Metastability of the Potts Ferromagnet on Random Regular Graphs

Amin Coja-Oghlan ✉
Faculty of Computer Science, TU Dortmund, Germany

Andreas Galanis ✉
Department of Computer Science, University of Oxford, UK

Leslie Ann Goldberg ✉
Department of Computer Science, University of Oxford, UK

Jean Bernoulli Ravelomanana ✉
Faculty of Computer Science, TU Dortmund, Germany

Daniel Štefankovič ✉
Department of Computer Science, Univerity of Rochester, NY, USA

Eric Vigoda ✉
Computer Science, University of California Santa Barbara, CA, USA

─── **Abstract** ───

We study the performance of Markov chains for the q-state ferromagnetic Potts model on random regular graphs. While the cases of the grid and the complete graph are by now well-understood, the case of random regular graphs has resisted a detailed analysis and, in fact, even analysing the properties of the Potts distribution has remained elusive. It is conjectured that the performance of Markov chains is dictated by metastability phenomena, i.e., the presence of "phases" (clusters) in the sample space where Markov chains with local update rules, such as the Glauber dynamics, are bound to take exponential time to escape, and therefore cause slow mixing. The phases that are believed to drive these metastability phenomena in the case of the Potts model emerge as local, rather than global, maxima of the so-called Bethe functional, and previous approaches of analysing these phases based on optimisation arguments fall short of the task.

Our first contribution is to detail the emergence of the metastable phases for the q-state Potts model on the d-regular random graph for all integers $q, d \geq 3$, and establish that for an interval of temperatures, delineated by the uniqueness and a broadcasting threshold on the d-regular tree, the two phases coexist. The proofs are based on a conceptual connection between spatial properties and the structure of the Potts distribution on the random regular graph, rather than complicated moment calculations. This significantly refines earlier results by Helmuth, Jenssen, and Perkins who had established phase coexistence for a small interval around the so-called ordered-disordered threshold (via different arguments) that applied for large q and $d \geq 5$.

Based on our new structural understanding of the model, we obtain various algorithmic consequences. We first complement recent fast mixing results for Glauber dynamics by Blanca and Gheissari below the uniqueness threshold, showing an exponential lower bound on the mixing time above the uniqueness threshold. Then, we obtain tight results even for the non-local and more elaborate Swendsen-Wang chain, where we establish slow mixing/metastability for the whole interval of temperatures where the chain is conjectured to mix slowly on the random regular graph. The key is to bound the conductance of the chains using a random graph "planting" argument combined with delicate bounds on random-graph percolation.

1 Introduction

1.1 Motivation

Spin systems on random graphs have turned out to be a source of extremely challenging problems at the junction of mathematical physics and combinatorics [36, 37]. Beyond the initial motivation of modelling disordered systems, applications have sprung up in areas as diverse as computational complexity, coding theory, machine learning and even screening for infectious diseases; e.g. [1, 14, 22, 34, 38, 40, 41]. Progress has been inspired largely by techniques from statistical physics, which to a significant extent still await a rigorous justification. The physicists' sophisticated but largely heuristic tool is the Belief Propagation message passing scheme in combination with a functional called the Bethe free energy [35]. Roughly speaking, the fixed points of Belief Propagation are conjectured to correspond to the "pure states" of the underlying distribution, with the Bethe functional gauging the relative weight of the different pure states. Yet at closer inspection matters are actually rather complicated. For instance, the system typically possesses spurious Belief Propagation fixed points without any actual combinatorial meaning, while other fixed points need not correspond to metastable states that attract dynamics such as the Glauber Markov chain [11, 15]. Generally, the mathematical understanding of the connection between Belief Propagation and dynamics leaves much to be desired.

In this paper we investigate the ferromagnetic Potts model on the random regular graph. Recall, for an integer $q \geq 3$ and real $\beta > 0$, the Potts model on a graph $G = (V, E)$ corresponds to a probability distribution $\mu_{G,\beta}$ over all possible configurations $[q]^V$, commonly referred to as the Boltzmann/Gibbs distribution; the weight of a configuration σ in the distribution is defined as $\mu_{G,\beta}(\sigma) = e^{\beta \mathcal{H}_G(\sigma)}/Z_\beta(G)$ where $\mathcal{H}_G(\sigma)$ is the number of edges that are monochromatic under σ, and $Z_\beta(G) = \sum_{\tau \in [q]^V} e^{\beta \mathcal{H}_G(\tau)}$ is the normalising factor of the distribution. In physics jargon, β corresponds to the so-called inverse-temperature of the model, $\mathcal{H}_G(\cdot)$ is known as the Hamiltonian, and $Z_\beta(\cdot)$ is the partition function. Note, since $\beta > 0$, the Boltzmann distribution assigns greater weight to configurations σ where many edges join vertices of the same colour; thus, the pairwise interactions between vertices are ferromagnetic.

The Potts model on the d-regular random graph has two distinctive features. First, the local geometry of the random regular graph is essentially deterministic. For any fixed radius ℓ, the depth-ℓ neighbourhood of all but a tiny number of vertices is just a d-regular tree. Second, the ferromagnetic nature of the model precludes replica symmetry breaking, a complex type of long-range correlations [35]. Given these, it is conjectured that the model on the random regular graph has a similar behaviour to that on the clique (the so-called mean field case), and there has already been some preliminary evidence of this correspondence [4, 20, 19, 22, 29]. On the clique, the phase transitions are driven by a battle between two subsets of configurations (phases): (i) the paramagnetic/disordered phase, consisting of configurations where every colour appears roughly equal number of times, and (ii) the ferromagnetic/ordered phase, where one of the colours appears more frequently than

the others. It is widely believed that these two phases also mark (qualitatively) the same type of phase transitions for the Potts model on the random regular graph, yet this has remained largely elusive.

The main reason that this behaviour is harder to establish on the random regular graph is that it has a non-trivial global geometry which makes both the analysis of the distribution and Markov chains significantly more involved (to say the least). In particular, the emergence of the metastable states in the distribution, which can be established by way of calculus in the mean-field case, is out of reach with single-handed analytical approaches in the random regular graph and it is therefore not surprising that it has resisted a detailed analysis so far. Likewise, the analysis of Markov chains is a far more complicated task since their evolution needs to be considered in terms of the graph geometry and therefore much harder to keep track of.

Our main contribution is to detail the emergence of the metastable states, viewed as fixed points of Belief Propagation on this model, and their connection with the dynamic evolution of the two most popular Markov chains, the Glauber dynamics and the Swensen-Wang chain. We prove that these natural fixed points, whose emergence is directly connected with the phase transitions of the model, have the combinatorial meaning in terms of both the pure state decomposition of the distribution and the Glauber dynamics that physics intuition predicts they should. The proofs avoid the complicated moment calculations and the associated complex optimistion arguments that have become a hallmark of the study of spin systems on random graphs [2]. Instead, building upon and extending ideas from [3, 16], we exploit a connection between spatial mixing properties on the d-regular tree and the Boltzmann distribution. Our metastability results for the Potts model significantly refine those appearing in the literature, especially those in [22, 29] which are more relevant to this work, see Section 1.6 for a more detailed discussion.

We expect that this approach might carry over to other examples, particularly other ferromagnetic models. Let us begin by recapitulating Belief Propagation.

1.2 Belief Propagation

Suppose that $n, d \geq 3$ are integers such that dn is even and let $\mathbb{G} = \mathbb{G}(n, d)$ be the random d-regular graph on the vertex set $[n] = \{1, \dots, n\}$. For an inverse temperature parameter $\beta > 0$ and an integer $q \geq 3$ we set out to investigate the Boltzmann distribution $\mu_{\mathbb{G},\beta}$; let us write $\boldsymbol{\sigma}_{\mathbb{G},\beta}$ for a configuration drawn from $\mu_{\mathbb{G},\beta}$.

A vital step toward understanding the Boltzmann distribution is to get a good handle on the partition function $Z_\beta(\mathbb{G})$. Indeed, according to the physicsts' cavity method, Belief Propagation actually solves both problems in one fell swoop [35]. To elaborate, with each edge $e = uv$ of $\mathbb{G}$, Belief Propagation associates two messages $\mu_{\mathbb{G},\beta,u \to v}, \mu_{\mathbb{G},\beta,v \to u}$, which are probability distributions on the set $[q]$ of colours. The message $\mu_{\mathbb{G},\beta,u \to v}(c)$ is defined as the marginal probability of v receiving colour c in a configuration drawn from the Potts model on the graph $\mathbb{G} - u$ obtained by removing u. The semantics of $\mu_{\mathbb{G},\beta,v \to u}$ is analogous.

Under the assumption that the colours of far apart vertices of $\mathbb{G}$ are asymptotically independent, one can heuristically derive a set of equations that links the various messages together. For a vertex v, let ∂v be the set of neighbours of v, and for an integer $\ell \geq 1$ let $\partial^\ell v$ be the set of vertices at distance precisely ℓ from v. The *Belief Propagation equations* read

$$\mu_{\mathbb{G},\beta,v \to u}(c) = \frac{\prod_{w \in \partial v \setminus \{u\}} 1 + (\mathrm{e}^\beta - 1)\mu_{\mathbb{G},\beta,w \to v}(c)}{\sum_{\chi \in [q]} \prod_{w \in \partial v \setminus \{u\}} 1 + (\mathrm{e}^\beta - 1)\mu_{\mathbb{G},\beta,w \to v}(\chi)} \quad (uv \in E(\mathbb{G}), \ c \in [q]). \quad (1)$$

The insight behind (1) is that once we remove v from the graph, its neighbours $w \neq u$ are typically far apart from one another because $\mathbb{G}$ contains only a negligible number of short cycles. Hence, we expect that in $\mathbb{G} - v$ the spins assigned to $w \in \partial v \setminus \{u\}$ are asymptotically independent. From this assumption it is straightforward to derive the sum-product-formula (1).

A few obvious issues spring to mind. First, for large β it is not actually true that far apart vertices decorrelate. This is because at low temperature there occur q different ferromagnetic pure states, one for each choice of the dominant colour. To break the symmetry between them one could introduce a weak external field that slighly boosts a specific colour or, more bluntly, confine oneself to a conditional distribution on subspace where a specific colour dominates. In the definition of the messages and in (1) we should thus replace the Boltzmann distribution by the conditional distribution $\mu_{\mathbb{G},\beta}(\,\cdot\, | \, S)$ for a suitable $S \subseteq [q]^n$. Second, even for the conditional measure we do not actually expect (1) to hold precisely. This is because for any finite n minute correlations between far apart vertices are bound to remain.

Nonetheless, precise solutions $(\mu_{v \to u})_{uv \in E(\mathbb{G})}$ to (1) are still meaningful. They correspond to stationary points of a functional called the *Bethe free energy*, which connects Belief Propagation with the problem of approximating the partition function [44]. Given a collection $(\mu_{u \to v})_{uv \in E(\mathbb{G})}$ of probability distributions on $[q]$, the Bethe functional reads

$$
\begin{aligned}
\mathcal{B}_{\mathbb{G},\beta}\left((\mu_{u \to v})_{uv \in E(\mathbb{G})}\right) = \frac{1}{n} \sum_{v \in V(\mathbb{G})} \log\left[\sum_{c \in [q]} \prod_{w \in \partial v} 1 + (\mathrm{e}^{\beta} - 1)\mu_{w \to v}(c)\right] \\
- \frac{1}{n} \sum_{vw \in E(\mathbb{G})} \log\left[1 + (\mathrm{e}^{\beta} - 1)\sum_{c \in [q]} \mu_{v \to w}(c)\mu_{w \to v}(c)\right].
\end{aligned}
\tag{2}
$$

According to the cavity method the maximum of $\mathcal{B}_{\mathbb{G},\beta}\left((\mu_{u \to v})_{uv \in E(\mathbb{G})}\right)$ over all solutions $(\mu_{u \to v})_{uv \in E(\mathbb{G})}$ to (1) should be asymptotically equal to $\log Z_{\beta}(\mathbb{G})$ with high probability.

In summary, physics lore holds that the solutions $(\mu_{u \to v})_{uv \in E(\mathbb{G})}$ to (1) are meaningful because they correspond to a decomposition of the phase space $[q]^n$ into pieces where long-range correlations are absent. Indeed, these "pure states" are expected to exhibit metastability, i.e., they trap dynamics such as the Glauber Markov chain for an exponential amount of time. Moreover, the relative probabilities of the pure states are expected to be governed by their respective Bethe free energy. In the following we undertake to investigate these claims rigorously.

Before proceeding, let us mention that ferromagnetic spin systems on random graphs have been among the first models for which predictions based on the cavity method could be verified rigorously. Following seminal work by Dembo and Montanari on the Ising model [18] vindicating the "replica symmetric ansatz", Dembo, Montanari and Sun [20] studied, among other things, the Gibbs unique phase of the Potts ferromagnet on the random regular graph, and Dembo, Montanari, Sly and Sun [20] established the free energy of the model for all β (and d even). More generally, Ruozzi [39] pointed out how graph covers [43] can be used to investigate the partition function of supermodular models, of which the Ising ferromagnet is an example. In addition, Barbier, Chan and Macris [4] proved that ferromagnetic spin systems on random graphs are generally replica symmetric in the sense that the multi-overlaps of samples from the Boltzmann distribution concentrate on deterministic values.

1.3 The ferromagnetic and the paramagnetic states

An obvious attempt at constructing solutions to the Belief Propagation equations is to choose identical messages $\mu_{u \to v}$ for all edges $uv \in E(\mathbb{G})$. Clearly, any solution $(\mu(c))_{c \in [q]}$ to the system

$$\mu(c) = \frac{(1 + (e^\beta - 1)\mu(c))^{d-1}}{\sum_{\chi \in [q]}(1 + (e^\beta - 1)\mu(\chi))^{d-1}} \qquad (c \in [q]) \qquad (3)$$

supplies such a "constant" solution to (1). Let $\mathcal{F}_{d,\beta}$ be the set of all solutions $(\mu(c))_{c \in [q]}$ to (3). The Bethe functional (2) then simplifies to

$$\mathcal{B}_{d,\beta}\big((\mu(c))_{c \in [q]}\big) = \log\left[\sum_{c \in [q]}\big(1 + (e^\beta - 1)\mu(c)\big)^d\right] - \frac{d}{2}\log\left[1 + (e^\beta - 1)\sum_{c \in [q]}\mu(c)^2\right]. \quad (4)$$

One obvious solution to (3) is the uniform distribution on $[q]$; we refer to that solution as paramagnetic/disordered and denote it by μ_{p}. Apart from μ_{p}, other solutions to (3) emerge as β increases for any $d \geq 3$. Specifically, let $\beta_u > 0$ be the supremum value of $\beta > 0$ where μ_{p} is the unique solution to (3).[1] Then, for $\beta = \beta_u$, one more solution μ_{f} emerges such that $\mu_{\mathrm{f}}(1) > \mu_{\mathrm{f}}(i) = \frac{1-\mu_{\mathrm{f}}(1)}{q-1}$ for $i = 2, \dots, q$, portending the emergence of a metastable state and, ultimately, a phase transition. In particular, for any $\beta > \beta_u$, a bit of calculus reveals there exist either one or two distinct solutions μ with $\mu(1) > \mu(i) = \frac{1-\mu(1)}{q-1}$ for $i = 2, \dots, q$; we denote by μ_{f} the solution of (3) which maximises the value $\mu(1)$ and refer to it as ferromagnetic/ordered. The value β_u is the so-called uniqueness threshold for the Potts model on the d-regular tree, see, e.g., [22] for a more detailed discussion and related pointers.

At the critical value

$$\beta_p = \max\left\{\beta \geq \beta_u : \mathcal{B}_{d,\beta}(\mu_{\mathrm{p}}) \geq \mathcal{B}_{d,\beta}(\mu_{\mathrm{f}})\right\} = \log\frac{q-2}{(q-1)^{1-2/d}-1}.$$

the ferromagnetic solution μ_{f} takes over from the paramagnetic solution μ_{p} as the global maximiser of the Bethe functional. For that reason, the threshold β_p is also known in the literature as the ordered-disordered threshold. Yet, up to the threshold

$$\beta_h = \log(1 + q/(d-2))$$

the paramagnetic solution remains a local maximiser of the Bethe free energy; later, in Section 2.2 we will see that β_h has a natural interpretation as a tree-broadcasting threshold (and is also a conjectured threshold for uniqueness in the random-cluster representation for the Potts model, see [28] for details).

The relevance of these critical values has been demonstrated in [22] (see also [19] for d even, and [29] for q large), where it was shown that $\frac{1}{n}\log Z_\beta(\mathbb{G})$ is asymptotically equal to $\max_\mu \mathcal{B}_{d,\beta}(\mu)$, the maximum ranging over μ satisfying (3). In particular, at the maximum it holds that $\mu = \mu_{\mathrm{p}}$ when $\beta < \beta_p$, $\mu = \mu_{\mathrm{f}}$ when $\beta > \beta_p$ and $\mu \in \{\mu_{\mathrm{p}}, \mu_{\mathrm{f}}\}$ when $\beta = \beta_p$.

---

[1] The value does not have a closed-form expression, but there is an equivalent formulation of it given by the equality $e^{\beta_u} = 1 + \inf_{y>1}\frac{(y-1)(y^{d-1}+q-1)}{y^{d-1}-y}$.

1.4 Slow mixing and metastability

To investigate the two BP solutions further and obtain connections to the dynamical evolution of the model, we need to look more closely how these two solutions $\mu_{\mathrm{p}}, \mu_{\mathrm{f}}$ manifest themselves in the random regular graph. To this end, we define for a given distribution μ on $[q]$ another distribution

$$\nu^\mu(c) = \frac{(1 + (\mathrm{e}^\beta - 1)\mu(c))^d}{\sum_{\chi \in [q]} (1 + (\mathrm{e}^\beta - 1)\mu(\chi))^d} \qquad\qquad (c \in [q]). \qquad\qquad (5)$$

Let $\nu_{\mathrm{f}} = \nu^{\mu_{\mathrm{f}}}$ and $\nu_{\mathrm{p}} = \nu^{\mu_{\mathrm{p}}}$ for brevity; of course $\nu_{\mathrm{p}} = \mu_{\mathrm{p}}$ is just the uniform distribution. The distributions ν_{f} and ν_{p} represent the expected Boltzmann marginals within the pure states corresponding to μ_{f} and μ_{p}. Indeed, the r.h.s. of (5) resembles that of (3) except that the exponents read d rather than $d - 1$. This means that we pass from messages, where we omit one specific endpoint of an edge from the graph, to actual marginals, where all d neighbours of a vertex are present. For small $\varepsilon > 0$, it will therefore be relevant to consider the sets of configurations

$$S_{\mathrm{f}}(\varepsilon) = \left\{ \sigma \in [q]^n : \sum_{c \in [q]} \left| |\sigma^{-1}(c)| - n\nu_{\mathrm{f}}(c) \right| < \varepsilon n \right\},$$

$$S_{\mathrm{p}}(\varepsilon) = \left\{ \sigma \in [q]^n : \sum_{c \in [q]} \left| |\sigma^{-1}(c)| - n\nu_{\mathrm{p}}(c) \right| < \varepsilon n \right\},$$

whose colour statistics are about $n\nu_{\mathrm{f}}$ and $n\nu_{\mathrm{p}}$, respectively; i.e., in S_{p}, all colours appear with roughly equal frequency, whereas in S_{f} colour 1 is favoured over the other $q - 1$ colours (which appear with roughly equal frequency).

We are now in position to state our main result for Glauber dynamics. Recall that, for a graph $G = (V, E)$, Glauber is initialised at a configuration $\sigma_0 \in [q]^V$; at each time step $t \geq 1$, Glauber draws a vertex uniformly at random and obtains a new configuration σ_t by updating the colour of the chosen vertex according to the conditional Boltzmann distribution given the colours of its neighbours. It is a well-known fact that Glauber converges in distribution to $\mu_{G,\beta}$; the mixing time of the chain is defined as the maximum number of steps t needed to get within total variation distance $\leq 1/4$ from $\mu_{G,\beta}$, where the maximum is over the choice of the initial configuration σ_0, i.e., the quantity $\max_{\sigma_0} \min\{t : d_{\mathrm{TV}}(\sigma_t, \mu_{G,\beta}) \leq 1/4\}$.

For metastability, we will consider Glauber launched from a random configuration from a subset $S \subseteq [q]^V$ of the state space. More precisely, let us denote by $\mu_{G,\beta,S} = \mu_{G,\beta}(\cdot \mid S)$ the conditional Boltzmann distribution on S. We call S a *metastable state for Glauber dynamics on G* if there exists $\delta > 0$ such that

$$\mathbb{P}\left[\min\{t : \sigma_t \notin S\} \leq \mathrm{e}^{\delta|V|} \mid \sigma_0 \sim \mu_{G,\beta,S} \right] \leq \mathrm{e}^{-\delta|V|}.$$

Hence, it will most likely take Glauber an exponential amount of time to escape from a metastable state.

▶ **Theorem 1.1.** *Let $d, q \geq 3$ be integers and $\beta > 0$ be real. Then, for all sufficiently small $\varepsilon > 0$, the following hold w.h.p. over the choice of $\mathbb{G} = \mathbb{G}(n, d)$.*
   (i) *If $\beta < \beta_h$, then $S_{\mathrm{p}}(\varepsilon)$ is a metastable state for Glauber dynamics on $\mathbb{G}$.*
   (ii) *If $\beta > \beta_u$, then $S_{\mathrm{f}}(\varepsilon)$ is a metastable state for Glauber dynamics on $\mathbb{G}$.*
Further, for $\beta > \beta_u$, the mixing time of Glauber is $\mathrm{e}^{\Omega(n)}$.

Thus, we can summarise the evolution of the Potts model as follows. For $\beta < \beta_u$ there is no ferromagnetic state. As β passes β_u, the ferromagnetic state S_{f} emerges first as a metastable state. Hence, if we launch Glauber from S_{f}, the dynamics will most likely remain

trapped in the ferromagnetic state for an exponential amount of time, even though the Boltzmann weight of the paramagnetic state is exponentially larger (as we shall see in the next section). At the point β_p the ferromagnetic state then takes over as the one dominating the Boltzmann distribution, but the paramagnetic state remains as a metastable state up to β_h. Note in particular that the two states coexist as metastable states throughout the interval (β_u, β_h).

The metastability for the Potts model manifests also in the evolution of the Swendsen-Wang (SW) chain, which is another popular and substantially more elaborate chain that makes non-local moves, based on the random-cluster representation of the model. For a graph $G = (V, E)$ and a configuration $\sigma \in [q]^V$, a single iteration of SW starting from σ consists of two steps.

- *Percolation step:* Let $M = M(\sigma)$ be the random edge-set obtained by adding (independently) each monochromatic edge under σ with probability $p = 1 - e^{-\beta}$.

- *Recolouring step:* Obtain the new $\sigma' \in [q]^V$ by assigning each component[2] of the graph (V, M) a uniformly random colour from $[q]$; for $v \in V$, we set σ'_v to be the colour assigned to v's component.

We define metastable states for SW dynamics analogously to above. The following theorem establishes the analogue of Theorem 1.1 for the non-local SW dynamics. Note here that SW might change the most-frequent colour due to recolouring step, so the metastability statement for the ferromagnetic phase needs to consider the set $S_f(\varepsilon)$ with its $q-1$ permutations.

▶ **Theorem 1.2.** *Let $d, q \geq 3$ be integers and $\beta > 0$ be real. Then, for all sufficiently small $\varepsilon > 0$, the following hold w.h.p. over the choice of $\mathbb{G} = \mathbb{G}(n, d)$.*

(i) *If $\beta < \beta_h$, then $S_p(\varepsilon)$ is a metastable state for SW dynamics on $\mathbb{G}$.*

(ii) *If $\beta > \beta_u$, then $S_f(\varepsilon)$ together with its $q-1$ permutations is a metastable state for SW dynamics on $\mathbb{G}$.*

Further, for $\beta \in (\beta_u, \beta_h)$, the mixing time of SW is $e^{\Omega(n)}$.

1.5 The relative weight of the metastable states

At the heart of obtaining the metastability results of the previous section is a refined understanding of the relative weight of the ferromagnetic and paramagnetic states. The following notion of non-reconstruction will be the key in our arguments; it captures the absence of long-range correlations within a set $S \subseteq [q]^n$, saying that, for any vertex v, a typical boundary configuration on $\sigma_{\partial^\ell v}$ chosen according to the conditional distribution on S does not impose a discernible bias on the colour of v (for large ℓ, n; recall, $\partial^\ell v$ is the set of all vertices at distance precisely ℓ from v). More precisely, let $\mu = \mu_{\mathbb{G}, \beta}$ and $\sigma \sim \mu$; the Boltzmann distribution exhibits *non-reconstruction given a subset* $S \subseteq [q]^n$ if for any vertex v it holds that

$$\lim_{\ell \to \infty} \limsup_{n \to \infty} \sum_{c \in [q]} \sum_{\tau \in S} \mathbb{E} \left[\mu(\tau \mid S) \times |\mu(\sigma_v = c \mid \sigma_{\partial^\ell v} = \tau_{\partial^\ell v}) - \mu(\sigma_v = c \mid S)| \right] = 0,$$

where the expectation is over the choice of the graph $\mathbb{G}$.

---

[2] Note, isolated vertices count as connected components.

▶ **Theorem 1.3.** *Let $d, q \geq 3$ be integers and $\beta > 0$ be real. The following hold for all sufficiently small $\varepsilon > 0$ as $n \to \infty$.*

(i) *For all $\beta < \beta_p$, $\mathbb{E}\left[\mu_{\mathbb{G}, \beta}(S_\mathrm{p})\right] \to 1$ and, if $\beta > \beta_u$, then $\mathbb{E}\left[\frac{1}{n} \log \mu_{\mathbb{G}, \beta}(S_\mathrm{f})\right] \to \mathcal{B}_{d, \beta}(\mu_\mathrm{f}) - \mathcal{B}_{d, \beta}(\mu_\mathrm{p})$.*

(ii) *For all $\beta > \beta_p$, $\mathbb{E}\left[\mu_{\mathbb{G}, \beta}(S_\mathrm{f})\right] \to 1/q$ and, if $\beta < \beta_h$, then $\mathbb{E}\left[\frac{1}{n} \log \mu_{\mathbb{G}, \beta}(S_\mathrm{p})\right] \to \mathcal{B}_{d, \beta}(\mu_\mathrm{p}) - \mathcal{B}_{d, \beta}(\mu_\mathrm{f})$.*

Furthermore, the Boltzmann distribution given S_p exhibits non-reconstruction if $\beta < \beta_h$ and the Boltzmann distribution given S_f exhibits non-reconstruction if $\beta > \beta_u$.

Theorem 1.3 shows that for $\beta < \beta_p$ the Boltzmann distribution is dominated by the paramagnetic state S_p for $\beta < \beta_p$. Nonetheless, at β_u the ferromagnetic state S_f and its $q - 1$ mirror images start to emerge. Their probability mass is determined by the Bethe free energy evaluated at μ_f. Further, as β passes β_p the ferromagnetic state takes over as the dominant state, with the paramagnetic state lingering on as a sub-dominant state up to β_h. Finally, both states S_p and S_f are free from long-range correlations both for the regime of β where they dominate and for those β where they are sub-dominant.

1.6 Discussion

Our slow mixing result for Glauber dynamics when $\beta > \beta_u$ (Theorem 1.1) significantly improves upon previous results of Bordewich, Greenhill and Patel [9] that applied to $\beta > \beta_u + \Theta_q(1)$. Similarly, our slow mixing result for Swendsen-Wang dynamics when $\beta \in (\beta_u, \beta_h)$ (Theorem 1.2) strengthens earlier results of Galanis, Štefankovič, Vigoda, Yang [22] which applied to $\beta = \beta_p$, and by Helmuth, Jenseen and Perkins [29] which applied for a small interval around β_p; both results applied only for q sufficiently large. To obtain our result for all integers $q, d \geq 3$, we need to carefully track how SW evolves on the random regular graph for configurations starting from the ferromagnetic and paramagnetic phases, by accounting for the percolation step via delicate arguments, whereas the approaches of [22, 29] side-stepped this analysis by considering the change in the number of monochromatic edges instead.

Our slow mixing results complement the recent fast mixing result of Blanca and Gheissari [6] for edge dynamics on the random d-regular graph that applies to all $\beta < \beta_u$. Roughly, edge dynamics is the analogue of Glauber dynamics for the random cluster representation of the Potts model (the random-cluster representation has nicer monotonicity properties). The result of [6] already implies a polynomial bound on the mixing time of SW when $\beta < \beta_u$ (due to comparison results by Ullrich that apply to general graphs [42]), and conversely our exponential lower bound on the mixing time of SW for $\beta \notin (\beta_u, \beta_h)$ implies an exponential lower bound on the mixing time of edge dynamics for $\beta \notin (\beta_u, \beta_h)$. The main open questions remaining are therefore showing whether Glauber dynamics for the Potts model mixes fast when $\beta \leq \beta_u$ and whether SW/edge-dynamics mixes fast when $\beta > \beta_h$. Extrapolating from the mean-field case (see discussion below), it is natural to conjecture that our slow mixing results are best-possible, i.e., for $\beta \leq \beta_u$, Glauber mixes rapidly and similarly, for $\beta \notin (\beta_u, \beta_h)$, SW mixes rapidly on the random regular graph.

Theorem 1.3, aside from being critical in establishing the aforementioned slow mixing and metastability results, is the first to establish for all $q, d \geq 3$ the coexistence of the ferromagnetic and paramagnetic phases for all β in the interval (β_u, β_h) and detail the logarithmic order of their relative weight in the same interval. Previous work in [22] showed coexistence for $\beta = \beta_p$ (for all $q, d \geq 3$) and [29] for β in a sub-interval of (β_u, β_h) around β_p (for large q and $d \geq 5$), see also footnote 3. We remark here that the approaches in [22, 29] establish more refined estimates on the deviations from the limiting value of the

log-partition function of the phases (in the corresponding regimes they apply), with [29] characterising in addition the limiting distribution using cluster-expansion methods. One can obtain analogous distributional characterisations for all $q, d \geq 3$ from our methods, once combined with the small subgraph conditioning method of [22]. It should be noted though the approach of [29] which goes through cluster expansion is more direct in that respect. We don't pursue such distributional results here since Theorem 1.3 is sufficient for our slow mixing results.

Together with Theorems 1.1 and 1.2, Theorem 1.3 delineates more firmly[3] the correspondence with the (simpler) mean-field case, the Potts model on the clique. In the mean-field case, there are qualitatively similar thresholds $\beta_u, \beta_p, \beta_h$ and the mixing time for Glauber and SW have been detailed for all β, even at criticality, see [7, 8, 25, 23, 17, 27, 31]. As mentioned earlier, the most tantalising question remaining open is to establish whether the fast mixing of SW for $\beta = \beta_u$ and $\beta \geq \beta_h$ in the mean-field case translates to the random regular graph as well. Another interesting direction is to extend our arguments to the random-cluster representation of the Potts model for all non-integer $q \geq 1$; note that the arguments of [5] and [29] do apply to non-integer q ($q \geq 1$ and q large, respectively). The proof of Theorem 1.3 relies on a truncated second moment computation, an argument that was applied to different models in [16, 13].

We further remark here that, from a worst-case perspective, it is known that sampling from the Potts model on d-regular graphs is #BIS-hard for $\beta > \beta_p$ [22], and we conjecture that the problem admits a poly-time approximation algorithm when $\beta < \beta_p$. However, even showing that Glauber mixes fast on any d-regular graph in the uniqueness regime $\beta < \beta_u$ is a major open problem, and Theorems 1.1 and 1.2 further demonstrate that getting an algorithm all the way to β_p will require using different techniques. On that front, progress has been made on the random regular graph: [29] obtained an algorithm for $d \geq 5$ and q large that applies to all β by sampling from each phase separately (using different tools), see also [10]. Moreover, for $\beta < \beta_p$, Efthymiou [21] gives an algorithm with weaker approximation guarantees but which applies to all $q, d \geq 3$ (see also [5]). In principle, and extrapolating again from the mean-field case, one could use Glauber/SW to sample from each phase on the random regular graph for all $q, d \geq 3$ and all β. Analysing such chains appears to be relatively far from the reach of current techniques even in the case of the random regular graph, let alone worst-case graphs. In the case of the Ising model however, the case $q = 2$, the analogue of this fast mixing question has recently been established for sufficiently large β in [26] on the random regular graph and the grid, exploiting certain monotonicity properties.

Finally, let us note that the case of the grid has qualitatively different behaviour than the mean-field and the random-regular case. There, the three critical points coincide and the behaviour at criticality depends on the value of q; the mixing time of Glauber and SW has largely been detailed, see [7, 33, 24].

---

[3] Note that the interval-behaviour on the random regular graph (and hence the correspondence with the mean-field case) is already implied to some extent by the interval-result of [29] (for q large and $d \geq 5$). Note however that the interval therein is contained strictly inside (β_u, β_h) and, in particular, its endpoints do not have the probabilistic interpretation of β_u, β_h. Nevertheless, [29] obtains various probabilistic properties of the metastable phases, including a stronger form of correlation decay than that of reconstruction that we consider here.

2   Overview

In this section we give an overview of the proofs of Theorems 1.1–1.3; for now, we will mostly work towards the proof of Theorem 1.3 which gives the main insights/tools that are needed to prove Theorems 1.1 and 1.2.

Fortunately, to prove Theorem 1.3, we do not need to start from first principles. Instead, we build upon the formula for the partition function $Z_\beta(\mathbb{G})$ and its proof via the second moment method from [22]. Additionally, we are going to seize upon facts about the non-reconstruction properties of the Potts model on the random d-regular tree, also from [22]. We will combine these tools with an auxiliary random graph model known as the planted model, which also plays a key role in the context of inference problems on random graphs [15].

Throughout most of the paper, instead of the simple random regular graph $\mathbb{G}$, we will work with the random d-regular multi-graph $\boldsymbol{G} = \boldsymbol{G}(n,d)$ drawn from the pairing model. Recall that $\boldsymbol{G}$ is obtained by creating d clones of each of the vertices from $[n]$, choosing a random perfect matching of the complete graph on $[n] \times [d]$ and subsequently contracting the vertices $\{i\} \times [d]$ into a single vertex i, for all $i \in [n]$. It is well-known that $\mathbb{G}$ is contiguous with respect to $\boldsymbol{G}$ [30], i.e., any property that holds w.h.p. for $\boldsymbol{G}$ also holds w.h.p. for $\mathbb{G}$.

The following notation will be handy. For a graph G and a configuration $\sigma \in [q]^{V(G)}$, define a probability distribution ν^σ on $[q]$ by letting

$$\nu^\sigma(s) = |\sigma^{-1}(s)|/n \qquad\qquad (s \in [q]).$$

In words, ν^σ is the empirical distribution of the colours under σ. Similarly, let $\rho^{G,\sigma} \in \mathcal{P}([q] \times [q])$ be the edge statistics of a given graph/colouring pair, i.e.,

$$\rho^{G,\sigma}(s,t) = \frac{1}{2|E(G)|} \sum_{u,v \in V(G)} \mathbf{1}\{uv \in E(G), \sigma_u = s, \sigma_v = t\}.$$

2.1   Moments and messages

The routine method for investigating the partition function and the Boltzmann distribution of random graphs is the method of moments [2]. The basic strategy is to calculate, one way or another, the first two moments $\mathbb{E}[Z_\beta(\boldsymbol{G})]$, $\mathbb{E}[Z_\beta(\boldsymbol{G})^2]$ of the partition function. Then we cross our fingers that the second moment is not much larger than the square of the first. It sometimes works. But potential pitfalls include a pronounced tendency of running into extremely challenging optimisation problems in the course of the second moment calculation and, worse, lottery effects that may foil the strategy altogether. While regular graphs in general and the Potts ferromagnet in particular are relatively tame specimens, these difficulties actually do arise once we set out to investigate metastable states. Drawing upon [3, 16] to sidestep these challenges, we develop a less computation-heavy proof strategy.

The starting point is the observation that the fixed points of (3) are intimately related to the moment calculation. This will not come as a surprise to experts, and indeed it was already noticed in [22]. To elaborate, let $\nu = (\nu(\sigma))_{\sigma \in [q]}$ be a probability distribution on the q colours. Moreover, let $\mathcal{R}(\nu)$ be the set of all symmetric matrices $(\rho(\sigma, \tau))_{\sigma, \tau \in [q]}$ with non-negative entries such that

$$\sum_{\tau \in [q]} \rho(\sigma, \tau) = \nu(\sigma) \qquad\qquad \text{for all } \sigma \in [q]. \qquad (6)$$

Relatively standard arguments (e.g., [12, Lemma 2.7]) show that the first moment satisfies

$$\lim_{n\to\infty}\frac{1}{n}\log\mathbb{E}[Z_\beta(\boldsymbol{G})] = \max_{\nu\in\mathcal{P}([q]),\rho\in R(\nu)} F_{d,\beta}(\nu,\rho), \qquad \text{where}$$

$$F_{d,\beta}(\nu,\rho) = (d-1)\sum_{\sigma\in[q]}\nu(\sigma)\log\nu(\sigma) - d\sum_{1\leq\sigma\leq\tau\leq q}\rho(\sigma,\tau)\log\rho(\sigma,\tau) + \frac{d\beta}{2}\sum_{\sigma\in[q]}\rho(\sigma,\sigma). \tag{7}$$

Thus, the first moment is governed by the maximum or maxima, as the case may be, of $F_{d,\beta}$. The function $F_{d,\beta}$ accounts for the contribution to $\mathbb{E}[Z_\beta(\boldsymbol{G})]$ coming from the set $\mathcal{Q}$ which consists of pairs (G,σ) with $\nu^\sigma = \nu + o(1)$ and $\rho^{G,\sigma} = \rho + o(1)$, i.e., the sum $\sum_{(G,\sigma)\in\mathcal{Q}}\mathbb{P}[\boldsymbol{G}=G]\mathrm{e}^{\beta\mathcal{H}_G(\sigma)}$ equals $\mathrm{e}^{nF_{d,\beta}(\nu,\rho)+o(n)}$.

We need to know that the maxima of $F_{d,b}$ are in one-to-one correspondence with the stable fixed points of (3). To be precise, a fixed point μ of (3) is *stable* if the Jacobian of (3) at μ has spectral radius strictly less than one. Let $\mathcal{F}_{d,\beta}^+$ be the set of all stable fixed points $\mu \in \mathcal{F}_{d,\beta}$. Moreover, let $\mathcal{F}_{d,\beta}^1$ be the set of all $\mu \in \mathcal{F}_{d,\beta}^+$ such that $\mu(1) = \max_{\sigma\in[q]}\mu(\sigma)$. In addition, let us call a local maximum (ν,ρ) of $F_{d,\beta}$ *stable* if there exist $\delta, c > 0$ such that

$$F_{d,\beta}(\nu',\rho') \leq F_{d,\beta}(\nu,\rho) - c\left(\|\nu-\nu'\|^2 + \|\rho-\rho'\|^2\right) \tag{8}$$

for all $\nu' \in \mathcal{P}([q])$ and $\rho' \in \mathcal{R}(\nu')$ such that $\|\nu-\nu'\| + \|\rho-\rho'\| < \delta$. Roughly, (8) provides that the Hessian of $F_{d,\beta}$ is negative definite on the subspace of all possible ν, ρ.

▶ **Lemma 2.2** ([22, Theorem 8]). *Suppose that $d, q \geq 3$ are integers and $\beta > 0$ is a real. The map $\mu \in \mathcal{P}([q]) \mapsto (\nu^\mu, \rho^\mu)$ defined by*

$$\nu^\mu(\sigma) = \frac{(1+(\mathrm{e}^\beta-1)\mu(\sigma))^d}{\sum_{\tau\in[q]}(1+(\mathrm{e}^\beta-1)\mu(\tau))^d}, \qquad \rho^\mu(\sigma,\tau) = \frac{\mathrm{e}^{\beta\mathbf{1}\{\sigma=\tau\}}\mu(\sigma)\mu(\tau)}{1+(\mathrm{e}^\beta-1)\sum_{s\in[q]}\mu(s)^2} \tag{9}$$

is a bijection from $\mathcal{F}_{d,\beta}^+$ to the set of stable local maxima of $F_{d,\beta}$. Moreover, for any fixed point μ we have $\mathcal{B}_{d,\beta}(\mu) = F_{d,\beta}(\nu^\mu, \rho^\mu)$.

For brevity, let $(\nu_\mathrm{p}, \rho_\mathrm{p}) = (\nu^{\mu_\mathrm{p}}, \rho^{\mu_\mathrm{p}})$ and $(\nu_\mathrm{f}, \rho_\mathrm{f}) = (\nu^{\mu_\mathrm{f}}, \rho^{\mu_\mathrm{f}})$. The following result characterises the stable fixed points $\mathcal{F}_{d,\beta}^1$.

▶ **Proposition 2.3** ([22, Theorem 4]). *Suppose that $d \geq 3, \beta > 0$.*
  (i) *If $\beta < \beta_u$, then (3) has a unique fixed point, namely the paramagnetic distribution ν_p on $[q]$. This fixed point is stable and thus $F_{d,\beta}$ attains its global maximum at $(\nu_\mathrm{p}, \rho_\mathrm{p})$.*
  (ii) *If $\beta_u < \beta < \beta_h$, then $\mathcal{F}_{d,\beta}^1$ contains two elements, namely the paramagnetic distribution ν_p and the ferromagnetic distribution ν_f; $(\nu_\mathrm{p}, \rho_\mathrm{p})$ is a global maximum of $F_{d,\beta}$ iff $\beta \leq \beta_p$, and $(\nu_\mathrm{f}, \rho_\mathrm{f})$ iff $\beta \geq \beta_p$.*
  (iii) *If $\beta > \beta_h$, then $\mathcal{F}_{d,\beta}^1$ contains precisely one element, namely the ferromagnetic distribution ν_f, and $(\nu_\mathrm{f}, \rho_\mathrm{f})$ is a global maximum of $F_{d,\beta}$.*

Like the first moment, the second moment boils down to an optimisation problem as well, albeit one of much higher dimension ($q^2 - 1$ rather than $q - 1$). Indeed, it is not difficult to derive the following approximation (once again, e.g., via [12, Lemma 2.7]). For a probability distribution $\nu \in \mathcal{P}([q])$ and a symmetric matrix $\rho \in \mathcal{R}(\nu)$ let $\mathcal{R}^\otimes(\rho)$ be the set of all tensors $r = (r(\sigma,\sigma',\tau,\tau'))_{\sigma,\sigma',\tau,\tau'\in[q]}$ such that $r(\sigma,\sigma',\tau,\tau') = r(\tau,\tau',\sigma,\sigma')$ for $\tau,\tau',\sigma,\sigma' \in [q]$ and

$$\sum_{\sigma',\tau'} r(\sigma,\sigma',\tau,\tau') = \sum_{\sigma',\tau'} r(\sigma',\sigma,\tau',\tau) = \rho(\sigma,\tau) \quad \text{for all } \sigma,\tau \in [q]. \tag{10}$$

Then, with $H(\cdot)$ denoting the entropy function, we have

$$\lim_{n \to \infty} \frac{1}{n} \log \mathbb{E}[(Z_\beta(\boldsymbol{G}))^2] = \max_{\nu, \rho \in \mathcal{R}(\nu), r \in \mathcal{R}^{\otimes}(\rho)} F_{d,\beta}^{\otimes}(\rho, r), \text{ where} \tag{11}$$

$$F_{d,\beta}^{\otimes}(\rho, r) = (d-1)H(\rho) + \frac{d}{2}H(r) + \frac{d\beta}{2} \sum_{\sigma, \sigma', \tau, \tau' \in [q]} \left(\mathbf{1}\{\sigma = \tau\} + \mathbf{1}\{\sigma' = \tau'\} \right) r(\sigma, \sigma', \tau, \tau').$$

A frontal assault on this optimisation problem is in general a daunting task due to the doubly-stochastic constraints in (10), i.e., the constraint $r \in \mathcal{R}^{\otimes}(\rho)$. But rather fortunately, to analyse the global maximum (over ν and ρ), these constraints can be relaxed, permitting an elegant translation of the problem to operator theory. In effect, the second moment computation can be reduced to a study of matrix norms. The result is summarised as follows.

▶ **Proposition 2.4** ([22, Theorem 7]). *For all $d, q \geq 3$ and $\beta > 0$ we have*

$$\max_{\nu, \rho \in \mathcal{R}(\nu), r \in \mathcal{R}^{\otimes}(\rho)} F_{d,\beta}^{\otimes}(\rho, r) = 2 \max_{\nu, \rho} F_{d,\beta}(\nu, \rho)$$

and thus $\mathbb{E}[Z_\beta(\boldsymbol{G})^2] = O(\mathbb{E}[Z_\beta(\boldsymbol{G})]^2)$.

Combining Lemma 2.2, Proposition 2.3 and Proposition 2.4, we obtain the following reformulation of [22, Theorem 7], which verifies that we obtain good approximations to the partition function by maximising the Bethe free energy on $\mathcal{F}_{d,\beta}$.

▶ **Theorem 2.5.** *For all integers $d, q \geq 3$ and real $\beta > 0$, we have* $\lim_{n \to \infty} n^{-1} \log Z_\beta(\mathbb{G}) = \max_{\mu \in \mathcal{F}_{d,\beta}} \mathcal{B}_{d,\beta}(\mu)$ *in probability.*

While the *global* maximisation of the function $F_{d,\beta}^{\otimes}$ and thus the proof of Theorem 2.5 boils down to matrix norm analysis, in order to prove Theorems 1.3 and 1.1 via the method of moments we would in addition need to get a good handle on all the *local* maxima. Unfortunately, we do not see a way to reduce this more refined question to operator norms (and it seems unlikely that one exists). Hence, it would seem that we should have to perform a fine-grained analysis of $F_{d,\beta}^{\otimes}$ after all. But luckily another path is open to us. Instead of proceeding analytically, we resort to probabilistic ideas. and we harness "quiet-planting" arguments with the notion of non-reconstruction on the Potts model on the d-regular tree. We review the latter in the next section.

2.2   Non-reconstruction on the regular tree

Let $\mathbb{T}_d$ be the infinite d-regular tree with root o. For a probability distribution $\mu \in \{\mu_{\mathrm{p}}, \mu_{\mathrm{f}}\}$ we define a broadcasting process $\boldsymbol{\sigma} = \boldsymbol{\sigma}_{d,\beta,\mu}$ on $\mathbb{T}_d$ as follows. Initially we draw the color $\boldsymbol{\sigma}_o$ of the root o from the distribution ν^μ. Subsequently, working our way down the levels of the tree, the color of a vertex v whose parent u has been coloured already is drawn from the distribution

$$\mathbb{P}\left[\boldsymbol{\sigma}_v = \sigma \mid \boldsymbol{\sigma}_u\right] = \frac{\mu(\sigma) \mathrm{e}^{\beta \mathbf{1}\{\sigma = \boldsymbol{\sigma}_u\}}}{\sum_{\tau \in [q]} \mu(\tau) \mathrm{e}^{\beta \mathbf{1}\{\tau = \boldsymbol{\sigma}_u\}}}.$$

Naturally, the colours of different vertices on the same level are pairwise independent, but not jointly since there is potentially some correlation with the root. Let $\partial^\ell o$ be the set of all vertices at distance precisely ℓ from o. We say that the broadcasting process has the *strong non-reconstruction property* if $\sum_{\tau \in [q]} \mathbb{E}\left[\left|\mathbb{P}\left[\boldsymbol{\sigma}_o = \tau \mid \boldsymbol{\sigma}_{\partial^\ell o}\right] - \mathbb{P}\left[\boldsymbol{\sigma}_o = \tau\right]\right|\right] = \mathrm{e}^{-\Omega(\ell)}$, where the expectation is over the random configuration $\boldsymbol{\sigma}_{\partial^\ell o}$ (distributed according to the broadcasting

process). In words, this says that the information about the spin of the root inferred from the spins of vertices at depth ℓ decays in the broadcasting process; the term "strong" refers that the decay is exponential with respect to the depth ℓ.

▶ **Proposition 2.6** ([22, Theorem 50]). *Let $d, q \geq 3$ be integers and $\beta > 0$ be real.*

  (i) *For $\beta < \beta_h$, the broadcasting process $\boldsymbol{\sigma}_{d,\beta,\mu_{\mathrm{p}}}$ has the strong non-reconstruction property.*

  (ii) *For $\beta > \beta_u$, the broadcasting process $\boldsymbol{\sigma}_{d,\beta,\mu_{\mathrm{f}}}$ has the strong non-reconstruction property.*

In order to prove Theorems 1.1–1.3 we will combine Proposition 2.6 with reweighted random graph models known as planted models. To be precise, we will consider two versions of planted models, a paramagnetic and a ferromagnetic one. Then we will deduce from Proposition 2.6 that the Boltzmann distribution of these planted models has the non-reconstruction property in a suitably defined sense. In combination with some general facts about Boltzmann distributions this will enable us to prove Theorems 1.1–1.3 without the need for complicated moment computations.

2.3 Second Moment via planting and non-reconstruction

We proceed to introduce the paramagnetic and the ferromagnetic version of the planted model. Roughly speaking, these are weighted versions of the common random regular graph $\mathbb{G}$ where the probability mass of a specific graph is proportional to the paramagnetic or ferromagnetic bit of the partition function. To be precise, for $\varepsilon > 0$, recall the subsets $S_{\mathrm{p}} = S_{\mathrm{p}}(\varepsilon), S_{\mathrm{f}} = S_{\mathrm{f}}(\varepsilon)$ of the configuration space $[q]^n$. Letting

$$Z_{\mathrm{f}}(G) = \sum_{\sigma \in S_{\mathrm{f}}} \mathrm{e}^{\beta \mathcal{H}_G(\sigma)} \quad \text{and} \quad Z_{\mathrm{p}}(G) = \sum_{\sigma \in S_{\mathrm{p}}} \mathrm{e}^{\beta \mathcal{H}_G(\sigma)}, \tag{12}$$

we define random graph models $\hat{\boldsymbol{G}}_{\mathrm{f}}, \hat{\boldsymbol{G}}_{\mathrm{p}}$ by

$$\mathbb{P}\left[\hat{\boldsymbol{G}}_{\mathrm{f}} = G\right] = \frac{Z_{\mathrm{f}}(G)\mathbb{P}\left[\boldsymbol{G} = G\right]}{\mathbb{E}[Z_{\mathrm{f}}(\boldsymbol{G})]}, \qquad \mathbb{P}\left[\hat{\boldsymbol{G}}_{\mathrm{p}} = G\right] = \frac{Z_{\mathrm{p}}(G)\mathbb{P}\left[\boldsymbol{G} = G\right]}{\mathbb{E}[Z_{\mathrm{p}}(\boldsymbol{G})]}. \tag{13}$$

Thus, $\hat{\boldsymbol{G}}_{\mathrm{f}}$ and $\hat{\boldsymbol{G}}_{\mathrm{p}}$ are d-regular random graphs on n vertices such that the probability that a specific graph G comes up is proportional to $Z_{\mathrm{f}}(G)$ and $Z_{\mathrm{p}}(G)$, respectively. Note, the expected value of $Z_{\mathrm{f}}(G)$ and $Z_{\mathrm{p}}(G)$ is captured by the function $F_{d,\beta}$, and we have (see Lemmas 3.2 and 3.3 in the full version)

$$\mathbb{E}[Z_{\mathrm{p}}(G)] = n^{O(1)} \exp(nF_{d,\beta}(\nu_{\mathrm{p}}, \rho_{\mathrm{p}})) \quad \text{and} \quad \mathbb{E}[Z_{\mathrm{f}}(G)] = n^{O(1)} \exp(nF_{d,\beta}(\nu_{\mathrm{f}}, \rho_{\mathrm{f}})). \tag{14}$$

The key ingredient to prove Theorem 1.3 is to quantify the overlap of two typical configurations in the conditional Boltzmann distributions (under S_{f} and S_{p}). To be precise, for a graph $G = (V, E)$, the overlap of two configurations $\sigma, \sigma' \in [q]^V$ is defined as the probability distribution $\nu(\sigma, \sigma') \in \mathcal{P}([q]^2)$ with

$$\nu_{c,c'}(\sigma, \sigma') = \frac{1}{n} \sum_{v \in V(G)} \mathbf{1}\left\{\sigma_v = c, \sigma'_v = c'\right\} \qquad (c, c' \in [q]).$$

For a graph G let $\boldsymbol{\sigma}_{G,\mathrm{f}}$ denote a sample from the conditional distribution $\mu_{G,\beta}(\cdot \mid S_{\mathrm{f}})$. and define $\boldsymbol{\sigma}_{G,\mathrm{p}}$ similarly for S_{p}. The following lemma studies the overlap for two configurations in the conditional distribution $\mu_{\hat{\boldsymbol{G}}_{\mathrm{p}},\beta}(\cdot \mid S_{\mathrm{p}})$, a similar lemma applies to the ferromagnetic phase S_{f}, see Lemma 3.9 in the full version.

▶ **Lemma 3.8.** *Let $d, q \geq 3$ be integers and $\beta < \beta_h$ be real. Let $\boldsymbol{\sigma}_{\hat{\boldsymbol{G}}_{\mathrm{p}},\mathrm{p}}, \boldsymbol{\sigma}'_{\hat{\boldsymbol{G}}_{\mathrm{p}},\mathrm{p}}$ be independent samples from $\mu_{\hat{\boldsymbol{G}}_{\mathrm{p}},\beta}(\,\cdot\mid S_{\mathrm{p}})$. Then $\mathbb{E}\big[d_{\mathrm{TV}}\big(\nu(\boldsymbol{\sigma}_{\hat{\boldsymbol{G}}_{\mathrm{p}},\mathrm{p}}, \boldsymbol{\sigma}'_{\hat{\boldsymbol{G}}_{\mathrm{p}},\mathrm{p}}), \nu_{\mathrm{p}} \otimes \nu_{\mathrm{p}}\big)\big] = o(1)$.*

To utilise Lemmas 3.8 and 3.9, we proceed to apply the second moment method to truncated versions of the paramagnetic and ferromagnetic partition functions $Z_{\mathrm{p}}, Z_{\mathrm{f}}$ where we expressly drop graphs that violate the overlap bounds from Lemmas 3.8. Thus, we introduce the events $\mathcal{E}_{\mathrm{p}} = \{G : \mathbb{E}\big[d_{\mathrm{TV}}(\nu(\boldsymbol{\sigma}_{G,\mathrm{p}}, \boldsymbol{\sigma}'_{G,\mathrm{p}}), \nu_{\mathrm{p}} \otimes \nu_{\mathrm{p}})\big] = o(1)\}$ and the analogous event $\mathcal{E}_{\mathrm{f}}$ for the ferromagnetic phase. Consider now the random variables

$$Y_{\mathrm{p}}(G) = Z_{\mathrm{p}}(G) \cdot \mathbf{1}\{G \in \mathcal{E}_{\mathrm{p}}\} \text{ and } Y_{\mathrm{f}}(G) = Z_{\mathrm{f}}(G) \cdot \mathbf{1}\{G \in \mathcal{E}_{\mathrm{f}}\}$$

Combining Lemma 3.8 with the Nishimori identity (17), we obtain

$$\frac{\mathbb{E}[Y_{\mathrm{p}}]}{\mathbb{E}[Z_{\mathrm{p}}]} = \mathbb{P}\left[\hat{\boldsymbol{G}}_{\mathrm{p}} \in \mathcal{E}_{\mathrm{p}}\right] \sim 1 \text{ and } \frac{\mathbb{E}[Y_{\mathrm{f}}]}{\mathbb{E}[Z_{\mathrm{f}}]} = \mathbb{P}\left[\hat{\boldsymbol{G}}_{\mathrm{f}} \in \mathcal{E}_{\mathrm{f}}\right] \sim 1$$

and thus $\mathbb{E}[Y_{\mathrm{p}}] \sim \mathbb{E}[Z_{\mathrm{p}}]$ and $\mathbb{E}[Y_{\mathrm{f}}] \sim \mathbb{E}[Z_{\mathrm{f}}]$. Crucially, estimating the second moments of these two random variables is a cinch because by construction we can avoid an explicit optimisation of the function $F_{d,\beta}^{\otimes}$ from (11). Indeed, because we drop graphs G whose overlaps stray far from the product measures $\nu_{\mathrm{p}} \otimes \nu_{\mathrm{p}}$ and $\nu_{\mathrm{f}} \otimes \nu_{\mathrm{f}}$, respectively, we basically just need to evaluate the function $F_{d,\beta}^{\otimes}$ at $\nu_{\mathrm{p}} \otimes \nu_{\mathrm{p}}$ and $\nu_{\mathrm{f}} \otimes \nu_{\mathrm{f}}$, which is a matter of relatively simple algebra (due to convexity arguments). We thus obtain the following.

▶ **Corollary 3.10.** *Let $d, q \geq 3$ be integers and $\beta > 0$ be real.*
  (i) *If $\beta < \beta_h$, then $\mathbb{E}[Y_{\mathrm{p}}(\boldsymbol{G})] \sim \mathbb{E}[Z_{\mathrm{p}}(\boldsymbol{G})]$ and $\mathbb{E}[Y_{\mathrm{p}}(\boldsymbol{G})^2] \leq \exp(o(n))\mathbb{E}[Z_{\mathrm{p}}(\boldsymbol{G})]^2$.*
  (ii) *If $\beta > \beta_u$, then $\mathbb{E}[Y_{\mathrm{f}}(\boldsymbol{G})] \sim \mathbb{E}[Z_{\mathrm{f}}(\boldsymbol{G})]$ and $\mathbb{E}[Y_{\mathrm{f}}(\boldsymbol{G})^2] \leq \exp(o(n))\mathbb{E}[Z_{\mathrm{f}}(\boldsymbol{G})]^2$.*
At this stage, one can combine Corollary 3.10 together with (14) to derive the first two parts of Theorem 1.3 (using also the results from Section 2.1).

3   Quiet planting and non-reconstruction

In this section we give an outline of the proof of Lemma 3.8, which was the main ingredient to carry out the second moment method of Section 2.3.

While the planted models defined in (13) are useful for the second-moment argument, working with them directly is rather unwieldy. Fortunately, there is a relatively simple way out using the so-called Nishimori identities; on the way, we will also introduce some of the ingredients that are used for the metastability/slow-mixing results.

To elaborate, we complement the definition (13) of the planted random graphs $\hat{\boldsymbol{G}}_{\mathrm{f}}, \hat{\boldsymbol{G}}_{\mathrm{p}}$ by also introducing a reweighted distribution on graphs for a specific configuration $\sigma \in [q]^n$. Specifically, we define a random graph $\hat{\boldsymbol{G}}(\sigma)$ by letting

$$\mathbb{P}\left[\hat{\boldsymbol{G}}(\sigma) = G\right] = \frac{\mathbb{P}[\boldsymbol{G} = G]\,e^{\beta\mathcal{H}_G(\sigma)}}{\mathbb{E}[e^{\beta\mathcal{H}_G(\sigma)}]}. \tag{15}$$

Furthermore, recalling the truncated partition functions $Z_{\mathrm{f}}, Z_{\mathrm{p}}$ from (12), we introduce reweighted random configurations $\hat{\boldsymbol{\sigma}}_{\mathrm{f}} = \hat{\boldsymbol{\sigma}}_{\mathrm{f}}(\varepsilon) \in [q]^n$ and $\hat{\boldsymbol{\sigma}}_{\mathrm{p}} = \hat{\boldsymbol{\sigma}}_{\mathrm{p}}(\varepsilon) \in [q]^n$ with distributions

$$\mathbb{P}\left[\hat{\boldsymbol{\sigma}}_{\mathrm{f}} = \sigma\right] = \frac{\mathbf{1}\{\sigma \in S_{\mathrm{f}}\}\,\mathbb{E}[e^{\beta\mathcal{H}_G(\sigma)}]}{\mathbb{E}[Z_{\mathrm{f}}(\boldsymbol{G})]}, \qquad \mathbb{P}\left[\hat{\boldsymbol{\sigma}}_{\mathrm{p}} = \sigma\right] = \frac{\mathbf{1}\{\sigma \in S_{\mathrm{p}}\}\,\mathbb{E}[e^{\beta\mathcal{H}_G(\sigma)}]}{\mathbb{E}[Z_{\mathrm{p}}(\boldsymbol{G})]}. \tag{16}$$

We have the following paramagnetic and ferromagnetic *Nishimori identities*. Nishimori identities were derived in [15] for a broad family of planted models which, however, does not include the planted ferromagnetic models $\hat{\boldsymbol{G}}_{\mathrm{p}}, \hat{\boldsymbol{G}}_{\mathrm{f}}$. Nonetheless, the (simple) proof of Proposition 3.1 is practically identical to that in [15] (and is given in the full version).

▶ **Proposition 3.1.** *We have the distributional equalities*

$$(\hat{\boldsymbol{G}}_{\mathrm{p}}, \boldsymbol{\sigma}_{\hat{\boldsymbol{G}}_{\mathrm{p}},\mathrm{p}}) \overset{d}{=} (\hat{\boldsymbol{G}}(\hat{\boldsymbol{\sigma}}_{\mathrm{p}}), \hat{\boldsymbol{\sigma}}_{\mathrm{p}}), \qquad\qquad\qquad (\hat{\boldsymbol{G}}_{\mathrm{f}}, \boldsymbol{\sigma}_{\hat{\boldsymbol{G}}_{\mathrm{f}},\mathrm{f}}) \overset{d}{=} (\hat{\boldsymbol{G}}(\hat{\boldsymbol{\sigma}}_{\mathrm{f}}), \hat{\boldsymbol{\sigma}}_{\mathrm{f}}). \qquad (17)$$

Proposition 3.1 paves the way for a more hands-on description of the planted models in (13). Indeed, the random graph models $(\hat{\boldsymbol{G}}(\hat{\boldsymbol{\sigma}}_{\mathrm{p}}), \hat{\boldsymbol{\sigma}}_{\mathrm{p}})$ and $(\hat{\boldsymbol{G}}(\hat{\boldsymbol{\sigma}}_{\mathrm{f}}), \hat{\boldsymbol{\sigma}}_{\mathrm{f}})$ are invariant under permutations of the vertices, so $\hat{\boldsymbol{\sigma}}_{\mathrm{p}}$ and $\hat{\boldsymbol{\sigma}}_{\mathrm{f}}$ are uniformly random given their colour statistics, and the random graphs $\hat{\boldsymbol{G}}(\hat{\boldsymbol{\sigma}}_{\mathrm{p}})$ and $\hat{\boldsymbol{G}}(\hat{\boldsymbol{\sigma}}_{\mathrm{f}})$ themselves are uniformly random and easy to sample given the planted assignment $\hat{\boldsymbol{\sigma}}_{\mathrm{p}}$ or $\hat{\boldsymbol{\sigma}}_{\mathrm{f}}$ and given the edge statistics $\rho^{\hat{\boldsymbol{G}}(\hat{\boldsymbol{\sigma}}_{\mathrm{p}}),\hat{\boldsymbol{\sigma}}_{\mathrm{p}}}$ and $\rho^{\hat{\boldsymbol{G}}(\hat{\boldsymbol{\sigma}}_{\mathrm{f}}),\hat{\boldsymbol{\sigma}}_{\mathrm{f}}}$. Moreover, because $(\nu_{\mathrm{p}}, \rho_{\mathrm{p}})$ and $(\nu_{\mathrm{f}}, \rho_{\mathrm{f}})$ are local maxima of the first moment function $F_{d,b}(\nu, \rho)$, a first moment argument based on (8) allows us to control the vertex-edge colour statistics very accurately, i.e., there exist $c, t_0 > 0$ such that for all $t \in [0, t_0]$

$$\mathbb{P}\left[d_{\mathrm{TV}}\left(\nu^{\hat{\boldsymbol{\sigma}}_{\mathrm{p}}}, \nu_{\mathrm{p}}\right) + d_{\mathrm{TV}}\left(\rho^{\hat{\boldsymbol{G}}(\hat{\boldsymbol{\sigma}}_{\mathrm{p}}),\hat{\boldsymbol{\sigma}}_{\mathrm{p}}}, \rho_{\mathrm{p}}\right) > t\right] \le n^{O(1)}\mathrm{e}^{-ct^2 n}, \qquad (18)$$

and similarly for the deviations from $(\nu_{\mathrm{f}}, \rho_{\mathrm{f}})$ in the ferromagnetic phase (see Lemmas 3.4 and 3.5 in the full version). At this point we have handy descriptions of the models $(\hat{\boldsymbol{G}}(\hat{\boldsymbol{\sigma}}_{\mathrm{p}}), \hat{\boldsymbol{\sigma}}_{\mathrm{p}})$ and $(\hat{\boldsymbol{G}}(\hat{\boldsymbol{\sigma}}_{\mathrm{f}}), \hat{\boldsymbol{\sigma}}_{\mathrm{f}})$, and therefore, via Proposition 3.1, $(\hat{\boldsymbol{G}}_{\mathrm{p}}, \boldsymbol{\sigma}_{\hat{\boldsymbol{G}}_{\mathrm{p}},\mathrm{p}})$ and $(\hat{\boldsymbol{G}}_{\mathrm{f}}, \boldsymbol{\sigma}_{\hat{\boldsymbol{G}}_{\mathrm{f}},\mathrm{f}})$.

We will next utilise the information on the distribution of $\hat{\boldsymbol{\sigma}}_{\mathrm{p}}, \rho^{\hat{\boldsymbol{G}}(\hat{\boldsymbol{\sigma}}_{\mathrm{p}}),\hat{\boldsymbol{\sigma}}_{\mathrm{p}}}$ to couple the distribution of the colouring produced by the tree broadcasting process and the colouring that $\hat{\boldsymbol{\sigma}}_{\mathrm{p}}$ induces on the neighbourhood of some particular vertex of $\hat{\boldsymbol{G}}(\hat{\boldsymbol{\sigma}}_{\mathrm{p}})$, say v. In particular, for $\ell = \lceil \log \log n \rceil$, the ℓ-neighbourhood of v is going to be tree-like, so conditional on the statistics $\nu^{\hat{\boldsymbol{\sigma}}_{\mathrm{p}}}, \rho^{\hat{\boldsymbol{G}}(\hat{\boldsymbol{\sigma}}_{\mathrm{p}}),\hat{\boldsymbol{\sigma}}_{\mathrm{p}}}$, an inductive coupling (see Lemma 3.6) shows that

$$d_{\mathrm{TV}}(\hat{\boldsymbol{\sigma}}_{\mathrm{p},\partial^\ell v}, \tau_{\partial^\ell o}) = d^\ell \left(d_{\mathrm{TV}}(\nu^{\hat{\boldsymbol{\sigma}}_{\mathrm{p}}}, \nu_{\mathrm{p}}) + d_{\mathrm{TV}}(\rho^{\hat{\boldsymbol{G}}(\hat{\boldsymbol{\sigma}}_{\mathrm{p}}),\hat{\boldsymbol{\sigma}}_{\mathrm{p}}}, \rho_{\mathrm{p}}) + n^{-0.99}\right).$$

From (18), it then follows that the last quantity is $o(n^{-1/5})$ with probability $1 - o(1/n)$. Hence, the colourings $\hat{\boldsymbol{\sigma}}_{\mathrm{p},\partial^\ell v}$ and $\tau_{\partial^\ell o}$ can be coupled such that both are identical with probability $1 - o(n^{-1/5})$. Consequently, from the tree broadcasting results of Proposition 2.6, we obtain that $\sum_{c \in [q]} \mathbb{E}\left|\nu_{\mathrm{p}}(c) - \mu_{\hat{\boldsymbol{G}}(\hat{\boldsymbol{\sigma}}_{\mathrm{p}}),\beta}(\boldsymbol{\sigma}_v = c \mid \boldsymbol{\sigma}_{\partial^\ell v} = \hat{\boldsymbol{\sigma}}_{\mathrm{p},\partial^\ell v})\right| < \ell^{-3}$ which translates via the Nishimori identity into

$$\sum_{c \in [q]} \mathbb{E}\left|\nu_{\mathrm{p}}(c) - \mu_{\hat{\boldsymbol{G}}_{\mathrm{p}},\beta}(\boldsymbol{\sigma}_v = c \mid \boldsymbol{\sigma}_{\partial^\ell v} = \boldsymbol{\sigma}_{\hat{\boldsymbol{G}}_{\mathrm{p}},\mathrm{p},\partial^\ell v})\right| < \ell^{-3}. \qquad (19)$$

Proof Sketch of Lemma 3.8. Due to the Nishimori identity (17), it suffices to prove that for a sample $\boldsymbol{\sigma}_{\hat{\boldsymbol{G}}(\hat{\boldsymbol{\sigma}}_{\mathrm{p}}),\mathrm{p}}$ from $\mu_{\hat{\boldsymbol{G}}(\hat{\boldsymbol{\sigma}}_{\mathrm{p}}),\beta}(\cdot \mid S_{\mathrm{p}})$ that

$$d_{\mathrm{TV}}\left(\nu(\hat{\boldsymbol{\sigma}}_{\mathrm{p}}, \boldsymbol{\sigma}_{\hat{\boldsymbol{G}}(\hat{\boldsymbol{\sigma}}_{\mathrm{p}}),\mathrm{p}}), \nu_{\mathrm{p}} \otimes \nu_{\mathrm{p}}\right) = o(1). \qquad (20)$$

To see (20), for colors $s, t \in [q]$, we consider the first and second moment of the number of vertices u with $\hat{\boldsymbol{\sigma}}_{\mathrm{p}}(u) = s$ and $\boldsymbol{\sigma}_{\hat{\boldsymbol{G}}(\hat{\boldsymbol{\sigma}}_{\mathrm{p}}),\mathrm{p}}(u) = t$. To facilitate the analysis of the second moment, it will be convenient to consider the following configuration $\boldsymbol{\sigma}'_{\hat{\boldsymbol{G}}(\hat{\boldsymbol{\sigma}}_{\mathrm{p}}),\mathrm{p}}$. Let $\boldsymbol{v}, \boldsymbol{w}$ be two random vertices such that $\hat{\boldsymbol{\sigma}}_{\mathrm{p}}(\boldsymbol{v}) = \hat{\boldsymbol{\sigma}}_{\mathrm{p}}(\boldsymbol{w}) = s$. Also let $\ell = \ell(n) = \lceil \log \log n \rceil$. Now, draw $\boldsymbol{\sigma}''_{\hat{\boldsymbol{G}}(\hat{\boldsymbol{\sigma}}_{\mathrm{p}}),\mathrm{p}}$ from $\mu_{\hat{\boldsymbol{G}}(\hat{\boldsymbol{\sigma}}_{\mathrm{p}}),\beta}(\cdot \mid S_{\mathrm{p}})$ and subsequently generate $\boldsymbol{\sigma}'_{\hat{\boldsymbol{G}}(\hat{\boldsymbol{\sigma}}_{\mathrm{p}}),\mathrm{p}}$ by re-sampling the colours of the vertices at distance less than ℓ from $\boldsymbol{v}, \boldsymbol{w}$ given the colours of the vertices at distance ℓ from $\boldsymbol{v}, \boldsymbol{w}$ and the event S_{p}. Then $\boldsymbol{\sigma}'_{\hat{\boldsymbol{G}}(\hat{\boldsymbol{\sigma}}_{\mathrm{p}}),\mathrm{p}}$ has distribution $\mu_{\hat{\boldsymbol{G}}(\hat{\boldsymbol{\sigma}}_{\mathrm{p}}),\beta}(\cdot \mid S_{\mathrm{p}})$. Moreover, since for two random vertices $\boldsymbol{v}, \boldsymbol{w}$ their ℓ-neighbourhoods are going to be disjoint w.h.p., the reconstruction property in (19) implies that w.h.p. for all $\chi, \chi' \in [q]$

$$\mathbb{P}\left[\boldsymbol{\sigma}'_{\hat{\boldsymbol{G}}_{\mathrm{p}},\mathrm{p}}(\boldsymbol{v}) = \chi, \boldsymbol{\sigma}'_{\hat{\boldsymbol{G}}_{\mathrm{p}},\mathrm{p}}(\boldsymbol{w}) = \chi' \mid \hat{\boldsymbol{\sigma}}_{\mathrm{p}}, \hat{\boldsymbol{G}}(\hat{\boldsymbol{\sigma}}_{\mathrm{p}}), \boldsymbol{v}, \boldsymbol{w}\right] = \nu_p(\chi)\nu_p(\chi') + o(1). \qquad (21)$$

Hence, for a colour $t \in [q]$ let $\boldsymbol{X}(s,t)$ be the number of vertices u with $\hat{\boldsymbol{\sigma}}_{\mathrm{p}}(u) = s$ and $\boldsymbol{\sigma}'_{\hat{\boldsymbol{G}}_{\mathrm{p,p}}}(v) = t$. Then (21) shows that w.h.p. $\mathbb{E}\left[\boldsymbol{X}(s,t) \mid \hat{\boldsymbol{\sigma}}_{\mathrm{p}}, \hat{\boldsymbol{G}}(\hat{\boldsymbol{\sigma}}_{\mathrm{p}})\right] \sim \frac{n}{q^2}$ and $\mathbb{E}\left[\boldsymbol{X}(s,t)^2 \mid \hat{\boldsymbol{\sigma}}_{\mathrm{p}}, \hat{\boldsymbol{G}}(\hat{\boldsymbol{\sigma}}_{\mathrm{p}})\right] \sim \frac{n^2}{q^4}$. Thus, (20) follows from Chebyshev's inequality. ◄

4 Metastability and Slow mixing

In this section, we prove Theorems 1.1 and 1.2. Recall from Section 1.3 the paramagnetic and ferromagnetic states $S_{\mathrm{p}}(\varepsilon)$ and $S_{\mathrm{f}}(\varepsilon)$ for $\varepsilon > 0$. For the purposes of this section we will need to be more systematic of keeping track the dependence of these phases on ε. In particular, we will use the more explicit notation $Z_{\mathrm{p}}^{\varepsilon}(G)$ and $Z_{\mathrm{f}}^{\varepsilon}(G)$ to denote the quantities $Z_{\mathrm{p}}(G)$ and $Z_{\mathrm{f}}(G)$, respectively, from (12). The following lemma, based on Theorem 1.3, reflects the fact that ν_{p} and ν_{f} are local maxima of the first-moment function $F_{d,\beta}$.

▶ **Lemma 4.1.** *Let $q, d \geq 3$ be integers and $\beta > 0$ be real. Then, for all sufficiently small constants $\varepsilon' > \varepsilon > 0$, there exists constant $\zeta > 0$ such that w.h.p. over $G \sim \boldsymbol{G}$, it holds that*
1. *If $\beta < \beta_h$, then $Z_{\mathrm{p}}^{\varepsilon}(G) \geq \mathrm{e}^{-n^{3/4}} \mathbb{E}[Z_{\mathrm{p}}^{\varepsilon}(\boldsymbol{G})]$ and $Z_{\mathrm{p}}^{\varepsilon'}(G) \leq (1 + \mathrm{e}^{-\zeta n}) Z_{\mathrm{p}}^{\varepsilon}(G)$.*
2. *If $\beta > \beta_u$, then $Z_{\mathrm{f}}^{\varepsilon}(G) \geq \mathrm{e}^{-n^{3/4}} \mathbb{E}[Z_{\mathrm{f}}^{\varepsilon}(\boldsymbol{G})]$ and $Z_{\mathrm{f}}^{\varepsilon'}(G) \leq (1 + \mathrm{e}^{-\zeta n}) Z_{\mathrm{f}}^{\varepsilon}(G)$.*

Theorem 1.1 will follow by way of a conductance argument. Let $G = (V,E)$ be a graph, and P be the transition matrix for the Glauber dynamics defined in Section 1.4. For a set $S \subseteq [q]^V$ define the *bottleneck ratio* of S to be $\Phi(S) = \frac{\sum_{\sigma \in S, \tau \notin S} \mu_{G,\beta}(\sigma) P(\sigma,\tau)}{\mu_{G,\beta}(S)}$. The following lemma provides a routine conductance bound (e.g., [32, Theorem 7.3]). For the sake of completeness the proof is included in the full version.

▶ **Lemma 4.2.** *Let $G = (V,E)$ be a graph. For any $S \subseteq [q]^V$ such that $\mu_G(S) > 0$ and any integer $t \geq 0$ we have $\|\mu_{G,S} P^t - \mu_{G,S}\|_{TV} \leq t\Phi(S)$.*

Proof of Theorem 1.1. We prove the statement for the pairing model $\boldsymbol{G}$, the result for $\mathbb{G}$ follows immediately by contiguity. Let $\varepsilon' > \varepsilon > 0$ and $\zeta > 0$ be small constants such that Lemma 4.1 applies, and let $G \sim \boldsymbol{G}$ be a graph satisfying the lemma. Set for convenience $\mu = \mu_{\boldsymbol{G},\beta}$; we consider first the metastability of $S_{\mathrm{f}}(\varepsilon)$ for $\beta > \beta_u$.

Since Glauber updates one vertex at a time it is impossible in one step to move from $\sigma \in S_{\mathrm{f}}(\varepsilon)$ to $\tau \in [q]^n \backslash S_{\mathrm{f}}(\varepsilon')$, i.e., $P(\sigma,\tau) = 0$, and therefore

$$\Phi\big(S_{\mathrm{f}}(\varepsilon)\big) = \frac{\sum_{\sigma \in S_{\mathrm{f}}(\varepsilon)} \sum_{\tau \notin S_{\mathrm{f}}(\varepsilon)} \mu(\sigma) P(\sigma,\tau)}{\mu\big(S_{\mathrm{f}}(\varepsilon)\big)} = \frac{\sum_{\sigma \in S_{\mathrm{f}}(\varepsilon)} \sum_{\tau \in S_{\mathrm{f}}(\varepsilon') \backslash S_{\mathrm{f}}(\varepsilon')} \mu(\sigma) P(\sigma,\tau)}{\mu\big(S_{\mathrm{f}}(\varepsilon)\big)}$$

By reversibility of Glauber, for any $\sigma, \tau \in [q]^n$ we have $\mu(\sigma) P(\sigma,\tau) = \mu(\tau) P(\tau,\sigma)$, and therefore the numerator is upper-bounded by $\mu\big(S_{\mathrm{f}}(\varepsilon') \backslash S_{\mathrm{f}}(\varepsilon)\big)$. Hence, $\Phi\big(S_{\mathrm{f}}(\varepsilon)\big) \leq \frac{\mu\big(S_{\mathrm{f}}(\varepsilon') \backslash S_{\mathrm{f}}(\varepsilon)\big)}{\mu\big(S_{\mathrm{f}}(\varepsilon)\big)} = \frac{Z_{\mathrm{f}}^{\varepsilon'}(G) - Z_{\mathrm{f}}^{\varepsilon}(G)}{Z_{\mathrm{f}}^{\varepsilon}(G)} \leq \mathrm{e}^{-\zeta n}$, where the last inequality follows from the fact that G satisfies Lemma 4.1. Lemma 4.2 now ensures that for all nonnegative integers $T \leq \mathrm{e}^{\zeta n/3}$

$$\big\|\mu\big(\cdot \mid S_{\mathrm{f}}(\varepsilon)\big) P^T - \mu\big(\cdot \mid S_{\mathrm{f}}(\varepsilon)\big)\big\|_{TV} \leq T \cdot \Phi(S_{\mathrm{f}}) \leq \mathrm{e}^{-2\zeta n/3}. \tag{22}$$

Now, consider the Glauber dynamics $(\sigma_t)_{t \geq 0}$ launched from σ_0 drawn from $\mu_{\mathbb{G},\beta,S_{\mathrm{f}}(\varepsilon)}$, and denote by $T_{\mathrm{f}} = \min\{t > 0 : \sigma_t \notin S_{\mathrm{f}}(\varepsilon)\}$ its escape time from $S_{\mathrm{f}}(\varepsilon)$. Observe that σ_t has the same distribution as $\mu(\cdot \mid S_{\mathrm{f}}(\varepsilon)) P^t$, so (22) implies that for all nonnegative integers $T \leq \mathrm{e}^{\zeta n/3}$ it holds that $\big|\mathbb{P}[\sigma_T \in S_{\mathrm{f}}(\varepsilon)] - 1\big| < \mathrm{e}^{-2\zeta n/3}$, or equivalently $\mathbb{P}[\sigma_T \notin S_{\mathrm{f}}(\varepsilon)] \leq \mathrm{e}^{-2\zeta n/3}$. By a

union bound over the values of T, we therefore obtain that $\mathbb{P}[T_f \leq e^{\zeta n/3}] \leq e^{-\zeta n/3}$, thus proving that $S_f(\varepsilon)$ is a metastable state for $\beta > \beta_u$. Analogous arguments show that $S_p(\varepsilon)$ is a metastable state for $\beta < \beta_h$.

The slow mixing of Glauber for $\beta > \beta_u$ follows from the metastability of $S_f(\varepsilon)$. In particular, from Theorem 1.3 we have that $\|\mu(\cdot \mid S_f(\varepsilon)) - \mu\| \geq 3/5$ and therefore, from (22), $\|\mu(\cdot \mid S_f(\varepsilon))P^T - \mu\| \geq 1/2$, yielding that the mixing time is $e^{\Omega(n)}$. ◄

The key and much more challenging ingredient to establish Theorem 1.2 is to bound the probability that Swendsen-Wang escapes $S_p(\varepsilon)$ and $S_f(\varepsilon)$. More precisely, for a graph G, a configuration $\sigma \in [q]^n$, and $S \subseteq [q]^n$, let $P^G_{SW}(\sigma \to S)$ denote the probability that after one step of SW on G starting from σ, we end up in a configuration in S.

The following proposition shows that for almost all pairs (G, σ) from the ferromagnetic planted distribution $(\hat{\boldsymbol{G}}(\hat{\boldsymbol{\sigma}}_f(\varepsilon)), \hat{\boldsymbol{\sigma}}_f(\varepsilon))$, the probability that SW leads to a configuration in the ferromagnetic phase, slightly enlarged, is $1 - e^{-\Omega(n)}$. Note here that SW might change the dominant colour due to recolouring step, so, for $\varepsilon > 0$, we consider the set of configurations $\tilde{S}_f(\varepsilon)$ that consists of the ferromagnetic phase $S_f(\varepsilon)$ together with its $q - 1$ permutations, and the probability that SW escapes from it, starting from a ferromagnetic state.

▶ **Proposition 4.4.** *Let $q, d \geq 3$ be integers and $\beta \in (\beta_u, \beta_h)$. Then, for all sufficiently small constants $\varepsilon' > \varepsilon > 0$, there exists constant $\eta > 0$ such that with probability $1 - e^{-\eta n}$ over the planted distribution $(G, \sigma) \sim (\hat{\boldsymbol{G}}(\hat{\boldsymbol{\sigma}}_f(\varepsilon)), \hat{\boldsymbol{\sigma}}_f(\varepsilon))$, it holds that $P^G_{SW}(\sigma \to \tilde{S}_f(\varepsilon')) \geq 1 - e^{-\eta n}$.*

An analogous Proposition 4.3 applies for the paramagnetic distribution $(\hat{\boldsymbol{G}}(\hat{\boldsymbol{\sigma}}_p(\varepsilon)), \hat{\boldsymbol{\sigma}}_p(\varepsilon))$. The proof of these Propositions requires a delicate analysis of the percolation step in SW since we need probability bounds that are exponentially close to 1. Especially for Proposition 4.4, the presence of a giant component (corresponding to the dominant colour) complicates the arguments significantly since we need to take into account the underlying vertex-edge colour statistics of $(\hat{\boldsymbol{G}}(\hat{\boldsymbol{\sigma}}_f(\varepsilon)), \hat{\boldsymbol{\sigma}}_f(\varepsilon)))$ studied in Section 3. Even with Propositions 4.3 and 4.4 at hand, concluding Theorem 1.2 requires a bit of work based on the planting ideas.

Proof Sketch of Theorem 1.2. We consider first the metastability for the ferromagnetic phase when $\beta > \beta_u$. Let $\varepsilon' > \varepsilon > 0$ and $\eta, \zeta > 0$ be small constants such that Lemma 4.1 and Proposition 4.4 apply. Let $\theta = \frac{1}{10} \min\{\eta, \zeta\}$.

Let $\mathcal{Q}$ be the set of d-regular (multi)graphs that satisfy both items in Lemma 4.1. Moreover, let $\mathcal{Q}'$ be the set of d-regular (multi)graphs G such that the set of configurations where SW has conceivable probability of escaping $\tilde{S}_f(\varepsilon')$ has small weight, i.e., the set

$$S_{\text{Bad}}(G) = \left\{ \sigma \in \tilde{S}_f(\varepsilon) \,\middle|\, P^G_{SW}(\sigma \to \tilde{S}_f(\varepsilon')) < 1 - e^{-\eta n} \right\}$$

has aggregate weight $Z_{\text{Bad}}(G) = \sum_{\sigma \in S_{\text{Bad}}(G)} e^{\beta \mathcal{H}(G)}$ less than $e^{-\theta n} Z^\varepsilon_f(G)$. For a d-regular graph G such that $G \in \mathcal{Q} \cap \mathcal{Q}'$, using arguments analogous to those for Glauber, we have that $\Phi_{SW}(\tilde{S}_f(\varepsilon)) \leq 10 e^{-\theta n}$. By arguments analogous to those in the proof of Theorem 1.1, we have that $\tilde{S}_f(\varepsilon)$ is a metastable state for graphs $G \in \mathcal{Q} \cap \mathcal{Q}'$. Therefore, to finish the metastability proof for the random graph, it suffices to show that $\mathbb{P}(\boldsymbol{G} \in \mathcal{Q} \cap \mathcal{Q}') = 1 - o(1)$.

To do this, let $\mathcal{G}(n, d)$ be the set of all multigraphs that can be obtained in the pairing model and $\Lambda_{d,\beta}(n) = \{(G, \sigma) \mid G \in \mathcal{G}(n, d), \ \sigma \in \tilde{S}_f(\varepsilon)\}$. Let $\mathcal{E}$ be the pairs $(G, \sigma) \in \Lambda_{d,\beta}(n)$ where one step of SW starting from G, σ stays within $\tilde{S}_f(\varepsilon')$ with probability $1 - e^{-\Omega(n)}$, more precisely $\mathcal{E} = \left\{ (G, \sigma) \in \Lambda_{d,\beta}(n) \,\middle|\, P^G_{SW}(\sigma \to \tilde{S}_f(\varepsilon')) \geq 1 - e^{-\eta n} \right\}$. The aggregate weight corresponding to pairs $(G, \sigma) \notin \mathcal{E}$ can be lower-bounded by

$$\sum_{(G,\sigma) \in \Lambda_{d,\beta} \setminus \mathcal{E}} e^{\beta \mathcal{H}_G(\sigma)} \geq \sum_{G \in \mathcal{Q} \setminus \mathcal{Q}'} \sum_{\sigma \in \Sigma_{\text{Bad}}(G)} e^{\beta \mathcal{H}_G(\sigma)} \geq e^{-\theta n} \sum_{G \in \mathcal{Q} \setminus \mathcal{Q}'} Z^\varepsilon_f(G).$$

For graphs $G \in \mathcal{Q}$ we have $Z_{\mathrm{f}}^{\varepsilon}(G) \geq \mathrm{e}^{-n^{3/4}}\mathbb{E}[Z_{\mathrm{f}}^{\varepsilon}(\boldsymbol{G})]$, and therefore

$$\sum_{(G,\sigma) \in \Lambda_{d,\beta}\backslash\mathcal{E}} \mathrm{e}^{\beta\mathcal{H}_G(\sigma)} \geq \mathrm{e}^{-(\theta n + n^{3/4})}\big|\mathcal{Q}\backslash\mathcal{Q}'\big| \; \mathbb{E}\big[Z_{\mathrm{f}}^{\varepsilon}(\boldsymbol{G})\big] = \mathrm{e}^{-(\theta n + n^{3/4})}\big|\mathcal{Q}\backslash\mathcal{Q}'\big| \frac{\sum_{(G,\sigma) \in \Lambda_{d,\beta}} \mathrm{e}^{\beta\mathcal{H}_G(\sigma)}}{\big|\mathcal{G}(n,d)\big|}$$

From the definition of $\big(\hat{\boldsymbol{G}}\big(\hat{\boldsymbol{\sigma}}_{\mathrm{f}}(\varepsilon)\big), \hat{\boldsymbol{\sigma}}_{\mathrm{f}}(\varepsilon)\big)$, cf. (15),(16), observe that

$$\frac{\sum_{(G,\sigma) \in \Lambda_{d,\beta}\backslash\mathcal{E}} \mathrm{e}^{\beta\mathcal{H}_G(\sigma)}}{\sum_{(G,\sigma) \in \Lambda_{d,\beta}} \mathrm{e}^{\beta\mathcal{H}_G(\sigma)}} = \mathbb{P}\big[\big(\hat{\boldsymbol{G}}\big(\hat{\boldsymbol{\sigma}}_{\mathrm{f}}(\varepsilon)\big), \hat{\boldsymbol{\sigma}}_{\mathrm{f}}(\varepsilon)\big) \in \Lambda_{d,\beta}\backslash\mathcal{E}\big] \leq \mathrm{e}^{-\eta n} \leq \mathrm{e}^{-10\theta n},$$

where the penultimate inequality follows from Proposition 4.4 and the last from the choice of θ. Combining the last two inequalities, we obtain $\mathbb{P}[\boldsymbol{G} \in \mathcal{Q}\backslash\mathcal{Q}'] = o(1)$. Since $\mathbb{P}[\boldsymbol{G} \in \mathcal{Q}] = 1 - o(1)$ from Lemma 4.1, it follows that $\mathbb{P}[\boldsymbol{G} \in \mathcal{Q} \cap \mathcal{Q}'] \geq \mathbb{P}[\boldsymbol{G} \in \mathcal{Q}] - \mathbb{P}[\boldsymbol{G} \in \mathcal{Q}\backslash\mathcal{Q}'] \geq 1 - o(1)$. This concludes the proof for the metastability of the ferromagnetic phase $\tilde{S}_{\mathrm{f}}(\varepsilon)$ when $\beta > \beta_u$.

A similar bottleneck-ratio argument shows that $S_{\mathrm{p}}(\varepsilon)$ is a metastable state for $\beta < \beta_h$. The slow mixing of SW for $\beta \in (\beta_u, \beta_h)$ follows from the metastability of $\tilde{S}_{\mathrm{f}}(\varepsilon)$ when $\beta \in (\beta_u, \beta_p]$ and the metastability of $S_{\mathrm{p}}(\varepsilon)$ when $\beta \in [\beta_p, \beta_h)$. In particular, let $S \in \{\tilde{S}_{\mathrm{f}}(\varepsilon), S_{\mathrm{p}}(\varepsilon)\}$ be such that $\big\|\mu\big(\cdot \mid S\big) - \mu\big\| \geq 1/2$, then Lemma 4.2 gives that for $T = \mathrm{e}^{\Omega(n)}$, it holds that $\big\|\mu\big(\cdot \mid S\big)P_{SW}^T - \mu\big\| \geq 1/2 - 1/10$, yielding that the mixing time is $\mathrm{e}^{\Omega(n)}$.   ◄

---- **References** ----

1   Emmanuel Abbe. Community detection and stochastic block models: Recent developments. *Journal of Maching Learning Research*, 18(1):6446–6531, 2017.

2   Dimitris Achlioptas, Assaf Naor, and Yuval Peres. Rigorous location of phase transitions in hard optimization problems. *Nature*, 435(7043):759–764, 2005.

3   Victor Bapst and Amin Coja-Oghlan. Harnessing the Bethe free energy. *Random Structures and Algorithms*, 49:694–741, 2016.

4   Jean Barbier, Chun Lam Clement Chan, and Nicolas Macris. Concentration of multi-overlaps for random dilute ferromagnetic spin models. *Journal of Statistical Physics*, 180:534–557, 2019.

5   Antonio Blanca, Andreas Galanis, Leslie Ann Goldberg, Daniel Štefankovič, Eric Vigoda, and Kuan Yang. Sampling in uniqueness from the Potts and random-cluster models on random regular graphs. *SIAM Journal on Discrete Mathematics*, 34(1):742–793, 2020.

6   Antonio Blanca and Reza Gheissari. Random-cluster dynamics on random regular graphs in tree uniqueness. *Communications in Mathematical Physics*, 386(2):1243–1287, 2021.

7   Antonio Blanca and Alistair Sinclair. Dynamics for the mean-field random-cluster model. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM)*, pages 528–543, 2015.

8   Antonio Blanca, Alistair Sinclair, and Xusheng Zhang. The critical mean-field Chayes-Machta dynamics. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM)*, 2021.

9   Magnus Bordewich, Catherine Greenhill, and Viresh Patel. Mixing of the Glauber dynamics for the ferromagnetic Potts model. *Random Structures and Algorithms*, 48(1):21–52, 2016.

10   Charlie Carlson, Ewan Davies, Nicolas Fraiman, Alexandra Kolla, Aditya Potukuchi, and Corrine Yap. Algorithms for the ferromagnetic Potts model on expanders. *arXiv preprint*, 2022. `arXiv:2204.01923`.

11   Amin Coja-Oghlan, Oliver Cooley, Mihyun Kang, Joon Lee, and Jean Bernoulli Ravelomanana. The sparse parity matrix. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 822–833, 2022.

12   Amin Coja-Oghlan, Charilaos Efthymiou, and Samuel Hetterich. On the chromatic number of random regular graphs. *Journal of Combinatorial Theory, Series B*, 116:367–439, 2016.

13   Amin Coja-Oghlan, Charilaos Efthymiou, Nor Jaafari, Mihyun Kang, and Tobias Kapetan-opoulos. Charting the replica symmetric phase. *Communications in Mathematical Physics*, 359(2):603–698, 2018.

14   Amin Coja-Oghlan, Oliver Gebhard, Max Hahn-Klimroth, and Philipp Loick. Optimal group testing. *Combinatorics, Probability and Computing*, 30(6):811–848, 2021.

15   Amin Coja-Oghlan, Florent Krzakala, Will Perkins, and Lenka Zdeborová. Information-theoretic thresholds from the cavity method. *Advances in Mathematics*, 333:694–795, 2018.

16   Amin Coja-Oghlan, Philipp Loick, Balázs F Mezei, and Gregory B Sorkin. The Ising antiferro-magnet and max cut on random regular graphs. *arXiv preprint*, 2020. `arXiv:2009.10483`.

17   Paul Cuff, Jian Ding, Oren Louidor, Eyal Lubetzky, Yuval Peres, and Allan Sly. Glauber dynamics for the mean-field Potts model. *Journal of Statistical Physics*, 149(3):432–477, 2012.

18   Amir Dembo and Andrea Montanari. Ising models on locally tree-like graphs. *The Annals of Applied Probability*, 20(2):565–592, 2010.

19   Amir Dembo, Andrea Montanari, Allan Sly, and Nike Sun. The replica symmetric solution for Potts models on *d*-regular graphs. *Communications in Mathematical Physics*, 327(2):551–575, 2014.

20   Amir Dembo, Andrea Montanari, and Nike Sun. Factor models on locally tree-like graph. *The Annals of Probability*, 41(6):4162–4213, 2013.

21   Charilaos Efthymiou. On sampling symmetric Gibbs distributions on sparse random graphs and hypergraphs. *arXiv preprint*, 2020. `arXiv:2007.07145`.

22   Andreas Galanis, Daniel Stefankovic, Eric Vigoda, and Linji Yang. Ferromagnetic Potts model: Refined #BIS-hardness and related results. *SIAM Journal of Computation*, 45(6):2004–2065, 2016.

23   Andreas Galanis, Daniel Štefankovič, and Eric Vigoda. Swendsen-Wang algorithm on the mean-field Potts model. *Random Structures and Algorithms*, 54(1):82–147, 2019.

24   Reza Gheissari and Eyal Lubetzky. Mixing times of critical two-dimensional Potts models. *Communications on Pure and Applied Mathematics*, 71(5):994–1046, 2018.

25   Reza Gheissari, Eyal Lubetzky, and Yuval Peres. Exponentially slow mixing in the mean-field Swendsen-Wang dynamics. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2018.

26   Reza Gheissari and Alistair Sinclair. Low-temperature Ising dynamics with random initializa-tions. *arXiv preprint*, 2021. `arXiv:2106.11296`.

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

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

29   Tyler Helmuth, Matthew Jenssen, and Will Perkins. Finite-size scaling, phase coexistence, and algorithms for the random cluster model on random graphs. *arXiv preprint*, 2020. `arXiv:2006.11580`.

30   Svante Janson, Andrzej Rucinski, and Tomasz Luczak. *Random graphs*. John Wiley and Sons, 2011.

31   Jungkyoung Lee. Energy landscape and metastability of curie–weiss–potts model. *Journal of Statistical Physics*, 187(1):1–46, 2022.

32   David A Levin and Yuval Peres. Markov chains and mixing times. *American Mathematical Society*, 2009.

33   Eyal Lubetzky and Allan Sly. Critical Ising on the square lattice mixes in polynomial time. *Communications in Mathematical Physics*, 313(3):815–836, 2012.

34   Marc Mézard. Mean-field message-passing equations in the Hopfield model and its generaliza-tions. *Physical Review E*, 95(2):022117, 2017.

35   Marc Mezard and Andrea Montanari. *Information, physics, and computation*. Oxford University Press, 2009.

36   Marc Mézard and Giorgio Parisi. The Bethe lattice spin glass revisited. *The European Physical Journal B-Condensed Matter and Complex Systems*, 20(2):217–233, 2001.

37   Marc Mézard and Giorgio Parisi. The cavity method at zero temperature. *Journal of Statistical Physics*, 111(1):1–34, 2003.

38   Tom Richardson and Ruediger Urbanke. *Modern coding theory*. Cambridge University Press, 2008.

39   Nicholas Ruozzi. The bethe partition function of log-supermodular graphical models. *Advances in Neural Information Processing Systems*, 25, 2012.

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

41   Allan Sly and Nike Sun. The computational hardness of counting in two-spin models on d-regular graphs. In *53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 361–369, 2012.

42   Mario Ullrich. Swendsen-Wang is faster than single-bond dynamics. *SIAM Journal on Discrete Mathematics*, 28(1):37–48, 2014.

43   Pascal O Vontobel. Counting in graph covers: A combinatorial characterization of the Bethe entropy function. *IEEE Transactions on Information Theory*, 59(9):6018–6048, 2013.

44   Jonathan S Yedidia, William T Freeman, and Yair Weiss. Understanding belief propagation and its generalizations. *Exploring artificial intelligence in the new millennium*, 8:239–269, 2003.