2012 IEEE 53rd Annual Symposium on Foundations of Computer Science

The computational hardness of counting in two-spin models on d-regular graphs

Allan Sly
Department of Statistics
University of California, Berkeley
Berkeley, California, United States
sly@stat.berkeley.edu

Nike Sun
Department of Statistics
Stanford University
Stanford, California, United States
nikesun@stanford.edu

***Abstract*—The class of two-spin systems contains several important models, including random independent sets and the Ising model of statistical physics. We show that for both the hard-core (independent set) model and the anti-ferromagnetic Ising model with arbitrary external field, it is NP-hard to approximate the partition function or approximately sample from the model on regular graphs when the model has non-uniqueness on the corresponding regular tree. Together with results of Jerrum–Sinclair, Weitz, and Sinclair–Srivastava–Thurley giving FPRAS's for all other two-spin systems except at the uniqueness threshold, this gives an almost complete classification of the computational complexity of two-spin systems on bounded-degree graphs.**

Our proof establishes that the normalized log-partition function of any two-spin system on bipartite locally tree-like graphs converges to a limiting "free energy density" which coincides with the (non-rigorous) Bethe prediction of statistical physics. We use this result to characterize the local structure of two-spin systems on locally tree-like bipartite expander graphs, which then become the basic gadgets in a randomized reduction to approximate MAX–CUT. Our approach is novel in that it makes no use of the second moment method employed in previous works on these questions.

***Keywords*-spin system, hard-core model, independent set, Ising model, Bethe free energy**

I. Introduction

Spin systems are stochastic models defined by local interactions on networks. The class of spin systems includes well-known combinatorial counting and constraint satisfaction problems. In this paper we classify the complexity of approximating the partition function for all homogeneous two-spin systems on bounded-degree graphs.

When interactions favor agreement of adjacent spins, the model is said to be *ferromagnetic*. Jerrum and Sinclair [1] gave a fully polynomial-time randomized approximation scheme (FPRAS) for approximating the partition function (the normalizing constant in the probability distribution) of the ferromagnetic Ising model, which covers all ferromagnetic two-spin systems. For anti-ferromagnetic systems such as the hard-core and anti-ferromagnetic Ising models, the complexity of approximating the partition function depends on the model parameters, and is known to be NP-hard when the interactions are sufficiently strong. Our first main result establishes that the computational transition for such models on d-regular graphs is located precisely at the uniqueness threshold (see Defn. I.6) for the corresponding model on the d-regular tree.

Theorem 1. For $d \geq 3$ and $\lambda > \lambda_c(d) = \frac{(d-1)^{d-1}}{(d-2)^d}$, unless NP $=$ RP there exists no FPRAS for the partition function of the hard-core model with fugacity λ on d-regular graphs.

The transition point $\lambda_c(d)$ is the uniqueness threshold for the hard-core model on the d-regular tree: it marks the point above which distant boundary conditions have a non-vanishing influence on the spin at the root. In a seminal paper [2], Weitz used computational tree methods to provide an FPTAS for the partition function of the hard-core model on graphs of maximum degree d at any $\lambda < \lambda_c(d)$. Together with Weitz's result, Thm. 1 completes the classification of the complexity of the hard-core model except at the threshold λ_c.

Previously it was shown that there is no FPRAS for the hard-core model at $\lambda d \geq 10000$ [3]. In the case of $\lambda = 1$ this was improved to $d \geq 25$ [4], [5], using random regular bipartite graphs as basic gadgets in a hardness reduction. Mossel et al. [6] showed that local MCMC algorithms are exponentially slow for $\lambda > \lambda_c(d)$, and conjectured that λ_c is in fact the threshold for existence of an FPRAS.

The first rigorous result establishing a computational transition at the uniqueness threshold appeared in [7], where hardness was shown for $\lambda_c(d) < \lambda < \lambda_c(d) + \epsilon(d)$ for some $\epsilon(d) > 0$. The proof relies on a detailed analysis of the hard-core model on random bipartite graphs, which are then used in a randomized reduction to MAX-CUT. More precisely the result of [7] gives hardness subject to a technical condition which was an artifact of a difficult second moment calculation from [6], and which could only be verified for $\lambda < \lambda_c(d)+\epsilon(d)$. Hardness was subsequently shown by Galanis et al. [8] for all $\lambda > \lambda_c(d)$ when $d \neq 4, 5$ by verifying the technical condition of [7].

In this paper we follow a different approach which is more conceptual and completely circumvents second moment method calculations. Moreover the same method of proof gives the analogous result for anti-ferromagnetic Ising models with arbitrary external field:

DOI 10.1109/FOCS.2012.56

Theorem 2. For $d \geq 3$, $B \in \mathbb{R}$ and $\beta < \beta_{c,\mathrm{af}}(B, d) < 0$, unless NP = RP there does not exist an FPRAS for the partition function of the anti-ferromagnetic Ising model with inverse temperature β and external field B on d-regular graphs.

Here $\beta_{c,\mathrm{af}}(B, d)$ denotes the uniqueness threshold for the anti-ferromagnetic Ising model with external field B on the d-regular tree. Extending the methods of Weitz [2], Sinclair et al. [9] (see also [10]) gave an FPTAS for the anti-ferromagnetic Ising model on d-regular graphs at inverse temperature $\beta > \beta_{c,\mathrm{af}}(B, d)$, so together with Thm. 2 this again establishes that the computational transition coincides with the tree uniqueness threshold.

The hard-core and anti-ferromagnetic Ising models together encompass all (non-degenerate) homogeneous two-spin systems on d-regular graphs (see §III-A). Thus, the results of [2], [1], [9] combined with Thms. 1 and 2 give a full classification of the computational complexity of approximating the partition function for (homogeneous) two-spin systems on d-regular graphs, except at the uniqueness thresholds $\lambda_c(d)$ and $\beta_{c,\mathrm{af}}(B, d)$.

In fact, we will show inapproximability in non-uniqueness regimes in a strong sense: not only does there not exist an FPRAS, but for any fixed choice of model parameters and d there exists $c > 0$ such that it is NP-hard even to approximate the partition function within a factor of e^{cn} on the class of d-regular graphs.

Independent results of Galanis–Štefankovič–Vigoda

In a simultaneous and independent work, Galanis, Štefankovič and Vigoda [11] established the result of Thm. 1, and Thm. 2 in the case of zero external field ($B = 0$). Their methods differ from ours: they analyze the second moment of the partition function on random bipartite d-regular graphs, and establish the condition necessary to apply the approach of [7]. They analyze a difficult optimization of a real function in several variables by relating the problem to certain tree recursions.

A. Reduction to MAX-CUT *via bipartite graphs*

Our proof is based on a detailed characterization (Thm. 5) of the local structure of anti-ferromagnetic two-spin systems on symmetric bipartite d-regular locally tree-like graphs. Specifically, we show that the joint distribution of all the spins in a large neighborhood of a typical vertex in the graph converges to a known (Gibbs) measure on the d-regular tree. Under the additional assumption that that the graph is an edge expander, when the model has non-uniqueness on the d-regular tree the spin distribution on the graph is divided into $+$ and $-$ phases where one or the other side of the graph has a linear number more vertices with $+$ spin.

Our main results Thms. 1 and 2 are then proved by a variation on the construction of [7], using the bipartite graphs in a randomized reduction to approximate MAX-CUT on 3-regular graphs, which is known to be NP-hard [12]. First, we use Thm. 5 to construct a symmetric bipartite d-regular locally tree-like graph G of large constant size such that, conditioned on the phase of the global configuration, spins at distant vertices are asymptotically independent with known marginals depending only on the side of the graph (Propn. III.2).

Given a 3-regular graph H on which we wish to approximate MAX-CUT, first we take a disjoint copy G_v of G for each vertex $v \in H$. After removing $3k$ edges from each G_v, for each edge $(u, v) \in H$ we add k edges joining each side of G_u to the corresponding side of G_v in such a way that the resulting graph H^G is d-regular.

The connections between gadgets do not substantially change the spin distributions inside them, and in particular the $\pm$ phases remain. The anti-ferromagnetic nature of the interaction, however, results in neighboring copies of G in H^G preferring to be in opposing phases. Using the asymptotic conditional independence result Propn. III.2 we can estimate the partition function for the model on H^G restricted to configurations of given phase on each copy of G within a factor of $e^{\epsilon|H|}$ (Lem. III.4). We find that the distribution is concentrated on configurations where the vector of phases gives a good cut of H, and the effect is strengthened as k is increased. Thus, for any $\epsilon > 0$, by taking k (hence G) to be sufficiently large a $(1 + \epsilon)$-approximation of MAX-CUT(H) can be determined from the partition function of the model on H^G, thereby completing the reduction.

Our reduction depends crucially on the detailed picture of the spin distribution developed in Thm. 5 and Propn. III.2. Using methods developed in [13], these results in turn are obtained as consequences of precise asymptotics for the partition function of two-spin models on bipartite d-regular graphs: we show that the log-partition function, normalized by the number of vertices in the graph, has an asymptotic value, the "free energy density," which is easily computed from the (non-rigorous) "Bethe prediction" of statistical physics (see §II). This is a result of independent interest, since lower bounds for partition functions on graphs have proved to be in general challenging. Our proof avoids the use of the second moment method and its heavy calculations and instead bounds the derivative of the partition function directly using properties of Gibbs measures on trees. Asymptotics for the partition function on general tree-like graphs were established for the ferromagnetic Ising model in [14], [15], [16], and for more general spin systems in uniqueness regimes in [16]. Our result for anti-ferromagnetic models is stated somewhat informally as follows; for the precise statement see Thm. 4.

Theorem 3. For any non-degenerate homogeneous two-spin model on bipartite d-regular locally tree-like graphs, the log-partition function normalized by the number of vertices has an asymptotic value which coincides with the Bethe free energy prediction.

We now formally introduce the models which we consider. We then define the notion of local (weak) convergence of graphs and give precise statements of our results on the partition function (Thm. 4) and local structure (Thm. 5) of these models on bipartite graphs.

B. Definition of spin systems

Let $G = (V, E)$ be a finite undirected graph, and $\mathscr{X}$ a finite alphabet of *spins*. A *spin system* or *spin model* on G is a probability measure on the space of *(spin) configurations* $\underline{\sigma} \in \mathscr{X}^V$ of form

$$\nu_G^{\underline{\psi}}(\underline{\sigma}) = \frac{1}{Z_G(\underline{\psi})} \prod_{(ij)\in E} \psi(\sigma_i, \sigma_j) \prod_{i\in V} \bar{\psi}(\sigma_i), \tag{I.1}$$

where ψ is a symmetric function $\mathscr{X}^2 \to \mathbb{R}_{\geq 0}$, $\bar{\psi}$ is a positive function $\mathscr{X} \to \mathbb{R}_{\geq 0}$, and $Z_G(\underline{\psi})$ is the normalizing constant, called the *partition function*. The pair $\underline{\psi} \equiv (\psi, \bar{\psi})$ is called a *specification* for the spin system (I.1).

In this paper we consider spin systems with an alphabet of size two; without loss $\mathscr{X} \equiv \{\pm 1\}$. The *Ising* model on G at inverse temperature β and external field B is given by

$$\nu_G^{\beta,B}(\underline{\sigma}) = \frac{1}{Z_G(\beta, B)} \prod_{(ij)\in E} e^{\beta\sigma_i\sigma_j} \prod_{i\in V} e^{B\sigma_i}. \tag{I.2}$$

The *hard-core* (or *independent set*) model on G at activity or fugacity λ is given by

$$\nu_G^{\lambda}(\underline{\sigma}) = \frac{1}{Z_G(\lambda)} \prod_{(ij)\in E} \mathbf{1}\{\bar{\sigma}_i\bar{\sigma}_j \neq 1\} \prod_{i\in V} \lambda^{\bar{\sigma}_i} \tag{I.3}$$

where $\bar{\sigma} \equiv \mathbf{1}\{\sigma = +1\} = (1+\sigma)/2$. Our definition (I.3) is trivially equivalent to the standard definition of the hard-core model which has spin 0 in place of -1, but we take $\mathscr{X} = \{\pm 1\}$ throughout to unify the notation.

C. Local convergence and the Bethe prediction

If G is any graph and v a vertex in G, write $B_t(v)$ for the subgraph induced by the vertices of G at graph distance at most t from v, and $\partial v \equiv B_1(v)\backslash\{v\}$ for the neighbors of v. We let $T \equiv (T, o)$ denote a general tree with root o, with $T^t \equiv B_t(o) \subseteq T$ the subtree of depth t. We also fix d throughout and write $\mathbb{T} \equiv (\mathbb{T}, o)$ for the rooted d-regular tree.

Definition I.1. Let $G_n = (V_n = [n], E_n)$ be a sequence of (random) finite undirected graphs, and let I_n denote a uniformly random vertex in V_n. The sequence G_n is said to *converge locally* to the d-regular tree $\mathbb{T}$ if for all $t \geq 0$, $B_t(I_n)$ converges to $\mathbb{T}^t$ in distribution with respect to the joint law $\mathbb{P}_n$ of (G_n, I_n): that is, $\lim_{n\to\infty} \mathbb{P}_n(B_t(I_n) \cong \mathbb{T}^t) = 1$ (where $\cong$ denotes graph isomorphism).

We write $\mathbb{E}_n$ for expectation with respect to $\mathbb{P}_n$ and impose the following integrability condition on the degree of I_n:

Definition I.2. The sequence G_n is *uniformly sparse* if the random variables $|\partial I_n|$ are uniformly integrable, that is, if

$$\lim_{L\to\infty} \limsup_{n\to\infty} \mathbb{E}_n[|\partial I_n| \mathbf{1}\{|\partial I_n| \geq L\}] = 0.$$

We assume throughout that G_n ($n \geq 1$) is a uniformly sparse graph sequence converging locally to the d-regular tree $\mathbb{T}$; this setting is hereafter denoted $G_n \to_{loc} \mathbb{T}$, and we write $Z_n \equiv Z_{G_n}(\underline{\psi})$. The *free energy density* for a specification $\underline{\psi}$ on G_n is defined by

$$\phi \equiv \lim_{n\to\infty} \phi_n \equiv \lim_{n\to\infty} \frac{1}{n} \mathbb{E}_n[\log Z_n], \tag{I.4}$$

provided the limit exists. For ferromagnetic spin systems on a broad class of locally tree-like graphs, heuristic methods from statistical physics yield an explicit (conjectural) formula for the value of ϕ, the so-called "Bethe prediction" Φ whose definition we recall in §II. For anti-ferromagnetic two-spin models, the Bethe prediction is well-defined only on graph sequences G_n which are nearly bipartite, in the following sense: let $\mathbb{T}_+$ denote the d-regular tree $\mathbb{T}$ with vertices colored $+1$ (black) or -1 (white) according to whether they are at even or odd distance from the root o; let $\mathbb{T}_-$ be $\mathbb{T}_+$ with the colors reversed. Let $\mathbf{T}$ be the random tree which equals $\mathbb{T}_+$ or $\mathbb{T}_-$ with equal probability; write $\mathbf{P}$ for the law of $\mathbf{T}$ and $\mathbf{E}$ for expectation with respect to $\mathbf{P}$.

Definition I.3. For $G_n \to_{loc} \mathbb{T}$, we say the G_n are *nearly bipartite*, and write $G_n \to_{loc} \mathbf{T}$ (equivalently $G_n \to_{loc} \mathbf{P}$), if there exists a (not necessarily proper) black-white coloring of G_n such that for all $t \geq 0$, $B_t(I_n) \to \mathbf{T}^t$ in distribution.

The precise statement of Thm. 3 is then as follows:

Theorem 4. Let $\underline{\psi}$ specify a non-degenerate homogeneous two-spin system.

(a) If $\underline{\psi}$ is ferromagnetic, then ϕ exists for any $G_n \to_{loc} \mathbb{T}$ and equals $\Phi_{\{\mathbb{T}\}}$ as defined by (II.2) (and given more explicitly by (II.4)).
(b) If $\underline{\psi}$ is anti-ferromagnetic, then ϕ exists for any $G_n \to_{loc} \mathbf{T}$ and equals $\Phi_{\{\mathbb{T}_\pm\}}$ as defined in (II.2) (and given more explicitly by (II.3)).

Remark I.4. Hereafter we treat $G_n \to_{loc} \mathbb{T}$ and $G_n \to_{loc} \mathbf{T}$ in a unified manner when possible by writing $G_n \to_{loc} \mathbb{P}_{\mathcal{T}}$ for $\mathbb{P}_{\mathcal{T}}$ the uniform measure on $\mathcal{T}$, which always denotes either $\{\mathbb{T}\}$ or $\{\mathbb{T}_\pm\}$. We write $\mathbb{E}_{\mathcal{T}}$ for expectation with respect to $\mathbb{P}_{\mathcal{T}}$.

D. Local structure of measures

Under some additional assumptions on G_n, Thm. 4, together with the arguments of [13], characterizes the asymptotic local structure of the spin systems $\nu_n \equiv \nu_{G_n}$.

For $G_n \to_{loc} \mathbf{T}$, let $\tau : V_n \to \{\pm\}$ denote the given black-white coloring of the vertices of G_n (hereafter writing $\pm$ as shorthand for ± 1). We say that G_n is *symmetric* if it is isomorphism-invariant to reversing the black-white coloring. For a spin configuration $\underline{\sigma} \in G_n$ we define the *phase* of $\underline{\sigma}$ to be

$$Y(\underline{\sigma}) \equiv \operatorname{sgn} \sum_i \tau_i \sigma_i, \quad \operatorname{sgn} x \equiv \mathbf{1}\{x \geq 0\} - \mathbf{1}\{x < 0\}.$$

Let $\nu_n^{\pm}$ denote the measure ν_n conditioned on the configurations of $\pm$ phase: that is,

$$\nu_n^{\pm}(\underline{\sigma}) \equiv \frac{1}{Z_n^{\pm}} \mathbf{1}\{Y(\underline{\sigma}) = \pm\} \prod_{(ij) \in E_n} \psi(\sigma_i, \sigma_j) \prod_{i \in V_n} \bar{\psi}(\sigma_i),$$

where $Z_n^{\pm}$ is the partition function restricted to the $\pm$ configurations. We will characterize the local structure of the measures $\nu_n^{\pm}$ on graph sequences satisfying an edge-expansion assumption, as follows:

Definition I.5. A graph $G = (V, E)$ is a $(\delta, \gamma, \lambda)$*-edge expander* if, for any set of vertices $S \subseteq V$ with $\delta|V| \leq |S| \leq \gamma|V|$, there are at least $\lambda|S|$ edges joining S to $V \backslash S$.

The measures $\nu_n^{\pm}$ will be related to Gibbs measures on the infinite tree. In particular, recall the definition of (Gibbs) uniqueness:

Definition I.6. For a rooted tree T, let $\mathscr{G}_T$ denote the set of Gibbs measures for the specification $\underline{\psi}$ on T. The specification is said to have *(Gibbs) uniqueness* (on T) if $|\mathscr{G}_T| = 1$.

Recalling Rmk. I.4, let $\mathscr{G}_{\mathcal{T}}$ denote the space of mappings $\nu : T \mapsto \nu(T)$, $T \in \mathcal{T}$ (with $\mathscr{G}_{\{\mathbb{T}\}} \hookrightarrow \mathscr{G}_{\{\mathbb{T}_{\pm}\}}$ in the obvious manner). When $\mathcal{T} = \{\mathbb{T}_{\pm}\}$ we write $\nu_{\pm}$ as shorthand for $\nu(\mathbb{T}_{\pm})$.

Definition I.7. An element $\nu \in \mathscr{G}_{\mathcal{T}}$ is *translation-invariant* if for $(T, o) \in \mathcal{T}$ and any vertex $x \in T$, the law on spin configurations of (T, x) induced by $\nu(T, o)$ coincides with $\nu(T, x)$.

(If $\mathcal{T} = \{\mathbb{T}\}$ then the preceding agrees with the usual definition of translation-invariance, whereas if $\mathcal{T} = \{\mathbb{T}_{\pm}\}$ then the projections $\nu(\mathbb{T}^{\pm})$ are *semi*-translation-invariant.)

For a two-spin model, let ν^+ (resp. ν^-) be the elements of $\mathscr{G}_{\mathcal{T}}$ defined by conditioning on all spins identically equal to 1 on the t-th level of black (resp. white) vertices and taking the weak limit as $t \to \infty$; the $\nu^{\pm}$ are translation-invariant. The projections $\mu^+ \equiv \nu_+^+ \equiv \nu^+(\mathbb{T}_+)$ and $\mu^- \equiv \nu_+^- \equiv \nu^-(\mathbb{T}_+)$, disregarding the black-white coloring on $\mathbb{T}_+$, are the extremal semi-translation-invariant Gibbs measures for the model on $\mathbb{T}$, and by symmetry

$$\mu^+ = \nu_-^- \equiv \nu^-(\mathbb{T}_-), \quad \mu^- = \nu_-^+ \equiv \nu^+(\mathbb{T}_-).$$

The model has uniqueness if and only if $\mu^+ = \mu^-$.

Definition I.8. For $G_n \sim \mathbb{P}_n$ a random graph sequence and ν_n any law on spin configurations $\underline{\sigma}_n$ of G_n, we say that $\mathbb{P}_n \otimes \nu_n$ *converges locally (weakly)* to $\mathbb{P}_{\mathcal{T}} \otimes \nu$ (for $\nu \in \mathscr{G}_{\mathcal{T}}$), and write $\mathbb{P}_n \otimes \nu_n \to_{loc} \mathbb{P}_{\mathcal{T}} \otimes \nu$, if it holds for all $t \geq 0$ that $(B_t(I_n), \underline{\sigma}_{B_t(I_n)})$ converges in distribution to $(T^t, \underline{\sigma}_t)$ where $T \sim \mathbb{P}_{\mathcal{T}}$ and $\underline{\sigma}_t$ is the restriction to T^t of $\underline{\sigma} \sim \nu(T)$.

Remark I.9. In [13, Defn. 2.3] three forms A, B, C of local convergence of measures are distinguished, with $C \Rightarrow B \Rightarrow A$. Our Defn. I.8 corresponds to the weakest form A: however, as explained in the proof of [13, Thm. 2.4 (II)], if the $(\nu(T))_{T \in \mathcal{T}}$ are *extremal* Gibbs measures then A, B, C are easily seen to be equivalent, so convergence in the sense of Defn. I.8 implies convergence in the *a priori* stronger sense of

$$\|\mathbb{P}_n[(B_t(I_n), \underline{\sigma}_{B_t(I_n)}) = \cdot] - \mathbb{P}_{\mathcal{T}}[(T^t, \underline{\sigma}_{T^t}) = \cdot]\|_{\text{TV}} \to 0.$$

Theorem 5. For any anti-ferromagnetic two-spin system on $G_n \to_{loc} \mathbf{T}$, the following hold:

(a) If the G_n are symmetric, then $\mathbb{P}_n \otimes \nu_n \to_{loc} \mathbf{P} \otimes [(\nu^+ + \nu^-)/2]$.

(b) If for all $\delta > 0$ the G_n are $(\delta, 1/2, \lambda_\delta)$-edge expanders for some $\lambda_\delta > 0$, then

$$\mathbb{P}_n \otimes \nu_n^{\pm} \to_{loc} \mathbf{P} \otimes \nu^{\pm}. \tag{I.5}$$

Further, with $\langle\,\rangle_\mu$ denoting expectation with respect to the Gibbs measure μ,

$$\frac{1}{n} Y(\underline{\sigma}) \sum_{i \in V} \tau_i \sigma_i \to \frac{1}{2}[\langle \sigma_o \rangle_{\mu^+} - \langle \sigma_o \rangle_{\mu^-}] \quad \text{in probability.} \tag{I.6}$$

Outline of the paper

In §II we review the Bethe prediction in the d-regular setting. In §III we prove the approximate conditional independence statement (Propn. III.2) and demonstrate the randomized reduction to MAX-CUT to prove our main results Thms. 1 and 2. The proofs of Thm. 3 (in the form Thm. 4) and Thm. 5 are given in the full version of this extended abstract, which may be found at `http://arxiv.org/abs/1203.2602`.

II. The Bethe prediction

In this section we review the Bethe prediction which gives the limiting free energy density ϕ for two-spin models on graph sequences $G_n \to_{loc} \mathbf{T}$ in Thm. 4. We refer to [17], [16] for more general background and references on the Bethe prediction, and here we describe only its specialization to the d-regular setting.

Given $\mathcal{T}$, let $\mathcal{T}_{\text{e}}$ denote the set of trees T rooted not at a vertex but at an oriented edge $x \to y$, obtained by distinguishing an oriented edge in $T \in \mathcal{T}$ and forgetting the root. Elements of $\mathcal{T}, \mathcal{T}_{\text{e}}$ are regarded modulo isomorphism: thus if $\mathcal{T} = \{\mathbb{T}\}$ then $\mathcal{T}_{\text{e}} = \{(\mathbb{T}, o \to j)\}$, and if $\mathcal{T} = \{\mathbb{T}_{\pm}\}$ then $\mathcal{T}_{\text{e}} = \{(\mathbb{T}_{\pm}, o \to j)\}$.

Let Δ denote the $(|\mathscr{X}|-1)$-dimensional simplex of probability measures on $\mathscr{X}$. A *message* is a mapping $h : \mathcal{T}_{\rm e} \to \Delta$; we write $\mathcal{H} \equiv \mathcal{H}(\mathcal{T})$ for the space of messages on $\mathcal{T}_{\rm e}$. For $T \in \mathcal{T}$, $x \to y$ in T, and $h \in \mathcal{H}$, write $h_{x\to y}$ for the image of $(T, x \to y) \in \mathcal{T}_{\rm e}$ under h, and define

$$\Phi_T(h) \equiv \Phi_T^{\rm vx}(h) - \Phi_T^{\rm e}(h)$$

where

$$\Phi_T^{\rm vx}(h) \equiv \log\bigg\{\sum_{\sigma_o}\bar\psi(\sigma_o)\prod_{j\in\partial o}\bigg(\sum_{\sigma_j}\psi(\sigma_o,\sigma_j)h_{j\to o}(\sigma_j)\bigg)\bigg\},$$

$$\Phi_T^{\rm e}(h) \equiv \frac{1}{2}\sum_{j\in\partial o}\log\bigg\{\sum_{\sigma_o,\sigma_j}\psi(\sigma_o,\sigma_j)h_{o\to j}(\sigma_o)h_{j\to o}(\sigma_j)\bigg\}.$$

The *Bethe free energy functional* on $\mathcal{H}(\mathcal{T})$ is defined by $\Phi_{\mathcal{T}}(h) \equiv \mathbb{E}_{\mathcal{T}}[\Phi_T(h)]$.

The *Bethe* or *belief propagation (BP) recursion* is the map

$$\mathsf{BP} \equiv \mathsf{BP}_{\mathcal{T}} : \mathcal{H}(\mathcal{T}) \to \mathcal{H}(\mathcal{T}),$$
$$(\mathsf{BP}h)_{x\to y}(\sigma) \equiv \bar{\mathsf{F}}[(h_{v\to x})_{v\in\partial x\setminus y}]$$

for $\bar{\mathsf{F}} : \Delta^{d-1} \to \Delta$ defined by

$$[\bar{\mathsf{F}}(\underline{h})](\sigma) \cong \bar\psi(\sigma)\prod_{j=1}^{d-1}\bigg\{\sum_{\sigma_j}\psi(\sigma,\sigma_j)h_j(\sigma_j)\bigg\}, \qquad \text{(II.1)}$$

where $\underline{h} \equiv (h_1,\ldots,h_{d-1}) \in \Delta^{d-1}$ and $\cong$ denotes equivalence up to a positive normalizing factor.

Definition II.1. For any homogeneous spin system on $G_n \to_{loc} \mathbb{P}_{\mathcal{T}}$, the *Bethe prediction* is that the free energy density ϕ of (I.4) exists and equals

$$\Phi \equiv \Phi_{\mathcal{T}} \equiv \sup_{h\in\mathcal{H}_\star}\Phi_{\mathcal{T}}(h) \qquad \text{(II.2)}$$

with $\mathcal{H}_\star \equiv \mathcal{H}_\star(\mathcal{T}) \subseteq \mathcal{H}(\mathcal{T})$ the set of all fixed points of $\mathsf{BP}_{\mathcal{T}}$.

For $h \in \Delta$ write $\mathsf{F}(h) \equiv \bar{\mathsf{F}}(h,\ldots,h)$: then $\mathcal{H}_\star(\{\mathbb{T}\})$ corresponds simply to the fixed points of F in simplex. For $h \in \mathcal{H}(\{\mathbb{T}_\pm\})$ we write $h_\pm \equiv h(\mathbb{T}_\pm, o \to j) \in \Delta$: then any $h \in \mathcal{H}_\star(\{\mathbb{T}_\pm\})$ must satisfy $h_\pm = \mathsf{F}(h_\mp)$, so $\mathcal{H}_\star(\{\mathbb{T}_\pm\})$ corresponds to the fixed points of the *double* recursion $\mathsf{F}^{(2)} \equiv \mathsf{F}\circ\mathsf{F}$.

In verifying the Bethe prediction we will identify the fixed points attaining the supremum in (II.2). In the anti-ferromagnetic case, with h^+ (resp. h^-) denoting the elements $h \in \mathcal{H}_\star(\{\mathbb{T}_\pm\})$ maximizing $h_+(+)$ (resp. $h_-(+)$), we will see that

$$\Phi_{\{\mathbb{T}_\pm\}} = \Phi_{\{\mathbb{T}_\pm\}}(h^+) = \Phi_{\{\mathbb{T}_\pm\}}(h^-). \qquad \text{(II.3)}$$

Explicitly, $h_+^+ = h_-^-$ (resp. $h_-^+ = h_+^-$) will be the fixed points of $\mathsf{F}^{(2)}$ giving maximal (resp. minimal) probability to spin $+$. The ferromagnetic case reduces to the Ising model: here, with $h^\pm$ denoting the elements of $\mathcal{H}_\star(\{\mathbb{T}\})$ maximizing $h_{o\to j}(\pm)$ on $\mathbb{T}$, we will see that

$$\Phi_{\{\mathbb{T}\}} = \Phi_{\{\mathbb{T}\}}(h^{\operatorname{sgn} B}). \qquad \text{(II.4)}$$

III. Computational hardness

In this section we prove the hardness results Thms. 1 and 2. In §III-A we show that for purposes of computing ϕ on d-regular locally tree-like graph sequences, all non-degenerate two-spin systems reduce to the Ising or hard-core models. In §III-B we construct and analyze the bipartite expander gadgets to be used in the reduction to MAX-CUT. We complete the reduction in §III-C, concluding the proof.

A. Reduction to Ising and hard-core on d-regular graphs

We now show that for the computation of the free energy density, all (non-degenerate) homogeneous two-spin models on graph sequences $G_n \to_{loc} \mathbf{T}$ reduce to either the Ising or hard-core model. Indeed, let $\underline{\psi} \equiv (\psi,\bar\psi)$ be a specification for a two-spin system with alphabet $\mathscr{X} = \{\pm\}$. If we define $\underline{\psi}'$ by $\psi'(\sigma,\sigma') \equiv \psi(\sigma,\sigma')\bar\psi(\sigma)^{1/d}\bar\psi(\sigma')^{1/d}$, and $\bar\psi'(\sigma) \equiv 1$, then

$$\frac{1}{n}\log Z_G(\underline{\psi}) - \frac{1}{n}\log Z_G(\underline{\psi}') = O(\mathbb{E}_n[|\partial I_n|\mathbf{1}\{|\partial I_n| \neq d\}]),$$

which for $G_n \to_{loc} \mathbb{T}$ tends to zero as $n \to \infty$ by uniform sparsity. Therefore we assume without loss $\bar\psi \equiv 1$, and consider the possibilities for ψ:

(1) If $\psi > 0$, then $\psi(\sigma,\sigma') = e^{B_0}e^{\beta\sigma\sigma'}e^{B\sigma/d}e^{B\sigma'/d}$ for β, B, B_0 defined by

$$\frac{\psi(+,+)}{\psi(-,-)} = e^{4B/d}, \quad \frac{\psi(+,+)\psi(-,-)}{\psi(+,-)^2} = e^{4\beta},$$
$$\psi(+,+)\psi(+,-)^2\psi(-,-) = e^{4B_0},$$

so $\phi_n - (d/2)B_0$ is asymptotically equal to the free energy density for the Ising model on G_n with parameters (β, B).

(2) If $\psi(+,-) = \psi(-,+) > 0$ and $\psi(-,-) > \psi(+,+) = 0$, then, recalling $\bar\sigma \equiv \mathbf{1}\{\sigma = +\}$, we have $\psi(\sigma,\sigma') = e^{B_0}\mathbf{1}\{\bar\sigma\bar\sigma' \neq 1\}\lambda^{\bar\sigma/d}\lambda^{\bar\sigma'/d}$ for B_0, λ defined by

$$\psi(-,-) \equiv e^{B_0}, \quad \frac{\psi(+,-)}{\psi(-,-)} \equiv \lambda^{1/d}.$$

Therefore $\phi_n - (d/2)B_0$ is asymptotically equal to the free energy density for the independent set model on G_n at fugacity λ.

The remaining two-spin models are degenerate, with free energy density which is easy to calculate:

(3) Suppose $\psi(+,-) = \psi(-,+) = 0$, so that $\psi(\sigma,\sigma')$ may be written as $\mathbf{1}\{\sigma = \sigma'\}e^{B_0}e^{B\sigma/d}e^{B\sigma'/d}$. Then

$$\phi_n = B_0\frac{\mathbb{E}_n[|E_n|]}{n} + B + \frac{1}{n}\mathbb{E}_n\bigg[\sum_{j=1}^{k(G_n)}\log(1+e^{-2B|C_j|})\bigg]$$

where the sum is taken over the connected components $C_1,\ldots,C_{k(G_n)}$ of G_n. We claim $\phi_n \to \phi = (d/2)B_0 + B$: we have $\liminf_{n\to\infty}(\phi_n - \phi) \geq 0$ (using uniform sparsity), and

$$\limsup_{n\to\infty}(\phi_n - \phi) \leq \limsup_{n\to\infty}\log 2\frac{\mathbb{E}_n[k(G_n)]}{n},$$

so it suffices to show $\mathbb{E}_n[k(G_n)]/n \to 0$. Indeed, if this fails then there exists $\epsilon > 0$ such that for infinitely many n, the event $\{k(G_n) \ge \epsilon n\}$ occurs with $\mathbb{P}_n$-probability at least ϵ. On this event, G_n has at least $\epsilon n/2$ components of size $\le 2/\epsilon$, so for $t > \log_k(2/\epsilon)$, $\limsup_{n\to\infty} \mathbb{P}_n(B_t(I_n) \not\cong \mathbb{T}^t) \ge \epsilon^2/2 > 0$, in contradiction of $G_n \to_{loc} \mathbb{T}$.

(4) Suppose instead $\psi(+,+) = \psi(-,-) = 0$ while $\psi(+,-) = \psi(-,+) > 0$. If the G_n are not exactly bipartite then $\phi_n = -\infty$. If they are exactly bipartite then
$$\phi_n = \log\psi(+,-)\frac{\mathbb{E}_n[|E_n|]}{n} + \log 2\frac{\mathbb{E}_n[k(G_n)]}{n},$$
and by the observation of (3) this converges to $\phi = (d/2)\log\psi(+,-)$.

B. Bipartite expander gadgets

In this section we construct the bipartite expander gadgets to be used in the reduction to MAX-CUT (Lem. III.1) and refine Thm. 5 to an approximate conditional independence statement for the gadgets (Propn. III.2). We conclude with the proof of our main results Thms. 1 and 2.

For any fixed positive integer k, G^k_{2n} will be a bipartite graph on $2n$ vertices with n even, defined as follows:

- Let H_n be a graph on n vertices of maximum degree d, generated by the configuration model as follows: take a uniformly random matching $\mathfrak{m}$ of $[dn]$, and put an edge (ij) in H_n for every edge $(i',j') \in \mathfrak{m}$ with $i' \in i+n\mathbb{Z}$, $j' \in j+n\mathbb{Z}$ (self-loops and multi-edges allowed).
- Take G_{2n} to be the bipartite double cover of H_n: the two parts of G_{2n} are $(i_+)_{i=1}^n$ and $(i_-)_{i=1}^n$, and we put two edges (i_+,j_-) and (j_+,i_-) in G_{2n} for every edge $(ij) \in H_n$ (multi-edges allowed).
- Choose k vertices $(i^\ell)_{\ell=1}^k$ uniformly at random from H_n, and for each ℓ choose $j^\ell \in \partial i^\ell$ uniformly at random. G^k_{2n} is the simple bipartite graph formed by deleting the edges $(i^\ell_\pm, j^\ell_\mp)$ from G_{2n} and merging any remaining multi-edges in the graph into single edges. Write $W^\pm \equiv \{i^\ell_\pm, j^\ell_\pm\}_{\ell=1}^k$ and $W \equiv W^+ \cup W^-$.

The graphs G_{2n} are d-regular with probability bounded away from zero as $n\to\infty$ (see e.g. [18, Ch. 9]). The following lemma gives their expansion property:

Lemma III.1. *Let k be fixed. For all $\delta > 0$ there exists $\lambda_\delta > 0$ such that the G^k_{2n} are $(\delta, 1/2, \lambda_\delta)$-edge expanders with high probability as $n \to \infty$.*

Proof: By stochastic domination we may assume $d = 3$. For $S \subset H_n$ with $|S| = m$, the probability that there are exactly j edges in H_n between S and its complement is
$$P_{j,m} = I_{j,m}\frac{\binom{3m}{j}\binom{3(n-m)}{j} j! M_{3(m-j)} M_{3(n-m-j)}}{M_{3n}},$$
where $I_{j,m}$ is the indicator that $m-j$ is even, and $M_\ell = (\ell-1)!! = \pi^{-1/2}\Gamma[(\ell+1)/2]2^{\ell/2}$ is the number of matchings on $[\ell]$ for ℓ even. By Stirling's approximation, if $\delta \le m/n \le 1-\delta$ and $j = \gamma n$, then
$$P_{j,m} = I_{j,m}\exp\Big\{-n\Big[\frac{3}{2}H(m/n) - \gamma\log\gamma + O_\delta(\gamma)\Big] + o_\delta(n)\Big\}$$
(where $H(p)$ denotes the binary entropy function $-p\log p - (1-p)\log(1-p)$). There are $\le e^{nH(m/n)}$ subsets of H_n of size m so there exists $\gamma_\delta > 0$ such that with probability at least $ne^{-nH(\delta)/4}$, all subsets of H_n of size between δn and $(1-\delta)n$ have expansion at least γ_δ.

We now show expansion for G^k_{2n}: since k does not change with n and the number of edges leaving any set of vertices decreases by at most a factor of 3 when multi-edges are merged into single edges, it suffices to show expansion for G_{2n}. Let $S_\pm$ be subsets of the $\pm$ sides of G_{2n} such that $S \equiv S_+ \cup S_-$ has size $\le n$. If the projection πS of S in H_n has size $\le (1-\delta)n$, then S has expansion at least $\gamma_\delta/2$. Suppose $|\pi S| \ge (1-\delta)n$: without loss $|S_+| \ge |S_-|$, so $|\pi S_+ \backslash \pi S_-| \ge (1/2-\delta)n$. If there are fewer than $\gamma|S|$ edges leaving S, then there must be at least $3(1/2-\delta)n - \gamma n$ edges between $\pi S_+\backslash\pi S_-$ and its complement in H_n. A similar analysis as above shows that for sufficiently small δ there exists $\gamma_\delta > 0$ such that the probability G_{2n} has such a set S is $\le e^{-n(\log 2)/4}$, and this concludes the proof. ∎

Recall that we use $W^\pm$ to denote the endpoints on the $\pm$ sides of the $2k$ edges deleted from G_{2n} in the formation of G^k_{2n}. Recall also the definitions of $\mu^\pm \in \mathscr{G}_{\mathbb{T}}$, and write $h^\pm \equiv h^{\mu^\pm}_{o\to j} \in \Delta$. For $h, h' \in \Delta$ define $h \otimes_\psi h' \in \Delta_{\mathscr{X}^2}$ by
$$(h\otimes_\psi h')(\sigma,\sigma') = \frac{h(\sigma)\psi(\sigma,\sigma')h(\sigma')}{z(h\otimes_\psi h')}, \qquad \text{(III.1)}$$
for $z(h\otimes_\psi h')$ the normalizing constant.

Proposition III.2. *The conditional measure $\nu^\pm_{G^k_{2n}}(\underline{\sigma}_W = \cdot)$ converges to the product measure*
$$Q^\pm_W(\underline{\sigma}) \equiv \prod_{w\in W^+} h^\pm(\sigma_w) \prod_{w\in W^-} h^\mp(\sigma_w).$$

Proof: Let B_t denote the union of the balls $B_t(w) \subseteq G_{2n}$ over $w \in \{i^\ell_\pm\}_{\ell=1}^k$; assume that B_t is a disjoint union of graphs isomorphic to $\mathbb{T}^t$ with internal boundary $S_t \equiv B_t\backslash B_{t-1}$, which is the case with high probability. For $\underline{\eta} \in \mathscr{X}^{S_t}$ let
$$\xi^\pm_{t,\ell,\underline{\eta}}(\cdot) \equiv \nu_{G^k_{2n}}(\sigma_{i^\ell_\pm} = \cdot \mid \underline{\sigma}_{S_t} = \underline{\eta}),$$
$$\zeta^\pm_{t,\ell,\underline{\eta}}(\cdot) \equiv \nu_{G^k_{2n}}(\sigma_{j^\ell_\pm} = \cdot \mid \underline{\sigma}_{S_t} = \underline{\eta}),$$
so that
$$\nu_{G_{2n}}[(\sigma_{i^\ell_+}, \sigma_{j^\ell_-}) = \cdot \mid \underline{\sigma}_{S_t} = \underline{\eta}] = \xi^+_{t,\ell,\underline{\eta}} \otimes_\psi \zeta^-_{t,\ell,\underline{\eta}}.$$
By Thm. 5, the conditional measures $\nu^+_{G_{2n}}(\underline{\sigma}_{B_t(i^\ell_+)} = \cdot)$ converge to μ^+. But by (I.6), $Y(\underline{\sigma})$ agrees with $Y_t(\underline{\sigma}) \equiv$ sgn $\sum_{i\in V\backslash B_t}\tau_i\sigma_i$ with high probability, so that convergence

also holds if we replace $\nu^{+}_{G_{2n}}$ *by* $\nu^{\pm t}_{G_{2n}}(\cdot) \equiv \nu_{G_{2n}}(\cdot \mid Y_t(\underline{\sigma}) = \pm)$. *In particular,*

$$\mathbb{E}_{2n}\Big[\Big\|\sum_{\underline{\eta}} \nu^{+t}_{G_{2n}}(\underline{\sigma}_{S_t} = \underline{\eta})\xi^{+}_{t,\ell,\underline{\eta}} \otimes_\psi \zeta^{-}_{t,\ell,\underline{\eta}} - h^{+} \otimes_\psi h^{-}\Big\|_{\mathrm{TV}}\Big]$$
$$= \mathbb{E}_{2n}\Big[\Big\|\langle \xi^{+}_{t,\ell,\underline{\sigma}_{S_t}} \otimes_\psi \zeta^{-}_{t,\ell,\underline{\sigma}_{S_t}}\rangle_{\nu^{+t}_{G_{2n}}} - h^{+} \otimes_\psi h^{-}\Big\|_{\mathrm{TV}}\Big]$$

tends to zero in the limit $n \to \infty$ *followed by* $t \to \infty$. *On the other hand, it is easily seen that* $(h \otimes_\psi h')(1,0)$ *is maximized by taking* $h(1)$ *and* $h'(0)$ *as large as possible. But in the limit* $t \to \infty$ *the values* $\xi^{\pm}_{t,\ell,\underline{\eta}}(1), \zeta^{\pm}_{t,\ell,\underline{\eta}}(1)$ *(with* $\underline{\eta}$ *arbitrary) are sandwiched between* $h^{\pm}(1)$, *so it must be that*

$$\mathbb{E}_{2n}\Big[\sum_{\sigma\in\{\pm\}} \Big\langle \|\xi^{\sigma}_{t,\ell,\underline{\sigma}_{S_t}} - h^{\sigma}\|_{\mathrm{TV}}\Big\rangle_{\nu^{+t}_{G_{2n}}}\Big] \to 0 \qquad \text{(III.2)}$$

in the limit $n \to \infty$ *followed by* $t \to \infty$.

We now claim that (III.2) *continues to hold after removal of the edges* $(i^{\ell}_{\pm}, j^{\ell}_{\mp})$. *Indeed,*

$$\frac{\nu^{+t}_{G_{2n}}(\underline{\sigma}_{S_t} = \underline{\eta})}{\nu^{+t}_{G^k_{2n}}(\underline{\sigma}_{S_t} = \underline{\eta})} = \frac{\left[\dfrac{Z^{+t}_{\mathrm{out}}(\underline{\eta}) Z_{\mathrm{in}}(\underline{\eta})}{\sum_{\underline{\eta}'} Z^{+t}_{\mathrm{out}}(\underline{\eta}') Z_{\mathrm{in}}(\underline{\eta}')}\right]}{\left[\dfrac{Z^{+t}_{\mathrm{out}}(\underline{\eta}) Z^{k}_{\mathrm{in}}(\underline{\eta})}{\sum_{\underline{\eta}'} Z^{+t}_{\mathrm{out}}(\underline{\eta}') Z^{k}_{\mathrm{in}}(\underline{\eta}')}\right]} \qquad \text{(III.3)}$$

where

$$Z^{\pm t}_{\mathrm{out}}(\underline{\eta}) \equiv Z_{G_{2n}\setminus B_{t-1}}[\{\underline{\sigma}_{G_{2n}\setminus B_{t-1}} : Y_t(\underline{\sigma}) = \pm,\, \underline{\sigma}_{S_t} = \underline{\eta}\}],$$
$$Z_{\mathrm{in}}(\underline{\eta}) \equiv Z_{B_t}[\{\underline{\sigma}_{B_t} : \underline{\sigma}_{S_t} = \underline{\eta}\}],$$
$$Z^{k}_{\mathrm{in}}(\underline{\eta}) \equiv Z_{B_t \cap G^k_{2n}}[\{\underline{\sigma}_{B_t} : \underline{\sigma}_{S_t} = \underline{\eta}\}].$$

Now note that for k *bounded and* t *large we have* $Z_{\mathrm{in}}(\underline{\eta}) \asymp Z^{k}_{\mathrm{in}}(\underline{\eta})$ *uniformly over* $\underline{\eta}$: *for Ising interactions at non-zero temperature this is obvious, while for the hard-core model*

$$\frac{Z_{\mathrm{in}}(\underline{\eta})}{Z^{k}_{\mathrm{in}}(\underline{\eta})} = \prod_{\ell=1}^{k}\Big\{[1 - \xi^{+}_{t,\ell,\underline{\eta}}(1)\zeta^{-}_{t,\ell,\underline{\eta}}(1)][1 - \xi^{-}_{t,\ell,\underline{\eta}}(1)\zeta^{+}_{t,\ell,\underline{\eta}}(1)]\Big\}$$

which for t *large is* $\asymp 1$ *uniformly over* $\underline{\eta}$. *Since the* $\xi^{\pm}_{t,\ell,\underline{\eta}}$ *and* $\zeta^{\pm}_{t,\ell,\underline{\eta}}$ *are* $\underline{\eta}$*-measurable, it follows from* (III.3) *that* (III.2) *continues to hold with* $\nu^{+t}_{G^k_{2n}}$ *in place of* $\nu^{+t}_{G_{2n}}$. *Since the spins* $(\sigma_w)_{w\in W}$ *are independent under* $\nu^{\pm t}_{G^k_{2n}}(\cdot \mid \underline{\sigma}_{S_t})$, *this further implies*

$$0 = \lim_{t\to\infty}\lim_{n\to\infty} \mathbb{E}_{2n}\Big[\|\nu^{+t}_{G^k_{2n}}(\underline{\sigma}_W = \cdot) - Q^{+}_W\|_{\mathrm{TV}}\Big]. \qquad \text{(III.4)}$$

Finally, by a similar argument as before $\lim_{n\to\infty} \nu_{G^k_{2n}}(Y(\underline{\sigma}) = Y_t(\underline{\sigma})) = 1$, *so* (III.4) *holds with* $\nu^{+}_{G^k_{2n}}$ *in place of* $\nu^{+t}_{G^k_{2n}}$ *which gives the result.* ■

C. *Randomized reduction to* MAX-CUT

We now demonstrate how to use Propn. III.2 to establish a randomized reduction from approximating the partition function to the problem of approximate MAX-CUT on 3-regular graphs, which is NP-hard [12]. We begin with the following easy observation:

Lemma III.3. *For anti-ferromagnetic two-spin models on* $G_n \to_{loc} \mathbf{T}$,

$$\lim_{n\to\infty} \mathbb{E}_n\Big[\nu_n\Big(\sum_{i\in V_n} \tau_i\sigma_i = 0\Big)\Big] = 0.$$

Proof: For the Ising model see [13, Lem. 4.1]. For the hard-core model, let A_n *denote the set of vertices* $i \in V_n$ *with* $B_2(i)$ *isomorphic to* $\mathbb{T}^2_+$, *the depth-two subtree of* $\mathbb{T}_+$; *then* A_n *is necessarily an independent set of black vertices. The probability that* $\sum_{i\in A_n} \tau_i\bar{\sigma}_i = \sum_{i\in A_n} \bar{\sigma}_i$ *takes value* j, *conditioned on all the spins* $(\bar{\sigma}_i)_{i\notin A_n}$, *is* $\mathbb{P}(X = j)$ *where* X *is a binomial random variable on* $N = |\{i \in A_n : \bar{\sigma}_{\partial i} \equiv 0\}|$ *number of trials with success probability* $\lambda/(1+\lambda)$. *If* $N \ge \epsilon n$ *then* $\mathbb{P}(X = j) = O(1/\sqrt{\epsilon n})$ *uniformly in* j *(e.g. by the Berry-Esséen theorem). If* $N < \epsilon n$ *then* $\sum_{i\in\partial A_n} \bar{\sigma}_i \ge (|A_n| - \epsilon n)/d$, *so*

$$\frac{1}{n}\sum_{i\in V_n} \tau_i\sigma_i = \frac{2}{n}\sum_{i\in V_n} \tau_i\bar{\sigma}_i - \frac{1}{n}\sum_{i\in V_n} \tau_i$$
$$< \epsilon - \frac{|A_n|/n - \epsilon}{d} + \frac{|V_n\setminus(A_n \cup \partial A_n)|}{n} - \frac{1}{n}\sum_{i\in V_n} \tau_i.$$

As $n \to \infty$ *the right-hand side tends in probability to* $[-1/2+\epsilon(d+1)]/d$, *which is negative for small* ϵ. *Combining the above observations concludes the proof for the hard-core model.* ■

Let H be a 3-regular graph on m vertices and construct the bipartite graph $G = G^{3k}_{2n}$ by the procedure described above. By Lem. III.3 and Propn. III.2, for any $\epsilon > 0$ there exists $n(\epsilon)$ large enough such that the following hold with positive probability:

(I) G^{3k}_{2n} was formed by removing $3k$ distinct edges from a d-regular graph G_{2n};
(II) $\nu_{G^{3k}_{2n}}(Y(\underline{\sigma}) = +) \le (1+\epsilon)/2$; and
(III) $\nu^{\pm}_{G^{3k}_{2n}}(\underline{\sigma}_W)/Q^{\pm}_W(\underline{\sigma}_W) \in [1-\epsilon, 1+\epsilon]$ for all $\underline{\sigma}_W$.

Consequently, for given ϵ we may find G^{3k}_{2n} satisfying properties (I)-(III) within finite time by deterministic search. We then construct from H and G a new graph H^G as follows:

- For each vertex $x \in H$ let G_x be a copy of G, and denote by $W^{\pm}_x$ the vertices of G_x corresponding to $W^{\pm}$ in G. Let $\widehat{H}^G$ be the disjoint union of the G_x, $x \in H$.
- For every edge $(x,y) \in H$, add $2k$ edges between W^{+}_x and W^{+}_y and similarly $2k$ edges between W^{-}_x and W^{-}_y. This can be done deterministically in such a way that the resulting graph, which we denote H^G, is d-regular.

We write a spin configuration on $\widehat{H}^G$ or H^G as $\underline{\sigma} \equiv (\underline{\sigma}_x)_{x\in H}$ where $\underline{\sigma}_x$ is the restriction of $\underline{\sigma}$ to G_x. We write $Y_x \equiv Y(\underline{\sigma}_x)$ for the phase of each $\underline{\sigma}_x$, and $\mathcal{Y}(\underline{\sigma}) \equiv (Y(\underline{\sigma}_x))_{x\in H} \in \{0,1\}^H$. Write $Z_{H^G}(\mathcal{Y})$ for the partition function for the two-spin model on H^G restricted to configurations of phase $\mathcal{Y}$, and define likewise $Z_{\widehat{H}^G}(\mathcal{Y})$.

Recalling (III.1), let

$$\Gamma \equiv z(h^+ \otimes_\psi h^+) z(h^- \otimes_\psi h^-), \quad \Theta \equiv z(h^+ \otimes_\psi h^-)^2,$$

and note that for anti-ferromagnetic two-spin models in non-uniqueness regimes, $\Theta > \Gamma$.

Lemma III.4. *For G satisfying properties* (I)-(III)*,*

$$[(1-\epsilon)/2]^m \le \frac{Z_{H^G}/Z_{\widehat{H}^G}}{\Gamma^{2k|E(H)|}(\Theta/\Gamma)^{2k\,\text{MAX-CUT}(H)}} \le (1+\epsilon)^m.$$

Proof: By (II)*,*

$$(1-\epsilon)^m \le 2^m \frac{Z_{\widehat{H}^G(\mathcal{Y})}}{Z_{\widehat{H}^G}} \le (1+\epsilon)^m \tag{III.5}$$

for all $\mathcal{Y} \in \{0,1\}^H$. By (III)*, the ratio*

$$\frac{Z_{H^G}(\mathcal{Y})}{Z_{\widehat{H}^G}(\mathcal{Y})} = \sum_{x\in H} \sum_{\underline{\sigma}_{W_x}} \nu_{G_x}^{Y_x}(\underline{\sigma}_{W_x}) \prod_{(ij)\in E(H^G)\setminus E(\widehat{H}^G)} \psi(\sigma_i, \sigma_j)$$

is within a $(1\pm\epsilon)^m$ factor of

$$\sum_{x\in H} \sum_{\underline{\sigma}_{W_x}} Q^{Y_x}(\underline{\sigma}_{W_x^+}) \prod_{(ij)\in E(H^G)\setminus E(\widehat{H}^G)} \psi(\sigma_i, \sigma_j),$$

which by direct calculation equals

$$\Gamma^{2k|E(H)|}(\Theta/\Gamma)^{2k\,\text{cut}(\mathcal{Y})}$$

where $\text{cut}(\mathcal{Y}) \equiv |\{(x,y) \in E(H) : Y_x \neq Y_y\}|$*, the number of edges crossing the cut of H induced by $\mathcal{Y}$. Combining with* (III.5) *gives*

$$\begin{aligned} Z_{H^G} &= \sum_{\mathcal{Y}} \frac{Z_{H^G}(\mathcal{Y})}{Z_{\widehat{H}^G}(\mathcal{Y})} Z_{\widehat{H}^G}(\mathcal{Y}) \\ &\le (1+\epsilon)^{2m} \Gamma^{2k|E(H)|}(\Theta/\Gamma)^{2k\,\text{MAX-CUT}(H)} Z_{\widehat{H}^G} \end{aligned}$$

and similarly

$$Z_{H^G} \ge 2^{-m}(1-\epsilon)^{2m}\Gamma^{2k|E(H)|}(\Theta/\Gamma)^{2k\,\text{MAX-CUT}(H)} Z_{\widehat{H}^G}.$$

Rearranging gives the stated result. ■

Using this lemma we now complete the reduction to approximate MAX-CUT:

Proof of Thms. 1 and 2: Let H be a 3-regular graph on m vertices, and note that the maximum cut of H is at least $3m/4$, the expected value of a random cut. Construct $\widehat{H}^G$, H^G as above. Since $\widehat{H}^G$ is a disjoint collection of constant-size graphs, its partition function can be computed in polynomial time. Suppose Z_{H^G} could be approximated within a factor of $e^{c|H^G|}$ in polynomial time for any $c > 0$: rearranging the result of Lem. III.4 gives

$$\begin{aligned} &\frac{\log\left(\dfrac{Z_{H^G}/Z_{\widehat{H}^G}}{\Gamma^{2k|E(H)|}(1+\epsilon)^m}\right)}{2k\log(\Theta/\Gamma)} \le \text{MAX-CUT}(H) \\ &\qquad \le \frac{\log\left(\dfrac{Z_{H^G}/Z_{\widehat{H}^G}}{\Gamma^{2k|E(H)|}[(1-\epsilon)/2]^m}\right)}{2k\log(\Theta/\Gamma)}, \end{aligned} \tag{III.6}$$

so within polynomial time one obtains upper and lower bounds for MAX-CUT(H) which differ by $O[(c|G|+1)m/k]$. Taking k large and c small then allows to compute MAX-CUT(H) up to an arbitrarily small multiplicative error: that is, we have completed the reduction to a PRAS for MAX-CUT on 3-regular graphs, in contradiction of the result of [12]. ■

ACKNOWLEDGMENT

We thank Andreas Galanis, Daniel Štefankovič, and Eric Vigoda for describing to us their methods and for sending us a draft of their paper. We thank Amir Dembo, David Gamarnik, Andrea Montanari, Alistair Sinclair, Piyush Srivastava, and David Wilson for helpful conversations. The research of A.S. is partially supported by an Alfred P. Sloan Research Fellowship. The research of N.S. is partially supported by a Department of Defense NDSEG Fellowship.

REFERENCES

[1] M. Jerrum and A. Sinclair, "Polynomial-time approximation algorithms for the Ising model," *SIAM J. Comput.*, vol. 22, no. 5, pp. 1087–1116, 1993. [Online]. Available: http://dx.doi.org/10.1137/0222066

[2] D. Weitz, "Counting independent sets up to the tree threshold," in *STOC'06: Proceedings of the 38th Annual ACM Symposium on Theory of Computing*. New York: ACM, 2006, pp. 140–149. [Online]. Available: http://dx.doi.org/10.1145/1132516.1132538

[3] M. Luby and E. Vigoda, "Fast convergence of the Glauber dynamics for sampling independent sets," *Random Structures Algorithms*, vol. 15, no. 3-4, pp. 229–241, 1999, statistical physics methods in discrete probability, combinatorics, and theoretical computer science (Princeton, NJ, 1997). [Online]. Available: http://dx.doi.org/10.1002/(SICI)1098-2418(199910/12)15:3/4¡229::AID-RSA3¿3.0.CO;2-X

[4] M. Dyer, A. Frieze, and M. Jerrum, "On counting independent sets in sparse graphs," in *40th Annual Symposium on Foundations of Computer Science (New York, 1999)*. IEEE Computer Soc., Los Alamitos, CA, 1999, pp. 210–217. [Online]. Available: http://dx.doi.org/10.1109/SFFCS.1999.814593

[5] ——, "On counting independent sets in sparse graphs," *SIAM J. Comput.*, vol. 31, no. 5, pp. 1527–1541 (electronic), 2002. [Online]. Available: http://dx.doi.org/10.1137/S0097539701383844

[6] E. Mossel, D. Weitz, and N. Wormald, "On the hardness of sampling independent sets beyond the tree threshold," *Probab. Theory Related Fields*, vol. 143, no. 3-4, pp. 401–439, 2009. [Online]. Available: http://dx.doi.org/10.1007/s00440-007-0131-9

[7] A. Sly, "Computational transition at the uniqueness threshold," *Foundations of Computer Science, Annual IEEE Symposium on*, vol. 0, pp. 287–296, 2010.

[8] A. Galanis, Q. Ge, D. Štefankovič, E. Vigoda, and L. Yang, "Improved inapproximability results for counting independent sets in the hard-core model," in *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, ser. Lecture Notes in Computer Science, L. Goldberg, K. Jansen, R. Ravi, and J. Rolim, Eds. Springer Berlin / Heidelberg, 2011, vol. 6845, pp. 567–578, 10.1007/978-3-642-22935-0_48. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-22935-0_48

[9] A. Sinclair, P. Srivastava, and M. Thurley, "Approximation algorithms for two-state anti-ferromagnetic spin systems on bounded degree graphs," Preprint, arXiv:1107.2368v3, 2011.

[10] L. Li, P. Lu, and Y. Yin, "Correlation decay up to uniqueness in spin systems," Preprint, arXiv:1111.7064v1, 2011.

[11] A. Galanis, D. Štefankovič, and E. Vigoda, "Inapproximability of the partition function for the antiferromagnetic Ising and hard-core models," Preprint, arXiv:1203.2226v2, 2012.

[12] P. Alimonti and V. Kann, "Hardness of approximating problems on cubic graphs," in *Algorithms and complexity (Rome, 1997)*, ser. Lecture Notes in Comput. Sci. Berlin: Springer, 1997, vol. 1203, pp. 288–298.

[13] A. Montanari, E. Mossel, and A. Sly, "The weak limit of Ising models on locally tree-like graphs," *Probability Theory and Related Fields*, vol. 152, pp. 31–51, 2012, 10.1007/s00440-010-0315-6. [Online]. Available: http://dx.doi.org/10.1007/s00440-010-0315-6

[14] A. Dembo and A. Montanari, "Ising models on locally tree-like graphs," *Ann. Appl. Probab.*, vol. 20, no. 2, pp. 565–592, 2010. [Online]. Available: http://dx.doi.org/10.1214/09-AAP627

[15] S. Dommers, C. Giardinà, and R. van der Hofstad, "Ising models on power-law random graphs," *J. Stat. Phys.*, vol. 141, no. 4, pp. 638–660, 2010. [Online]. Available: http://dx.doi.org/10.1007/s10955-010-0067-9

[16] A. Dembo, A. Montanari, and N. Sun, "Factor models on locally tree-like graphs," Preprint, arXiv:1110.4821v2, 2011.

[17] A. Dembo and A. Montanari, "Gibbs measures and phase transitions on sparse random graphs," *Braz. J. Probab. Stat.*, vol. 24, no. 2, pp. 137–211, 2010. [Online]. Available: http://dx.doi.org/10.1214/09-BJPS027

[18] S. Janson, T. Łuczak, and A. Rucinski, *Random graphs*, ser. Wiley-Interscience Series in Discrete Mathematics and Optimization. Wiley-Interscience, New York, 2000.