Assume that for any arbitrarily small constant $\delta > 0$, there is an oracle approx_δ such that, for any J with $\lambda_{\max}(J) - \lambda_{\min}(J) \leq \gamma$, we have that, when $F = \text{approx}_\delta(J)$, $|F - \log(Z(J))| \leq \delta m$. We will show how to obtain a PTAS for MAXCUT on 3-regular graphs, i.e., approximate MAXCUT on 3-regular graphs within an arbitrarily small factor.

Let H be a 3-regular graph H on m vertices, an instance of MAXCUT. The maximum cut of H is at least the expected value of a random cut which is equal to $3m/4$. We then construct H^G and \hat{H}^G as above. Observe that $Z(\hat{H}^G)$ can be computed in poly-time since \hat{H}^G is a disjoint collection of constant-size gadget graphs. Moreover, by Claim 6, H^G is an instance of $\text{SPECTRALISING}(\gamma)$. So, we can use the oracle approx_δ on H^G , which will give us an output F_H with the guarantee

$$|F_H - \log Z_{H^G}| \leq \delta mn.$$

Lemma 7 implies that

$$\frac{3 \log \left(\frac{Z_{H^G} / Z_{\hat{H}^G}}{A^{3mt/2} (1+4\epsilon)^m} \right)}{t \log(B/A)} \leq \text{maxcut}(H) \leq \frac{3 \log \left(\frac{2^m Z_{H^G} / Z_{\hat{H}^G}}{A^{3mt/2} (1-4\epsilon)^m} \right)}{t \log(B/A)}.$$

Thus, by using the output F_H we can compute upper and lower bounds for the maximum cut value, which differ by $O((\delta n + 1)m/t)$. Since $m \leq 4/3 \max\text{cut}(H)$, to show the desired PTAS for MAXCUT, it only remains to show that the quantity $R = (\delta n + 1)/t$ can be made arbitrarily small, say less than some target value ζ , where $\zeta > 0$ is an arbitrary constant. We first take t to be sufficiently large, so that $1/t < \zeta/2$ is sufficiently small. This makes n to be large, but still a constant, and hence n/t is a constant. So, by taking δ small enough, we will have $\delta n/t < \zeta/2$, making $R < \zeta$ as desired.

This yields the desired PTAS. Since MAXCUT is APX-hard [2], we conclude that it is NP-hard to approximate Z_J within some exponential factor, as wanted. ◀

3.2 Proof of Theorem 3

For integers $d, n \geq 3$ with dn even, let $G_{n,d}$ be a d -regular graph chosen uniformly at random among all such graphs with vertex set $V = \{1, 2, \dots, n\}$. Let $S \subseteq [n]$ be an arbitrary subset of the vertices of size t . Consider the Ising distribution μ_J with $J = \beta A$ where A is the adjacency matrix of G and $\beta > \frac{1}{2} \ln(1 + \frac{2}{d-2})$.

The range of β corresponds to the so-called non-uniqueness regime on the d -regular tree; roughly, this implies that on the d -regular tree of height h , when we condition the leaves to be $+$ and take the limit $h \rightarrow \infty$, the marginal probability that the root is plus converges to some value $q^+ > 1/2$. Similarly, when we condition the leaves to be $-$, the marginal probability that the root is plus converges to some value $q^- < 1/2$.³

It is well-known by now [13, 42] that this behaviour on the tree manifests itself on the random d -regular graph, roughly because of the tree-like neighborhoods in the latter. To make this more precise in our setting, analogously to Section 2, for a subset $S \subseteq V$, define the phase $Y_S(\sigma)$ of a configuration $\sigma \in \{-1, +1\}^V$ to be $+$ if $\sum_{i \in V \setminus S} \sigma_i \geq 0$, and $-$ otherwise. We also define the product measures Q_S^\pm on S analogously to (2), using now the values of q^+, q^- as defined above (see also Footnote 3). Then, the following lemma captures the main properties of the gadget that we need.

► **Lemma 8.** *Let $d \geq 3$ be an integer and $\beta > \frac{1}{2} \ln(1 + \frac{2}{d-2})$. Then, for any real $\epsilon > 0$ and integer $t \geq 1$, for all sufficiently large integers $n = n(t, \epsilon)$ with $n - t$ odd, the following holds with probability $1 - \epsilon$ over the choice of $G \sim G_{n,d}$. Let $S \subseteq V$ be a subset of vertices with $|S| = t$.*

Consider the Ising model with interaction matrix $J = \beta A$ where A is the adjacency matrix of G . Then:

1. $\lambda_{\max}(J) - \lambda_{\min}(J) \leq \beta(d + 2\sqrt{d-1}) + \epsilon$.
2. *The phases appear with the same probability, i.e., $\Pr_{\sigma \sim \mu_J}[Y_\sigma = +] = \Pr_{\sigma \sim \mu_J}[Y_\sigma = -] = 1/2$.*
3. *Conditioned on the phase, the joint distribution of the spins in S is approximately given by the product distribution Q_S^\pm , i.e.,*

$$\text{for any } \tau \in \{-1, +1\}^S, \text{ it holds that } \Pr_{\sigma \sim \mu_J}[\sigma_S = \tau \mid Y_\sigma = \pm] = (1 \pm \epsilon)Q_S^\pm(\tau).$$

Proof. The first item is Friedman's result [20], see also [8]. The second item is by symmetry of the configuration space (since n is odd). The third item follows by [42, Theorem 2.4], see also [13, Theorem 2.7] and [47, Proposition 4.2] for related results. Technically, there is a bit of work to translate the results here, we give the details in the full version. ◀

³ To define q^+, q^- more explicitly, for $\beta > \frac{1}{2} \ln(1 + \frac{2}{d-2})$, let $\tilde{q}^+ > 1 > \tilde{q}^- > 0$ be the solutions of $x = \left(\frac{\exp(2\beta)x+1}{x+\exp(2\beta)}\right)^{d-1}$. Then, q^+, q^- are defined from $\frac{q^+}{1-q^+} = \tilde{q}^+ \frac{\exp(2\beta)\tilde{q}^++1}{\tilde{q}^++\exp(2\beta)}$ and $\frac{q^-}{1-q^-} = \tilde{q}^- \frac{\exp(2\beta)\tilde{q}^-+1}{\tilde{q}^-+\exp(2\beta)}$, see also [21, Section 3].

► **Remark 9.** We use the gadget of Lemma 8 for some large but otherwise constant value of n . So, we can find a d -regular graph G satisfying Items 1-3 of Lemma 8 in deterministic time.

We are now ready to prove Theorem 3, which we restate here for convenience.

► **Theorem 3.** Fix any integer $d \geq 4$ and real $\gamma > \frac{1}{2} \ln(1 + \frac{2}{d-3})(d-1 + 2\sqrt{d-2})$. Then, it is NP-hard to approximate $\text{BOUNDEDSPECTRALISING}(d, \gamma)$, even within an exponential factor 2^{cN} for some constant $c = c(\gamma) > 0$.

Proof. Let

$$\beta_{d-1} := \frac{1}{2} \ln\left(1 + \frac{2}{d-3}\right), \quad \lambda_{d-1} := d-1 + 2\sqrt{d-2} \tag{13}$$

and set $\beta = \beta_{d-1} + \eta$, $\lambda = \lambda_{d-1} + \eta$ where $\eta > 0$ is a small constant so that $\beta\lambda + 2\eta < \gamma$ (note that such an η exists since $\gamma > \beta_{d-1}\lambda_{d-1}$).

Assume that we are given a 3-regular instance H of MAXCUT with m vertices. Let G be a $(d-1)$ -regular gadget with n vertices for some sufficiently large n , i.e., G satisfies Items 1-3 of Lemma 8 for degree $d-1$ and $\beta = \beta_{d-1} + \eta$, see also Remark 9. So, according to Item 1 there, the interaction matrix J_G corresponding to G satisfies $\lambda_{\max}(J_G) - \lambda_{\min}(J_G) \leq \beta\lambda$.

Using G , the construction of the graph H^G is identical to that of Section 3.1, i.e., we have a distinct copy of G for each node of H and, for each pair of neighbouring nodes of H , we add a matching of size $t/3$ between the corresponding gadgets using the vertices in S . Note that H^G has maximum degree d , so the interaction matrix of H^G , denoted by J henceforth, has at most d non-zero entries per row.

The weight of an edge inside the gadget is $w_+ = \beta > 0$ and the weight of the edges that connect two gadgets is $w_- = -\eta < 0$ (antiferromagnetic connections). Analogously to Claim 5, the symmetric matrix J can be written as $D + E \in \mathbb{R}^{nm \times nm}$, where (i) D is a block diagonal matrix with the matrix in each block being the $n \times n$ adjacency matrix of G scaled by w_+ , and (ii) E contains in each row exactly one non-zero element of magnitude w_- . The same argument as in the proof of Claim 6 gives that

$$\begin{aligned} |\lambda_{\max}(J) - \lambda_{\min}(J)| &\leq |\lambda_{\max}(J) - \lambda_{\max}(D)| + |\lambda_{\max}(D) - \lambda_{\min}(D)| + |\lambda_{\min}(D) - \lambda_{\min}(J)| \\ &\leq 2\eta + \beta\lambda < \gamma. \end{aligned}$$

This establishes that H^G is a valid instance of $\text{BOUNDEDSPECTRALISING}(d, \gamma)$.

Now, using Item 3 of Lemma 8, we obtain the exact same estimate as in Lemma 7 (with the same expressions for the constants A, B modulo the new values of w_+ and w_-), and therefore the same argument used in the proof of Theorem 1 applies verbatim to show NP-hardness of approximating the partition function within an arbitrarily small exponential factor. ◀

► **Remark 10.** Note that we could make the graph H^G to be d -regular for any integer $d \geq 3$ by taking the gadget G to be a random d -regular graph with a matching of size t removed (and using the endpoints of the matching as the set S of terminals); this more refined construction has been used for example in the hardness results of [46, 47, 21]. While one can show the analogue of Items 2 and 3 with minor modifications (analogously to what was done in the proof of Lemma 8), the proof of Item 1 for this modified gadget seems to require more careful adaptation of the proofs in [20, 8]. It is nevertheless reasonable to expect that the same bound on the range of the eigenvalues as stated currently in Item 1 will still apply; provided this is indeed the case, one can improve slightly the parameters of Theorem 3 to $d \geq 3$ and $\gamma > \beta_d \lambda_d$, where β_d, λ_d are as in (13).

References

- 1 Vedat Levi Alev and Lap Chi Lau. Improved analysis of higher order random walks and applications. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2020, pages 1198–1211, 2020.
- 2 Paola Alimonti and Viggo Kann. Hardness of approximating problems on cubic graphs. In *Proceedings of the Third Italian Conference on Algorithms and Complexity*, CIAC '97, pages 288–298, 1997.
- 3 Nima Anari, Vishesh Jain, Frederic Koehler, Huy Tuan Pham, and Thuy-Duong Vuong. Entropic independence in high-dimensional expanders: Modified log-sobolev inequalities for fractionally log-concave polynomials and the ising model. *arXiv preprint*, 10:32–42, 2021. [arXiv:2106.04105](https://arxiv.org/abs/2106.04105).
- 4 Nima Anari, Kuikui Liu, and Shayan Oveis Gharan. Spectral independence in high-dimensional expanders and applications to the hardcore model. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1319–1330, 2020.
- 5 Afonso S Bandeira, Dmitriy Kunisky, and Alexander S Wein. Computational hardness of certifying bounds on constrained pca problems. In *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*, volume 151, page 78. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020.
- 6 Ivona Bezáková, Andreas Galanis, Leslie Ann Goldberg, and Daniel Štefankovič. Fast sampling via spectral independence beyond bounded-degree graphs. *ACM Trans. Algorithms*, 20(1), January 2024.
- 7 Rajendra Bhatia. *Perturbation bounds for matrix eigenvalues*. SIAM, 2007.
- 8 Charles Bordenave. A new proof of friedman’s second eigenvalue theorem and its extension to random lifts. In *Annales Scientifiques de l’École Normale Supérieure*, volume 4, pages 1393–1439, 2020.
- 9 Guy Bresler. Efficiently learning Ising models on arbitrary graphs. In *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing*, STOC '15, pages 771–782, 2015.
- 10 Yuansi Chen and Ronen Eldan. Localization schemes: A framework for proving mixing bounds for markov chains. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 110–122. IEEE, 2022.
- 11 Zongchen Chen, Kuikui Liu, and Eric Vigoda. Optimal mixing of Glauber dynamics: entropy factorization via high-dimensional expansion. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2021, pages 1537–1550, 2021.
- 12 Yuval Dagan, Constantinos Daskalakis, Nishanth Dikkala, and Anthimos Vardis Kandiros. Learning Ising models from one or multiple samples. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2021, pages 161–168, 2021.
- 13 Amir Dembo and Andrea Montanari. Ising models on locally tree-like graphs. *The Annals of Applied Probability*, 20(2):565–592, 2010.
- 14 Jian Ding, Eyal Lubetzky, and Yuval Peres. The mixing time evolution of glauber dynamics for the mean-field ising model. *Communications in Mathematical Physics*, 289(2):725–764, 2009.
- 15 Martin Dyer, Leslie Ann Goldberg, Catherine Greenhill, and Mark Jerrum. The relative complexity of approximate counting problems. *Algorithmica*, 38:471–500, 2004.
- 16 Charilaos Efthymiou and Kostas Zampetakis. On sampling diluted spin glasses using glauber dynamics. *arXiv preprint*, 2024. [arXiv:2403.08921](https://arxiv.org/abs/2403.08921).
- 17 Ahmed El Alaoui, Andrea Montanari, and Mark Sellke. Sampling from the sherrington-kirkpatrick gibbs measure via algorithmic stochastic localization. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 323–334. IEEE, 2022.
- 18 Ronen Eldan, Frederic Koehler, and Ofer Zeitouni. A spectral condition for spectral gap: fast mixing in high-temperature ising models. *Probability theory and related fields*, 182(3):1035–1051, 2022.

- 19 Weiming Feng, Heng Guo, and Jiaheng Wang. Swendsen-Wang dynamics for the ferromagnetic Ising model with external fields. *Information and Computation*, 294:105066, 2023.
- 20 Joel Friedman. *A proof of Alon's second eigenvalue conjecture and related problems*. American Mathematical Soc., 2008.
- 21 Andreas Galanis, Daniel Štefankovič, and Eric Vigoda. Inapproximability of the partition function for the antiferromagnetic Ising and hard-core models. *Combinatorics, Probability and Computing*, 25(4):500–559, 2016.
- 22 Antoine Gerschenfeld and Andrea Montanari. Reconstruction for models on random graphs. In *2007 48th Annual IEEE Symposium on Foundations of Computer Science*, pages 194–204, 2007.
- 23 Leslie Ann Goldberg and Mark Jerrum. Approximating the partition function of the ferromagnetic Potts model. *J. ACM*, 59(5), 2012.
- 24 Heng Guo and Mark Jerrum. Random cluster dynamics for the Ising model is rapidly mixing. *The Annals of Applied Probability*, 28(2):1292–1313, 2018.
- 25 Linus Hamilton, Frederic Koehler, and Ankur Moitra. Information theoretic properties of Markov random fields, and their algorithmic applications. In *Advances in Neural Information Processing Systems*, volume 30, 2017.
- 26 Brice Huang, Andrea Montanari, and Huy Tuan Pham. Sampling from spherical spin glasses in total variation via algorithmic stochastic localization. *arXiv preprint*, 2024. [arXiv:2404.15651](https://arxiv.org/abs/2404.15651).
- 27 Vishesh Jain, Frederic Koehler, and Andrej Risteski. Mean-field approximation, convex hierarchies, and the optimality of correlation rounding: a unified perspective. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 1226–1236, 2019.
- 28 Mark Jerrum and Alistair Sinclair. Polynomial-time approximation algorithms for the Ising model. *SIAM Journal on computing*, 22(5):1087–1116, 1993.
- 29 Mark R. Jerrum, Leslie G. Valiant, and Vijay V. Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theoretical Computer Science*, 43:169–188, 1986.
- 30 Adam Klivans and Raghu Meka. Learning graphical models using multiplicative weights. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 343–354, 2017.
- 31 Frederic Koehler, Holden Lee, and Andrej Risteski. Sampling approximately low-rank Ising models: MCMC meets variational methods. In *Conference on Learning Theory*, pages 4945–4988. PMLR, 2022.
- 32 Dmitriy Kunisky. Optimality of Glauber dynamics for general-purpose Ising model sampling and free energy approximation. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 5013–5028, 2024.
- 33 Steffen L Lauritzen. *Graphical models*, volume 17. Clarendon Press, 1996.
- 34 David A Levin, Malwina J Luczak, and Yuval Peres. Glauber dynamics for the mean-field Ising model: cut-off, critical power law, and metastability. *Probability Theory and Related Fields*, 146:223–265, 2010.
- 35 Liang Li, Pinyan Lu, and Yitong Yin. Correlation decay up to uniqueness in spin systems. In *Proceedings of the twenty-fourth annual ACM-SIAM symposium on Discrete algorithms*, pages 67–84, 2013.
- 36 Kuikui Liu, Sidhanth Mohanty, Amit Rajaraman, and David X Wu. Fast mixing in sparse random Ising models. *arXiv preprint*, 2024. [arXiv:2405.06616](https://arxiv.org/abs/2405.06616).
- 37 Eyal Lubetzky and Allan Sly. Critical Ising on the square lattice mixes in polynomial time. *Communications in Mathematical Physics*, 313(3):815–836, 2012.
- 38 Fabio Martinelli and Enzo Olivieri. Approach to equilibrium of Glauber dynamics in the one phase region: I. The attractive case. *Communications in Mathematical Physics*, 161(3):447–486, 1994.

70:14 On Sampling from Ising Models with Spectral Constraints

- 39 Fabio Martinelli and Enzo Olivieri. Approach to equilibrium of Glauber dynamics in the one phase region: II. The general case. *Communications in Mathematical Physics*, 161(3):487–514, 1994.
- 40 Marc Mezard and Andrea Montanari. *Information, physics, and computation*. Oxford University Press, 2009.
- 41 Marc Mézard, Giorgio Parisi, and Miguel Angel Virasoro. *Spin glass theory and beyond: An Introduction to the Replica Method and Its Applications*, volume 9. World Scientific Publishing Company, 1987.
- 42 Andrea Montanari, Elchanan Mossel, and Allan Sly. The weak limit of Ising models on locally tree-like graphs. *Probability Theory and Related Fields*, 152:31–51, 2012.
- 43 Elchanan Mossel and Allan Sly. Exact thresholds for Ising–Gibbs samplers on general graphs. *The Annals of Probability*, 41(1), 2013.
- 44 Andrej Risteski. How to calculate partition functions using convex programming hierarchies: provable bounds for variational methods. In *Conference on Learning Theory*, pages 1402–1416. PMLR, 2016.
- 45 Alistair Sinclair, Piyush Srivastava, and Marc Thurley. Approximation algorithms for two-state anti-ferromagnetic spin systems on bounded degree graphs. *Journal of Statistical Physics*, 155(4):666–686, 2014.
- 46 Allan Sly. Computational transition at the uniqueness threshold. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 287–296. IEEE, 2010.
- 47 Allan Sly and Nike Sun. The computational hardness of counting in two-spin models on d -regular graphs. In *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*, pages 361–369. IEEE, 2012.
- 48 Michel Talagrand. *Mean field models for spin glasses: Volume I: Basic examples*, volume 54. Springer Science & Business Media, 2010.