Fast and Slow Mixing of the Kawasaki Dynamics on Bounded-Degree Graphs

Aiya Kuchukova
School of Mathematics, Georgia Institute of Technology, Atlanta, GA, USA

Marcus Pappik
Hasso Plattner Institute, University of Potsdam, Germany

Will Perkins
School of Computer Science, Georgia Institute of Technology, Atlanta, GA, USA

Corrine Yap
School of Mathematics, Georgia Institute of Technology, Atlanta, GA, USA

Abstract

We study the worst-case mixing time of the global Kawasaki dynamics for the fixed-magnetization Ising model on the class of graphs of maximum degree Δ. Proving a conjecture of Carlson, Davies, Kolla, and Perkins, we show that below the tree uniqueness threshold, the Kawasaki dynamics mix rapidly for all magnetizations. Disproving a conjecture of Carlson, Davies, Kolla, and Perkins, we show that the regime of fast mixing does not extend throughout the regime of tractability for this model: there is a range of parameters for which there exist efficient sampling algorithms for the fixed-magnetization Ising model on max-degree Δ graphs, but the Kawasaki dynamics can take exponential time to mix. Our techniques involve showing spectral independence in the fixed-magnetization Ising model and proving a sharp threshold for the existence of multiple metastable states in the Ising model with external field on random regular graphs.

2012 ACM Subject Classification Mathematics of computing → Markov-chain Monte Carlo methods; Theory of computation → Randomness, geometry and discrete structures

Keywords and phrases ferromagnetic Ising model, fixed-magnetization Ising model, Kawasaki dynamics, Glauber dynamics, mixing time

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2024.56

Category RANDOM

Related Version *Full Version*: https://arxiv.org/abs/2405.06209 [36]

Funding *Marcus Pappik*: supported by the HPI Research School on Data Science and Engineering.
Will Perkins: supported in part by NSF grant CCF-2309708.
Corrine Yap: supported in part by NIH grant R01GM126554.

Acknowledgements The authors thank Zongchen Chen and Marcus Michelen for very helpful discussions.

1 Introduction

The Ising model on a finite graph $G = (V, E)$ is the following probability distribution on $\Omega = \{+1, -1\}^V$:

$$\mu_{G,\beta,\lambda}(\sigma) = \frac{\lambda^{|\sigma|^+} e^{\beta m_G(\sigma)}}{Z_G(\beta, \lambda)} \tag{1}$$

where $|\sigma|^+ = |\{\sigma^{-1}(+1)\}|$ is the number of vertices assigned a $+1$ spin under σ which we call the *size* of σ, and $m_G(\sigma)$ is the number of monochromatic edges in G under the 2-coloring given by $\sigma \in \Omega$. The measure $\mu_{G,\beta,\lambda}$ is called the *Gibbs measure* on G with *inverse temperature*

$\beta \geq 0$ and *external field* $\lambda \geq 0$. The normalizing constant $Z_G(\beta, \lambda) = \sum_{\sigma \in \Omega} \lambda^{|\sigma|^+} e^{\beta m_G(\sigma)}$ is the *partition function* of the Ising model. Throughout this paper, we focus on the *ferromagnetic* case, $\beta \geq 0$, in which agreeing spins on edges are preferred.

Spin models on graphs are the source of many interesting computational problems. Questions about the tractability of approximate counting (estimating the partition function) and approximate sampling (from the Gibbs distribution) are studied extensively.

In the case of the ferromagnetic Ising model, Jerrum and Sinclair [35] showed that there is a polynomial-time approximation algorithm on all graphs at all temperatures, and Randall and Wilson [42] gave an efficient sampling algorithm.

In other cases, such as the anti-ferromagnetic Ising model ($\beta < 0$) and the hard-core model of weighted independent sets, approximate counting and sampling can be computationally hard (e.g., no polynomial-time algorithm exists unless NP=RP). For the class $\mathcal{G}_\Delta$ of graphs of maximum degree Δ, these two models exhibit *computational thresholds*: as the activity or external field parameter λ varies, there is a sharp threshold between tractability (efficient approximate counting and sampling) and intractability (NP-hardness) [28, 46–48]. Moreover, the critical value $\lambda_c = \lambda_c(\Delta, \beta)$ is the phase transition point of the corresponding model on the infinite Δ-regular tree $\mathbb{T}_\Delta$ (more precisely, it is the threshold for the *uniqueness of Gibbs measure* on $\mathbb{T}_\Delta$, a notion which we discuss shortly). Thus there is a remarkable connection between computational thresholds and statistical physics phase transitions. Even further, the threshold λ_c has recently been shown to be a *dynamical threshold*: it is the threshold for rapid mixing of the Glauber dynamics, a natural Markov chain for sampling from spin models like the Ising or hard-core models, on graphs in $\mathcal{G}_\Delta$ [2, 15, 41, 48]. So in these cases, three different thresholds (computational, dynamical, uniqueness on the tree) coincide.

A very similar picture has emerged for the model of a uniformly random independent set of a given size. For the class of graphs $\mathcal{G}_\Delta$, there is a critical density $\alpha_c(\Delta)$ so that if $\alpha < \alpha_c$, there are efficient algorithms to approximately count and sample independent sets of density α, while if $\alpha > \alpha_c$ no such algorithms exist unless NP=RP [22]. Jain, Michelen, Pham, and Vuong [33] recently proved that this computational threshold α_c also marks the dynamical threshold – for $\alpha < \alpha_c$, the natural "down-up" random walk on independent sets of a given size mixes rapidly. The threshold $\alpha_c(\Delta)$ is closely connected to a uniqueness threshold on the tree: it is the smallest expected density of an independent set in the hard-core model on $G \in \mathcal{G}_\Delta$ at activity $\lambda_c(\Delta)$.

Returning to the ferromagnetic Ising model ($\beta \geq 0$), the picture is fundamentally different and not completely understood. While there is no computational threshold (there are efficient algorithms for all parameters) one can still ask about the relationship between uniqueness and dynamical thresholds. The natural dynamics in this setting are the *Glauber dynamics*, a Markov chain on the state space Ω with stationary distribution $\mu_{G,\beta,\lambda}$ which at each step chooses a uniformly random vertex and updates its spin according to the conditional distribution given the spins of its neighbors. For the case $\lambda = 1$ ("no external field") the dynamical threshold has been identified, and it coincides with the uniqueness threshold. For $\Delta \geq 3$, let the *critical inverse temperature* of the Ising model on $\mathbb{T}_\Delta$ be denoted by

$$\beta_u(\Delta) := \ln\left(\frac{\Delta}{\Delta - 2}\right).$$

The value $\beta_u(\Delta)$ is the Gibbs uniqueness threshold for the Ising model (with $\lambda = 1$) on $\mathbb{T}_\Delta$ (see e.g. [6] and below in Section 2.1 for a precise definition). Mossel and Sly [40] proved that for $0 \leq \beta < \beta_u$ and any λ, the Glauber dynamics are rapidly mixing for any $G \in \mathcal{G}_\Delta$. This threshold in β is sharp due to the analysis of the random Δ-regular graph in [23, 31]: for $\beta > \beta_u$ and $\lambda = 1$, the Glauber dynamics for the Ising model take exponential time to mix.

For general $\lambda \geq 0$, in the regime $\beta > \beta_u$, the threshold landscape is not as well understood. Note that the model is symmetric around $\lambda = 1$ by swapping the role of $+$ and $-$ spins and so for each threshold, its inverse is also a threshold; for clarity we will define thresholds for the case $\lambda \geq 1$. Let $\lambda_u(\Delta, \beta)$ be the Gibbs uniqueness threshold of the ferromagnetic Ising model on $\mathbb{T}_\Delta$; that is, λ_u is the smallest $\lambda_0 \geq 1$ so that there is a unique Gibbs measure for the Ising model on $\mathbb{T}_\Delta$ with inverse temperature β and external field λ, for all $\lambda > \lambda_0$ (again see [6] and Section 2.1 for details). The value of λ_u can be given implicitly as the solution to an equation involving Δ, β, and λ. Unlike in the above mentioned examples, while λ_u marks a phase transition on the tree, it does not mark a computational transition (since sampling from the ferromagnetic Ising model is tractable on all graphs and all parameters) and it has not been established as a dynamical threshold (though this also has not been ruled out). Below in Theorem 2 we show that the worst-case mixing time of Glauber dynamics over $\mathcal{G}_\Delta$ is exponential when $|\log \lambda| < \log \lambda_u$.

The complementary result (fast mixing of the Glauber dynamics for $G \in \mathcal{G}_\Delta$ when $|\log \lambda| > \log \lambda_u$) is not known to hold. Instead, sufficient conditions for fast mixing have been given that require λ to be somewhat larger than λ_u. An interesting insight is that upper bounds on the dynamical threshold are often connected to zero-freeness of the map $\lambda \mapsto Z_G(\beta, \lambda)$ considered as a complex polynomial. Throughout this paper, we particularly focus on the *analytic threshold* $\lambda_a(\Delta, \beta)$, defined by the following requirement: for all $G \in \mathcal{G}_\Delta$, every compact $D \subset (\lambda_a(\Delta, \beta), \infty)$ and every partial spin assignment $\tau_U : U \to \{-1, +1\}$, $U \subset V$ it holds that $Z_G^{\tau_U}(\beta, \lambda)$ (the partition function restricted to configurations that are consistent with τ_U) is non-zero for all λ in some uniform complex neighborhood of D. A formal definition of λ_a is given in Section 2.5. In contrast to the uniqueness threshold, $\lambda_a(\Delta, \beta)$ has not been determined. It is known that $\lambda_a(\Delta, \beta) \geq \lambda_u(\Delta, \beta)$ and the best known upper bound is

$$\lambda_a(\Delta, \beta) \leq \min\left\{\frac{(\Delta-2)e^{2\beta} - \Delta}{e^{\beta(2-\Delta)}}, e^{\beta\Delta}\right\} =: \bar{\lambda}_a\,. \tag{2}$$

The first expression in the minimum of (2) was proven by Shao and Sun [44], and the second bound of $e^{\beta\Delta}$ (which is smaller than the first expression for $\Delta \geq 4$ and β large enough) was proven by Shao and Ye [45].

It turns out that this analytic threshold λ_a is closely related to the dynamical threshold. More precisely, Chen, Liu, and Vigoda [17] proved that the first bound in (2) can be used to define a regime in which the ferromagnetic Ising model satisfies ℓ_∞-independence (see Section 2.4), a stronger version of spectral independence that implies rapid mixing of Glauber dynamics. Their derivation of the threshold used techniques similar to those of Shao and Sun [44] which resulted in coinciding bounds, but a more systematic connection was provided by Chen, Liu and Vigoda in [16]. They showed that for a broad class of spin systems, sufficiently strong zero-freeness assumptions imply ℓ_∞-independence. With small adjustments, we use their technique to argue that the ferromagnetic Ising model satisfies ℓ_∞-independence for all $|\log \lambda| > \log \lambda_a(\Delta, \beta)$ (see Theorem 22).

The main focus of this paper is on dynamical thresholds of the *fixed-magnetization* Ising model with inverse temperature β and magnetization η. The magnetization (per vertex) of an Ising configuration σ is $\eta(\sigma) := \frac{\sum_{v \in V(G)} \sigma_v}{|V(G)|}$. A configuration σ of magnetization η has size (number of $+1$ spins) exactly $k = \lfloor n\frac{\eta+1}{2} \rfloor$. We denote by Ω_k the configurations of size k.

The fixed-magnetization Ising model with inverse temperature $\beta \geq 0$ and magnetization $\eta \in [-1, 1]$ is then a probability distribution defined similarly to (1) but on Ω_k, where $k = \lfloor n\frac{\eta+1}{2} \rfloor$, as

$$\hat{\mu}_{G,\beta,\eta}(\sigma) = \frac{e^{\beta m_G(\sigma)}}{\hat{Z}_{G,\eta}(\beta)},$$

where

$$\hat{Z}_{G,\eta}(\beta) = \sum_{\sigma \in \Omega_k} e^{\beta m_G(\sigma)}$$

is the fixed-magnetization partition function. Here we use floors to avoid restricting to values of η where $n\frac{\eta+1}{2}$ is an integer. The distribution $\hat{\mu}_{G,\beta,\eta}$ is exactly that of $\mu_{G,\beta,\lambda}$ conditioned on the event $\{\sigma \in \Omega_k\}$. Note that the external field plays no role in the fixed-magnetization model since $\lambda^{|\sigma|^+}$ is constant on Ω_k.

In statistical physics, the fixed-magnetization Ising model is the *canonical ensemble* while the Ising model is the *grand canonical ensemble*. The fixed-magnetization model on lattices is studied in, e.g., [13, 24], where interesting geometric behavior is described; the behavior of the Kawasaki dynamics (the natural analogue of Glauber dynamics) on $\mathbb{Z}^d$ has been studied extensively in, e.g., [9–11, 38]. Here we focus on dynamical behavior over the class of all graphs of maximum degree Δ.

To understand algorithmic and dynamical thresholds in the fixed-magnetization Ising model, we need to define some further parameters. The mean magnetization of the + measure on $\mathbb{T}_\Delta$ (explained in detail in Section 2.1) is

$$\eta^+_{\Delta,\beta,\lambda} := \tanh\left(L^* + \operatorname{artanh}(\tanh(L^*)\tanh(\beta/2))\right)$$

where L^* is the largest solution to

$$L = \log(\lambda) + (\Delta - 1)\operatorname{artanh}(\tanh(L)\tanh(\beta/2)) .$$

We are specifically interested in the following three quantities:

$$\eta_c(\Delta, \beta) = \eta^+_{\Delta,\beta,1} \qquad \eta_u(\Delta, \beta) = \eta^+_{\Delta,\beta,\lambda_u} \qquad \eta_a(\Delta, \beta) = \eta^+_{\Delta,\beta,\lambda_a} .$$

For $\beta > \beta_u$, we have $0 < \eta_c < \eta_u \leq \eta_a$. It is not known if the last inequality is strict or not (just as it is not known if $\lambda_a = \lambda_u$).

Carlson, Davies, Kolla, and Perkins [12] showed recently that the fixed-magnetization Ising model exhibits quite different algorithmic behavior than the Ising model: it exhibits a computational threshold. In particular, for $\beta < \beta_u$ and any η, as well as for $\beta > \beta_u$ and $|\eta| > \eta_c$, there are efficient approximate counting and sampling algorithms for the Ising model at fixed mean magnetization η on $\mathcal{G}_\Delta$, while for $\beta > \beta_u$ and $|\eta| < \eta_c$, there are no such algorithms unless NP=RP. Thus β_u and η_c mark the computational threshold in the fixed-magnetization Ising model.

Here we study dynamical thresholds for the fixed-magnetization Ising model on $\mathcal{G}_\Delta$. Given a distribution, one candidate for an efficient approximate sampling algorithm is a Markov chain whose stationary distribution is our target distribution, but the efficiency of this algorithm depends on the mixing time. Recall that the mixing time of a Markov chain is the number of steps, in the worst-case over initial distribution, required for a Markov chain to reach 1/4 total variation distance of its stationary distribution (see Section 2.4 for a formal definition). As mentioned above, the natural dynamics associated to the fixed-magnetization

Ising model are the *Kawasaki dynamics*, which is a reversible Markov chain on Ω_k. At each step of the chain, a $+1$ vertex and a -1 vertex are chosen uniformly at random and have their spins swapped with a probability depending on the ratio of the Ising probabilities of the two configurations. This is sometimes referred to as the *global* Kawasaki dynamics, whereas the *local* Kawasaki dynamics restrict to swapping spins of neighboring vertices.

Our main contributions concern the *mixing time* of the Kawasaki dynamics. Taking $\|\mu - \nu\|_{\mathrm{TV}} := \sup_{A \in \mathcal{A}} |\mu(A) - \nu(A)|$ to be the total variation distance between probability distributions μ and ν on a probability space $(\Omega, \mathcal{A})$, the mixing time of a Markov chain on Ω that has transition matrix P and stationary distribution π is

$$\tau_{mix} := \inf \left\{ t : \max_{x \in \Omega} \|P^t(x, \cdot) - \pi\|_{\mathrm{TV}} \leq \frac{1}{4} \right\} .$$

Resolving one conjecture of Carlson, Davies, Kolla, and Perkins and disproving another (part (i) and (ii) respectively of [12, Conjecture 1]), we establish thresholds in the mean magnetization for fast and slow mixing of the Kawasaki dynamics on $\mathcal{G}_\Delta$.

▶ **Theorem 1.** *For the Kawasaki dynamics, the following two statements hold:*

(1) *If $0 \leq \beta < \beta_u$ or if $\beta > \beta_u$ and $|\eta| > \eta_a$, then the Kawasaki dynamics for $\hat{\mu}_{G,\beta,\eta}$ have mixing time $O(|V(G)|^2)$ for all $G \in \mathcal{G}_\Delta$.*

(2) *There exists a sequence of graphs $G_n \in \mathcal{G}_\Delta$ with $|V(G_n)| \to \infty$ such that for $\beta > \beta_u$ and $|\eta| < \eta_u$, the Kawasaki dynamics for $\hat{\mu}_{G,\beta,\eta}$ have mixing time* $\exp(\Omega(|V(G_n)|))$ *on G.*

Fast mixing of the dynamics for all η when $\beta < \beta_u$ was conjectured in [12]. The slow mixing for some $\eta > \eta_c$ disproves the conjecture from [12] asserting the coincidence of the algorithmic and dynamical thresholds. If it were established that $\lambda_a(\Delta, \beta) = \lambda_u(\Delta, \beta)$ then Theorem 1 would give the sharp dynamical threshold for the fixed-magnetization model. It is an interesting question to understand the dynamical threshold in both the Ising model and fixed-magnetization Ising model if instead it holds that $\lambda_u < \lambda_a$.

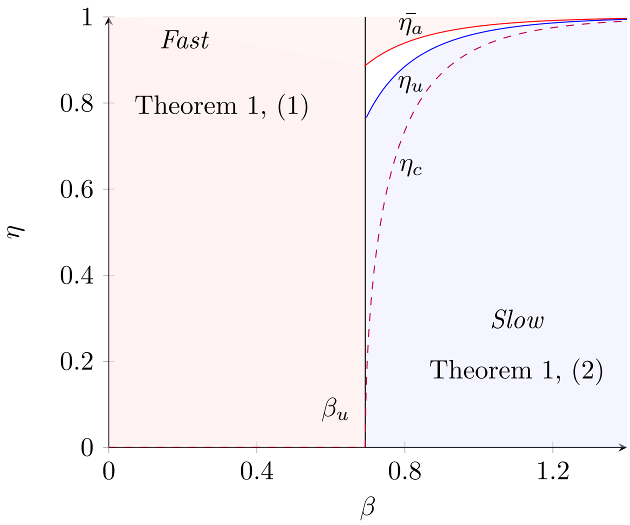

Figure 1 Sketch of the phase space for the fixed-magnetization model on $\mathcal{G}_\Delta$ when $\Delta = 4$, where $\bar{\eta_a} = \eta_{\Delta,\beta,\bar{\lambda}_a}$.

A diagram of the computational and dynamical thresholds for the fixed-magnetization Ising model is given in Figure 1.

Towards the proof of Theorem 1,(2), we establish that the Glauber dynamics for the Ising model on the random Δ-regular graph takes exponential time to mix when $\beta > \beta_u$ and λ is in the non-uniqueness regime for $\mathbb{T}_\Delta$.

▶ **Theorem 2.** *Fix $\Delta \geq 3$, $\beta > \beta_u(\Delta)$, and $|\log \lambda| < \log \lambda_u(\Delta, \beta)$. Let G be a uniformly random Δ-regular graph on n vertices. Then with high probability as $n \to \infty$, the mixing time of the Glauber dynamics for the Ising model on G is $e^{\Theta(n)}$.*

This theorem complements the result of Can, van der Hofstad, and Kumagai [8] showing that when $|\log \lambda| > \log \lambda_u$, with high probability over the random regular graph the mixing time of the Glauber dynamics is $O(n \log n)$; they conjectured that the mixing time is exponential when $|\log \lambda| < \log \lambda_u$, which Theorem 2 confirms.

Theorem 2 also fills in more of the picture for dynamical thresholds in the Ising model on graphs in $\mathcal{G}_\Delta$; see Figure 2.

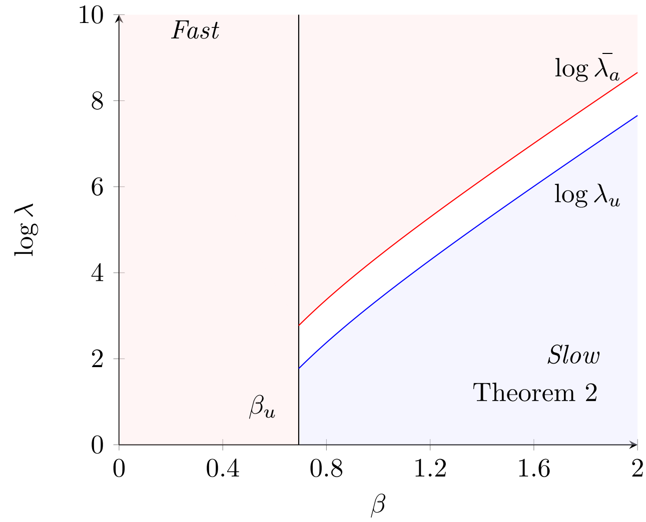

Figure 2 Sketch of the phase space for the Ising model Glauber dynamics on $\mathcal{G}_\Delta$ when $\Delta = 4$.

Before we give an overview of our proof techniques, we state some open questions. Our first question is concerned with the relation between the analytic threshold and the uniqueness threshold for the Ising model.

▶ **Question 3.** *Does $\lambda_a(\Delta, \beta) = \lambda_u(\Delta, \beta)$?*

If the answer is yes, then by the results above we would have a complete characterization of the dynamical thresholds in the Ising and fixed magnetization Ising models on $\mathcal{G}_\Delta$.

Next we conjecture the following improvement of part (1) of Theorem 1.

▶ **Conjecture 4.** *If $0 \leq \beta < \beta_u$ or if $\beta > \beta_u$ and $|\eta| > \eta_a$, then the Kawasaki dynamics for $\hat{\mu}_{G,\beta,\eta}$ are optimally mixing: the mixing time is in $O(|V(G)| \cdot \log(|V(G)|))$ for all $G \in \mathcal{G}_\Delta$.*

The analogous statement for independent sets is proved in [33] by proving a log-Sobolev inequality for the down-up walk with constant $\Omega(1/n)$.

While we focus on global Kawasaki dynamics in this paper, we suggest that our results also apply to the local dynamics. Note that for studying local Kawasaki dynamics, it makes sense to assume that G is connected. In this case, we believe that a Markov chain comparison argument as in [26] can be used to show that the mixing times of the local and global dynamics only differ by a polynomial factor. While our slow mixing result for the global dynamics uses identical copies of disjoint random graphs, our arguments should still apply if they are connected with a sparse set of edges. As a consequence, both slow and rapid mixing from Theorem 1 would carry over. A full proof of this is left for future work.

1.1 Overview of Techniques

The proofs of Theorems 1 and 2 involve several different ingredients, including local central limit theorems, spectral independence, and first- and second-moment methods for spin models on random graph. We give an overview of the techniques here.

1.1.1 Fast Mixing

At a high level, the proof of Theorem 1, (1) follows the strategy used by Jain, Michelen, Pham, and Vuong [33] to show fast mixing for the down-up walk on independent sets of density less than $\alpha_c(\Delta)$.

In order to derive an upper bound on the mixing time of the Kawasaki dynamics for the fixed-magnetization Ising model, we prove that the spectral gap of the associated transition matrix is bounded below by $\Omega(1/n)$. To achieve this, we study a related down-up Ising walk on Ω_k while arguing that the respective spectral gaps of the Kawasaki dynamics and the down-up walk are within a constant factor of each other. This allows us to make use of recent literature that relates the spectral gap of a down-up walk to spectral independence [1, 2, 14].

Informally speaking, spectral independence captures the idea that for most pairs of vertices $v, w \in V$, the spins assigned to v and w by a random configuration from $\hat{\mu}_{G,\beta,\eta}$ are almost independent. While spectral independence for the Ising model has been studied before by Chen, Liu, and Vigoda [17], no comparable result exists for the fixed-magnetization model. To derive the required spectral independence property, we follow an approach introduced in [33] to analyze the down-up walk for fixed-size independent sets. The idea is to choose λ such that a random configuration from $\mu_{G,\beta,\lambda}$ has expected magnetization per vertex close to η. We then view $\hat{\mu}_{G,\beta,\eta}$ as $\mu_{G,\beta,\lambda}$ conditioned on the desired magnetization.

We use this perspective to show that $\hat{\mu}_{G,\beta,\eta}$ satisfies ℓ_∞-independence as follows:

(1) An extremal combinatorics result on the magnetization of the Ising model from [12] shows that for any $G \in \mathcal{G}_\Delta$, the value of λ that achieves expected magnetization η satisfies $|\log \lambda| > \log \lambda_a$ if $|\eta| > \eta_a$. This allows us to use an approach by Chen, Liu, and Vigoda [16] to derive $O(1)$-ℓ_∞-independence for the Ising model for all such λ based on our zero-freeness assumption.

(2) We next show that the probability under $\mu_{G,\beta,\lambda}$ of drawing a configuration with exactly the correct magnetization is sufficiently large, and that this probability does not change significantly after conditioning on the spin of a vertex. For the former, a lower bound of $\Theta(1/\sqrt{n})$ can be derived from existing local central limit theorems for the expected number of $+1$ spins [12]. For the latter, we perform a similar analysis to [33] and use an Edgeworth expansion to prove that conditioning on the spin of a vertex changes this probability by at most $O(n^{-3/2})$. For both results it is crucial that the Ising model satisfies sufficiently strong zero-freeness assumptions for all considered λ.

The above discussion indicates how we obtain spectral independence for $\hat{\mu}_{G,\beta,\eta}$. The bulk of our work comes from leveraging this to derive a lower bound on the spectral gap of the down-up walk. This requires us to prove that spectral independence also holds when an arbitrary vertex set $U \subset V$ with $|U| < k$ is fixed (or *pinned*) to have spin $+1$. Such pinnings interfere with the proof strategy above for several reasons. First of all, pinning vertices to $+1$

decreases the λ that we need to choose to obtain the desired magnetization η. In particular, if we aim for $\eta > \eta_a$, this might cause the required value of λ to leave the regime in which zero-freeness (and ℓ_∞-independence) for the Ising model is guaranteed. We circumvent this by observing that the Kawasaki dynamics is symmetric under swapping $+1$ and -1 spins. Hence, it suffices to consider $\eta < -\eta_a$, and an application of the FKG inequality ensures that we only need to consider $\lambda < 1/\lambda_a(\Delta, \beta)$ for all relevant pinnings.

The second difficulty is that once the number of free vertices $k - |U|$ becomes sub-linear in n, both the local central limit theorem and the Edgeworth expansion can fail. Similar to [33], we solve this issue by using the localization framework by Chen and Eldan [14], which allows us to factorize the spectral gap of the down-up walk into the spectral gaps of two Markov chains that are easier to analyze. The first chain is a generalization of the down-up walk that updates $\Theta(n)$ vertices in each step, and we can analyze its spectral gap based on the spectral independence result described above using the local-to-global framework for local spectral expanders [1, 2, 15, 17]. The second walk is a simple down-up walk but with a set of vertices $U \subset V$ pinned to $+1$. In particular, we need to show that there is some $\alpha > 0$ (depending on β and Δ) such that for $k - |U| \leq \alpha n$, the spectral gap of such a pinned down-up walk is bounded below by $\Omega(1/n)$.

For bounding the spectral gap of the pinned walk, we use a coupling argument. Specifically, we construct a suitable metric on the state space such that the distance between two coupled copies of the Markov chain contracts in expectation in each step. For the independent set model studied in [33], such a contracting coupling is well known, appearing in the original "path coupling" paper of Bubley and Dyer [7]. In contrast, for the fixed-magnetization Ising model, no such result exists, and the default choice of coupling (sometimes called the identity coupling) and metric (the number of vertices on which both configurations differ) does not exhibit the desired contraction. Roughly speaking, this is because the ferromagnetism can cause certain types of disagreements to increase the probability that new disagreements are created. We overcome this problem by studying a refined metric, which assigns different weights to "good" and "bad" disagreements in a way that guarantees that distances under this new metric decrease in expectation under the coupling, thus establishing the desired bound on the spectral gap.

1.1.2 Slow Mixing

For the slow mixing results, we leverage the connection between the Ising model on the infinite tree $\mathbb{T}_\Delta$ and the behavior of the model on a uniformly random Δ-regular graph. In the relevant range of parameters ($\beta > \beta_u$, $1 < \lambda < \lambda_u$) there are two distinct Ising Gibbs measures on $\mathbb{T}_\Delta$, the "plus measure" and the "minus measure." On the random graph these two Gibbs measures manifest themselves as a dominant and subdominant metastable state: sets of configurations for which the Glauber dynamics take exponential time to escape from. The existence of multiple metastable states immediately shows slow mixing of the Glauber dynamics (Theorem 2), and we then use this to construct a graph on which the Kawasaki dynamics is slow mixing, proving Theorem 1,(2).

To do this, we exhibit the existence of a bottleneck in the state space of the model on a Δ-regular graph H constructed as the disjoint union of several copies of a random Δ-regular graph. We define two different subsets of configurations of the fixed-magnetization Ising model on H: in the set of configurations S_1, each copy of the random graph comprising H has magnetization η; in the set S_2, some copies have magnetization approximately $\eta_+ > \eta$ and some copies have magnetization approximately $\eta_- < \eta$ (chosen in such a way that their average is η). We then show that a third set S_3 separates S_1 and S_2 (under single-step updates of the Kawasaki dynamics) and carries exponentially less probability mass in the fixed-magnetization Ising model than either S_1 or S_2. Via a standard conductance argument this proves exponentially slow mixing of the Kawasaki dynamics.

Bounds on the weights of the sets S_1, S_2, and S_3 will follow from the existence of the metastable states on the random graph. One metastable state consists of configurations with magnetization close to $\eta^+_{\Delta,\beta,\lambda}$ and the other consists of configurations with magnetization close to $\eta^-_{\Delta,\beta,\lambda}$. That is, the two metastable states are in correspondence with the two distinct extremal Gibbs measures on $\mathbb{T}_\Delta$ (which is why $\lambda < \lambda_u$ is crucial).

Identifying the metastable states follows from determining which states (organized according to their magnetizations) contribute significantly to the partition function $Z_G(\beta, \lambda)$ of the Ising model on the random Δ-regular graph. A first guess about how much each state contributes to $Z_G(\beta, \lambda)$ would be to take the expected contribution. The exponential order of this expectation is captured by a function $f_{\Delta,\beta,\lambda}(\eta)$. From [29], we know that the critical points of this function correspond to fixed points of a recursion on $\mathbb{T}_\Delta$, and that the second-moment method can be used to lower bound the contribution of the state with magnetization η, where η is the maximum of $f_{\Delta,\beta,\lambda}(\eta)$. This suffices to determine the dominant state of the Ising model on the random graph (as was done in much greater generality by Dembo and Montanari in [23]).

To identify subdominant metastable states, however, we need to analyze the contribution of states with magnetization η when η is a local maximum of $f_{\Delta,\beta,\lambda}(\eta)$. For this we follow the approach of [19] utilizing non-reconstruction in planted models. While their setting is the q-state Potts model for $q \geq 3$, many of their results can be translated to our context of the external-field Ising model. We discuss their techniques in greater detail in Section 3.2 and in the full paper [36].

When we construct the graph H as the union of random graphs, we also must understand how the behavior of the fixed-magnetization Ising model relates to that of the Ising model. To do this, we give a new and simple argument in Section 3.2 to bound the probability of hitting a given magnetization in the Ising model.

Interestingly, while the graph on which we show slow mixing is the union of random regular graphs, the behavior of the Kawasaki dynamics on a single copy of the random regular graph can be very different. Recently, Bauerschmidt, Bodineau, and Dagallier [4] (see also [5]) showed that the local Kawasaki dynamics for the fixed-magnetization Ising model mixes in time $O(n \log^6 n)$ on random Δ-regular graphs at all magnetizations when $\beta < 1/(8\sqrt{\Delta - 1})$. In particular, when Δ is sufficiently large this regime of fast mixing includes parameters outside the tree uniqueness phase, i.e. inside the range of parameters for which we prove exponentially-slow mixing in the worst case over graphs in $\mathcal{G}_\Delta$.

1.2 Outline

In Section 2, we collect preliminary results that will be used in our proofs. In Section 3 we give a more detailed overview on our main steps for proving Theorem 1 and Theorem 2. In particular, in Section 3.1, we discuss our fast-mixing result, Theorem 1,(1), and in Section 3.2 we discuss our slow-mixing results, Theorem 1,(2) and Theorem 2. All proofs and more details can be found in the full version of the paper [36].

2 Preliminaries

Throughout the paper and unless otherwise stated, we will make the following assumptions: $\Delta \geq 3$ is fixed, $\beta \geq 0$, $G = (V, E) \in \mathcal{G}_\Delta$, and $n = |V|$.

We will often switch between notation of η for the magnetization per vertex and $k = \lfloor \frac{\eta+1}{2} n \rfloor$ for the number of $+1$ spins in such a configuration. We will thus abuse notation and write $\hat{\mu}_{G,\beta,k}$ for $\hat{\mu}_{G,\beta,\eta}$ and $\hat{Z}_{G,k}(\beta)$ for $\hat{Z}_{G,\eta}(\beta)$ when it makes things more clear. We will also on occasion drop G and β from the subscripts of our Gibbs measure notation as well as the subscripts and argument of our partition function notation when G and β do not play a role in the proofs.

2.1 Ising Model on the Infinite Tree

Let $\mathbb{T}_\Delta$ denote the infinite Δ-regular tree. Since it has infinitely many vertices, one cannot define the Ising model on $\mathbb{T}_\Delta$ via (1). Instead, the Dobrushin-Lanford-Ruelle equations can be used to define "infinite-volume Gibbs measures" for the Ising model and other spin models on infinite graphs. This approach says that a probability measure μ on $\{\pm 1\}^{V(\mathbb{T}_\Delta)}$ is a Gibbs measure for the Ising model at inverse temperature β and external field λ if the conditional measure on any finite set of vertices given a configuration on the complement is the Ising model defined by (1) with the appropriate boundary conditions. See [30] for more details.

A main question about Gibbs measures on infinite graphs is whether for a given specification of parameters (i.e. β and λ in the Ising case) and a given infinite graph G there is a unique Gibbs measure or multiple distinct Gibbs measures. The transition between uniqueness and non-uniqueness as a parameter varies marks a phase transition.

Understanding uniqueness and non-uniqueness of the Ising model on $\mathbb{T}_\Delta$ is relatively simple because of monotonicity and the FKG inequality. There are two extreme infinite-volume Gibbs measures in the sense of maximizing or minimizing the probability that a fixed vertex of $\mathbb{T}_\Delta$ gets a $+1$ spin: the "$+$ measure" on $\mathbb{T}_\Delta$ is the Gibbs measure realized by taking a weak limit of finite-volume Gibbs measures on depth N truncations of $\mathbb{T}_\Delta$ with boundary vertices assigned $+1$ spins; the "$-$ measure" is the weak limit of finite-volume measures with boundary vertices receiving -1 spins.

The quantities $\eta^+_{\Delta,\beta,\lambda}$ and $\eta^-_{\Delta,\beta,\lambda}$ are the respective expectations of σ_v (for any fixed v in $\mathbb{T}_\Delta$) under these two Gibbs measures. The quantities can be calculated as solutions to fixed point equations (see e.g. [6]), giving

$$\eta^+_{\Delta,\beta,\lambda} = \tanh\left(L^* + \operatorname{artanh}(\tanh(L^*)\tanh(\beta/2))\right)$$

where L^* is the largest solution to

$$L = \log(\lambda) + (\Delta - 1)\operatorname{artanh}(\tanh(L)\tanh(\beta/2))\,.$$

The following proposition summarizes information about $\eta^+_{\Delta,\beta,\lambda}$, $\eta^-_{\Delta,\beta,\lambda}$ and Gibbs uniqueness that we will use (all follow from the results in [6]).

▶ **Proposition 5.** *Fix $\Delta \geq 3$.*

- *There is uniqueness of Gibbs measure for the Ising model with parameters β, λ on $\mathbb{T}_\Delta$ if and only if $\eta^+_{\Delta,\beta,\lambda} = \eta^-_{\Delta,\beta,\lambda}$.*
- *For $\beta \leq \beta_u(\Delta) = \ln\left(\frac{\Delta}{\Delta-2}\right)$, there is uniqueness for all λ.*
- *For $\beta > \beta_u(\Delta)$ there is $\lambda_u > 1$ so that there is uniqueness if and only if $|\log\lambda| > \log\lambda_u$.*
- *$\eta^+_{\Delta,\beta,\lambda}$ is continuous and strictly increasing in λ on the interval $[1,\infty)$. In particular, recall that $\eta_c(\Delta,\beta) = \eta^+_{\Delta,\beta,1}$ and $\eta_u(\Delta,\beta) = \eta^+_{\Delta,\beta,\lambda_u}$; then for every $\eta \in [\eta_c, \eta_u]$ there is $\lambda \in [1, \lambda_u]$ so that $\eta^+_{\Delta,\beta,\lambda} = \eta$.*

Finally, it will be important to bound the expected magnetization in the Ising model for given β, λ and any $G \in \mathcal{G}_\Delta$. The bound is an extremal result proved in [12].

▶ **Theorem 6** ([12, Theorem 3]). *For $G \in \mathcal{G}_\Delta$, $\lambda \geq 1$, and $\beta \geq 0$,*

$$\mathbb{E}_{\sigma \sim \mu_{G,\beta,\lambda}}[\eta(\sigma)] \leq \eta^+_{\Delta,\beta,\lambda}.$$

2.2 Pinned Models

For the fast-mixing argument, we will frequently consider pinned versions of our models, meaning conditioned on some subset of vertices having been assigned a particular spin. For $U \subset V$, we call a function $\tau_U : U \to \{+1, -1\}$ a *pinning* on U. We write $\Omega^{\tau_U} = \{\sigma \in \Omega \mid \forall u \in U : \sigma(u) = \tau_U(u)\}$ for the set of Ising configurations on G that agree with τ_U on U. The *Ising partition function with pinning* τ_U is defined as

$$Z_G^{\tau_U}(\beta, \lambda) = \sum_{\sigma \in \Omega^{\tau_U}} \lambda^{|\sigma|^+} e^{\beta m_G(\sigma)},$$

and the *Ising model under pinning* τ_U is defined by Gibbs measure

$$\mu^{\tau_U}_{G,\beta,\lambda}(\sigma) = \frac{\mathbb{1}_{\sigma \in \Omega^{\tau_U}} \lambda^{|\sigma|^+} e^{\beta m_G(\sigma)}}{Z_G^{\tau_U}(\beta, \lambda)}.$$

Note that for $\lambda > 0$, it holds that $\mu^{\tau_U}_{\beta,\lambda}$ is a well-defined probability distribution with support Ω^{τ_U}. We allow for the case $U = \emptyset$, which is equivalent to the unpinned Ising model. Often, τ_U will be the constant $+1$ function on U, in which case we write Ω^U, Z_G^U and $\mu^U_{\beta,\lambda}$.

Analogously to the Ising model, we will also impose pinnings on the fixed-magnetization model. To this end, set $\Omega_k^{\tau_U} = \{\sigma \in \Omega_k : \forall u \in U : \sigma(u) = \tau_U(u)\}$ and define the *fixed-magnetization partition function with pinning* τ_U as

$$\hat{Z}^{\tau_U}_{G,k}(\beta) = \sum_{\sigma \in \Omega_k^{\tau_U}} e^{\beta m_G(\sigma)}.$$

The *fixed-magnetization Ising model under pinning* τ_U is a probability measure with support $\Omega_k^{\tau_U}$ defined by

$$\hat{\mu}^{\tau_U}_{G,\beta,k}(\sigma) = \frac{\mathbb{1}_{\sigma \in \Omega_k^{\tau_U}} e^{\beta m_G(\sigma)}}{\hat{Z}^{\tau_U}_{G,k}(\beta)}.$$

Throughout the paper, we assume $|\tau_U|^+ \leq k$ so that the expression above is well-defined. As with the Ising model, we write Ω_k^U, $\hat{Z}^U_{G,k}$ and $\hat{\mu}^U_{G,\beta,k}$ when τ_U is the constant $+1$ function.

2.3 Kawasaki Dynamics, Down-up Walk, and Glauber Dynamics

Here we formally define the three Markov chains that we will analyze. Our main object of study is the Kawasaki dynamics for the fixed-magnetization Ising model. For this, we fix a size k where $1 \leq k \leq |V| - 1$.

▶ **Definition 7** (Kawasaki dynamics). *The Kawasaki dynamics on Ω_k is a Markov chain $\mathcal{K}_{\beta,k} = (X_t)_{t \geq 0}$ given by the following update rule:*

1. *Pick $u \in X_t^{-1}(+1)$ and $w \in X_t^{-1}(-1)$ uniformly at random, and set $X \in \Omega_k$ such that $X(v) = X_t(w), X(w) = X_t(v)$, and $X(u) = X_t(u)$ for $u \neq v, w$.*
2. *Set $X_{t+1} = X$ with probability $\min\left\{1, \frac{\hat{\mu}_{G,\beta,k}(X)}{\hat{\mu}_{G,\beta,k}(X_t)}\right\}$, and set $X_{t+1} = X_t$ otherwise.*

In other words, the Kawasaki dynamics chooses two vertices with opposite spins and swaps their spins with probability proportional to the change in monochromatic edges.

For proving fast mixing of the Kawasaki dynamics, we use the down-up walk on the +1 spins as a proxy for our analysis. Here we will also need to consider the Markov chain under plus pinnings.

▶ **Definition 8** (Down-up walk with plus pinnings). *For $U \subset V$ and with $|U| < k$ we define the* +1-down-up walk *on Ω_k^U as a Markov chain $\mathcal{P}_{\beta,k}^U = (Y_t)_{t\geq 0}$, given by the following update rule:*

1. *Pick $v \in Y_t^{-1}(+1) \setminus U$ uniformly at random and set $W = Y_t^{-1}(+1) \setminus \{v\}$.*
2. *Draw Y_{t+1} from $\hat{\mu}_{G,\beta,k}^W$.*

We write $\mathcal{P}_{\beta,k}$ if $U = \emptyset$.

The following observation is easy to check.

▶ **Observation 9.** *$\mathcal{K}_{\beta,k}$ and $\mathcal{P}_{\beta,k}$ are ergodic and reversible with respect to $\hat{\mu}_{\beta,k}$. Moreover, there is a constant $C \geq 1$ that only depends on Δ and β such that for all $\sigma_1 \neq \sigma_2$*

$$\frac{1}{C} \cdot \mathcal{P}_{\beta,k}(\sigma_1, \sigma_2) \leq \mathcal{K}_{\beta,k}(\sigma_1, \sigma_2) \leq C \cdot \mathcal{P}_{\beta,k}(\sigma_1, \sigma_2).$$

Lastly, we also consider the Glauber dynamics for the Ising model.

▶ **Definition 10** (Glauber dynamics). *The Glauber dynamics on Ω is a Markov chain $(X_t)_{t\geq 0}$, given by the following update rule:*

1. *Pick $v \in V(G)$ uniformly at random.*
2. *For $u \neq v$, set $X_{t+1}(u) = X_t(u)$, and sample $X_{t+1}(v)$ from the marginal distribution at v conditioned on $X_{t+1}(N(v))$.*

2.4 Mixing Times

Our goal in analyzing the Kawasaki dynamics is to understand the *mixing time* of this Markov chain. Given two probability distributions μ and ν on probability space $(\Omega, \mathcal{A})$, let

$$\|\mu - \nu\|_{\mathrm{TV}} := \sup_{A \in \mathcal{A}} |\mu(A) - \nu(A)|$$

be the *total variation distance* between μ and ν. For a Markov chain on Ω with transition matrix P and unique stationary distribution π, we may then define

$$d(t) := \max_{x \in \Omega} \|P^t(x, \cdot) - \pi\|_{\mathrm{TV}}.$$

▶ **Definition 11.** *The* mixing time *is*

$$\tau_{mix} = \inf\left\{t : d(t) \leq \frac{1}{4}\right\}.$$

See, e.g., [37] for background on Markov chains and mixing times. We use several different techniques to analyze the mixing time of the Kawasaki dynamics, which we now describe.

2.4.1 Upper Bounds on Mixing Time

A common way to upper-bound the mixing time of a reversible Markov chain P is by lower-bounding its spectral gap, which can be defined via the following variational characterization.

▶ **Definition 12.** *Let P be a transition matrix that is reversible with respect to π. We denote by $\mathfrak{gap}(P)$ the* spectral gap *(or Poincaré constant) of P, which is defined as the largest constant γ such that $\gamma \mathrm{Var}_\pi(f) \leq \mathcal{E}_P(f,f)$ for any function $f : \Omega \to \mathbb{R}$, where $\mathrm{Var}_\pi(f)$ is the variance of f with respect to π and $\mathcal{E}_P$ is the Dirichlet form of P, given by*

$$\mathcal{E}_P(f,g) = \frac{1}{2} \sum_{x,y \in \Omega} (f(x) - f(y))(g(x) - g(y))P(x,y)\pi(x) \qquad f, g : \Omega \to \mathbb{R}.$$

Using this characterization of the spectral gap, we have the following observation.

▶ **Observation 13.** *Suppose P_1 and P_2 are transition matrices that are both reversible with respect to π. If there are constants $\alpha_1, \alpha_2 > 0$ such that $\alpha_1 \cdot P_1(x,y) \leq P_2(x,y) \leq \alpha_2 \cdot P_1(x,y)$ for all $x \neq y$, then $\alpha_1 \cdot \mathfrak{gap}(P_1) \leq \mathfrak{gap}(P_2) \leq \alpha_2 \cdot \mathfrak{gap}(P_1)$.*

On account of Observation 9, this allows us to study the spectral gap of the down-up walk $\mathcal{P}_{\beta,k}$ instead of the Kawasaki dynamics $\mathcal{K}_{\beta,k}$.

An upper bound on the mixing time of an ergodic, reversible Markov chain with transition matrix P can be obtained from its spectral gap via the following standard relationship (see [37, Theorem 12.4]):

$$\tau_{\mathrm{mix}} \leq \mathfrak{gap}(P)^{-1} \cdot \log\left(\frac{4}{\min_{x \in \Omega} \pi(x)}\right).$$

There are various ways to obtain bounds on the spectral gap of a Markov chain, one of which is to construct a contracting coupling. For a transition matrix P, we say that a Markov chain $(X_t, Y_t)_{t \geq 0}$ on $\Omega \times \Omega$ is a *coupling* of P with itself if each of the marginal processes $(X_t)_{t \geq 0}$ and $(Y_t)_{t \geq 0}$ is a Markov chain with transition matrix P. We use this notion to bound the spectral gap.

▶ **Theorem 14** ([37, Theorem 13.1])**.** *Suppose Ω is finite and let $(X_t, Y_t)_{t \geq 0}$ be a coupling of P with itself. If there is a constant $c > 0$ and a function $\rho : \Omega \times \Omega \to \mathbb{R}_{\geq 0}$ such that $\rho(x,y) = 0$ if and only if $x = y$, and for all $t \in \mathbb{Z}_{\geq 0}$ it holds that*

$$\mathbb{E}[\rho(X_{t+1}, Y_{t+1}) \mid X_t, Y_t] \leq (1 - c)\rho(X_t, Y_t),$$

then the spectral gap of P is at least c.

We will use Theorem 14 to show that the down-up walk $\mathcal{P}^U_{\beta,k}$ has a spectral gap of $\Omega(1/k)$ whenever $k - |U| \leq \alpha n$ for some α depending on Δ and β. In particular, by the symmetry of the Kawasaki dynamics under swapping all spins, this proves a spectral gap of $\Omega(1/k)$ for $\mathcal{K}_{\beta,k}$ if $k \leq \alpha n$ or $k \geq (1 - \alpha)n$, but it does not cover the full regime of Theorem 1,(1).

To prove the full result of Theorem 1,(1), we prove that $\hat{\mu}^U_{\beta,k}$ satisfies spectral independence for suitable $k \in \mathbb{N}$ and sets $U \subset V$. Spectral independence is a property of the stationary distribution π of a Markov chain, and it was recently used to bound the spectral gap and prove rapid mixing of various chains [1,2,14,15,17,33]. For the following discussion of spectral independence, we restrict ourselves to distributions on $\Omega = 2^V$ where V is some finite set (e.g., the vertices of a graph). Note that this encompasses both the fixed-magnetization Ising model as well as the Ising model, by associating $S \in \Omega$ with the Ising configuration that maps all vertices in S to $+1$. In this setting, we adopt the following notation: for a distribution π on Ω, a subset S drawn from π, and $v \in V$, we write $\pi(v)$ for the probability that $v \in S$ and $\pi(\overline{v})$ for the probability that $v \notin S$. We extend this to conditional probabilities, writing for example $\pi(v \mid \overline{u})$ for the probability that $v \in S$ given $u \notin S$.

▶ **Definition 15.** *The* influence matrix *of a distribution π on 2^V is the matrix $M_\pi \in \mathbb{R}^{n\times n}$ with entries*

$$M_\pi[u,v] = \begin{cases} 0 & \text{if } \pi(u) = 0 \\ \pi(v \mid u) - \pi(v) & \text{otherwise} \end{cases}$$

Using this definition of M_π, the definition of spectral independence of π is as follows.

▶ **Definition 16.** *A probability distribution π on 2^V is called C-*spectrally independent *(for $C \geq 0$) if the largest eigenvalue of M_π is at most C.*

Since directly bounding the largest eigenvalue of M_π is usually challenging, a common approach is to bound the ℓ_∞-norm of M_π instead. This leads to the stronger notion of ℓ_∞-independence.

▶ **Definition 17.** *A probability distribution π on 2^V is C-ℓ_∞-*independent *(for $C \geq 0$) if*

$$\|M_\pi\|_\infty := \max_{u\in V} \sum_{v\in V} |M_\pi[u,v]|$$

is at most C.

▶ Remark 18. There are various definitions of the pairwise influence matrix in the literature [2,15,17]. For spin systems with two possible states for each vertex (such as the Ising model), pairwise influence is commonly defined as $M_\pi[u,v] = \pi(v \mid u) - \pi(v \mid \overline{u})$. However, note that switching between the two definitions only changes the spectral radius by some constant factor, provided that $\pi(v)$ is uniformly bounded away from 0 and 1. Since this is the case for the Ising model, given that $\lambda > 0$, existing spectral independence results such as [17] carry over to our definition. Moreover, Definition 15 is arguably more natural for canonical ensembles, such as the fixed-magnetization Ising model, as it relates more directly to local spectral expansion of the associated simplicial complex (see [36] for details).

There are different ways to derive bounds on the spectral gap of a Markov chain from spectral independence. The most popular approach is the use of *local-to-global theorems*, which are applicable whenever the Markov chain in question can be represented as a down-up walk on a suitable weighted simplicial complex [1,2,15,17]. Local-to-global theorems allow us to express the spectral gap of the down-up walk in terms of spectral gaps of local walks on the complex, which can then be related to the spectrum of the pairwise influence matrix.

A more recent framework was introduced by Chen and Eldan [14] and uses *localization schemes*. A localization scheme maps a probability distribution π on Ω to a localization process – a random sequence of probability measures that interpolates between π and a random Dirac measure. Via the localization process, a localization scheme gives rise to a Markov chain with stationary distribution π. Provided that the localization process exhibits a property called "approximate conservation of variance," this can be used to bound the spectral gap of the associated Markov chain. For a broad class of localization schemes, approximate conservation of variance follows if all measures along the localization process exhibit spectral independence. Since we are studying the fixed-magnetization Ising model, we are particularly concerned with distributions π on Ω_k. In this setting, the canonical choice for a localization scheme is the subset simplicial-complex localization (see [14, Example 5]), and the natural associated Markov chain is the down-up walk $\mathcal{P}_{\beta,k}$.

The main difficulty in applying the above frameworks in our setting is that they usually assume $O(1)$-spectral independence of the pinned distributions $\hat{\mu}^U_{\beta,k}$ for all $U \subset V$ with $0 \leq |U| \leq k-1$. Unfortunately, we will not be able to derive spectral independence for all such U. Moreover, for the localization framework, it is not clear if the subset simplicial-complex localization allows us to derive approximate conservation of variance from spectral independence. To overcome these difficulties, we use an argument similar to that of Jain, Michelen, Pham and Vuong [33]. We combine the techniques above as follows: first, we use a localization scheme to show that for any $\ell \leq k-1$, the spectral gap of $\mathcal{P}_{\beta,k}$ is bounded below by the product of the spectral gap of the pinned down-up walk $\mathcal{P}^U_{\beta,k}$ for any $U \subset V$ with $|U| = \ell$ and the spectral gap of the (k, ℓ)-down-up walk, a modified version of $\mathcal{P}_{\beta,k}$ that resamples $k-\ell$ plus spins in each step. Choosing ℓ such that $k-\ell \leq \alpha n$ for some suitable $\alpha > 0$, we can use a coupling argument as discussed before to show that $\mathfrak{gap}(\mathcal{P}^U_{\beta,k}) \in \Omega(1/k)$ for every $U \subset V$ with $|U| = \ell$. To lower-bound the spectral gap of the (k, ℓ)-down-up walk, we use a local-to-global theorem by Chen, Liu and Vigoda [15]. This only requires us to show that $\hat{\mu}^W_{\beta,k}$ satisfies $O(1)$-spectral independence for all $W \subset V$ with $k - |W| \geq \alpha' n$ for some $0 < \alpha' < \alpha$. The range of k for which we can show this $O(1)$-spectral independence leads to the magnetization range given in Theorem 1,(1).

2.4.2 Lower Bounds on Mixing Time

To prove slow mixing, we exhibit the existence of a bottleneck in the state space, a set of configurations which separates two parts of the state space and carries an exponentially smaller probability in the stationary distribution than either of the two parts. The following lemma captures a simple form of this argument, often phrased in terms of *conductance*, for proving lower bounds on the mixing times of Markov chains.

▶ **Lemma 19.** *Let $(X_t)_{t\geq 0}$ be a Markov chain on the state space Ω with stationary distribution π. Suppose there exists disjoint sets $S_1, S_2, S_3 \subset \Omega$ so that the following hold:*

- *For the chain to pass from S_2 to S_1 it must pass through S_3;*
- *$\pi(S_1) \geq \pi(S_2)$*
- *$\pi(S_3) \leq e^{-\Omega(n)}\pi(S_2)$.*

Then the mixing time of the chain (X_t) is $\exp(\Omega(n))$.

The statement is an immediate corollary of, e.g., [25, Claim 2.3].

To prove Theorem 2, we define S_1, S_2, S_3 to be sets of configurations with certain magnetizations. S_1 will be those configurations whose magnetization per vertex is close to that of the plus measure on $\mathbb{T}_\Delta$ (when $\lambda > 1$); S_2 will be those whose magnetization per vertex is close to that of the minus measure; and S_3 will be configurations whose magnetization is just larger than that of S_2.

To prove Theorem 1,(2), we consider a graph H made up of disjoint copies of a random regular graph. We define S_1 to be the set of configurations with magnetization η on each copy; S_2 will be a set of configurations with magnetization η_+ on some copies and η_- on others, for $\eta_- < \eta < \eta_+$, such that the overall magnetization is η. Again S_3 will be a neighborhood of S_2. In both cases, the main work will be in verifying the conditions of Lemma 19.

2.5 Thresholds for Zero-Freeness and Spectral Independence

The definition of $\lambda_a(\Delta, \beta)$ is based on viewing the Ising partition function as a polynomial in the (complex) variable λ. We write $\mathcal{N}(z, \delta)$ for the open ball of radius $\delta > 0$ around $z \in \mathbb{C}$.

▶ **Definition 20** (Absolute zero-freeness). *Given $\beta \geq 0$, $\Delta \in \mathbb{N}$, $\lambda > 0$ and $\delta > 0$, we say that the Ising model is* absolutely δ-zero-free at activity λ *if for all graphs $G \in \mathcal{G}_\Delta$, all pinnings τ_U with $U \subseteq V$ and all $\lambda' \in \mathcal{N}(\lambda, \delta)$ it holds that $Z_G^{\tau_U}(\beta, \lambda') \neq 0$.*

We now define $\lambda_a(\Delta, \beta)$ as follows.

▶ **Definition 21.** *For $\Delta \in \mathbb{N}$ and $\beta \geq \beta_u(\Delta)$ we set $\lambda_a(\Delta, \beta)$ to be the smallest $\lambda_a \geq 1$ such that for every compact set $D \subset (\lambda_a, \infty)$ there is some $\delta > 0$ such that for all $\lambda \in D$ the Ising model is absolutely δ-zero-free at λ.*

An important implication of absolute zero-freeness is given in the following theorem. Its proof follows a similar argument to those in [16] while using the ferromagnetism of the model and Montel's theorem (see [49]) to avoid the requirement of multivariate zero-freeness. The proof can be found in the full version of the paper [36].

▶ **Theorem 22.** *Fix $\beta \geq 0$ and $\Delta \in \mathbb{N}$. Let $D \subset \mathbb{R}_{>0}$ be compact and assume there is some $\delta > 0$ such that the ferromagnetic Ising model is absolutely δ-zero-free at every $\lambda \in D$. Then, there is some constant $C > 0$, only depending on D, λ, β and Δ, such that for all $\lambda \in D$, $G \in \mathcal{G}_\Delta$ and all pinnings τ_U it holds that $\hat{\mu}_{G,\beta,\lambda}^{\tau_U}$ is C-ℓ_∞-independent.*

3 Main Statements and Proof Structure

We briefly state the most important steps for showing Theorem 1. All proofs and intermediate steps are omitted and can be found in the full version of the paper [36].

3.1 Rapid Mixing

We start with discussing our proof of the rapid mixing result in part (1) of Theorem 1. The structure of the entire proof is illustrated in Figure 3.

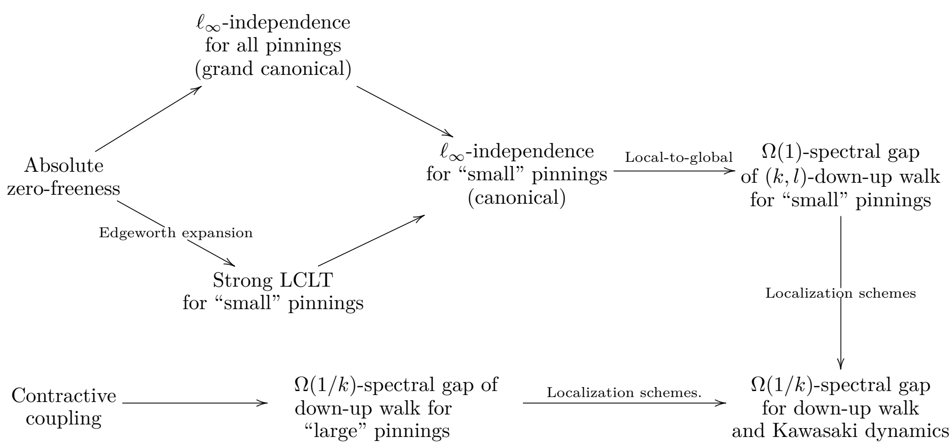

Figure 3 The structure of the rapid mixing proof.

All results in this subsection are given in the context of the following assumptions.

▶ **Condition 23.**

1. *Let $\beta \geq 0$, and let $D \subset \mathbb{R}_{>0}$ be compact such that there is some $\delta > 0$ for which the Ising model is absolutely δ-zero-free for all $\lambda \in D$. Further, let $\lambda \in D$.*
2. *Let $\alpha \in [0, 1)$, let $U \subset V$ with $|U| \leq \alpha n$ and let τ_U be a pinning of U.*
3. *Let $\sigma \sim \mu_{\beta,\lambda}^{\tau_U}$ and let $X = |\sigma|^+$.*

Our first step is to show that zero-freeness implies a strengthened version of a local central limit theorem for X via Edgeworth expansion. Using similar arguments as Jain, Michelen, Pham and Vuong [33] for the hard-core model, we obtain the following result.

▶ **Theorem 24.** *Suppose Condition 23 holds. Let $d \in \mathbb{N}$ and let $\ell \in \mathbb{R}$ such that $\mathbb{E}[X] + \ell \in \mathbb{Z}_{\geq 0}$. Set $s = \sqrt{\mathrm{Var}(X)}$ and $\beta_j = \frac{\kappa_j(X)}{j! s^j}$ for all $j \in \mathbb{N}$, and write $H_k(\cdot)$ for the k^{th} Hermite polynomial. It holds that*

$$\mu_{\beta,\lambda}^{\tau_U}(X - \mathbb{E}[X] = \ell) = \frac{e^{-\frac{\ell^2}{2s^2}}}{\sqrt{2\pi}s}\left(1 + \sum_{r\geq 3} H_r(\ell/s) \sum_{k_3,\ldots,k_{2d+1}} \prod_{j=3}^{2d+1} \frac{\beta_j^{k_j}}{k_j!}\right) + O\left(n^{-d}\right)$$

where the inner sum is over tuples $k_3, \ldots, k_{2d+1} \in \mathbb{Z}_{\geq 0}$ such that $\sum_j k_j \cdot j = r$ and $\sum_j k_j \cdot \frac{j-2}{2} \leq d$, and the implied constants depend only on Δ, β, δ, D, d and α.

Our next ingredient is to use zero-freeness to obtain a stability result for the cumulants of X under adding vertices to the pinning. Writing $\kappa_j(X)$ for the jth cumulant of X, we have the following statement.

▶ **Lemma 25.** *Suppose Condition 23 holds. Let $v \in V \setminus U$, and let $\tau_U, +_v$ denote the pinning on $U \cup \{v\}$ that maps v to $+1$ and all other vertices $u \in U$ to $\tau_U(u)$. Let $X^+ = |\sigma'|^+$ for $\sigma' \sim \mu_{\beta,\lambda}^{\tau_U,+_v}$. For all $j \in \mathbb{N}$ it holds that $|\kappa_j(X^+) - \kappa_j(X)| = O(1)$ with implied constants only depending on Δ, β, δ, D and j.*

The analog of Lemma 25 for the hard-core model was proven in [33]. However, their arguments are tailored to the hard-core model and do not apply in our setting. Instead, we provide a more general argument based on an application of Montel's theorem that is inspired by [43].

Using Theorem 24 and Lemma 25, we get the following stability result for the probability of having exactly k vertices assigned to $+1$.

▶ **Lemma 26.** *Suppose Condition 23 holds and assume further that $|U| + 1 \leq \alpha n$. Let $k \in \mathbb{Z}_{\geq 0}$ be such that $|\mathbb{E}[X] - k| \leq L$ for some $L \in \mathbb{R}_{\geq 0}$. For all $v \in V \setminus U$ it holds that*

$$\mu_{\beta,\lambda}^{\tau_U}(X = k) = \Theta(n^{-1/2}), \tag{3}$$

$$\left|\mu_{\beta,\lambda}^{\tau_U}(X = k) - \mu_{\beta,\lambda}^{\tau_U}(X = k \mid \sigma(v) = +1)\right| = O(n^{-3/2}) \tag{4}$$

with implied constants depending only on Δ, β, δ, D, L and α.

Next, recall that by Theorem 22 zero-freeness implies ℓ_∞-independence for the ferromagnetic Ising model. Combining this with Lemma 26 for a suitable λ, we get the following ℓ_∞-independence result for the fixed magnetization model.

▶ **Theorem 27.** *Assume $0 \leq \beta < \beta_u(\Delta)$ and $\gamma \in (0, 1/2]$, or $\beta \geq \beta_u(\Delta)$ and $\gamma \in (0, \frac{1-\eta_a}{2})$ for $\eta_a = \eta_a(\Delta, \beta)$. For all $k \coloneqq \gamma n \in \mathbb{N}$, all $\alpha \in [0, \gamma)$ and $U \subset V$ with $|U| \leq \alpha n$ it holds that $\hat{\mu}_{\beta,k}^U$ is C-ℓ_∞-independent for a constant C depending only on Δ, β, γ and α.*

Using Theorem 27, we can apply a local-to-global theorem from [15] to show that for every $k - \ell \in \Theta(n)$ the spectral gap of the (k, ℓ)-down-up walk is in $\Omega(1)$. However, to get the desired spectral gap for $\mathcal{P}_{\beta,k}$ (and $\mathcal{K}_{\beta,k}$), we require one last ingredient, which is to show that the spectral gap of the pinned down-up walk $\mathcal{P}^U_{\beta,k}$ is in $\Omega(1/n)$ whenever $k = \gamma n$ and $U \subset V$ are such that $k - |U|$ is small enough.

In the setting of fixed-size independent sets studied in [33], such a result was previously known due to Theorem 14 and a path coupling by Bubley and Dyer [7]. In contrast, a straightforward application of path coupling with the Hamming metric does not work in our setting. Instead, we introduce a modified metric on the state space that takes into account how likely a disagreement is to spread, which allows us to prove the following result.

▶ **Lemma 28.** *Let $G \in \mathcal{G}_\Delta$ with n sufficiently large. There exists some $\alpha = \alpha(\Delta, \beta) > 0$ such that for all $0 < k \leq n/2$ and all $U \subset V$ with $0 < k - |U| \leq \alpha n$ it holds that $\mathcal{P}^U_{\beta,k}$ has spectral gap $\Omega(1/k)$ with constants depending on β and Δ.*

We can now proceed to sketch our proof of the rapid mixing part of Theorem 1. We first note that the Kawasaki dynamics Markov chain is invariant under swapping all spins (i.e. the mapping $\sigma \mapsto -\sigma$), allowing us to focus on $k \leq n/2$ (or equivalently the magnetization regime $\eta \leq 0$). Moreover, by Observation 9 it suffices to prove the desired spectral gap for the down-up walk $\mathcal{P}_{\beta,k}$ for the respective values of k. Using a localization schmeme, we argue that the spectral gap of $\mathcal{P}_{\beta,k}$ is bounded below by the product of $\inf_{U \in \binom{V}{\ell}} \mathfrak{gap}(\mathcal{P}^U_{\beta,k})$ and the spectral gap of the (k, ℓ)-down-up walk. By Lemma 28, we know that $\inf_{U \in \binom{V}{\ell}} \mathfrak{gap}(\mathcal{P}^U_{\beta,k}) \in \Omega(1/k)$ whenever ℓ is such that $k - \ell \leq \alpha n$ for some $\alpha = \alpha(\Delta, \beta) > 0$. Moreover, by Theorem 27 and a local-to-global theorem from [15], we can derive a $\Omega(1)$ spectral gap for the (k, ℓ)-down-up walk. Combining both concludes our rapid mixing proof.

3.2 Metastability and Slow Mixing

In this section we prove slow-mixing results for both the Ising Glauber dynamics and fixed magnetization Kawasaki dynamics when $\beta > \beta_u(\Delta)$ and $|\log \lambda| < \log \lambda_u$ and $|\eta| < \eta_u$ respectively. The structure of the proof is illustrated below in Figure 4.

$$\begin{tikzcd}
\text{Non-uniqueness on } \mathbb{T}_\Delta \arrow[r] & \begin{array}{c}\text{Metastability on}\\ \text{random graphs}\end{array} \arrow[r] \arrow[d] & \begin{array}{c}\text{Slow mixing of}\\ \text{Kawasaki dynamics}\end{array} \\
 & \begin{array}{c}\text{Slow mixing of}\\ \text{Glauber dynamics}\end{array} &
\end{tikzcd}$$

Figure 4 The structure of the slow mixing proof.

Note. As in the previous section, both perspectives of fixed magnetization per vertex η and fixed size k will be useful in our arguments. We will use $Z_{G,\eta}(\beta, \lambda)$ (where we sometimes drop the parameters β and λ for convenience) to denote the contribution to the Ising model partition function $Z_G(\beta, \lambda)$ from configurations of magnetization η. The notation $Z_{G,k}(\beta, \lambda)$ will mean the contributions to $Z_G(\beta, \lambda)$ from configurations of size k. When $k = \lfloor n\frac{\eta+1}{2} \rfloor$, we have $Z_{G,\eta} = Z_{G,k}$ and will use the notations interchangeably.

Our goal is to understand how configurations of different magnetizations typically contribute to the partition function $Z_G(\beta, \lambda)$ when G is a random Δ-regular graph. To start, we shift to a slightly different model called the *configuration model*, which we will denote $\mathbf{G}$. To generate a graph from this model for a given Δ and n, take Δ copies of $[n]$ and a uniformly random perfect matching on the Δn vertices, and then identify the copies corresponding to the same vertex. This gives a random Δ-regular multigraph, and it is well-known that properties holding with high probability for the configuration model also hold with high probability for the uniform random Δ-regular graph when Δ is constant [34].

We say the model has multiple *metastable states* if the function $\lim_{n\to\infty} \frac{1}{n}\mathbb{E}\log Z_{G,\eta}(\beta,\lambda)$ has more than one local maximum as η varies. A first attempt at understanding this phenomenon would be to look at the first moment, and understand the local maxima of

$$f_{\Delta,\beta,\lambda}(\eta) := \lim_{n\to\infty} \frac{1}{n} \log \mathbb{E} Z_{G,\eta}(\beta,\lambda) \tag{5}$$

as a function of η (with the crucial distinction between the two functions being the interchange of the expectation and logarithm).

Using computations similar to those found in [18, 19, 29], we can derive an expression for $f_{\Delta,\beta,\lambda}(\eta)$. We then proceed by studying the the maxima of $f_{\Delta,\beta,\lambda}(\eta)$ as a one-variable function with respect to η. By a result in [29] (following [27, 41]), we know that the critical points of $f_{\Delta,\beta,\lambda}(\eta)$ correspond exactly to fixed points of the *tree recursion* for the Ising model on $\mathbb{T}_\Delta$, which are the solutions to the equation

$$R = \frac{\lambda(Re^\beta + 1)^{\Delta-1}}{(R + e^\beta)^{\Delta-1}}\,. \tag{6}$$

▶ **Theorem 29** ([29, Theorem 9, Lemma 11])**.** *There is a 1-to-1 correspondence between the fixed points of the tree recursion given in* (6) *and the critical points of $f_{\Delta,\beta,\lambda}(\eta)$. Moreover, the stable fixed points of the tree recursion given in* (6) *are in 1-to-1 correspondence with Hessian local maxima of $f_{\Delta,\beta,\lambda}(\eta)$.*

Recall that a fixed point is *stable* if the absolute value of the derivative at that point is less than 1. A local maximum is a *Hessian local maximum* if the Hessian is negative definite at that point. In particular, as our functions are univariate (after fixing Δ, β, λ), this is simply saying that the second derivative is negative which implies the existence of a local maximum.

For the above theorem to be useful, we need to understand the solutions of (6).

▶ **Proposition 30.** *For $\beta > \beta_u$, the following hold:*

(1) *If $|\log\lambda| > \log\lambda_u$, then* (6) *has a unique fixed point. It is stable and hence corresponds to the global maximizer of $f_{\Delta,\beta,\lambda}$. This maximizer is $\eta^+_{\Delta,\beta,\lambda} = \eta^-_{\Delta,\beta,\lambda}$.*

(2) *If $|\log\lambda| = \log\lambda_u$, then* (6) *has two distinct fixed points, one of which is stable and corresponds to the global maximizer of $f_{\Delta,\beta,\lambda}$. The other corresponds to an inflection point of $f_{\Delta,\beta,\lambda}$.*

(3) *If $|\log\lambda| < \log\lambda_u$, then* (6) *has three distinct fixed points. The largest and the smallest are both stable, corresponding to the only two local maxima of $f_{\Delta,\beta,\lambda}$. When $\lambda > 1$, $\eta^+_{\Delta,\beta,\lambda}$ is the unique global maximizer; when $\lambda < 1$, $\eta^-_{\Delta,\beta,\lambda}$ is the unique global maximizer; when $\lambda = 1$ then $\eta^+_{\Delta,\beta,\lambda}, \eta^-_{\Delta,\beta,\lambda}$ are both global maximizers.*

Portions of this statement have been shown in, for example, [29, 30, 32], and we give a complete proof in [36]. An illustration of $f_{\Delta,\beta,\lambda}(\eta)$ is given in Figure 5; the left plot appears for $\lambda > \lambda_u$ (Case 1 above) and the right plot appears for $1 < \lambda < \lambda_u$ (Case 3 above).

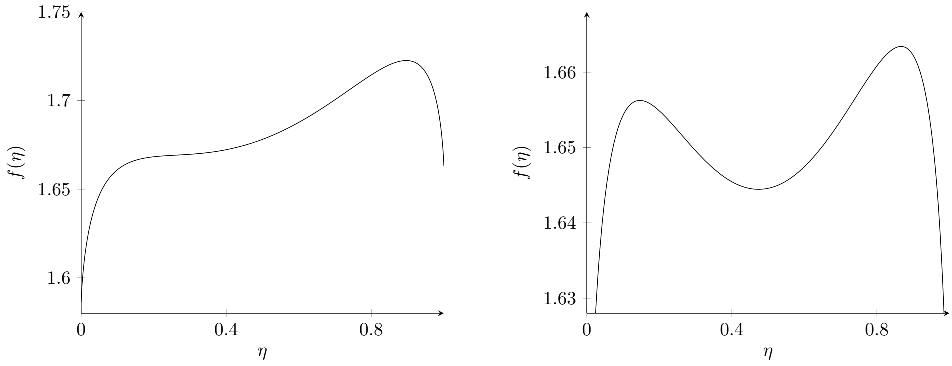

Figure 5 Sketch of the function $f_{\Delta,\beta,\lambda}(\eta)$ for $\Delta = 4$, $\beta = \ln(2) + 0.1$, and (left) $\lambda = 1.08$, (right) $\lambda = 1.01$.

While the behavior described in part (3) suggests metastability, Proposition 30 is only about the expected partition function, and we will need to show that multiple local maxima exist with high probability over the random graph. This will involve showing a lower bound on the partition function at the two local maxima and and upper bound everywhere else.

Via Markov's inequality, the next statement gives a high probability approximate upper bound on $Z_{G,\eta}(\lambda)$.

▶ **Lemma 31.** *Fix $\beta \geq 0, \lambda > 0$. With probability $1 - o(1)$ over the random Δ-regular graph G on n vertices, it holds for every η that*

$$Z_{G,\eta}(\lambda) \leq n^2 \cdot \mathbb{E} Z_{G,\eta}(\lambda)\,.$$

We further prove lower bounds on $Z_{G,\eta}$ for values of η which are local maxima. For a global maximum, this was proved in [29] via the second moment method.

▶ **Theorem 32** ([29, Theorem 8])**.** *Fix $\lambda > 0$ and suppose that η is a global maximizer of $f_{\Delta,\beta,\lambda}$. With probability $1 - o(1)$ over the random Δ-regular graph G on n vertices,*

$$Z_{G,\eta}(\lambda) \geq \frac{1}{n}\mathbb{E} Z_{G,\eta}(\lambda)\,.$$

We prove the following corresponding statement for the local maximizers.

▶ **Proposition 33.** *Fix $\lambda > 0$ and suppose that η is a local maximizer of $f_{\Delta,\beta,\lambda}$. For any $\zeta > 0$, with probability $1 - o(1)$ over the random Δ-regular graph G on n vertices,*

$$Z_{G,\eta}(\lambda) \geq e^{-\zeta n}\mathbb{E}[Z_{G,\eta}(\lambda)]\,.$$

The proof of Proposition 33 follows the template of Coja-Oghlan, Galanis, Goldberg, Ravelomanana, Štefankovič, and Vigoda [19] in proving metastability in the zero-field ferromagnetic Potts model (which in turn used ideas from [3,21]). The argument involves various techniques such as studying the *planted model*, *Nishimori identities* [20], and *non-reconstruction of broadcasting processes* [19, 29, 39], and it is presented in the full paper [36]. We can now sketch the proofs of our slow mixing results.

Slow mixing of Glauber Dynamics

We start with sketching our proof of Theorem 2. Let $\beta > \beta_u(\Delta), \lambda \in [1, \lambda_u)$, and $G \sim \mathbf{G}$. Let $\eta = \eta^+_{\Delta,\beta,\lambda}$, the mean magnetization of the root of $\mathbb{T}_\Delta$ under the $+$ boundary conditions with external field λ, and let $\eta_- = \eta^-_{\Delta,\beta,\lambda}$, the same but under the $-$ boundary conditions.

As η and η_- are global and local maximizers of $f_{\Delta,\beta,\lambda}$, there are $\epsilon > 0$ and $\delta > 0$ so that:
1. $\mathbb{E}[Z_{G,\eta'}(\lambda)] \le e^{-\delta n}\mathbb{E}[Z_{G,\eta}(\lambda)]$ for all η' such that $|\eta' - \eta| > \epsilon$.
2. $\mathbb{E}[Z_{G,\eta'}(\lambda)] \le e^{-\delta n}\mathbb{E}[Z_{G,\eta_-}(\lambda)]$ for all η' such that $|\eta' - \eta_-| \in (\epsilon, 2\epsilon)$.

Next, we sketch how we construct the configuration sets S_1, S_2, S_3 for applying Lemma 19, where we assume here for simplicity that the magnetization η can actually be realized on G. For $\epsilon > 0$ as above, we set:

S_1 : configurations with magnetization η

S_2 : configurations with magnetization in $[\eta_- - \epsilon, \eta_- + \epsilon]$

S_3 : configurations with magnetization in $[\eta_- - 2\epsilon, \eta_- - \epsilon) \cup (\eta_- + \epsilon, \eta_- + 2\epsilon]$.

First, note that the Glauber dynamics starting in S_2 must pass through S_3 to reach S_1. Abbreviating $\mu_{G,\beta,\lambda}$ as μ, we can use Lemma 31,Theorem 32 and Property 1 from above to show that $\mu(S_2) < \mu(S_1)$ a.a.s. over G. Similarly, using Proposition 30, Property 2 and Proposition 33 yields $\mu(S_3) \le e^{-\Omega(n)}\mu(S_2)$ a.a.s. Hence, applying Lemma 19, we conclude that the mixing time of Glauber dynamics on G is $\exp(\Omega(n))$.

Slow Mixing of the Kawasaki Dynamics

We proceed with sketching the proof of part (2) of Theorem 1. Let $\beta > \beta_u(\Delta)$. We consider a graph H consisting of m identical copies $G_1, G_2, \dots G_m$ of a random Δ-regular graph G from $\mathbf{G}$, where is m is determined later based on η. We will separately consider the cases of $|\eta| \in (\eta_c, \eta_u)$ and $|\eta| \le \eta_c$, and assume without loss of generality that $\eta > 0$.

We start with the case $\eta \in (\eta_c, \eta_u)$. By Proposition 5, there exists $\lambda_\eta \in (1, \lambda_u)$ such that $\eta = \eta^+_{\Delta,\beta,\lambda_\eta}$. For $\lambda_+ \in (\lambda_\eta, \lambda_u)$, set $\eta_+ = \eta^+_{\Delta,\beta,\lambda_+}$ and $\eta_- = \eta^-_{\Delta,\beta,\lambda_+}$. In particular, note that we may choose λ_+ such that there are $m, \ell \in \mathbb{N}$ with $\ell < m$ and $m\eta = \ell\eta_+ + (m-\ell)\eta_-$, where m is used for constructing H. Further, observe that η is the global maximizer of $f_{\Delta,\beta,\lambda_\eta}$ and that η_+ and η_- are the global and local maximizers of $f_{\Delta,\beta,\lambda_+}$. Hence, there are $\epsilon > 0$ and $\delta > 0$ so that:
1. $\mathbb{E}[Z_{G,\eta'}(\lambda_\eta)] \le e^{-\delta n}\mathbb{E}[Z_{G,\eta}(\lambda_\eta)]$ for all η' such that $|\eta' - \eta| > \epsilon$.
2. $\mathbb{E}[Z_{G,\eta'}(\lambda_+)] \le e^{-\delta n}\mathbb{E}[Z_{G,\eta_+}(\lambda_+)]$ for all η' such that $|\eta' - \eta_+| \in (\epsilon, 2\epsilon)$.
3. $\mathbb{E}[Z_{G,\eta'}(\lambda_+)] \le e^{-\delta n}\mathbb{E}[Z_{G,\eta_-}(\lambda_+)]$ for all η' such that $|\eta' - \eta_-| \in (\epsilon, 2\epsilon)$.

As for proving slow mixing of Glauber dynamics, we aim for applying Lemma 19. To sketch the construction of S_1, S_2, S_3, we again assume here for simplicity that a magnetization of η can be realized on each subgraph G_i. Given a configuration, we write η_{G_i} for the magnetization on subgraph G_i. We then take the following subsets of configurations on H with overall magnetization η:

S_1 : $\eta_{G_i} = \eta$ for all $1 \le i \le m$,

S_2 : $\eta_{G_i} \in [\eta_+ - \epsilon, \eta_+ + \epsilon]$ for all $i \le \ell$ and $\eta_{G_i} \in [\eta_- - \epsilon, \eta_- + \epsilon]$ for all $i > \ell$,

S_3 : $\eta_{G_i} \in [\eta_+ - 2\epsilon, \eta_+ + \epsilon]$ for all $i \le \ell$ and $\eta_{G_i} \in [\eta_- - \epsilon, \eta_- + 2\epsilon]$ for all $i > \ell$, and there exists $i \le \ell$ with $\eta_{G_i} \in [\eta_+ - 2\epsilon, \eta_+ - \epsilon]$ or $i > \ell$ with $\eta_{G_i} \in [\eta_- + \epsilon, \eta_- + 2\epsilon]$.

Note that the Kawasaki dynamics have to pass through S_3 to get from S_2 to S_1. Moreover, abbreviating $\hat{\mu}_{H,\beta,k}$ as $\hat{\mu}$, we can use Theorem 32, Lemma 31 and Property 1 to show that $\hat{\mu}(S_1) \ge \hat{\mu}(S_2)$, and we can use Lemma 31, Properties 2 and 3, Theorem 32 and Proposition 33 to show that $\hat{\mu}(S_3) \le e^{-\Theta(n)}\hat{\mu}(S_2)$ a.s.s. Hence, applying Lemma 19, we conclude that the mixing time of Kawasaki dynamics on H is $\exp(\Omega(n))$.

In the case that $0 < \eta \leq \eta_c$, we require a slightly different argument since we cannot apply Proposition 5 to η. Instead, we argue that for all $\eta \in (0, \eta_c]$ we can choose $\delta' > 0$ sufficiently small such that for all $\eta_+ \in (\eta_c, \eta_c + \delta')$ and $\eta_- = \eta^-_{\Delta,\beta,\lambda_{\eta_+}}$ it holds that $\eta_- < \eta < \eta_+$. In particular, we may choose η_+ such that $m\eta = \ell\eta_+ + (m - \ell)\eta_-$ for some $m, \ell \in \mathbb{N}, \ell < m$. We then define S_1, S_2, S_3 (again with some slight simplification here) by

$S_1 \colon \eta_{G_i} \in [\eta_- - \epsilon, \eta_- + \epsilon]$ for all $i \leq m - \ell$ and $\eta_{G_i} \in [\eta_+ - \epsilon, \eta_+ + \epsilon]$ else,
$S_2 \colon \eta_{G_i} \in [\eta_+ - \epsilon, \eta_+ + \epsilon]$ for all $i \leq \ell$ and $\eta_{G_i} \in [\eta_- - \epsilon, \eta_- + \epsilon]$ else,
$S_3 \colon \eta_{G_i} \in [\eta_+ - 2\epsilon, \eta_+ + \epsilon]$ for all $i \leq \ell$ and $\eta_{G_i} \in [\eta_- - \epsilon, \eta_- + 2\epsilon]$ else, and there exists $i \leq \ell$ with $\eta_{G_i} \in [\eta_+ - 2\epsilon, \eta_+ - \epsilon]$ or $i > \ell$ with $\eta_{G_i} \in [\eta_- + \epsilon, \eta_- + 2\epsilon]$.

By symmetry, we have $\hat{\mu}(S_1) = \hat{\mu}(S_2)$ and by the same arguments as before it holds that $\hat{\mu}(S_3) \leq e^{-\Theta(n)}\hat{\mu}(S_2)$. Applying Lemma 19 then gives the desired result.

References

1 Vedat Levi Alev and Lap Chi Lau. Improved analysis of higher order random walks and applications. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 1198–1211, 2020. doi:10.1145/3357713.3384317.

2 Nima Anari, Kuikui Liu, and Shayan Oveis Gharan. Spectral independence in high-dimensional expanders and applications to the hardcore model. *SIAM Journal on Computing*, 0(0):FOCS20–1, 2021. doi:10.1137/20M1367696.

3 Victor Bapst and Amin Coja-Oghlan. Harnessing the Bethe free energy. *Random structures & algorithms*, 49(4):694–741, 2016. doi:10.1002/rsa.20692.

4 Roland Bauerschmidt, Thierry Bodineau, and Benoit Dagallier. Kawasaki dynamics beyond the uniqueness threshold. *arXiv preprint arXiv:2310.04609*, 2023. doi:10.48550/arXiv.2310.04609.

5 Roland Bauerschmidt, Thierry Bodineau, and Benoit Dagallier. Stochastic dynamics and the Polchinski equation: an introduction. *arXiv preprint arXiv:2307.07619*, 2023. doi:10.48550/arXiv.2307.07619.

6 Rodney J Baxter. *Exactly solved models in statistical mechanics*. Elsevier, 2016.

7 Russ Bubley and Martin Dyer. Path coupling: A technique for proving rapid mixing in Markov chains. In *Proceedings 38th Annual Symposium on Foundations of Computer Science*, pages 223–231. IEEE, 1997. doi:10.1109/SFCS.1997.646111.

8 Van Hao Can, Remco van der Hofstad, and Takashi Kumagai. Glauber dynamics for Ising models on random regular graphs: cut-off and metastability. *ALEA*, 18(1):1441–1482, 2021. doi:10.30757/ALEA.v18-52.

9 N Cancrini and F Martinelli. On the spectral gap of Kawasaki dynamics under a mixing condition revisited. *Journal of Mathematical Physics*, 41(3):1391–1423, 2000. doi:10.1063/1.533192.

10 N Cancrini, F Martinelli, and C Roberto. The logarithmic Sobolev constant of Kawasaki dynamics under a mixing condition revisited. *Annales de l'Institut Henri Poincare (B) Probability and Statistics*, 38(4):385–436, 2002. doi:10.1016/S0246-0203(01)01096-2.

11 Nicoletta Cancrini, F Cesi, and F Martinelli. The spectral gap for the Kawasaki dynamics at low temperature. *Journal of statistical physics*, 95:215–271, 1999. doi:10.1023/A:1004581512343.

12 Charlie Carlson, Ewan Davies, Alexandra Kolla, and Will Perkins. Computational thresholds for the fixed-magnetization Ising model. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1459–1472, 2022. doi:10.1145/3519935.3520003.

13 Raphaël Cerf and Ágoston Pisztora. On the Wulff crystal in the Ising model. *Annals of probability*, pages 947–1017, 2000. doi:10.1214/aop/1019160324.

14 Yuansi Chen and Ronen Eldan. Localization schemes: A framework for proving mixing bounds for Markov chains. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 110–122. IEEE, 2022. `doi:10.1109/FOCS54457.2022.00018`.

15 Zongchen Chen, Kuikui Liu, and Eric Vigoda. Optimal mixing of Glauber dynamics: Entropy factorization via high-dimensional expansion. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 1537–1550, 2021. `doi:10.1145/3406325.3451035`.

16 Zongchen Chen, Kuikui Liu, and Eric Vigoda. Spectral independence via stability and applications to holant-type problems. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 149–160. IEEE, 2022. `doi:10.1109/FOCS52979.2021.00023`.

17 Zongchen Chen, Kuikui Liu, and Eric Vigoda. Rapid mixing of Glauber dynamics up to uniqueness via contraction. *SIAM Journal on Computing*, 52(1):196–237, 2023. `doi:10.1137/20M136685X`.

18 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. `doi:10.1016/j.jctb.2015.09.006`.

19 Amin Coja-Oghlan, Andreas Galanis, Leslie Ann Goldberg, Jean Bernoulli Ravelomanana, Daniel Štefankovič, and Eric Vigoda. Metastability of the Potts ferromagnet on random regular graphs. *Communications in Mathematical Physics*, pages 1–41, 2023. `doi:10.1007/s00220-023-04644-6`.

20 Amin Coja-Oghlan, Florent Krzakala, Will Perkins, and Lenka Zdeborová. Information-theoretic thresholds from the cavity method. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 146–157, 2017. `doi:10.1145/3055399.3055420`.

21 Amin Coja-Oghlan, Philipp Loick, Balázs F Mezei, and Gregory B Sorkin. The Ising antiferromagnet and max cut on random regular graphs. *SIAM Journal on Discrete Mathematics*, 36(2):1306–1342, 2022. `doi:10.1137/20M137999X`.

22 Ewan Davies and Will Perkins. Approximately counting independent sets of a given size in bounded-degree graphs. *SIAM Journal on Computing*, 52(2):618–640, 2023. `doi:10.1137/21M1466220`.

23 Amir Dembo and Andrea Montanari. Ising models on locally tree-like graphs. *Ann. Appl. Probab.*, 20(1):565–592, 2010. `doi:10.1214/09-AAP627`.

24 Roland Lvovich Dobrushin, Roman Koteckỳ, and Senya Shlosman. *Wulff construction: a global shape from local interaction*, volume 104. American Mathematical Society Providence, 1992.

25 Martin Dyer, Alan Frieze, and Mark Jerrum. On counting independent sets in sparse graphs. *SIAM Journal on Computing*, 31(5):1527–1541, 2002. `doi:10.1137/S0097539701383844`.

26 Martin Dyer, Leslie Ann Goldberg, Mark Jerrum, and Russell Martin. Markov chain comparison. *Probability Surveys*, 3(none):89–111, 2006. `doi:10.1214/154957806000000041`.

27 Andreas Galanis, Qi Ge, Daniel Štefankovič, Eric Vigoda, and Linji Yang. Improved inapproximability results for counting independent sets in the hard-core model. *Random Structures & Algorithms*, 45(1):78–110, 2014. `doi:10.1002/rsa.20479`.

28 Andreas Galanis, Daniel Štefankovič, and Eric Vigoda. Inapproximability of the partition function for the antiferromagnetic Ising and hard-core models. *Combinatorics, Probability and Computing*, 25(4):500–559, 2016. `doi:10.1017/S0963548315000401`.

29 Andreas Galanis, Daniel Štefankovič, Eric Vigoda, and Linji Yang. Ferromagnetic Potts model: Refined #BIS-hardness and related results. *SIAM Journal on Computing*, 45(6):2004–2065, 2016. `doi:10.1137/140997580`.

30 Hans-Otto Georgii. *Gibbs measures and phase transitions*, volume 9. Walter de Gruyter, 2011.

31 Antoine Gerschenfeld and Andrea Montanari. Reconstruction for models on random graphs. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07)*, pages 194–204. IEEE, 2007. `doi:10.1109/FOCS.2007.58`.

32 Heng Guo and Pinyan Lu. Uniqueness, spatial mixing, and approximation for ferromagnetic 2-spin systems. *ACM Transactions on Computation Theory (TOCT)*, 10(4):1–25, 2018. doi:10.1145/3265025.

33 Vishesh Jain, Marcus Michelen, Huy Tuan Pham, and Thuy-Duong Vuong. Optimal mixing of the down-up walk on independent sets of a given size. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1665–1681. IEEE, 2023. doi:10.1109/FOCS57990.2023.00101.

34 Svante Janson, Tomasz Łuczak, and Andrzej Rucinski. *Random graphs*. John Wiley & Sons, 2011.

35 Mark Jerrum and Alistair Sinclair. Polynomial-time approximation algorithms for the Ising model. *SIAM Journal on computing*, 22(5):1087–1116, 1993. doi:10.1137/0222066.

36 Aiya Kuchukova, Marcus Pappik, Will Perkins, and Corrine Yap. Fast and slow mixing of the kawasaki dynamics on bounded-degree graphs. *arXiv preprint arXiv:2405.06209*, 2024. doi:10.48550/arXiv.2405.06209.

37 David A Levin and Yuval Peres. *Markov chains and mixing times*, volume 107. American Mathematical Soc., 2017.

38 Sheng Lin Lu and Horng-Tzer Yau. Spectral gap and logarithmic Sobolev inequality for Kawasaki and Glauber dynamics. *Communications in Mathematical Physics*, 156(2):399–433, 1993. doi:10.1007/BF02098489.

39 Fabio Martinelli, Alistair Sinclair, and Dror Weitz. Fast mixing for independent sets, colorings, and other models on trees. *Random Structures & Algorithms*, 31(2):134–172, 2007. doi:10.1002/rsa.20132.

40 Elchanan Mossel and Allan Sly. Exact thresholds for Ising–Gibbs samplers on general graphs. *The Annals of Probability*, 41(1):294–328, 2013. doi:10.1214/11-AOP737.

41 Elchanan Mossel, Dror Weitz, and Nicholas Wormald. On the hardness of sampling independent sets beyond the tree threshold. *Probability Theory and Related Fields*, 143(3-4):401–439, 2009. doi:10.1007/s00440-007-0131-9.

42 Dana Randall and David Wilson. Sampling spin configurations of an Ising system. In *Symposium on Discrete Algorithms: Proceedings of the tenth annual ACM-SIAM symposium on Discrete algorithms*, pages 959–960, 1999. doi:10.1145/314500.314945.

43 Guus Regts. Absence of zeros implies strong spatial mixing. *Probability Theory and Related Fields*, pages 1–21, 2023. doi:10.1007/s00440-023-01190-z.

44 Shuai Shao and Yuxin Sun. Contraction: A unified perspective of correlation decay and zero-freeness of 2-spin systems. *Journal of Statistical Physics*, 185:1–25, 2021. doi:10.1007/s10955-021-02831-0.

45 Shuai Shao and Xiaowei Ye. From zero-freeness to strong spatial mixing via a christoffel-darboux type identity. *arXiv preprint arXiv:2401.09317*, 2024. doi:10.48550/arXiv.2401.09317.

46 Allan Sly. Computational transition at the uniqueness threshold. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 287–296. IEEE, 2010. doi:10.1109/FOCS.2010.34.

47 Allan Sly and Nike Sun. Counting in two-spin models on d-regular graphs. *Annals of Probability*, 42(6):2383–2416, 2014. doi:10.1214/13-AOP888.

48 Dror Weitz. Counting independent sets up to the tree threshold. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 140–149, 2006. doi:10.1145/1132516.1132538.

49 Lawrence Zalcman. Normal families: new perspectives. *Bulletin of the American Mathematical Society*, 35(3):215–230, 1998. doi:10.1090/S0273-0979-98-00755-1.