

Inapproximability for Antiferromagnetic Spin Systems in the Tree Non-Uniqueness Region^{*}

Andreas Galanis
Georgia Tech
agalanis@cc.gatech.edu.

Daniel Štefankovič
University of Rochester
stefanko@cs.rochester.edu.

Eric Vigoda
Georgia Tech
vigoda@cc.gatech.edu.

ABSTRACT

A remarkable connection has been established for antiferromagnetic 2-spin systems, including the Ising and hard-core models, showing that the computational complexity of approximating the partition function for graphs with maximum degree Δ undergoes a phase transition that coincides with the statistical physics uniqueness/non-uniqueness phase transition on the infinite Δ -regular tree. Despite this clear picture for 2-spin systems, there is little known for multi-spin systems. We present the first analog of the above inapproximability results for multi-spin systems.

The main difficulty in previous inapproximability results was analyzing the behavior of the model on random Δ -regular bipartite graphs, which served as the gadget in the reduction. To this end one needs to understand the moments of the partition function. Our key contribution is connecting: (i) induced matrix norms, (ii) maxima of the expectation of the partition function, and (iii) attractive fixed points of the associated tree recursions (belief propagation). The view through matrix norms allows a simple and generic analysis of the second moment for any spin system on random Δ -regular bipartite graphs. This yields concentration results for any spin system in which one can analyze the maxima of the first moment. The connection to fixed points of the tree recursions enables an analysis of the maxima of the first moment for specific models of interest.

For k -colorings we prove that for even k , in the tree non-uniqueness region (specifically for semi-translation invariant measures which corresponds to $k < \Delta$) it is NP-hard, unless NP=RP, to approximate the number of colorings for triangle-free Δ -regular graphs. Our proof extends to the antiferromagnetic Potts model, and, in fact, to every antiferromagnetic model under a mild condition.

1. INTRODUCTION

1.1 Background

Spin systems are a general framework from statistical physics that captures classical physics models, including the Ising and Potts models, and models of particular combinatorial interest, including k -colorings and the hard-core lattice gas model defined on independent sets. We define these combinatorial models more precisely before presenting the context of our results.

The hard-core lattice gas model is an example of a 2-spin system. For a graph $G = (V, E)$, configurations of the model are the set Ω of independent sets of G . The model is parameterized by an activity $\lambda > 0$, and a configuration $\sigma \in \Omega$ is assigned weight $w(\sigma) = \lambda^{|\sigma|}$. The Gibbs distribution is $\mu(\sigma) = w(\sigma)/Z$ where the normalizing factor is known as the partition function and is defined as $Z = \sum_{\tau \in \Omega} w(\tau)$. In the hard-core model the spins correspond to occupied/unoccupied. Multi-spin systems are models with more than 2 spins, an example being the k -colorings problem. In the colorings problem, for a graph $G = (V, E)$, configurations are the set Ω of assignments of a set of k colors to vertices so that neighboring vertices receive different colors. The Gibbs distribution is the uniform distribution over Ω , and in this case the partition function $Z = |\Omega|$ is the number of k -colorings in G .

The hard-core model and colorings are examples of antiferromagnetic systems – neighboring vertices “prefer” to have different spins. In contrast, in ferromagnetic systems neighboring spins tend to align. We defer the formal definitions of ferromagnetic and antiferromagnetic systems to Section 6 where we discuss how our results extend to general spin systems.

The focus of this paper is the computational complexity of computing the partition function. Exact computation of the partition function is typically #P-complete, even for very restricted classes of graphs [Gre00]. Hence our focus is on the existence of a fully-polynomial approximation scheme

– either a deterministic FPTAS or randomized FPRAS – for estimating the partition function. For any spin system, (approximate) sampling from the Gibbs distribution implies an FPRAS for estimating the partition function, and hence our hardness results also apply to the associated sampling problem.

The computational complexity of approximating the partition function on Δ -regular graphs is now well-understood for 2-spin systems, such as the Ising and hard-core models. For ferromagnetic 2-spin systems, there is an FPRAS for estimating the partition function [GJP03, JS93] for all parameters (for graphs of maximum degree Δ , the picture is more complicated, see [LLZ14]). The picture is more intricate (and fascinating) for antiferromagnetic 2-spin systems. We will detail the picture after introducing the statistical physics notion of a phase transition.

Let $\mathbb{T}_{\Delta,\ell}$ denote the complete Δ -regular tree of depth ℓ with root r . The question of interest is whether or not we can fix a configuration on the leaves of $\mathbb{T}_{\Delta,\ell}$ so that the root is influenced by this boundary configuration in the limit $\ell \rightarrow \infty$. For the example of colorings, fix a coloring σ_ℓ of the leaves (such that there is at least one coloring of the rest of the tree that is consistent with σ_ℓ). Look at a random coloring of the tree $\mathbb{T}_{\Delta,\ell}$ conditioned on the leaves having coloring σ_ℓ . For all sequences (σ_ℓ) of fixed leaf colorings, if in the limit $\ell \rightarrow \infty$, the marginal at the root is uniform over the k colors, then we say uniqueness holds, and otherwise we say non-uniqueness holds. (The terminology comes from statistical physics where the focus is on the set of infinite-volume Gibbs measures.)

For the hard-core model the critical activity is $\lambda_c(\Delta) = (\Delta - 1)^{\Delta-1}/(\Delta - 2)^\Delta$ [Kel91]. Weitz [Wei06] presented an FPTAS for estimating the partition function in the tree uniqueness region (i.e., when $\lambda < \lambda_c(\Delta)$). On the other side, Sly [Sly10] proved that, unless NP=RP, it is NP-hard to obtain an FPRAS for Δ -regular graphs in the tree non-uniqueness region (i.e., when $\lambda > \lambda_c(\Delta)$). These results were extended to all 2-spin antiferromagnetic models, for the positive side see [SST12, LLY13] and for the negative side see [SS12, GSV12]. For 2-spin antiferromagnetic models, this establishes a beautiful picture connecting the computational complexity of approximating the partition function to statistical physics phase transitions in the infinite tree.

1.2 Main Results

The picture for multi-spin systems (systems with $q > 2$ possible spins for vertices) is much less clear; the above approaches for 2-spin systems do not extend to multi-spin models in a straightforward manner. We aim to establish the analog of the above inapproximability results for the colorings problem, namely, NP-hardness in the tree non-uniqueness region. Our techniques and results generalize to a broad class of antiferromagnetic spin systems.

For the colorings problem, even understanding the uniqueness threshold is challenging. Jonasson [Jon02] established uniqueness when $k \geq \Delta + 1$, and it is easy to show non-uniqueness when $k \leq \Delta$ since a fixed coloring on the leaves can “freeze” the internal coloring. For 2-spin systems uniqueness can be characterized by the existence of multiple solutions of a certain system of equations (8), called tree recursions, see Section 6 for additional explanation. In statistical physics terminology the solutions to these equations correspond to semi-translation invariant measures on the infinite

tree \mathbb{T}_Δ . For colorings the uniqueness threshold and the semi-translation invariant uniqueness threshold no longer coincide. In particular, Brightwell and Winkler [BW02] established, for semi-translation invariant measures, uniqueness when $k \geq \Delta$ and non-uniqueness when $k < \Delta$.

We prove, for even k , that it is NP-hard to approximate the number of colorings (in other words, NP-hard to approximate the partition function) when there is non-uniqueness of semi-translation invariant Gibbs measures on \mathbb{T}_Δ , i.e., when $k < \Delta$. Moreover, our result proves hardness for the class of triangle-free Δ -regular graphs. Hence, our result is particularly interesting in the region $k = \Omega(\Delta/\log \Delta)$ since a seminal result of Johansson [Joh96, MR02] shows that all triangle-free graphs are colorable with $O(\Delta/\log \Delta)$ colors. His proof, which uses the nibble method and the Lovász Local Lemma, can be made algorithmic using the constructive proof of [MT10]. For general graphs with maximum degree Δ , the interesting region is $k = \Delta - O(\sqrt{\Delta})$, since Molloy and Reed [MR01] showed, for sufficiently large constant Δ , a polynomial-time algorithm to determine if a graph with maximum degree Δ is k -colorable when $k \geq \Delta - \sqrt{\Delta} + 3$. We note that most parts of the proof extend to the odd k case as well, modulo a technical condition described in the end of Section 4.

Here is the formal statement of our inapproximability result for colorings.

THEOREM 1. *For all even $k \geq 3$, all $\Delta \geq 3$, for the k -colorings problem, when $k < \Delta$, unless NP=RP, there is no FPRAS that approximates the partition function for triangle-free Δ -regular graphs. Moreover, there exists $\varepsilon = \varepsilon(k, \Delta)$ such that, unless NP=RP, one cannot approximate the partition function within a factor $2^{\varepsilon n}$ for triangle-free Δ -regular graphs (where n is the number of vertices).*

Our result also extends to the anti-ferromagnetic Potts model. In the q -state Potts model there is a parameter $B > 0$ which corresponds to the “temperature” and controls the strength of the interactions along an edge. For a graph $G = (V, E)$, the set Ω of configurations are assignments $\sigma : V \rightarrow [q]$. Each configuration has a weight $w(\sigma) = B^{m(\sigma)}$ where $m(\sigma)$ is the number of monochromatic edges in σ . The Gibbs distribution is $\mu(\sigma) = w(\sigma)/Z$ where $Z = \sum_{\tau \in \Omega} w(\tau)$ is the partition function. The case $B > 1$ is the ferromagnetic Potts model, and $B < 1$ is the antiferromagnetic Potts model. Colorings corresponds to the $B = 0$ case, and the Ising model is the $q = 2$ case.

The uniqueness/non-uniqueness threshold for the infinite tree \mathbb{T}_Δ is not known for the antiferromagnetic Potts model. We prove that the uniqueness/non-uniqueness threshold for semi-translation invariant Gibbs measures on \mathbb{T}_Δ occurs at $B_c(\Delta) = \frac{\Delta-q}{\Delta}$. We believe this threshold coincides with the uniqueness/non-uniqueness threshold, unlike in the case of colorings. We prove, for even q , that approximating the partition function is NP-hard in the non-uniqueness region for semi-translation invariant measures.

THEOREM 2. *For all even $q \geq 3$, all $\Delta \geq 3$, for the anti-ferromagnetic q -state Potts model, for all $B < \frac{\Delta-q}{\Delta}$, unless NP=RP, there is no FPRAS that approximates the partition function for triangle-free Δ -regular graphs. Moreover, there exists $\varepsilon = \varepsilon(q, \Delta)$ such that, unless NP=RP, one cannot approximate the partition function within a factor $2^{\varepsilon n}$ for triangle-free Δ -regular graphs (where n is the number of vertices).*

In fact, we obtain inapproximability of the partition function for any antiferromagnetic model when there is non-uniqueness of semi-translation invariant measures on \mathbb{T}_Δ and mild additional conditions. Our results for general models are stated in Section 6.

1.3 Proof Approach

The key gadget in the inapproximability results for 2-spin models is a random Δ -regular bipartite graph. The rough idea for the hard-core model is that in the tree non-uniqueness region, on a random Δ -regular bipartite graph, an independent set from the Gibbs distribution is “unbalanced” with high probability (the fraction of occupied vertices in the two parts of the bipartition differ by a constant). To analyze random regular bipartite graphs, the original inapproximability result of Sly [Sly10] relied on a second moment analysis of Mossel et al. [MWW09], which Sly called a technical tour-de-force. The optimization at the heart of that analysis was difficult enough that his result only held for λ close to the uniqueness threshold.

We present a new approach for the associated optimization problem which is at the heart of the second moment analysis. Our approach yields a simple, short analysis that holds for *any model* on random Δ -regular bipartite graphs. The key idea is to define a new function Φ , which is represented as an induced matrix norm, and has the same critical points as the first moment. We can then use the fact that induced matrix norms are multiplicative over tensor product to analyze the second moment.

We present the main components of our general method to analyze the second moment in the next section. We then present the main ideas in the reduction in Section 5 and our results for general antiferromagnetic models in Section 6.

2. SECOND MOMENT ANALYSIS

A general q -spin system is specified by a symmetric $q \times q$ interaction matrix $\mathbf{B} = (B_{ij})_{i,j \in [q]}$ with non-negative entries, which specify the strength of the interaction between the spins. For example, the interaction matrix for the Potts model has off-diagonal entries equal to 1 and its diagonal entries equal to B . For a finite undirected graph $G = (V, E)$, a q -spin system is a probability distribution μ_G over the space Ω_G of all *configurations*, i.e., spin assignments $\sigma : V \rightarrow [q]$. The weight of a configuration $\sigma \in \Omega_G$ is the product of neighboring spin interactions, that is,

$$w_G(\sigma) = \prod_{(u,v) \in E} B_{\sigma(u)\sigma(v)}.$$

The Gibbs distribution μ_G is defined as $\mu_G(\sigma) = w_G(\sigma)/Z_G$ where the partition function Z_G is $Z_G = \sum_{\sigma \in \Omega_G} w_G(\sigma)$. We drop the subscript G when the graph under consideration is clear.

2.1 Basic Definitions: Matrix Norms

We will reformulate the maxima of the first and second moments in terms of matrix norms. We first recall the basic definitions regarding matrix norms. The usual vector norms are denoted as:

$$\|\mathbf{x}\|_p = \left(\sum_{i=1}^n x_i^p \right)^{1/p}.$$

We will use the subordinate matrix norm (also known as the induced matrix norm) which will be denoted as $\|\cdot\|_{p \rightarrow q}$ and

is defined as:

$$\|\mathbf{A}\|_{p \rightarrow q} = \max_{\|\mathbf{x}\|_p=1} \|\mathbf{A}\mathbf{x}\|_q.$$

Note that if \mathbf{A} has non-negative entries then one can restrict the maximization to \mathbf{x} with non-negative entries. A well-known example of an induced norm is the spectral norm $\|\cdot\|_{2 \rightarrow 2}$.

2.2 First and Second Moments

For antiferromagnetic models on a random Δ -regular bipartite graph $G = (V, E)$ with bipartition $V = V_1 \cup V_2$, the goal is to understand the Gibbs distribution μ_G by looking at the distribution of spin values in V_1 and V_2 . Let $n = |V_1| = |V_2|$. For a configuration $\sigma : V \rightarrow [q]$, we shall denote the set of vertices assigned spin i by $\sigma^{-1}(i)$. Denote by Δ_q the simplex

$$\Delta_q = \{(x_1, x_2, \dots, x_q) \in \mathbb{R}^q \mid \sum_{i=1}^q x_i = 1 \text{ and } x_i \geq 0 \text{ for } i = 1, \dots, q\}.$$

For $\alpha, \beta \in \Delta_q$, let

$$\Sigma^{\alpha, \beta} = \{\sigma : V \rightarrow \{1, \dots, q\} \mid |\sigma^{-1}(i) \cap V_1| = \alpha_i n, \\ |\sigma^{-1}(i) \cap V_2| = \beta_i n \text{ for } i = 1, \dots, q\},$$

that is, configurations in $\Sigma^{\alpha, \beta}$ assign $\alpha_i n$ and $\beta_i n$ vertices in V_1 and V_2 the spin value i , respectively. We will be interested in the total weight $Z_G^{\alpha, \beta}$ of configurations in $\Sigma^{\alpha, \beta}$, namely

$$Z_G^{\alpha, \beta} = \sum_{\sigma \in \Sigma^{\alpha, \beta}} w(\sigma).$$

We study $Z_G^{\alpha, \beta}$ by looking at the moments $\mathbf{E}_G[Z_G^{\alpha, \beta}]$ and $\mathbf{E}_G[(Z_G^{\alpha, \beta})^2]$, where the expectation is over the distribution of the random Δ -regular bipartite graph, from hereon denoted by \mathcal{G} .

For $\alpha, \beta \in \Delta_q$, denote the leading term of the first and second moments as:

$$\Psi_1(\alpha, \beta) = \Psi_1^{\mathbf{B}}(\alpha, \beta) := \lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbf{E}_G[Z_G^{\alpha, \beta}], \\ \Psi_2(\alpha, \beta) = \Psi_2^{\mathbf{B}}(\alpha, \beta) := \lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbf{E}_G \left[\left(Z_G^{\alpha, \beta} \right)^2 \right].$$

We will refer to the pairs α, β that maximize the first moment as *dominant phases*. In the uniqueness region there is a unique dominant phase and it has $\alpha = \beta$. In contrast, for 2-spin antiferromagnetic models and for colorings in the semi-translation non-uniqueness region, the dominant phases have $\alpha \neq \beta$, and one expects this would hold for all antiferromagnetic models.

Our main technical result relates the second moment to the first moment, for any model on random bipartite regular graphs.

THEOREM 3. *For any spin system, for all $\Delta \geq 3$,*

$$\max_{\alpha, \beta} \Psi_2(\alpha, \beta) = 2 \max_{\alpha, \beta} \Psi_1(\alpha, \beta).$$

Crucially, Theorem 3 implies that $\Psi_2(\alpha, \beta) = 2\Psi_1(\alpha, \beta)$ for dominant phases, which is key for our arguments, since it will eventually allow us to prove strong concentration properties for the random variables $Z_G^{\alpha, \beta}$. We do this in two stages, first establishing the second moment (see [GSV13, Section J]), and then using the small graph conditioning

method (see [GSV13, Section G.1]). These concentration properties are described thoroughly in Section 5, where we also describe how they are used in the reduction.

2.3 Reformulating the First Moment in Terms of Matrix Norms

A key component in the analysis of the second moment is the following function Φ . Let $p = \Delta/(\Delta - 1)$. For non-negative $\mathbf{r} = (R_1, \dots, R_q)$, $\mathbf{c} = (C_1, \dots, C_q)$, define $\Phi(\mathbf{r}, \mathbf{c})$ by:

$$\exp(\Phi(\mathbf{r}, \mathbf{c})/\Delta) = \frac{\mathbf{r}^\top \mathbf{B} \mathbf{c}}{\|\mathbf{r}\|_p \|\mathbf{c}\|_p}.$$

We will show that the critical points of Φ and Ψ_1 match in the sense that there is a one-to-one correspondence between them and their values are equal at the corresponding critical points. The full statement is contained in [GSV13, Theorem 18], but the important element for the current discussion is captured in the following lemma:

LEMMA 4.

$$\max_{\alpha, \beta \in \Delta_q} \Psi_1(\alpha, \beta) = \max_{\mathbf{r}, \mathbf{c}} \Phi(\mathbf{r}, \mathbf{c}).$$

Therefore, to determine the dominant phases of Ψ_1 it suffices to study Φ . The maximum of Φ can be compactly expressed in terms of matrix norms as follows:

$$\begin{aligned} \max_{\mathbf{r}, \mathbf{c}} \exp(\Phi(\mathbf{r}, \mathbf{c})/\Delta) &= \max_{\mathbf{c}} \max_{\mathbf{r}} \frac{\mathbf{r}^\top \mathbf{B} \mathbf{c}}{\|\mathbf{r}\|_p \|\mathbf{c}\|_p} \\ &= \max_{\mathbf{c}} \frac{\|\mathbf{B} \mathbf{c}\|_\Delta}{\|\mathbf{c}\|_p} \\ &= \|\mathbf{B}\|_{p \rightarrow \Delta}, \end{aligned} \quad (1)$$

where the second equality follows from matrix norm duality.

Hence, the dominant phases of Ψ_1 can be expressed in terms of matrix norms:

$$\max_{\alpha, \beta \in \Delta_q} \exp(\Psi_1(\alpha, \beta)/\Delta) = \|\mathbf{B}\|_{\frac{\Delta}{\Delta-1} \rightarrow \Delta}. \quad (2)$$

2.4 Analyzing the Second Moment: Proof of Theorem 3

To analyze the second moment function Ψ_2 we will reduce it to the first moment optimization in the following manner. The key observation is that the associated optimization for the second moment is equivalent to a first moment optimization of a “paired-spin” model which is specified by the tensor product of the original interaction matrix with itself. This property enables us to relate the maximum for the second moment calculations with the maximum of the first moment calculations.

PROOF OF THEOREM 3. The second moment considers a pair of configurations, say σ and σ' , which are constrained to have a given phase α for V_1 and β for V_2 , where $V = V_1 \cup V_2$. We capture this constraint using a pair of vectors γ, δ corresponding to the overlap between σ and σ' , in particular, γ_{ij} (and δ_{ij}) is the number of vertices in V_1 (and V_2 , respectively) with spin i in σ and spin j in σ' .

Recall, $\Psi_1^{\mathbf{B}}$ indicates the dependence of the function Ψ_1 on the interaction matrix \mathbf{B} ; to simplify the notation we will drop the exponent if it is \mathbf{B} . We have (see [GSV13, Remark 3, Section I] for more details on this connection)

$$\Psi_2(\alpha, \beta) = \max_{\gamma, \delta} \Psi_1^{\mathbf{B} \otimes \mathbf{B}}(\gamma, \delta), \quad (3)$$

where the optimization in (3) is constrained to γ and δ such that

$$\sum_i \gamma_{ik} = \alpha_k, \sum_k \gamma_{ik} = \alpha_i, \sum_j \delta_{j\ell} = \beta_\ell \text{ and } \sum_\ell \delta_{j\ell} = \beta_j. \quad (4)$$

Ignoring the four constraints in (4) can only increase the value of (3) and hence

$$\begin{aligned} \max_{\alpha, \beta} \exp(\Psi_2(\alpha, \beta)/\Delta) &\leq \max_{\gamma, \delta} \exp(\Psi_1^{\mathbf{B} \otimes \mathbf{B}}(\gamma, \delta)/\Delta) \\ &= \|\mathbf{B} \otimes \mathbf{B}\|_{\frac{\Delta}{\Delta-1} \rightarrow \Delta}. \end{aligned} \quad (5)$$

The key fact we now use is that for induced norms $\|\cdot\|_{p \rightarrow q}$ with $p \leq q$ it holds (c.f., [Ben77, Proposition 10.1]) that:

$$\|\mathbf{B} \otimes \mathbf{B}\|_{p \rightarrow q} = \|\mathbf{B}\|_{p \rightarrow q} \|\mathbf{B}\|_{p \rightarrow q}. \quad (6)$$

Therefore,

$$\max_{\alpha, \beta} \Psi_2(\alpha, \beta) \leq 2\Delta \log \|\mathbf{B}\|_{\frac{\Delta}{\Delta-1} \rightarrow \Delta} = 2 \max_{\alpha, \beta} \Psi_1(\alpha, \beta). \quad (7)$$

To complete the proof of Theorem 3 it just remains to prove the reverse inequality, which follows from the fact that $\mathbf{E}[X^2] \geq \mathbf{E}[X]^2$. \square

3. TREE RECURSIONS AND CRITICAL POINTS OF THE FIRST MOMENT

The second moment results of the previous section will be used to establish that the Gibbs distribution is concentrated at the global maxima of $\Psi_1(\alpha, \beta)$. To simplify the analysis of the local maxima of Ψ_1 we connect them to attractive fixpoints of the associated tree recursions.

The fixpoints of the following tree recursions correspond to semi-translation invariant Gibbs measures:

$$\hat{R}_i \propto \left(\sum_{j=1}^q B_{ij} C_j \right)^{\Delta-1} \quad \text{and} \quad \hat{C}_j \propto \left(\sum_{i=1}^q B_{ij} R_i \right)^{\Delta-1}. \quad (8)$$

The fixpoints are those R_i 's and C_j 's such that $\hat{R}_i \propto R_i$ and $\hat{C}_j \propto C_j$, for all $i, j \in [q]$. The fixpoints of the tree recursions correspond to critical points of Ψ_1 , as was first observed in [MWW09], see [GSV13, Section D.2] for a derivation in our setting.

We call a fixpoint x of a function f a *Jacobian attractive fixpoint* if the Jacobian of f at x has spectral radius less than 1. We say that a critical point α, β is a *Hessian local maximum* if the Hessian of Ψ_1 at α, β is negative definite. (Note this is a sufficient condition for α, β to be a local maximum.)

We prove the following theorem in [GSV13, Section E], see [GSV13, Theorem 18] for a more detailed version of the below connection.

THEOREM 5. *Jacobian attractive fixpoints of tree recursions (8) (viewed as a function $(R_1, \dots, R_q, C_1, \dots, C_q) \mapsto (\hat{R}_1, \dots, \hat{R}_q, \hat{C}_1, \dots, \hat{C}_q)$) correspond to Hessian local maxima of Ψ_1 .*

Theorem 5 is important for analyzing the global maxima of Ψ_1 for colorings and antiferromagnetic Potts model (see Section 4). Moreover, it yields a simple formulation of the Hessian matrix of Ψ_1 at dominant phases, which is crucial to

compute the moments' asymptotics and apply the small subgraph conditioning method for general models (see [GSV13, Sections G & K]).

It is interesting to expand a bit more on the connection of Theorem 5 and the small subgraph conditioning method. Recall, the scope of the latter is to obtain almost sure results (over the choice of the random graph) using a variance analysis of the partition function. The variance analysis attributes fluctuations from the expected value of the partition function to the presence of the short cycles in the random graph. Effectively, the random graphs are locally tree-like and the small subgraph conditioning method is, in a certain sense, a way to capture the deviation from the tree structure. In this context, Theorem 5 can be used to quantify the following: fluctuations due to cycles are controlled by the eigenvalues of the stability matrix of the tree recursions. Precisely, Theorem 5 supplies us with bounds on the related eigenvalues which, together with the small subgraph conditioning method, yield that the fluctuation of $Z_G^{\alpha, \beta}$ by the presence of a cycle of length 2ℓ is captured by the disagreement percolation probability along a path of length 2ℓ in the Gibbs measure specified by the dominant phase (α, β) .

4. DOMINANT PHASES FOR ANTIFERROMAGNETIC POTTS AND COLORINGS MODELS

To obtain Theorems 1 and 2, we need to figure out the dominant phases in a random bipartite regular graph for the antiferromagnetic Potts and colorings models. Recall, the interaction matrix \mathbf{B} for the Potts model is completely determined by a parameter B , which is equal to $\exp(-\beta)$ where β is the inverse temperature in the standard notation for the Potts model. The antiferromagnetic regime corresponds to $0 < B < 1$. The coloring model is the zero temperature limit of the Potts model and corresponds to the particular case $B = 0$ in what follows. We should note that in Statistical Physics terms, the arguments of this section are closely related to the phase diagrams of the models.

Recall from Section 2.2 that a dominant phase is a global maximum of Ψ_1 . As we noted in Section 3, the critical points of Ψ_1 are given by fixpoints of the tree recursions, which for the Potts model are positive solutions to the following system of equations:

$$R_i \propto \left(BC_i + \sum_{j \neq i} C_j \right)^{\Delta-1}, \quad C_j \propto \left(BR_j + \sum_{i \neq j} R_i \right)^{\Delta-1}, \quad (9)$$

Given a fixpoint of the tree recursions (9), we will classify whether it is a Hessian local maximum of Ψ_1 using Theorem 5.

We begin our considerations by examining the fixpoints (9) which are translation invariant, i.e., satisfy $R_i \propto C_i$ for every $i \in [q]$. The following simple lemma implies that there is a unique translation invariant fixpoint.

LEMMA 6. *Let $0 \leq B < 1$ and $\Delta \geq 3$. Then a solution of (9) satisfies $R_i \propto C_i$ for every $i \in [q]$ iff $R_1 = \dots = R_q$ and $C_1 = \dots = C_q$.*

PROOF. By the symmetries of the model, we may assume an arbitrary ordering of the R_i 's. Since $0 \leq B < 1$, equations (9) imply the reverse ordering of the C_i 's. Thus, if $R_i \propto C_i$ for every $i \in [q]$, it must be the case that the ordering is trivial, i.e. $R_1 = \dots = R_q$ and $C_1 = \dots = C_q$. \square

We next identify the regime (in terms of B) where the translation invariant fixpoint corresponds to a Hessian local maximum of Ψ_1 . In this regime, Theorem 5 asserts that the translation invariant fixpoint is also Jacobian stable, hinting that uniqueness of the Gibbs measure might hold. While this rule of thumb is false in general (a counterexample is given by the *ferromagnetic* Potts model), for the colorings model it does capture the uniqueness regime, while for the antiferromagnetic Potts model, it captures the conjectured uniqueness regime. More precisely, we have the following.

LEMMA 7. *For $q < \Delta$, in the regime $0 \leq B < \frac{\Delta-q}{\Delta}$, the translation invariant fixpoint is Jacobian unstable. If $q \geq \Delta+1$ or $B > \frac{\Delta-q}{\Delta} > 0$, the translation invariant fixpoint is Jacobian stable.*

PROOF. Let \mathbf{A} be the matrix with off-diagonal entries equal to $1/(B+q-1)$ and diagonal entries equal to $B/(B+q-1)$. The translation invariant fixpoint is Jacobian stable iff the eigenvalues of \mathbf{A} other than its largest eigenvalue (which is equal to 1) have magnitude less than $1/(\Delta-1)$, see [GSV13, Lemma 45] for more details. This translates into whether $(1-B)/(B+q-1) < 1/(\Delta-1)$, yielding the second part of the lemma. Similarly, for the first part, one needs to check that $(1-B)/(B+q-1) > 1/(\Delta-1)$. \square

A much more technical argument, an extension of an argument in [BW02] for colorings, gives the following refinement of Lemma 7. In the regime where the translation invariant fixpoint is Jacobian stable, it is also the unique fixpoint of (9) (for colorings this also holds when the number of colors is equal to Δ), see [GSV13, Section H.5.1]. Thus, the second part of Lemma 7 a posteriori captures the range of temperatures where the global maximum of Ψ_1 is achieved by the translation invariant phase. In contrast, for the range of temperatures in the first part of Lemma 7, multiple semi-translation invariant fixpoints exist.

Let us now focus on $q < \Delta$ and $0 \leq B < \frac{\Delta-q}{\Delta}$. In this regime, analyzing the dominant phases is more intricate. The reason is that the number of fixpoints of (9) is a function of q . Moreover, the number of Jacobian stable fixpoints varies as a function of q, Δ, B . To explain the situation more concretely, we first need the following structural statement for the solutions of equations (9), namely that solutions of (9) are supported on at most 3 values for the R_i 's and similarly for the C_i 's. The lemma and its proof can be found as [GSV13, Lemma 47].

LEMMA 8. *Let $(R_1, \dots, R_q, C_1, \dots, C_q)$ be a positive solution of the system (9). Let t_R be the number of values on which the R_i 's are supported and define similarly t_C . Then $t_R = t_C$ and $t_R, t_C \leq 3$.*

We denote by (q_1, q_2, q_3) the multiplicities of the values that a fixpoint is supported. For example $(q, 0, 0)$ corresponds to the translation invariant fixpoint. Less trivially, for q even, the values of the R_i 's for a fixpoint of type $(q/2, q/2, 0)$ are supported on two distinct values each appearing with multiplicity $q/2$, and similarly for the C_j 's. To obtain the dominant phases, we have to compare the values of Ψ_1 at the different fixpoints, which is a difficult and unpleasant task.

To circumvent this difficulty, we instead turn to the function Φ , which by Lemma 4, can be used to study the dominant phases of Ψ_1 . Finding the maxima of Φ is itself a

difficult problem, since it involves computing $\|\mathbf{B}\|_{\frac{\Delta-1}{\Delta} \rightarrow \Delta}$. However, the extra structure allows for a natural relaxation scheme (allowing the multiplicities q_1, q_2, q_3 to be reals with sum equal to q), which can be penetrated analytically for all values of B, q, Δ . For q even, we can tie this relaxation back to the norm and hence to the global maxima of Ψ_1 .

We thus obtain the following lemma, which identifies the dominant phases for even q .

LEMMA 9. *For $0 \leq B < \frac{\Delta-1}{\Delta}$ and even $q \geq 3$, the maximum of Ψ_1 is attained at fixpoints of type $(q/2, q/2, 0)$. Moreover, there is a unique fixpoint of type $(q/2, q/2, 0)$ (up to permutations of the colors).*

For odd q , we are unable to establish whether the global maxima of Ψ_1 are fixpoints of type $(\lfloor q/2 \rfloor, \lfloor q/2 \rfloor, 1)$ or of type $(\lfloor q/2 \rfloor, \lceil q/2 \rceil, 0)$.

5. REDUCTION

In this section, we prove our inapproximability results. We start by reviewing the main components of the reduction for 2-spin systems (as carried out in [Sly10, SS12]) and in particular the hard-core model. This will allow us to isolate the parts of the argument which do not extend to the multi-spin case and motivate our reduction scheme. To simplify the presentation, we shall focus on the colorings model (k even), but the same ideas can be generalized to the Potts model and (with more technical effort) to arbitrary antiferromagnetic multi-spin models.

The basic gadget in the reduction is a bipartite random graph, which we denote by G . The sides of the bipartition have an equal number of vertices, and the sides are labelled with $+$ and $-$. Most vertices in G have degree Δ but there is also a small number of degree $\Delta - 1$ vertices (to allow to make connections between gadgets without creating degree $\Delta + 1$ vertices). For $s = \{+, -\}$, let the vertices in the s -side be $U^s \cup W^s$ where the vertices in $U = U^+ \cup U^-$ have degree Δ and the vertices in $W = W^+ \cup W^-$ have degree $\Delta - 1$. The phase of an independent set I is $+$ (resp. $-$) if I has more vertices in U^+ (resp. U^-). Note that the phase depends only on the spins of the ‘‘large’’ portion of the graph, i.e., the spins of vertices in U .

In non-uniqueness regimes, the gadget G has two important properties, both of which can be obtained by building on the second moment argument we outlined earlier. First, the phase of a random independent set I is equal to $+$ or $-$ with probability roughly equal to $1/2$. Second, conditioned on the phase of a random independent set I , the spins of the vertices in W are approximately independent, i.e., the marginal distribution on W is close to a product distribution. In this product distribution if the phase is $+$ (resp. $-$), a vertex in W^+ is in I with probability p^+ (resp. p^-), while a vertex in W^- is in I with probability p^- (resp. p^+). The values p^\pm correspond to maxima of the function Ψ_1 and, crucially (as we shall demonstrate shortly), they satisfy $p^+ \neq p^-$.

Using the second moment analysis of Section 2 and in particular Theorem 3, we can prove that an analogous phenomenon takes place for the k -colorings model in the semi-translation non-uniqueness regime. The main difference is that, instead of two phases, the number of phases is equal to the number of maximizers of the function Ψ_1 . For k even, the phase of a coloring is determined by the dominant set

of $k/2$ colors on U^+ , i.e., the $k/2$ colors with largest frequencies among vertices of U^+ . Each of the $\binom{k}{k/2}$ phases appears with roughly equal probability and given the phase, the marginal distribution on W is close to a product distribution, which we now describe. We can compute explicit values $a = a(k, \Delta), b = b(k, \Delta)$ such that for a phase $T \in \binom{[k]}{k/2}$ the probability mass function \mathbf{x} of a vertex in W^+ has its i -th entry equal to a if $i \in T$ and equal to b if $i \notin T$. Similarly, the probability mass function \mathbf{y} of a vertex in W^- has its i -th entry equal to b if $i \in T$ and equal to a if $i \notin T$. We note here that the values of a, b correspond to the values of $R_1, \dots, R_q, C_1, \dots, C_q$ of the $(q/2, q/2, 0)$ fixpoint of Lemma 9.

Let \mathcal{Q} be the union of the pairs (\mathbf{x}, \mathbf{y}) over all phases. Hereafter, we will identify the phases with elements of \mathcal{Q} . Note that if $(\mathbf{x}, \mathbf{y}) \in \mathcal{Q}$, then $(\mathbf{y}, \mathbf{x}) \in \mathcal{Q}$ as well. We also denote by \mathcal{Q}' the union of unordered elements of \mathcal{Q} . Elements of \mathcal{Q}' are called unordered phases (we use \mathbf{p} to denote unordered phases). Given a phase $\mathbf{p} = \{\mathbf{x}, \mathbf{y}\}$ an ordering of the pair will be called ‘‘assigning spin to the phase’’. The two ordered phases corresponding to the unordered phase \mathbf{p} will be denoted by \mathbf{p}^+ and \mathbf{p}^- .

The conditional independence property is crucial. It allows to quantify the effect of using vertices of W as terminals to make connections between copies of the gadget G . For example, consider the following type of connection, which we refer to as parallel. Let $v^+ \in W^+, v^- \in W^-$ and consider two copies of the gadget G , say G_1, G_2 . For $i = 1, 2$ denote by v_i^+, v_i^- the images of v^+, v^- in G_i . Now add the edges (v_1^+, v_2^+) and (v_1^-, v_2^-) and denote the final graph by G_{12} . Thus, a parallel connection corresponds to joining the $+, +$ and $-, -$ sides of two copies of the gadget.

Clearly, random colorings of G_{12} can be generated by first generating random colorings of G_1, G_2 and keeping the resulting coloring if v_1^\pm, v_2^\pm have different colors. We thus have that the partition function of G_{12} is equal to $(Z_G)^2$ times the probability that v_1^\pm, v_2^\pm have different colors in random colorings of G_1, G_2 . The latter quantity can easily be computed if we condition on the phases $(\mathbf{x}_1, \mathbf{y}_1), (\mathbf{x}_2, \mathbf{y}_2)$ of the colorings in G_1, G_2 , and this is equal to $(1 - \mathbf{x}_1^\top \mathbf{x}_2)(1 - \mathbf{y}_1^\top \mathbf{y}_2)$.

By taking logarithms, we can assume a parallel connection between gadgets with phases $(\mathbf{x}_1, \mathbf{y}_1)$ and $(\mathbf{x}_2, \mathbf{y}_2)$ incurs an (additive) weight

$$w_p((\mathbf{x}_1, \mathbf{y}_1), (\mathbf{x}_2, \mathbf{y}_2)) = \ln(1 - \mathbf{x}_1^\top \mathbf{x}_2) + \ln(1 - \mathbf{y}_1^\top \mathbf{y}_2).$$

In the hard-core model, parallel connections are sufficient to give hardness. In this case, we have that $\mathcal{Q}' = \{\mathbf{p}\}$ and $\mathcal{Q} = \{\mathbf{p}^+, \mathbf{p}^-\}$ and the respective function $w_p(\cdot, \cdot)$ satisfies

$$w_p(\mathbf{p}^+, \mathbf{p}^+) = w_p(\mathbf{p}^-, \mathbf{p}^-) < w_p(\mathbf{p}^+, \mathbf{p}^-). \quad (10)$$

Thus, in this case, $w_p(\cdot, \cdot)$ takes only two values and neighboring gadgets prefer to have different phases. Now assume that H is an instance of MAX-CUT and replace each vertex in H by a copy of the gadget G , while for each edge of H , connect the respective gadgets in parallel. The partition function of the final graph is dominated from phase assignments which correspond to large cuts in H . This intuition is the basis of the reduction in [Sly10, SS12].

For the colorings model, reducing from MAX-CUT poses an extra challenge. While for every unordered phase \mathbf{p} equation (10) continues to hold, a short calculation shows that the optimal configuration for a triangle of gadgets connected in parallel is to give all three gadgets different phases. To

bypass this entanglement, we need to introduce some sort of ferromagnetism in the reduction to enforce gadgets corresponding to vertices of H to use a single (unordered) phase. To achieve this, we use *symmetric* connections, which correspond to having not only $(+, +)$, $(-, -)$ connections of the gadgets, but also $(+, -)$ and $(-, +)$. Thus, a symmetric connection whose endpoints have phases $(\mathbf{x}_1, \mathbf{y}_1)$, $(\mathbf{x}_2, \mathbf{y}_2)$ incurs (additive) weight

$$\begin{aligned} w_s((\mathbf{x}_1, \mathbf{y}_1), (\mathbf{x}_2, \mathbf{y}_2)) \\ = w_p((\mathbf{x}_1, \mathbf{y}_1), (\mathbf{x}_2, \mathbf{y}_2)) + w_p((\mathbf{x}_1, \mathbf{y}_1), (\mathbf{y}_2, \mathbf{x}_2)). \end{aligned}$$

Symmetric connections will allow us to enforce a single unordered phase to all gadgets, while parallel connections will allow us to recover a maximum-cut partition. To have some modularity in our construction, rather than reducing from MAX-CUT directly, we use the following “phase labeling problem”.

COLORINGS PHASE LABELING PROBLEM(\mathcal{Q}):

INPUT: undirected edge-weighted multigraph $H = (V, E)$ and a partition of the edges $\{E_p, E_s\}$.

OUTPUT: $\text{MAXLWT}(H) := \max_{\mathcal{Y}} \text{LWT}_H(\mathcal{Y})$, where the maximization is over all possible phase labelings $\mathcal{Y} : V \rightarrow \mathcal{Q}$ and

$$\text{LWT}_H(\mathcal{Y}) := \sum_{\{u,v\} \in E_s} w_s(\mathcal{Y}(u), \mathcal{Y}(v)) + \sum_{\{u,v\} \in E_p} w_p(\mathcal{Y}(u), \mathcal{Y}(v)).$$

Edges in E_p (resp. E_s) correspond to parallel (resp. symmetric) connections and we shall refer to them as parallel (resp. symmetric) edges. The arguments in [SS12], which we sketched earlier, can easily be adapted to show that an algorithm for approximating the partition function to an arbitrarily small exponential factor yields a PRAS for the phase labeling problem, see [GSV13, Lemma 11] and its proof in [GSV13, Section C]. It then remains to prove that a PRAS for the phase labeling problem yields a PRAS for MAX-CUT on 3-regular graphs.

Our reduction relies on the following gadget which “prefers” the unordered phase of two distinguished vertices u and v to agree. For a phase assignment \mathcal{Y} with ordered phases, we denote by \mathcal{Y}' the respective phase assignment with unordered phases.

LEMMA 10. *A constant sized gadget J_1 with two distinguished vertices u, v can be constructed with the following property: all edges of J_1 are symmetric and the following is true,*

$$\max_{\mathcal{Y}; \mathcal{Y}'(u)=\mathcal{Y}'(v)} \text{LWT}_{J_1}(\mathcal{Y}) > \varepsilon_1 + \max_{\mathcal{Y}; \mathcal{Y}'(u) \neq \mathcal{Y}'(v)} \text{LWT}_{J_1}(\mathcal{Y}), \quad (11)$$

where $\varepsilon_1 > 0$ is a constant depending only on k and Δ .

We give the proof of the critical Lemma 10 after the (simpler) proof of the following.

LEMMA 11. *A PRAS for COLORINGS PHASE LABELING PROBLEM(\mathcal{Q}) yields a PRAS for MAX-CUT on 3-regular graphs.*

PROOF. Let ε_1 be as in Lemma 10 and

$$t := 2 \lceil (\max_{\mathbf{p}_1, \mathbf{p}_2} w_p(\mathbf{p}_1, \mathbf{p}_2) - \min_{\mathbf{p}_1, \mathbf{p}_2} w_p(\mathbf{p}_1, \mathbf{p}_2)) / \varepsilon_1 \rceil.$$

Given a 3-regular instance $H = (V, E)$ of MAX-CUT, we first declare all edges of H to be parallel. Moreover, for every edge (u', v') of H , take t copies of gadget J_1 from

Lemma 10, identify (merge) their u vertices with u' , and identify (merge) their v vertices with v' . Let H' be the final graph.

To find the optimal phase labeling of H' , we may focus on the phase assignment restricted to vertices in H , since each gadget J_1 can be independently set to its optimal value conditioned on the phases for its distinguished vertices u and v . We claim that

$$\text{MAXLWT}(H') = C_1 \text{MAXCUT}(H) + (C_2 + C_3 t)|E|, \quad (12)$$

for constants C_1, C_2, C_3 to be specified later (depending only on k, Δ). Using the trivial bound $\text{MAXCUT}(H) \geq |E|/2 = 3|V|/4$, the lemma follows easily from (12). We thus focus on proving (12).

The key idea is that for any phase labeling $\mathcal{Y} : V \rightarrow \mathcal{Q}$, changing the unordered phases of vertices in H to the same unordered phase $\mathbf{p} \in \mathcal{Q}'$, while keeping the spins, can only increase the weight of the labeling. Indeed, for $(u, v) \in E$ such that $\mathcal{Y}'(u) = \mathcal{Y}'(v)$, no change in the weight of the labeling occurs, using (11). For $(u, v) \in E$ such that $\mathcal{Y}'(u) \neq \mathcal{Y}'(v)$, the potential (weight) loss from the parallel edge (u, v) is compensated by the gain on the t copies of J_1 by (11) and the choice of t .

For phase labelings which assign vertices of H the same unordered phase \mathbf{p} , to attain the maximum weight for a phase labeling, we only need to choose the spins, in order to maximize the contribution from parallel edges (the edges of H). The same argument we discussed for the hard-core model, (10) yields that the optimal choice of spins to the phases induces a maximum-cut partition of H . For such a spin assignment, the contribution from parallel edges is $C_1 \text{MAXCUT}(H) + C_2|E|$, where

$$C_1 := w_p(\mathbf{p}^+, \mathbf{p}^-) - w_p(\mathbf{p}^-, \mathbf{p}^-) \text{ and } C_2 := w_p(\mathbf{p}^-, \mathbf{p}^-).$$

The contribution from symmetric edges is $C_3 t|E|$, where $C_3 := \max_{\mathcal{Y}; \mathcal{Y}'(u)=\mathcal{Y}'(v)=\mathbf{p}} \text{LWT}_{J_1}(\mathcal{Y})$. This proves (12). \square

PROOF OF LEMMA 10. Let $\mathcal{Q}' := \{\mathbf{p}_1, \dots, \mathbf{p}_{Q'}\}$ and $\mathbf{p}_i := \{\mathbf{x}_i, \mathbf{y}_i\}$ for $i \in [Q']$. Denote by K the multigraph on Q' vertices $b_1, b_2, \dots, b_{Q'}$ with the following symmetric edges: self-loop on b_i for $i \in [Q']$ and two edges between b_i and b_j for every $i, j \in [Q']$ with $i \neq j$. We first prove that the optimal phase assignments \mathcal{Y} of K are those which assign each vertex b_i a distinct phase from \mathcal{Q}' (note that the spin of the phase does not matter since all edges of K are symmetric). The desired gadget J_1 will be constructed afterwards.

Let \mathcal{Y} be a phase labeling of K and s_i be the number of vertices assigned phase \mathbf{p}_i . Let \mathbf{s} be the vector $(s_1, \dots, s_{Q'})^\top$. Note that $\mathbf{1}^\top \mathbf{s} = Q'$, where $\mathbf{1}$ is the all one vector with dimension Q' . Then

$$\text{LWT}_K(\mathcal{Y}) = \sum_{i,j \in [Q']} s_i s_j w_s(\mathbf{p}_i, \mathbf{p}_j) = \mathbf{s}^\top \mathbf{A} \mathbf{s},$$

where \mathbf{A} is the $Q' \times Q'$ matrix whose (i, j) entry equals $w_s(\mathbf{p}_i, \mathbf{p}_j)$. Note that \mathbf{A} is symmetric and $\mathbf{1}$ is an eigenvector of \mathbf{A} (because of the transitive symmetry of phases). Moreover, if we let $\mathbf{s}' = \mathbf{s} - \mathbf{1}$, then $\mathbf{1}^\top \mathbf{s}' = 0$. It follows that

$$\mathbf{s}^\top \mathbf{A} \mathbf{s} = \mathbf{1}^\top \mathbf{A} \mathbf{1} + (\mathbf{s}')^\top \mathbf{A} \mathbf{s}'. \quad (13)$$

If \mathbf{A} is negative definite, equation (13) shows that the all ones labeling is better than any other labeling. Hence the result will follow if we prove that \mathbf{A} is negative definite.

Let $\mathbf{z}_1, \dots, \mathbf{z}_Q := \mathbf{x}_1, \dots, \mathbf{x}_{Q'}, \mathbf{y}_1, \dots, \mathbf{y}_{Q'}$ and let $\hat{\mathbf{A}}$ be the $Q \times Q$ matrix whose ij -entry is $\ln(1 - \mathbf{z}_i^\top \mathbf{z}_j)$. Using the definition of the weights $w_s(\cdot, \cdot)$, it is easy to check that for any vector \mathbf{s} it holds that

$$\mathbf{s}^\top \mathbf{A} \mathbf{s} = (\mathbf{s}, \mathbf{s})^\top \hat{\mathbf{A}}(\mathbf{s}, \mathbf{s}),$$

so it suffices to prove that $\hat{\mathbf{A}}$ is negative definite. We will show here that $\hat{\mathbf{A}}$ is negative semi-definite; the proof that $\hat{\mathbf{A}}$ is regular (and hence negative definite) is trickier and is given in the proof of the more general [GSV13, Lemma 10]. Note that the entries of $\hat{\mathbf{A}}$ are obtained by applying $z \mapsto \ln(1 - z)$ to each entry of the Gram matrix of the vectors $\mathbf{z}_1, \dots, \mathbf{z}_Q$. Since for $|z| < 1$ we have $\ln(1 - z) = -z - z^2/2 - z^3/3 - \dots$, by Schur's product theorem (see Corollary 7.5.9 in [HJ90]) we obtain that $\hat{\mathbf{A}}$ is negative semi-definite, as desired.

To construct the gadget J_1 , we overlay two copies of K as follows. Let K_u (resp. K_v) be a copy of K , where the image of $b_{Q'}$ is renamed to u (resp. v). Overlay K_u, K_v by identifying the images of $b_1, \dots, b_{Q'-1}$ in the two copies. Thus, the resulting graph J_1 has two self loops on b_i for $i \in [Q' - 1]$, four edges between b_i and b_j for every $i, j \in [Q' - 1]$ with $i \neq j$, two edges between u and b_i for $i \in [Q' - 1]$, two edges between v and b_i for $i \in [Q' - 1]$ and a self loop on u, v .

Note that for every phase labeling \mathcal{Y} of J_1 , we have

$$\text{LWT}_{J_1}(\mathcal{Y}) = \text{LWT}_{K_u}(\mathcal{Y}) + \text{LWT}_{K_v}(\mathcal{Y})$$

and hence

$$\text{MAXLWT}(J_1) \leq 2\text{MAXLWT}(K).$$

Using that the optimal phase labelings for K are those which assign each vertex a distinct phase from Q' , we obtain that the inequality holds at equality for those (and only those) phase labelings which assign u, v a common phase $\mathbf{p} \in Q'$ and vertices $b_1, \dots, b_{Q'-1}$ a distinct phase from $Q' - \{\mathbf{p}\}$. This yields the ε_1 in the statement of the lemma. Note that ε_1 depends only on Q' , which in turn is completely determined by k, Δ . \square

6. GENERAL RESULTS

The inapproximability results for colorings can be extended to general antiferromagnetic models on bounded degree graphs. The key concept in the general theorem is again the existence of long range correlations, in the form of semi-translational non-uniqueness on the infinite regular tree.

We use the following definition of antiferromagnetic models, which is in terms of the signature of the interaction matrix \mathbf{B} , i.e., the signs of its eigenvalues. Recall that the interaction matrix \mathbf{B} is symmetric and hence its eigenvalues are real. Moreover, it is simple to see that the matrices \mathbf{B} which correspond to non-degenerate models should be irreducible. The Perron-Frobenius theorem then implies that one of the eigenvalues with the largest magnitude is positive. When all the other eigenvalues are negative, we show in [GSV13, Section A] (see [GSV13, Corollary 9]) that neighboring spins prefer to be different and hence we call such models antiferromagnetic. This notion of antiferromagnetism extends naturally the definition for 2-spin models (see [GJP03, LLY13, SS12]) and also captures antiferromagnetism in the Potts model. A more thorough discussion is given in [GSV13, Section A].

Proceeding to the general hardness results, we have already displayed that the second moment argument for random bipartite graphs covers arbitrary models (general interaction matrix \mathbf{B}). There are two properties of the gadget G in Section 5 we need to ensure. First, the symmetry breaking between the two sides of the graph, i.e., a typical configuration should have different color frequencies on the two sides of the graph. Second, to be able to quantify the interaction between neighboring gadgets, we need to ensure that phases appear with roughly equal probability and that given the phase, the spins of the vertices with degree $\Delta - 1$ are approximately independent.

The first property can be guaranteed by the absence of translation invariant phases, i.e., maximizers of the function Ψ_1 of the form (\mathbf{x}, \mathbf{x}) . Since in uniqueness regimes the only maximizer is translation invariant, this can be the case only in non-uniqueness regimes. The second property is subtler and relates to the concentration properties of the random variable $Z_G^{\alpha, \beta}$. The required concentration is sufficiently strong when the maximizers α, β (viewed as unordered pairs) of the function Ψ_1 are (i) Hessian maxima, i.e., the Hessian matrix of Ψ_1 is negative definite when evaluated at (α, β) , and (ii) permutation symmetric, i.e., obtainable from one another by a suitable permutation of the set of spins. We clarify that the permutations must be automorphisms of the interaction matrix \mathbf{B} . For example, note that the maxima for the colorings model are permutation symmetric.

Given these assumptions, the reduction of Section 5 (with some extra work) can be adapted to give the following general inapproximability result.

THEOREM 12. *Let $q \geq 2, \Delta \geq 3$. For an antiferromagnetic q -spin system with interaction matrix \mathbf{B} , if the dominant semi-translation invariant Gibbs measures on the tree \mathbb{T}_Δ are permutation symmetric and all of them are Hessian dominant and not translation invariant then, unless $\text{NP}=\text{RP}$, there is no FPRAS for approximating the partition function for triangle free Δ -regular graphs. Moreover, there exists $\varepsilon = \varepsilon(q, \Delta)$ such that, unless $\text{NP}=\text{RP}$, one cannot approximate the partition function within a factor $2^{\varepsilon n}$ for triangle-free Δ -regular graphs (where n is the number of vertices).*

We note that the inapproximability results for antiferromagnetic 2-spin systems [Sly10, SS12, GSV12] and Theorems 1 and 2 follow as corollaries of Theorem 12. We briefly sketch the details for 2-spin systems. Recall that for 2-spin systems, non-uniqueness coincides with the existence of multiple semi-translation invariant Gibbs measures on the tree. Moreover, in non-uniqueness, the maximizers of Ψ_1 are exactly two pairs (α, β) and (β, α) with $\alpha \neq \beta$. Thus, they do not correspond to translation invariant measures and satisfy the permutation symmetric property. Finally, it can also be verified that they are Hessian dominant and hence the hypotheses of Theorem 12 are satisfied.

7. REFERENCES

- [AK97] P. Alimonti and V. Kann. Hardness of approximating problems on cubic graphs. In *Proceedings of Algorithms and Complexity, Third Italian Conference, (CIAC)*, pages 288–298, 1997.

- [Ben77] G. Bennett. Schur multipliers. *Duke Math. J.*, 44(3):603–639, 1977.
- [BW02] G. R. Brightwell and P. Winkler. Random colorings of a Cayley tree. In *Contemporary combinatorics*, volume 10 of *Bolyai Soc. Math. Stud.*, pages 247–276. János Bolyai Math. Soc., Budapest, 2002.
- [Bru81] N. G. de Bruijn. *Asymptotic Methods in Analysis*, Dover, New York, 1981.
- [DMSS12] A. Dembo, A. Montanari, A. Sly, and N. Sun. The replica symmetric solution for Potts models on d -regular graphs. *arXiv:1207.5500 [math.PR]*, 2012. Preprint is available from the arXiv at: <http://arxiv.org/abs/1207.5500>.
- [GSV13] A. Galanis, D. Štefankovič, and E. Vigoda. Inapproximability for Antiferromagnetic Spin Systems in the Tree Non-Uniqueness Region. This is the full version of this paper which is available from the arXiv at: <http://arxiv.org/abs/1305.2902>
- [GSV12] A. Galanis, D. Štefankovič, and E. Vigoda. Inapproximability of the partition function for the antiferromagnetic Ising and hard-core models. *CoRR*, abs/1203.2226, 2012. Preprint is available from the arXiv at: <http://arxiv.org/abs/1203.2226>.
- [Geo11] H.-O. Georgii. *Gibbs measures and phase transitions*, volume 9 of *de Gruyter Studies in Mathematics*. Walter de Gruyter & Co., Berlin, second edition, 2011.
- [GJP03] L. A. Goldberg, M. Jerrum, and M. Paterson. The computational complexity of two-state spin systems. *Random Struct. Algorithms*, 23(2):133–154, 2003.
- [Gre00] C. Greenhill. The complexity of counting colourings and independent sets in sparse graphs and hypergraphs. *Comput. Complex.*, 9(1):52–72, 2000.
- [HJ90] R. A. Horn and C. R. Johnson. *Matrix analysis*. Cambridge University Press, Cambridge, 1990.
- [Jan95] S. Janson. Random regular graphs: Asymptotic distributions and contiguity. *Combinatorics, Probability & Computing*, 4:369–405, 1995.
- [JLR11] S. Janson, T. Łuczak, and A. Ruciński. *Random Graphs*. Wiley Series in Discrete Mathematics and Optimization. Wiley, 2011.
- [JS93] M. Jerrum and A. Sinclair. Polynomial-time Approximation Algorithms for the Ising Model. *SIAM Journal on Computing*, 22(5):1087–1116, 1993.
- [Joh96] A. Johansson. Asymptotic choice number for triangle free graphs. Technical Report 91-5, DIMACS, 1996.
- [Jon02] J. Jonasson. Uniqueness of uniform random colorings of regular trees. *Statistics & Probability Letters*, 57:243–248, 2002.
- [Kel91] F. P. Kelly. Loss networks. *Ann. Appl. Probab.*, 1(3):319–378, 1991.
- [LLY13] L. Li, P. Lu, and Y. Yin. Correlation Decay up to Uniqueness in Spin Systems. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 67–84, 2013.
- [LLZ14] J. Liu, P. Lu, and C. Zhang, The Complexity of Ferromagnetic Two-spin Systems with External Fields. *CoRR*, abs/1402.4346, 2014. Preprint is available from the arXiv at: <http://arxiv.org/abs/1402.4346>.
- [Lue04] D.G. Luenberger. *Linear and Nonlinear Programming*. Springer-Verlag, New York, second edition, 2004.
- [MR01] M. Molloy and B. A. Reed. Colouring graphs when the number of colours is nearly the maximum degree. In *Proceedings on 33rd Annual ACM Symposium on Theory of Computing (STOC)*, pages 462–470, 2001.
- [MR02] M. Molloy and B. Reed. *Graph colouring and the probabilistic method*, volume 23 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 2002.
- [MT10] R. A. Moser and G. Tardos. A constructive proof of the general Lovász local lemma. *J. ACM*, 57(2), 2010.
- [MWW09] E. Mossel, D. Weitz, and N. Wormald. On the hardness of sampling independent sets beyond the tree threshold. *Probab. Theory Related Fields*, 143(3-4):401–439, 2009.
- [RW94] R. W. Robinson and N. C. Wormald. Almost all regular graphs are Hamiltonian. *Random Struct. Algorithms*, 5(2):363–374, 1994.
- [SST12] A. Sinclair, P. Srivastava, and M. Thurley. Approximation algorithms for two-state anti-ferromagnetic spin systems on bounded degree graphs. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 941–953, 2012.
- [Sly10] A. Sly. Computational transition at the uniqueness threshold. In *Proceedings of the 51th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 287–296, 2010.
- [SS12] A. Sly and N. Sun. The computational hardness of counting in two-spin models on d -regular graphs. In *Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 361–369, 2012.
- [Wei06] D. Weitz. Counting independent sets up to the tree threshold. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC)*, pages 140–149, 2006.