Ferromagnetic Potts Model: Refined #BIS-hardness and Related Results*

Andreas Galanis†1, Daniel Štefankovič‡2, Eric Vigoda§3, and Linji Yang4

1 **University of Oxford, Wolfson Building, Parks Road, Oxford, OX1 3QD, UK**
andreas.galanis@cs.ox.ac.uk
2 **University of Rochester, Rochester, NY, 14627, USA**
stefanko@cs.rochester.edu
3 **Georgia Institute of Technology, Atlanta, GA, 30332, USA**
vigoda@cc.gatech.edu
4 **Facebook, Inc., Menlo Park, CA, 94025, USA**
ljyang@gatech.edu

Abstract

Recent results establish for the hard-core model (and more generally for 2-spin antiferromagnetic systems) that the computational complexity of approximating the partition function on graphs of maximum degree Δ undergoes a phase transition that coincides with the uniqueness/non-uniqueness phase transition on the infinite Δ-regular tree. For the ferromagnetic Potts model we investigate whether analogous hardness results hold. Goldberg and Jerrum showed that approximating the partition function of the ferromagnetic Potts model is at least as hard as approximating the number of independent sets in bipartite graphs, so-called #BIS-hardness. We improve this hardness result by establishing it for bipartite graphs of maximum degree Δ. To this end, we first present a detailed picture for the phase diagram for the infinite Δ-regular tree, giving a refined picture of its first-order phase transition and establishing the critical temperature for the coexistence of the disordered and ordered phases. We then prove for all temperatures below this critical temperature (corresponding to the region where the ordered phase "dominates") that it is #BIS-hard to approximate the partition function on bipartite graphs of maximum degree Δ.

The #BIS-hardness result uses random bipartite regular graphs as a gadget in the reduction. The analysis of these random graphs relies on recent results establishing connections between the maxima of the expectation of their partition function, attractive fixpoints of the associated tree recursions, and induced matrix norms. In this paper we extend these connections to random regular graphs for all ferromagnetic models. Using these connections, we establish the Bethe prediction for every ferromagnetic spin system on random regular graphs, which says roughly that the expectation of the log of the partition function Z is the same as the log of the expectation of Z. As a further consequence of our results, we prove for the ferromagnetic Potts model that the Swendsen-Wang algorithm is torpidly mixing (i. e., exponentially slow convergence to its stationary distribution) on random Δ-regular graphs at the critical temperature for sufficiently large q.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems

Keywords and phrases Ferromagnetic Potts model, approximate counting, spin systems, phase transition, random regular graphs

Digital Object Identifier 10.4230/LIPIcs.APPROX-RANDOM.2014.677

* Full version [15].

† The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) ERC grant agreement no. 334828. The paper reflects only the authors' views and not the views of the ERC or the European Commission. The European Union is not liable for any use that may be made of the information contained therein.
‡ Research supported in part by NSF grant CCF-1318374.
§ Research supported in part by NSF grant CCF-1217458.

17th Int'l Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX'14) / 18th Int'l Workshop on Randomization and Computation (RANDOM'14).
Editors: Klaus Jansen, José Rolim, Nikhil Devanur, and Cristopher Moore; pp. 677–691
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Background

1.1 Spin Systems

We study the ferromagnetic Potts model and present tools which are useful for any ferromagnetic spin system on random regular graphs. Hence we begin with a general definition of a spin system.

A spin system is defined, for an n-vertex graph $G = (V, E)$ and integer $q \geq 2$, on the space Ω of configurations σ which are assignments $\sigma : V \to [q]$. The model is characterized by its energy or Hamiltonian $H(\sigma)$ which is a function of the spin assignments to the vertices. In the classical examples of the Ising ($q = 2$) and Potts ($q \geq 3$) models without external field, the Hamiltonian $H(\sigma)$ is the number of monochromatic edges in σ. Each configuration has a weight $w(\sigma) = \exp(-\beta H(\sigma))$ for a parameter β corresponding to the "inverse temperature" which controls the strength of edge interactions.

The *Gibbs distribution* is defined as $\mu(\sigma) = w(\sigma)/Z$ where $Z = Z_G(\mathbf{B}) = \sum_\sigma w(\sigma)$ is the *partition function*. In our general setup, a *specification* of a q-state spin model is defined by a symmetric $q \times q$ interaction matrix $\mathbf{B}$ with non-negative entries. The weight of a configuration in this general setup is given by:

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

Many of our results also apply to models with arbitrary external fields since we will work with Δ-regular graphs and in this case the external field can be incorporated into the interaction matrix.

The Ising ($q = 2$) and Potts ($q > 2$) models have interaction matrices with diagonal entries $B := \exp(-\beta)$ and off-diagonal entries 1. The models are called ferromagnetic if $B > 1$ since then neighboring spins prefer to align and antiferromagnetic if $B < 1$. The hard-core model is an example of a 2-spin antiferromagnetic system, its interaction matrix is defined so that Ω is the set of independent sets of G and configuration $\sigma \in \Omega$ has weight $w(\sigma) = \lambda^{|\sigma|}$ for activity $\lambda > 0$.

We are not aware of a general definition of ferromagnetic and antiferromagnetic models. We use the following notions which generalize the analogous notions for 2-spin and for the Potts model. The ferromagnetic definition captures that neighboring spins preferring to align (see [15, Observation 1] in the full version of this paper). To avoid degenerate cases, we assume throughout this paper that $\mathbf{B}$ is ergodic, that is, irreducible and aperiodic, see [15, Section 1.2] in the full version for a detailed discussion. Hence, by the Perron-Frobenius theorem (and since $\mathbf{B}$ is non-negative) the eigenvalue of $\mathbf{B}$ with the largest magnitude is positive.

▶ **Definition 1.** A model is called **ferromagnetic** if $\mathbf{B}$ is positive definite. Equivalently we have that all of its eigenvalues are positive and also that

$$\mathbf{B} = \hat{\mathbf{B}}^T \hat{\mathbf{B}}, \tag{1}$$

for some $q \times q$ matrix $\hat{\mathbf{B}}$.

In contrast to the above notion of a ferromagnetic system, in [17] a model is called **antiferromagnetic** if all of the eigenvalues of $\mathbf{B}$ are negative except for the largest (which, as noted above, is positive).

1.2 Known Connections to Phase Transitions

Exact computation of the partition function is #P-complete, even for very restricted classes of graphs [23]. Hence we focus on whether there is a fully-polynomial (randomized or deterministic) approximation scheme, a so-called FPRAS or FPTAS.

One of our goals in this paper is to refine our understanding of connections between approximating the partition function on graphs of maximum degree Δ with phase transitions on the infinite Δ-regular tree $\mathbb{T}_\Delta$. A phase transition of particular interest in the infinite tree $\mathbb{T}_\Delta$ is the uniqueness/non-uniqueness threshold. Roughly speaking, in the uniqueness phase, if one fixes a so-called "boundary condition" which is a configuration σ_ℓ (for instance, an independent set in the hard-core model) on the vertices distance ℓ from the root, then in the Gibbs distribution conditioned on this configuration, is the root "unbiased"? Specifically, for all sequences (σ_ℓ) of boundary conditions, in the limit $\ell \to \infty$, does the root have the same marginal distribution? If so, there is a unique Gibbs measure on the infinite tree and hence we say the model is in the uniqueness region. If there are sequences of boundary conditions which influence the root in the limit then we say the model is in the non-uniqueness region.

For 2-spin antiferromagnetic spin systems, it was shown that there is an FPTAS for estimating the partition function for graphs of maximum degree Δ when the infinite tree $\mathbb{T}_\Delta$ is in the uniqueness phase [28]. On the other side, unless NP=RP, there is no FPRAS for the partition function for Δ-regular graphs when $\mathbb{T}_\Delta$ is in the non-uniqueness phase [36] (see also [16]). Recently, an analogous NP-hardness result was shown for approximating the number of k-colorings on triangle-free Δ-regular graph for even k when $k < \Delta$. In contrast to the above inapproximability results for antiferromagnetic systems, for the Ising model with or without external field [26] and for 2-spin ferromagnetic spin systems without external field [22] there is an FPRAS for all graphs. The situation for ferromagnetic multi-spin models, the ferromagnetic Potts being the most prominent example, is more intricate.

#BIS refers to the problem of computing the number of independent sets in bipartite graphs. A series of results has presented evidence that there is unlikely to be a polynomial-time algorithm for #BIS, since a number of unsolved counting problems have be shown to be #BIS-hard (for example, see [13, 2, 7]). The growing anecdotal evidence for #BIS-hardness suggests that the problem is intractable, though weaker than NP-hardness. More recently it was shown in [6] that for antiferromagnetic 2-spin models it is #BIS-hard to approximate the partition function on bipartite graphs of maximum degree Δ when the parameters of the model lie in the non-uniqueness region of the infinite Δ-regular tree $\mathbb{T}_\Delta$.

2 Results for the Potts Model

2.1 #BIS-hardness for the Potts model

Goldberg and Jerrum [20] showed that approximating the partition function of the ferromagnetic Potts model is #BIS-hard, hence it appears likely that the ferromagnetic Potts model is inapproximable for general graphs. We refine this #BIS-hardness result for the ferromagnetic Potts model. We prove that approximating the partition function for the ferromagnetic Potts model on *bipartite* graphs of maximum degree Δ is #BIS-hard for temperatures above the appropriate phase transition point in the infinite tree $\mathbb{T}_\Delta$. The appropriate phase transition in the Potts model is not the uniqueness/non-uniqueness threshold, but rather it is the ordered/disordered phase transition which occurs at $B = \mathfrak{B}_o$ as explained in the next section.

Formally, we study the following problem.

Name. #BipFerroPotts(q, B, Δ).
Instance. A bipartite graph G with maximum degree Δ.
Output. The partition function for the q-state Potts model on G.

We use the notion of approximation-preserving reductions, denoted as $\leq_{AP}$, formally defined in [13]. We can now formally state our main result.

▶ **Theorem 2.** *For all $q \geq 3$, all $\Delta \geq 3$, for the ferromagnetic q-state Potts model, for any $B > \mathfrak{B}_o$,*

$$\#\text{BIS} \leq_{AP} \#\text{BipFerroPotts}(q, B, \Delta),$$

where $\mathfrak{B}_o$ is given by (4)*.*

2.2 Potts Model Phase Diagram

To understand the critical point $\mathfrak{B}_o$ we need to delve into the nature of the phase transition in the ferromagnetic Potts model on the infinite Δ-regular tree $\mathbb{T}_\Delta$. We focus on how the phase transition manifests on a random Δ-regular graph.

For a configuration $\sigma \in \Omega$, denote the set of vertices assigned spin i by $\sigma^{-1}(i)$. Let $\triangle_q$ denote the $(q-1)$-simplex:

$$\triangle_t = \{(x_1, x_2, \ldots, x_t) \in \mathbb{R}^t \mid \textstyle\sum_{i=1}^t x_i = 1 \text{ and } x_i \geq 0 \text{ for } i = 1, \ldots, t\}.$$

We refer to $\boldsymbol{\alpha} \in \triangle_q$ as a *phase*. For a phase $\boldsymbol{\alpha}$, denote the set of configurations with frequencies of colors given by $\boldsymbol{\alpha}$ as:

$$\Sigma^{\boldsymbol{\alpha}} = \left\{\sigma : V \to \{1, \ldots, q\} \,\middle|\, |\sigma^{-1}(i)| = \lfloor \alpha_i n \rfloor \text{ for } i = 1, \ldots, q\right\},$$

and denote the partition function restricted to these configurations by:

$$Z_G^{\boldsymbol{\alpha}} = \textstyle\sum_{\sigma \in \Sigma^{\boldsymbol{\alpha}}} w_G(\sigma).$$

Let $\mathcal{G}$ denote the uniform distribution over Δ-regular graphs. Denote the exponent of the first moment as:

$$\Psi_1(\boldsymbol{\alpha}) := \Psi_1^{\mathbf{B}}(\boldsymbol{\alpha}) := \lim_{n \to \infty} \frac{1}{n} \log \mathbf{E}_{\mathcal{G}}\big[Z_G^{\boldsymbol{\alpha}}\big]. \tag{2}$$

The expression for Ψ_1 can be found in the full version of this paper, see [15, Section 4].

Those $\boldsymbol{\alpha}$ which are global maxima of Ψ_1 we refer to as *dominant* phases. We will see in Section 3.2 that the candidates for dominant phases correspond to stable fixpoints of the so-called tree recursions. There will be two phases of particular interest; we refer to these phases as the disordered phase and the ordered phase. The disordered phase is the uniform vector $\boldsymbol{\alpha} = (1/q, \ldots, 1/q)$. The ordered phase refers to a phase with one color dominating in the following sense: one coordinate is equal to $a > 1/q$ and the other $q-1$ coordinates are equal to $(1-a)/(q-1)$. Due to the symmetry of the Potts model, when the ordered phase dominates, in fact, the q symmetric ordered phases dominate. These ordered phases have a specific $a = a(B)$ which corresponds to a fixpoint of the tree recursions. The exact definition of this marginal a is not important at this stage, and hence we defer its explicit definition to a more detailed discussion which can be found in the full version of this paper, see [15, Section 8].

One of the difficulties for the Potts model is that the nature of the uniqueness/non-uniqueness phase transition on $\mathbb{T}_\Delta$ is inherently different from that of the Ising model. The ferromagnetic Ising model undergoes a second-order phase transition on $\mathbb{T}_\Delta$ which manifests itself on random Δ-regular graphs in the following manner. In the uniqueness region the disordered phase dominates, and in the non-uniqueness region the 2 ordered phases dominate.

In contrast, the ferromagnetic Potts model undergoes a first-order phase transition at the critical activity $\mathfrak{B}_u$. For $B < \mathfrak{B}_u$ there is a unique Gibbs measure on $\mathbb{T}_\Delta$. For $B \geq \mathfrak{B}_u$ there are multiple Gibbs measures on $\mathbb{T}_\Delta$, however there is a second critical activity $\mathfrak{B}_o$ corresponding to the disordered/ordered phase transition: for $B \leq \mathfrak{B}_o$ the disordered phase dominates, and for $B \geq \mathfrak{B}_o$ the ordered phases dominate (and at the critical point $\mathfrak{B}_o$ all of these $q+1$ phases dominate).

We present a detailed picture of the phase diagram for the ferromagnetic Potts model. Previously, Häggström [24] established the uniqueness threshold $\mathfrak{B}_u$ by studying percolation in the random cluster representation. In addition, Dembo et al. [11, 12] studied the ferromagnetic Potts model (including the case with an external field) and proved that for $B > \mathfrak{B}_u$, either the disordered or the q ordered phases are dominant, but they did not establish the precise regions where each phase dominates. For the simpler case of the complete graph (known as the Curie-Weiss model), [9] detailed the phase diagram.

Häggström [24] established that the uniqueness/non-uniqueness threshold for the infinite tree $\mathbb{T}_\Delta$ occurs at $\mathfrak{B}_u$ which is the unique value of B for which the following polynomial has a double root in $(0, 1)$:

$$(q-1)x^{\Delta} + (2-B-q)\,x^{\Delta-1} + Bx - 1. \tag{3}$$

The disordered phase is dominant in the uniqueness region and continues to dominate until the following activity (which was considered by Peruggi et al. [33]):

$$\mathfrak{B}_o := \frac{q-2}{(q-1)^{(1-2/\Delta)}-1}. \tag{4}$$

Finally, Häggström [24] considers the following activity $\mathfrak{B}_{rc}$, which he conjectures is a (second) threshold for uniqueness of the random-cluster model, defined as:

$$\mathfrak{B}_{rc} := 1 + \frac{q}{\Delta-2}.$$

Note, $\mathfrak{B}_u < \mathfrak{B}_o < \mathfrak{B}_{rc}$.

We prove the following picture for the phase diagram for the ferromagnetic Potts model (the proof can be found in the full version [15, Section 8]). Note, to prove that a function has a local maximum at a critical point, a standard approach is to show that its Hessian is negative definite. We often need this stronger condition in our proofs, hence we call such a critical point a *Hessian local maximum*. Moreover, those dominant phases $\boldsymbol{\alpha}$ where the Hessian of Ψ_1 is negative definite are called *Hessian dominant* phases. Note that dominant phases always exist but a dominant phase can fail to be Hessian (when some eigenvalue of the underlying Hessian is equal to zero). In Section 3.2, we give an alternative formulation of the Hessian condition in terms of the local stability of fixpoints of the tree recursions.

▶ **Theorem 3.** *For all $q \geq 3$ and $\Delta \geq 3$, for the ferromagnetic Potts model the following holds at activity B:*

$B < \mathfrak{B}_u$**:** *There is a unique infinite-volume Gibbs measure on $\mathbb{T}_\Delta$. The disordered phase is Hessian dominant phase, and there are no other local maxima of Ψ_1.*

$\mathfrak{B}_u < B < \mathfrak{B}_{rc}$*: The local maxima of Ψ_1 are the disordered phase $\mathbf{u}$ and the q ordered phases (the ordered phases are permutations of each other). All of these $q+1$ phases are Hessian local maxima. Moreover:*

$\mathfrak{B}_u < B < \mathfrak{B}_o$: *The disordered phase is Hessian dominant.*
$B = \mathfrak{B}_o$: *Both the disordered phase and the ordered phases are Hessian dominant.*
$\mathfrak{B}_o < B < \mathfrak{B}_{rc}$: *The ordered phases are Hessian dominant.*

$B \geq \mathfrak{B}_{rc}$: *The q ordered phases (which are permutations of each other) are Hessian dominant. For $B > \mathfrak{B}_{rc}$ there are no other local maxima of Ψ_1.*

2.3 Swendsen-Wang Algorithm

An algorithm of particular interest for the ferromagnetic Potts model is the Swendsen-Wang algorithm. The Swendsen-Wang algorithm is an ergodic Markov chain whose stationarity distribution is the Gibbs distribution. It utilizes the random-cluster representation to overcome potential "bottlenecks" for rapid mixing that are expected to arise in the non-uniqueness region. As a consequence of the above picture for the phase diagram on the infinite tree $\mathbb{T}_\Delta$ and our tools for analyzing random regular graphs, we can prove torpid mixing of the Swendsen-Wang algorithm at activities near the disordered/ordered phase transition point $\mathfrak{B}_o$. (Torpid mixing means that the mixing time is exponentially slow.)

The Swendsen-Wang algorithm utilizes the random cluster representation of the Potts model to potentially overcome bottlenecks that obstruct the simpler Glauber dynamics. It is formally defined as follows. From a configuration $X_t \in \Omega$:

- Let M be the set of monochromatic edges in X_t.
- For each edge $e \in M$, delete it with probability $1/B$. Let M' denote the set of monochromatic edges that were not deleted.
- In the graph (V, M'), for each connected component, choose a color uniformly at random from $[q]$ and assign all vertices in that component the chosen color. Let X_{t+1} denote the resulting spin configuration.

There are few results establishing rapid mixing of the Swendsen-Wang algorithm beyond what is known for the Glauber dynamics, see [37] for recent progress showing rapid mixing on the 2-dimensional lattice. However, there are several results establishing torpid mixing of the Swendsen-Wang algorithm at a critical value for the q-state ferromagnetic Potts model: on the complete graph ($q \geq 3$) [21], on Erdös-Rényi random graphs ($q \geq 3$) [8], and on the d-dimensional integer lattice $\mathbb{Z}^d$ (q sufficiently large) [3, 4].

Using our detailed picture of the phase diagram of the ferromagnetic Potts model and our generic second moment analysis for ferromagnetic models on random regular graphs which we explain in a moment, we establish torpid mixing on random Δ-regular graphs at the phase coexistence point $\mathfrak{B}_o$.

▶ **Theorem 4.** *For all $\Delta \geq 3$ and $q \geq 2\Delta/\log\Delta$, with probability $1-o(1)$ over the choice of a random Δ-regular graph, for the ferromagnetic Potts model with $B = \mathfrak{B}_o$, the Swendsen-Wang algorithm has mixing time $\exp(\Omega(n))$.*

3 Results for Ferromagnetic Models

3.1 Second Moment and Bethe Prediction Results

We analyze the Gibbs distribution on random Δ-regular graphs using second moment arguments. The challenging aspect of the second moment is determining the phase that dominates, as we will describe more precisely momentarily. In a straightforward analysis of the second moment, this reduces to an optimization problem over q^4 variables for a complicated expression. Even for $q = 2$ tackling this requires significant effort (see, for example, [32] for the hard-core model).

In a recent paper [17] we analyzed antiferromagnetic systems on *bipartite* random Δ-regular graphs, to use as gadgets for inapproximability results. In that work we presented a new approach for simplifying the analysis of the second moment for antiferromagnetic models using the theory of matrix norms. In this paper we extend that approach using the theory of matrix norms to analyze the second moment for random Δ-regular graphs (non-bipartite) for ferromagnetic systems. We obtain a short, elegant proof that the exponential order of the second moment is twice the exponential order of the first moment.

Denote the leading term of the second moment as

$$\Psi_2(\boldsymbol{\alpha}) := \Psi_2^{\mathbf{B}}(\boldsymbol{\alpha}) := \lim_{n\to\infty} \frac{1}{n} \log \mathbf{E}_{\mathcal{G}}\big[\,(Z_G^{\boldsymbol{\alpha}})^2\,\big]. \tag{5}$$

Our main technical result is the analysis of the second moment. In particular, we will relate the maximum of the second moment to the maximum of the first moment. To analyze the second moment we need to determine the phase $\boldsymbol{\alpha}$ that maximizes Ψ_2. We first show how to reexpress the maximum of Ψ_1 in a form that can be readily expressed in terms of matrix norms. The details are given in [15, Section 5.1] of the full version of this paper. Then, using the Cholesky decomposition of the interaction matrix $\mathbf{B}$ and properties of matrix norms we show that the maximum of Ψ_2 equals the value of a function at a tensor product of the dominant phases of the first moment. From there, we obtain the following theorem, whose proof can be found in [15, Section 5.2] of the full version.

▶ **Theorem 5.** *For a ferromagnetic model with interaction matrix* $\mathbf{B}$*,*

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

In particular, for dominant $\boldsymbol{\alpha}$*,* $\Psi_2(\boldsymbol{\alpha}) = 2\Psi_1(\boldsymbol{\alpha})$*.*

Combining Theorem 5 with an elaborate variance analysis known as the small subgraph conditioning method allows us to prove concentration for $Z_G^{\boldsymbol{\alpha}}$ (see Lemma 10). In particular, we verify the so-called *Bethe prediction* for general ferromagnetic models on random Δ-regular graphs, which is captured in our setting by equation (6) in the following theorem.

▶ **Theorem 6.** *Let* $\mathbf{B}$ *specify a ferromagnetic model. Then, if there exists a Hessian dominant phase, it holds that*

$$\lim_{n\to\infty} \frac{1}{n} \mathbf{E}_{\mathcal{G}}[\log Z_G] = \lim_{n\to\infty} \frac{1}{n} \log \mathbf{E}_{\mathcal{G}}[Z_G]. \tag{6}$$

Note that for a ferromagnetic model the interaction matrix $\mathbf{B}$ is positive definite and hence the entries on the diagonal are all positive. Thus Z_G is always positive for every graph G.

Theorem 6 can be extended to general models (not necessarily ferromagnetic) on random Δ-regular graphs under the stronger assumption that there is a unique semi-translation invariant Gibbs measure on $\mathbb{T}_\Delta$. In this setting, one also obtains the analogue of Theorem 5 and as a consequence concentration for $Z_G^{\boldsymbol{\alpha}}$ for the (unique) dominant phase $\boldsymbol{\alpha}$, which can be used to verify (in complete analogy) equation (6), see [15, Section 11.2] in the full version for details and a more thorough discussion.

3.2 Connection to Tree Recursions

As a consequence of Theorem 5, to analyze ferromagnetic models on random regular graphs, one only needs to analyze the first moment. To simplify the analysis of the first moment, we establish the following connection to the so-called tree recursions. An analogous connection was established in [17] for antiferromagnetic models on random bipartite Δ-regular graphs.

A key concept are the following recursions corresponding to the partition function on trees, and hence we refer to them as the (depth one) *tree recursions*:

$$\widehat{R}_i \propto \Big(\sum_{j=1}^{q} B_{ij} R_j\Big)^{\Delta-1} \tag{7}$$

The fixpoints of the tree recursions are those $\mathbf{R} = (R_1, \dots, R_q)$ such that: $\widehat{R}_i \propto R_i$ for all $i \in [q]$. We refer to a fixpoint $\mathbf{R}$ of the tree recursions as *Jacobian attractive* if the Jacobian at $\mathbf{R}$ has spectral radius less than 1. We prove the following theorem detailing the connections between the tree recursions and the critical points of the partition function for random regular graphs.

▶ **Theorem 7.** *Assume that the model is ferromagnetic. Jacobian attractive fixpoints of the (depth one) tree recursions are in one-to-one correspondence with the Hessian local maxima of* Ψ_1.

The above connection fails for antiferromagnetic models, e.g., for the antiferromagnetic Potts model the uniform fixpoint is a global maximum but it is not a stable fixpoint of the tree recursions for small enough temperature. (In fact, for antiferromagnetic models every solution of the tree recursions is a local maximum, see [15, Remark 3] in the full version.)

Using the above connection we establish the detailed picture for the dominant phases of the ferromagnetic Potts model as stated in Theorem 3.

3.3 Organization

In the following section we prove Theorem 2 showing #BIS-hardness for the Potts model in the ordered region. We also give the proof of Theorem 6 for the Bethe prediction in ferromagnetic models on random Δ-regular graphs in Section 5. The proofs of Theorems 3, 4, 5, 7 are given in the full version of this paper [15]. Specifically, in [15, Section 5] we analyze the second moment and thereby prove Theorem 5 for ferromagnetic models. In [15, Section 10] we prove Theorem 4 establishing torpid mixing of the Swendsen-Wang algorithm at the critical value $B = \mathfrak{B}_o$. We prove the connection between Jacobian attractive fixpoints of the tree recursions and the Hessian local maxima of Ψ_1 in [15, Section 6] and hence obtain Theorem 7. We then use this connection to prove Theorem 3 detailing the phase diagram in [15, Section 8].

4 #BIS-hardness for Potts

We first give a rough description of our reduction. We will construct a gadget G which is a balanced, bipartite graph on $(2 + o(1))n$ vertices. There will be $m' = O(n^{1/8})$ vertices on each side of G which will have degree $\Delta - 1$, the remainder have degree Δ. The key is that G behaves similarly to a random bipartite Δ-regular graph. Hence, the q ordered phases will dominate (for B above $\mathfrak{B}_o$). We will take an instance H for #FerroPotts(q,B) where H has m' vertices. We then replace each vertex in H by a gadget G. Then we will use the degree $\Delta - 1$ vertices in these gadgets to encode the edges of H, while preserving bipartiteness. The resulting graph H^G will have bounded degree Δ and the Potts model on H^G will "simulate" the Potts model on H.

The gadget G is defined by two parameters θ, ψ where $0 < \theta, \psi < 1/8$. The gadget is identical to that used by Sly [35]. The construction of the gadget G has two parts. First construct the following bipartite graph $\overline{G}$ with vertex set $V^+ \cup V^-$. For $s \in \{+, -\}$, $|V^s| = n + m'$ where m' will be defined precisely later. Take Δ random perfect matchings between V^+ and V^-. Then remove a matching of size m' from one of the Δ matchings. Call this graph $\overline{G}$. In the second stage, for each side of $\overline{G}$, partition the degree $\Delta - 1$ vertices into n^θ equal sized sets and attach to each set a $(\Delta-1)$-ary tree of depth ℓ where $\ell = \lfloor \psi \log_{\Delta-1} n \rfloor$. (Use the vertices of $\overline{G}$ as the leaves of these trees.) Hence each side contains n^θ trees of size n^ψ. (More precisely, $(\Delta - 1)^{\lfloor \theta \log_{\Delta-1} n \rfloor}$ trees of size $(\Delta - 1)^{\lfloor \psi \log_{\Delta-1} n \rfloor}$.) This defines the gadget G. For $s \in \{+, -\}$, let R^s denote the roots of the trees on side s. Notice that the roots R^s have degree $\Delta - 1$ and these will be used to encode the edges of H as described above. Note that $m' = (\Delta - 1)^{\lfloor \theta \log_{\Delta-1} n \rfloor + \lfloor \psi \log_{\Delta-1} n \rfloor}$ and $m' = O(n^{1/8})$. Finally, let $U^+ \cup U^-$ denote the vertices of degree Δ in the initial graph $\overline{G}$ and $W^+ \cup W^-$ denote the vertices of degree $\Delta - 1$ in $\overline{G}$.

Denote by $G = (V, E)$ the final graph. Recall, for a configuration $\sigma \in \Omega$, the set of vertices assigned spin i is denoted by $\sigma^{-1}(i)$. The phase of a configuration $\sigma : V \to [q]$ is defined as the dominant spin among vertices in $U = U^+ \cup U^-$:

$$Y(\sigma) := \arg\max_{i \in [q]} |\sigma^{-1}(i) \cap U|,$$

where ties are broken with an arbitrary deterministic criterion (e.g., the lowest index).

The gadget G behaves like a random bipartite Δ-regular graph because $m' \ll n$, as we will detail in the upcoming Lemma 8. Hence, since $B > \mathfrak{B}_o$, Theorem 3 implies that the q ordered phases are dominant. Therefore, we will get that for a sample σ from the Gibbs distribution, the phase of σ will be (close to) uniformly distributed over these q ordered phases. Let phase i refer to the ordered phase where spin i is the majority. Once we condition on the phase for the vertices in U, say it is phase i, then each of the roots, roughly independently, will have spin i with probability $\approx p$ and spin $j \neq i$ with probability $\approx (1 - p)/(q - 1)$ where p is the probability that the root of the infinite $(\Delta - 1)$-ary tree has spin i in the ordered phase i. [1] Hence, for each of the q possible phases, we define the following product distribution on the configurations $\sigma_R : R \to [q]$. For $i \in [q]$, let

$$Q^i_R(\sigma_R) = p^{|\sigma_R^{-1}(i)|} \left(\frac{1-p}{q-1}\right)^{|R \setminus \sigma_R^{-1}(i)|}.$$

The following lemma gives the precise formulation of the aforementioned properties of the gadget and is proved using methods in [35]. The proof is given in [15, Section 9.1] of the full version.

▶ **Lemma 8.** *For every $q, \Delta \geq 3$ and $B > \mathfrak{B}_o$, there exist constants $\theta, \psi > 0$ such that the graph G satisfies the following with probability $1 - o(1)$ over the choice of the graph:*

---

[1] The ordered phase $\boldsymbol{\alpha} = (a, (1 - a)/(q - 1), \dots, (1 - a)/(q - 1))$ specifies the marginal probabilities for the root of the infinite Δ-regular tree. To account for the root having degree $\Delta - 1$ one obtains that:

$$p = \frac{a^{(\Delta-1)/\Delta}}{(a/(1-a))^{(\Delta-1)/\Delta} + (q-1)^{1/\Delta}}.$$

1. *The phases occur with roughly equal probability, so that for every phase $i \in [q]$, we have*

$$\Big|\mu_G\big(Y(\sigma) = i\big) - \frac{1}{q}\Big| \le n^{-2\theta}.$$

2. *Conditioned on the phase i, the spins of vertices in R are approximately independent, that is,*

$$\max_{\sigma_R}\Big|\frac{\mu_G\big(\sigma_R \,|\, Y = i\big)}{Q^i_R(\sigma_R)} - 1\Big| \le n^{-2\theta}.$$

With Lemma 8 at hand, we can now formally state the reduction that we sketched earlier. Let $B > \mathfrak{B}_o$. Let H be a graph on n' vertices, where $n' \le n^{\theta/4}$ and θ is as in Lemma 8. Assuming an FPRAS for the ferromagnetic Potts model on max degree Δ graphs and temperature B, we will show that we can approximate $Z_H(B^*)$, the partition function of H in the ferromagnetic Potts model with temperature B^*, where B^* will be determined shortly.

To do this, we first construct a graph H^G. First, take $|H|$ disconnected copies of the gadget G in Lemma 8 and identify each copy with a vertex $v \in H$. Denote by $\hat{H}^G$ the resulting graph, G_v the copy of the gadget associated to the vertex v in H and by R_v^+, R_v^-, R_v the images of R^+, R^-, R in the gadget G_v, respectively. Finally, we denote by R_H the set of vertices $\cup_v R_v$. We next add the edges of H in $\hat{H}^G$. To do this, fix an arbitrary orientation of the edges of H. For each oriented edge (u, v) of H, we add an edge between one vertex in R_u^+ and one vertex in R_v^-, using mutually distinct vertices for distinct edges of H. The resulting graph will be denoted by H^G. Note that H^G is bipartite and has maximum degree Δ.

For a graph H and activity $B \ge 1$, recall that $Z_H(B)$ is the partition function for the ferromagnetic Potts model at activity B on the graph H. We have the following connection:

▶ **Lemma 9.** *Let $\Delta, q \ge 3$ and $B > \mathfrak{B}_o$. There exists B^* such that the following holds*

$$\big(1 - O(n^{-\theta})\big)\frac{q^{n'} Z_{H^G}(B)}{C_H\big(Z_G(B)\big)^{n'}} \le Z_H(B^*) \le \big(1 + O(n^{-\theta})\big)\frac{q^{n'} Z_{H^G}(B)}{C_H\big(Z_G(B)\big)^{n'}},$$

where $C_H = D^{|E(H)|}$ and $D = 1 + (B-1)\Big(\frac{2p(1-p)}{(q-1)^2} + (q-2)\frac{(1-p)^2}{(q-1)^2}\Big)$.

Using Lemma 9 we can now prove that for all $\Delta \ge 3$, all $B > \mathfrak{B}_o$, it is #BIS-hard to approximate the partition function for the ferromagnetic Potts model on bipartite graphs of maximum degree Δ.

Proof of Theorem 2. Goldberg and Jerrum [20] showed that for every B it is #BIS-hard to approximate the partition function of the ferromagnetic Potts on all graphs. Fix $\Delta, q \ge 3$ and $B > \mathfrak{B}_o$ for which we intend to prove Theorem 2. Let B^* be defined by Lemma 9. We first show that an FPRAS for approximating the partition function with activity B on graphs with maximum degree Δ implies an FPRAS for approximating the partition function with activity B^* on all graphs. It will then be clear that our reduction is in fact approximation-preserving and hence the theorem will be proven.

Suppose that there exists an FPRAS for approximating the partition function with activity B on graphs with maximum degree Δ. Take an input instance H for which we would like to estimate the partition function of the Potts model at activity B^*. First generate a random gadget G using the construction defined earlier. This graph G satisfies the properties in Lemma 8 with probability $1 - o(1)$. Approximate the partition function of G at activity B within a multiplicative factor $1 \pm \varepsilon/10n'$ using our presumed FPRAS. Also, using the presumed FPRAS approximate the partition function of H^G at activity B within a multiplicative factor $1 \pm \varepsilon/2$. The bounds for $Z_H(B^*)$ in Lemma 9 are then within a factor $1 \pm \varepsilon$ for sufficiently large n, giving an FPRAS for approximating the partition function at activity B^*. This, together with the result of [20], implies an FPRAS for counting independent sets in bipartite graphs. ◀

Proof of Lemma 9. Recall that $\hat{H}^G$ are the disconnected copies of the gadgets, as defined in the construction of H^G. Note, $Z_{\hat{H}^G}(B) = \big(Z_G(B)\big)^{n'}$. Hence to prove the lemma it suffices to analyze $\frac{Z_{H^G}(B)}{Z_{\hat{H}^G}(B)}$.

For a configuration σ on H^G, for each $v \in H$, let $Y_v(\sigma)$ denote the phase of σ on G_v. Denote the vector of these phases by $\mathcal{Y}(\sigma) = (Y_v(\sigma))_{v\in H} \in [q]^H$, we refer to $\mathcal{Y}(\sigma)$ as the phase vector for σ.

For $\mathcal{U} \in [q]^H$, let $\Omega_{\mathcal{U}}$ denote the set of configurations σ on H^G where $\mathcal{Y}(\sigma) = \mathcal{U}$. Let $Z_{H^G}(\mathcal{U})$ be the partition function of H^G restricted to configurations $\sigma \in \Omega_{\mathcal{U}}$, that is,

$$Z_{H^G}(\mathcal{U}) = \sum_{\sigma\in\Omega_{\mathcal{U}}} B^{m(\sigma)},$$

where for a configuration σ, $m(\sigma)$ is the number of monochromatic edges under σ. We may view $\mathcal{U}$ as an assignment $V(H) \to [q]$ where $V(H)$ are the vertices in the graph H. Hence, we can consider the number of monochromatic edges in the graph H under the assignment $\mathcal{U}$, which we denote by $m(\mathcal{U})$. Recall the goal is to analyze $\frac{Z_{H^G}(B)}{Z_{\hat{H}^G}(B)}$. To this end we will analyze $\frac{Z_{H^G}(\mathcal{U})}{Z_{\hat{H}^G}(\mathcal{U})}$ for every $\mathcal{U}$ and then we will use that every $\mathcal{U}$ is (close to) equally likely in $\hat{H}^G$ which will follow from Property 1 in Lemma 8. Notice that once we fix an assignment to all of the roots in R_H then the gadgets G_v are independent of each other. Hence we have that:

$$\frac{Z_{H^G}(\mathcal{U})}{Z_{\hat{H}^G}(\mathcal{U})} = \sum_{\sigma_{R_H}} \mu_{\hat{H}^G}(\sigma_{R_H} \,|\, \mathcal{Y}(\sigma) = \mathcal{U}) \prod_{(u,v)\in E(H^G)\setminus E(\hat{H}^G)} B^{\mathbf{1}\{\sigma_{R_H}(u)=\sigma_{R_H}(v)\}}.$$

Note that $\mu_{\hat{H}^G}(\sigma_{R_H} \,|\, \mathcal{Y}(\sigma) = \mathcal{U}) = \big(1 + O(n^{-\theta})\big) \prod_{v\in V(H)} Q^{\mathcal{U}_v}_{R_v}(\sigma_{R_v})$ since $\hat{H}^G$ is a union of disconnected copies of G and in each copy of G we have Property 2 of Lemma 8. It follows that

$$\begin{aligned}\frac{Z_{H^G}(\mathcal{U})}{Z_{\hat{H}^G}(\mathcal{U})} &= \big(1 + O(n^{-\theta})\big) \sum_{\sigma_{R_H}} \prod_{v\in V(H)} Q^{\mathcal{U}_v}_{R_v}(\sigma_{R_v}) \prod_{(u,v)\in E(H^G)\setminus E(\hat{H}^G)} B^{\mathbf{1}\{\sigma_{R_H}(u)=\sigma_{R_H}(v)\}} \\ &= \big(1 + O(n^{-\theta})\big) A^{m(\mathcal{U})} D^{|E(H)|-m(\mathcal{U})},\end{aligned}$$

where A (resp. D) is the expected weight of an edge for two gadgets which have the same (resp. different) phases. Simple calculations show that

$$A = 1 + (B-1)\Big(p^2 + \frac{(1-p)^2}{q-1}\Big), \quad D = 1 + (B-1)\Big(\frac{2p(1-p)}{(q-1)^2} + (q-2)\frac{(1-p)^2}{(q-1)^2}\Big).$$

Letting $B^* = A/D$ and $C_H = D^{|E(H)|}$, we obtain

$$\frac{Z_{H^G}(\mathcal{U})}{Z_{\hat{H}^G}(\mathcal{U})} = \big(1 + O(n^{-\theta})\big)(B^*)^{m(\mathcal{U})} C_H. \tag{8}$$

Property 1 in Lemma 8 gives that for every $\mathcal{U}$ it holds that

$$\big(1 - O(n^{-\theta})\big) q^{-n'} \le \Big(\frac{1}{q} - n^{-2\theta}\Big)^{n'} \le \frac{Z_{\hat{H}^G}(\mathcal{U})}{Z_{\hat{H}^G}} \le \Big(\frac{1}{q} + n^{-2\theta}\Big)^{n'} \le \big(1 + O(n^{-\theta})\big) q^{-n'}. \tag{9}$$

We also have

$$Z_{H^G}(B) = \sum_{\mathcal{U}} Z_{H^G}(\mathcal{U}) = \sum_{\mathcal{U}} \frac{Z_{H^G}(\mathcal{U})}{Z_{\hat{H}^G}(\mathcal{U})} Z_{\hat{H}^G}(\mathcal{U}) = Z_{\hat{H}^G} \sum_{\mathcal{U}} \frac{Z_{H^G}(\mathcal{U})}{Z_{\hat{H}^G}(\mathcal{U})} \frac{Z_{\hat{H}^G}(\mathcal{U})}{Z_{\hat{H}^G}}. \tag{10}$$

Using the estimates (8), (9) in (10), we obtain

$$\big(1 - O(n^{-\theta})\big) q^{-n'} C_H Z_H(B^*) \leq \frac{Z_{H^G}(B)}{Z_{\hat{H}^G}(B)} \leq \big(1 + O(n^{-\theta})\big) q^{-n'} C_H Z_H(B^*).$$

The result follows after observing that $Z_{\hat{H}^G}(B) = \big(Z_G(B)\big)^{n'}$ and rearranging the inequality. ◀

5 Bethe Prediction for Ferromagnetic Models on Random Δ-regular Graphs

5.1 Small Subraph Conditioning Method

By Theorem 5, we have that for the random variable $Z_G^{\boldsymbol{\alpha}}$, when $\boldsymbol{\alpha}$ is a global maximizer of Ψ_1, the exponential order of its second moment is twice the exponential order of its first moment. This is not sufficient however to obtain high probability results, since it turns out that, in the limit $n \to \infty$, the ratio of the second moment to the square of the first moment converges to a constant greater than 1. Hence, the second moment method fails to give statements that hold with high probability over a uniform random Δ-regular graph. More specifically, to obtain our results we need sharp lower bounds on the partition function which hold for almost all Δ-regular graphs. In the setting we described, the second moment method only implies the existence of a graph which satisfies the desired bounds and even there in a not sufficiently strong form.

For random Δ-regular graph ensembles, the standard way to circumvent this failure is to use the small subgraph conditioning method of Robinson and Wormald [34]. While the method is quite technical, its application is relatively streamlined when employed in the right framework. The method was first used for the analysis of spin systems in the work of [32] for the hard-core model and subsequently in [35], [16]. In [17], we extended the approach to q-spin models for all $q \geq 2$, where the major technical obstacle was the computation of certain determinants which arise in the computation of the moments' asymptotics. While the arguments there are for random *bipartite* Δ-regular graphs, the approach extends in a straightforward manner to random Δ-regular graphs.

We defer the details of the application of the method in the present setting to the full version of the paper, see [15, Section 11.1]. We state here the following lemma which is the final outcome of the method.

▶ **Lemma 10.** *For every ferromagnetic model* $\mathbf{B}$*, if* $\boldsymbol{\alpha}$ *is a Hessian dominant phase (c.f. Section 3.2) with probability* $1 - o(1)$ *over the choice of the graph* $G \sim \mathcal{G}(n, \Delta)$*, it holds that* $Z_G^{\boldsymbol{\alpha}} \geq \frac{1}{n}\mathbf{E}\big[Z_G^{\boldsymbol{\alpha}}\big]$.

5.2 Proof of Theorem 6

Using Lemma 10, the proof of Theorem 6 is straightforward.

Proof of Theorem 6. Let $\boldsymbol{\alpha}$ be a Hessian dominant phase, whose existence is guaranteed by the assumptions. By Lemma 10, with probability $1 - o(1)$ over the choice of the graph, we have $Z_G^{\boldsymbol{\alpha}} \geq \frac{1}{n}\mathbf{E}\big[Z_G^{\boldsymbol{\alpha}}\big]$, which implies $\frac{1}{n} \log Z_G \geq \Psi_1(\boldsymbol{\alpha}) + o(1)$.

Moreover, since the model is ferromagnetic, for Δ-regular graphs G with n vertices, $\frac{1}{n}\log Z_G \geq C$ for some constant $C > -\infty$ (explicitly, one can take $C := \frac{\Delta}{2}\log\max_{i\in[q]} B_{ii}$, see the remarks after Theorem 6). We thus obtain

$$\liminf_{n\to\infty} \frac{1}{n}\mathbf{E}_{\mathcal{G}}[\log Z_G] \geq \liminf_{n\to\infty} \big[(1-o(1))\Psi_1(\boldsymbol{\alpha}) + o(1)C\big] = \Psi_1(\boldsymbol{\alpha}).$$

By Jensen's inequality, we also have

$$\limsup_{n\to\infty} \frac{1}{n}\mathbf{E}_{\mathcal{G}}[\log Z_G] \leq \lim_{n\to\infty} \frac{1}{n}\log \mathbf{E}_{\mathcal{G}}[Z_G].$$

All that remains to show is that $\frac{1}{n}\log \mathbf{E}_{\mathcal{G}}[Z_G] = \Psi_1(\boldsymbol{\alpha}) + o(1)$. This is straightforward; if we decompose Z_G as $Z_G = \sum_{\boldsymbol{\alpha}'} Z_G^{\boldsymbol{\alpha}'}$, we obtain $\exp(o(n))\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha}}] \geq \mathbf{E}_{\mathcal{G}}[Z_G] \geq \mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha}}]$. Note the $\exp(o(n))$ is there to allow for dominant phases which are not Hessian.

This concludes the proof. ◀

References

1. G. Bennett. Schur multipliers. *Duke Mathematical Journal*, 44(3):603–639, 1977.
2. A. A. Bulatov, M. Dyer, L. A. Goldberg, M. Jerrum, and C. McQuillan. The expressibility of functions on the boolean domain, with applications to counting CSPs. *Journal of the ACM*, 60(5):Article No. 32, October 2013.
3. C. Borgs, J. T. Chayes, A. Frieze, J. H. Kim, P. Tetali, E. Vigoda, and V. H. Vu. Torpid mixing of some Monte Carlo Markov chain algorithms in statistical physics. In *Proceedings of the 40th IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 218–229, 1999.
4. C. Borgs, J. T. Chayes, and P. Tetali. Tight bounds for mixing of the Swendsen-Wang algorithm at the Potts transition point. *Probability Theory and Related Fields*, 152(3-4):509–557, 2012.
5. 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.
6. J.-Y. Cai, A. Galanis, L. A. Goldberg, H. Guo, M. Jerrum, D. Štefankovič, and E. Vigoda. #BIS-hardness for 2-spin systems on bipartite bounded degree graphs in the tree non-uniqueness region. Preprint available from the arXiv at: http://arxiv.org/abs/1311.4451
7. X. Chen, M. E. Dyer, L. A. Goldberg, M. Jerrum, P. Lu, C. McQuillan, and D. Richerby. The complexity of approximating conservative counting CSPs. In *Proceedings of the Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 148–159, 2013.
8. C. Cooper and A. M. Frieze. Mixing properties of the Swendsen-Wang process on classes of graphs. *Random Structures and Algorithms*, 15(3-4):242–261, 1999.
9. M. Costeniuc, R. S. Ellis, and H. Touchette. Complete analysis of phase transitions and ensemble equivalence for the Curie-Weiss-Potts model. *Journal of Mathematical Physics*, 46(6):paper 063301, 2005.
10. A. Dembo and A. Montanari. Ising models on locally tree-like graphs. *The Annals of Applied Probability*, 20(2):565–592, 2010.
11. A. Dembo, A. Montanari, A. Sly, and N. Sun. The replica symmetric solution for Potts models on d-regular graphs. *Communications in Mathematical Physics* (to appear). Preprint is available from the arXiv at: http://arxiv.org/abs/1207.5500.
12. A. Dembo, A. Montanari, and N. Sun. Factor models on locally tree-like graphs. *The Annals of Applied Probability* (to appear). Preprint is available from the arXiv at: http://arxiv.org/abs/1110.4821.

13 M. E. Dyer, L. A. Goldberg, C. S. Greenhill, and M. Jerrum. The relative complexity of approximate counting problems. *Algorithmica*, 38(3):471–500, 2003.

14 A. Galanis, Q. Ge, D. Štefankovič, E. Vigoda, and L. Yang. Improved inapproximability results for counting independent sets in the hard-core model. *Random Struct. Algorithms* (to appear). Preprint is available from the arXiv at: http://arxiv.org/abs/1105.5131

15 A. Galanis, D. Štefankovič, E. Vigoda, and L. Yang. Ferromagnetic Potts model: Refined #BIS-hardness and related results. Full version of this paper is available from the arXiv at: http://arxiv.org/abs/1311.4839

16 A. Galanis, D. Štefankovič, and E. Vigoda. Inapproximability of the partition function for the antiferromagnetic Ising and hard-core models. Preprint is available from the arXiv at: http://arxiv.org/abs/1203.2226

17 A. Galanis, D. Štefankovič, and E. Vigoda. Inapproximability for antiferromagnetic spin systems in the tree non-uniqueness region. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing (STOC)*, 823–831, 2014. Full version is available from the arXiv at: http://arxiv.org/abs/1305.2902

18 H.-O. Georgii. *Gibbs measures and phase transitions*, volume 9 of *de Gruyter Studies in Mathematics*. Walter de Gruyter & Co., Berlin, second edition, 2011.

19 A. Gerschenfeld and A. Montanari. Reconstruction for models on random graphs. In *Proceedings of the 48th Annual Symposium on Foundations of Computer Science* (FOCS), 194-204, 2007.

20 L. A. Goldberg and M. Jerrum. Approximating the partition function of the ferromagnetic Potts model. *Journal of the ACM*, 59(5):Article No. 25, 2012.

21 V. K. Gore and M. R. Jerrum. The Swendsen-Wang process does not always mix rapidly. *Journal of Statistical Physics*, 97(1-2):67–86, 1999.

22 L. A. Goldberg, M. Jerrum, and M. Paterson. The computational complexity of two-state spin systems. *Random Struct. Algorithms*, 23(2):133–154, 2003.

23 C. Greenhill. The complexity of counting colourings and independent sets in sparse graphs and hypergraphs. *Comput. Complex.*, 9(1):52–72, 2000.

24 O. Häggström. The random-cluster model on a homogeneous tree. *Probability Theory and Related Fields*, 104(2):231–253, 1996.

25 S. Janson. Random regular graphs: Asymptotic distributions and contiguity. *Combinatorics, Probability & Computing*, 4:369–405, 1995.

26 M. Jerrum and A. Sinclair. Polynomial-time approximation algorithms for the Ising model. *SIAM Journal on Computing*, 22(5):1087–1116, 1993.

27 F. P. Kelly. Loss networks. *The Annals of Applied Probability*, 1(3):319–378, 1991.

28 L. Li, P. Lu, and Y. Yin. Correlation decay up to uniqueness in spin systems. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 47–66. 2013.

29 F. Martinelli, A. Sinclair, and D. Weitz. Fast mixing for independent sets, colorings and other models on trees. In *Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 456–465. 2004.

30 R. Montenegro and P. Tetali. Mathematical aspects of mixing times in Markov chains. *Foundations and Trends in Theoretical Computer Science*, 1(3):237–354, 2006.

31 E. Mossel and A. Sly. Exact thresholds for Ising-Gibbs samplers on general graphs. *Annals of Probability*, 41(1):294–328, 2013.

32 E. Mossel, D. Weitz, and N. Wormald. On the hardness of sampling independent sets beyond the tree threshold. *Probability Theory and Related Fields*, 143(3-4):401–439, 2009.

33 F. Peruggi, F. Di Liberto, and G. Monroy. Phase diagrams of the q-state Potts model on Bethe lattices. *Physica A*, 141(1):151–186, 1987.

34 R. W. Robinson and N. C. Wormald. Almost all regular graphs are Hamiltonian. *Random Structures and Algorithms*, 5(2):363–374, 1994.

35 A. Sly. Computational transition at the uniqueness threshold. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 287–296, 2010.

36 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.

37 M. Ullrich. Rapid mixing of Swendsen-Wang dynamics in two dimensions. Ph.D. Thesis, Universität Jena, Germany, 2012. The thesis is available from the arXiv at: http://arxiv.org/abs/1212.4908

38 L. G. Valiant. The complexity of enumeration and reliability problems. *SIAM J. Computing*, 8(3):410–421, 1979.

39 D. Weitz. Counting independent sets up to the tree threshold. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC)*, 140–149, 2006.