Phase Transition for the Mixing Time of the Glauber Dynamics for Coloring Regular Trees

Prasad Tetali [*]   Juan C. Vera [†]   Eric Vigoda [‡]   Linji Yang[‡]

Abstract

We prove that the mixing time of the Glauber dynamics for random k-colorings of the complete tree with branching factor b undergoes a phase transition at $k = b(1+o_b(1))/\ln b$. Our main result shows nearly sharp bounds on the mixing time of the dynamics on the complete tree with n vertices for $k = Cb/\ln b$ colors with constant C. For $C \geq 1$ we prove the mixing time is $O(n^{1+o_b(1)} \ln^2 n)$. On the other side, for $C < 1$ the mixing time experiences a slowing down, in particular, we prove it is $O(n^{1/C+o_b(1)} \ln^2 n)$ and $\Omega(n^{1/C-o_b(1)})$. The critical point $C = 1$ is interesting since it coincides (at least up to first order) to the so-called reconstruction threshold which was recently established by Sly. The reconstruction threshold has been of considerable interest recently since it appears to have close connections to the efficiency of certain local algorithms, and this work was inspired by our attempt to understand these connections in this particular setting.

1 Introduction.

There has been considerable interest in recent years in understanding the mixing time of Markov chains arising from single-site updates (known as Glauber dynamics) for sampling spin systems on finite graphs. The Glauber dynamics is well-studied both for its computational purposes, most immediately its use in Markov chain Monte Carlo (MCMC) algorithms, and for its physical motivation as a model of how physical systems reach equilibrium. Several works in this topic focus on exploring the dynamical and spatial connections between the mixing time and equilibrium properties of the spin system. A notable example of such equilibrium properties is the uniqueness of the infinite volume Gibbs measure, which very roughly speaking corresponds to the influence of

---
[*]School of Mathematics and School of Computer Science, Georgia Institute of Technology, Atlanta GA 30332. Email: tetali@math.gatech.edu. Research supported in part by NSF grant DMS-0701043 and CCF-0910584.

[†]Department of Management Sciences, University of Waterloo, Waterloo ON. Email: jvera@waterloo.edu.

[‡]School of Computer Science, Georgia Institute of Technology, Atlanta GA 30332. Email: {vigoda,ljyang}@gatech.edu. Research supported in part by NSF grant CCF-0830298 and CCF-0910584.

a worst-case boundary condition. Recently a related weaker notion known as the reconstruction threshold has been the focus of considerable study. Reconstruction considers the influence of a "typical" boundary condition (we define it more precisely momentarily).

Much of the recent interest in reconstruction stems from its conjectured connections to the efficiency of local algorithms on trees and tree-like graphs, such as sparse random graphs. The Glauber dynamics is one particular example of such a local algorithm, another important example is belief propagation algorithms. The recent work of Achlioptas and Coja-Oghlan [1] gives strong evidence for the "algorithmic barriers" that arise in the reconstruction phase for several constraint satisfaction problems, including colorings, on sparse random graphs. In this paper we show the mixing time of the Glauber dynamics for random colorings of the complete tree undergoes a phase transition, and the critical point appears to coincide with the reconstruction threshold.

We study the heat-bath version of the Glauber dynamics on the complete tree with branching factor b for the case of (proper vertex) k-colorings. Proper colorings correspond in the physics community to the zero-temperature limit of the anti-ferromagnetic Potts model, and the infinite complete tree is known as the Bethe lattice. Let $\mathcal{C} = \{1, 2, \ldots, k\}$ denote the set of k colors, and $T_\ell = (V, E)$ denote the complete tree with branching factor b, height ℓ and n vertices. We are looking at the set Ω of proper vertex k-colorings which are assignments $\sigma : V \to \mathcal{C}$ such that for all $(v, w) \in E$ we have $\sigma(v) \neq \sigma(w)$. The Glauber dynamics for colorings is a Markov chain (X_t) whose state space is Ω and transitions $X_t \to X_{t+1}$ are defined as follows:

- Choose a vertex v uniformly at random.

- For all $w \neq v$ set $X_{t+1}(w) = X_t(w)$.

- Choose $X_{t+1}(v)$ uniformly at random from its set of available colors $\mathcal{C} \setminus X_t(N(v))$ where $N(v)$ denotes the neighbors of v.

For the complete tree, when $k \geq 3$ the dynamics is ergodic where the unique stationary distribution is the uniform distribution over Ω. The mixing time is

the number of steps, from the worst initial state, to reach within variation distance $\leq 1/2e$ of the stationary distribution. We also consider the relaxation time which is the inverse of the spectral gap of the transition matrix. We formally define these notions in Section 3.

For general graphs of maximum degree b, the Glauber dynamics is ergodic when $k \geq b + 2$ and the best result for arbitrary graphs proves $O(n^2)$ mixing time when $k > 11b/6$ [31]. There are a variety of improvements for classes of graphs with high degree or girth, see [11] for a survey, and recently, Mossel and Sly [27] proved polynomial mixing time for sparse random graphs $G(n, d/n)$ for constant $d > 1$ for some constant number of colors polynomially depending on d.

There are two phase transitions of primary interest in the tree T_ℓ – uniqueness and reconstruction. These phase transitions are realized by analyzing the influence of the boundary condition, which in the case of tree corresponds to fixing the coloring of the leaves. We say uniqueness holds if for all boundary conditions, if we consider the uniform distribution conditional on the boundary condition, the influence at the root decays in the limit $\ell \to \infty$ (i.e., the root is uniformly distributed over the set $\mathcal{C}$ in the limit). Jonasson [17] established that the uniqueness threshold is at $k = b + 2$. When $k \leq b + 1$ it is not hard to see that there are boundary conditions which, in fact, "freeze" the root, moreover, the Glauber dynamics is not ergodic in the case even when $k = b + 2$. Martinelli et al [25] analyzed the Glauber dynamics on the tree T_ℓ with a fixed boundary condition. They proved $O(n \log n)$ mixing time when $k \geq b + 3$ for any boundary condition.

The reconstruction threshold corresponds to the influence of a random boundary condition. In particular, we first choose a random coloring of T_ℓ, the colors of the leaves are fixed, and we rechoose a random coloring for the internal tree from this conditional distribution. Reconstruction is said to hold if the leaves have a non-vanishing (as $\ell \to \infty$) influence on the root in expectation. We refer to the reconstruction threshold as the critical point for the transition between the reconstruction and non-reconstruction phases. It was recently established by Sly that the reconstruction threshold occurs at $k = b(1 + o(1))/\ln b$ [4, 29].

A general connection between reconstruction and the convergence time of the Glauber dynamics was shown by Berger et al [3] who showed, for general spin systems, that $O(n)$ relaxation time on the complete tree (without boundary conditions) implies non-reconstruction. A new work of Ding et al [7] gives very sharp bounds on the mixing time of the Glauber dynamics for the Ising model on the complete tree, and illustrates it undergoes a phase transition at the recon-

struction threshold. For the case of colorings, recently Hayes et al [14] proved polynomial mixing time of the Glauber dynamics for any planar graph with maximum degree b when $k > 100b/\ln b$. Subsequently, improved results were established for the tree. In particular, Goldberg et al [12] proved the mixing time is $n^{\Omega(b/(k \ln b))}$ for the complete tree with branching factor b, and Lucier et al [21] proved the mixing time is $n^{O(1 + b/(k \ln b))}$ for any tree with maximum degree b and the number of colors $k \geq 4$. In a following paper, Lucier et al [22] further prove the same upper bound for the case when $k = 3$.

Our goal is to understand the relationship between the reconstruction threshold and the mixing time. Thus we want to establish a more precise picture than provided by the results of [12] and [21]. Our main result provides (nearly) sharp bounds on the mixing time and relaxation time of the Glauber dynamics for the complete tree, establishing a phase transition at the critical point $k = b(1 + o_b(1))/\ln b$. Our proofs build upon the approaches used by [12] and [21].

THEOREM 1.1. *For all $C > 0$, there exists b_0 such that, for all $b > b_0$, for $k = Cb/\ln b$, the Glauber dynamics on the complete tree T on n vertices with branching factor b and height $H = \lfloor \log_b n \rfloor$ satisfies the following:*

1. **For $C \geq 1$:**

$$\Omega\left(\frac{n \ln n}{b \operatorname{poly}(\log b)}\right) \leq T_{\mathrm{mix}} \leq O(n^{1 + o_b(1)} \ln^2 n)$$
$$\Omega(n) \leq T_{\mathrm{relax}} \leq O(n^{1 + o_b(1)})$$

2. **For $C < 1$:**

$$\Omega(n^{1/C - o_b(1)}) \leq T_{\mathrm{mix}} \leq O(n^{1/C + o_b(1)} \ln^2 n)$$
$$\Omega(n^{1/C - o_b(1)}) \leq T_{\mathrm{relax}} \leq O(n^{1/C + o_b(1)})$$

where the $o_b(1)$ functions are $O(\ln \ln b/\ln b)$ for the upper bounds, $b^{1 - 1/C}/C$ for the lower bounds when $1/2 < C < 1$ and exactly zero for the lower bounds when $0 < C \leq 1/2$. The constants in the $\Omega(\cdot)$ and $O(\cdot)$ are universal constants.

REMARK 1.1. *When $C \geq 1$, the lower bound of the mixing time is proved by Hayes and Sinclair [13] in a more general setting, and for the particular case of the heat-bath version of the Glauber dynamics on the complete tree, we believe it can be improved to $\Omega(n \ln n/\operatorname{poly}(\log b))$ by the same proof. The lower bound of the relaxation time simply follows from the fact that the probability of selecting a specific vertex to recolor in one step of the dynamics is $1/n$. Note, the results of Berger et al [3] imply a lower bound of $T_{\mathrm{relax}} \geq \omega(n)$ for the case $C < 1$ since reconstruction holds in this region.*

Our result extends to more general k and b, thereby refining the general picture provided by [12] and [21].

THEOREM 1.2. *There exists b_0 such that, for all k, b satisfying $b/(k \ln b) > 2$ and $b > b_0$, the Glauber dynamics on the complete tree of n vertices with branching factor b satisfies the following:*

$$\Omega(n^{b/(k \ln b)}) \leq T_{\text{mix}} \leq O(n^{b/(k \ln b)+\gamma} \ln^2 n)$$
$$\Omega(n^{b/(k \ln b)}) \leq T_{\text{relax}} \leq O(n^{b/(k \ln b)+\gamma}),$$

where

$$\gamma = \gamma(b) = 1 - \frac{\ln k}{\ln b} + \frac{\ln \ln b}{\ln b} + \frac{O(1)}{\ln b}$$

is at most a small constant.

REMARK 1.2. *The constants in the $\Omega(\cdot)$ and $O(\cdot)$ of Theorem 1.2 are universal constants. Also, note that when $k = b^\alpha$ for constant $\alpha < 1$, then $\lim_{b \to \infty} \gamma = 1 - \alpha$, and when k is constant, then $\lim_{b \to \infty} \gamma = 1$.*

2 Proof Overview.

We now give an outline of the proofs of Theorem 1.1. Readers can refer to Section 3 for the definitions and background materials.

2.1 Upper bounds. We first sketch the proof approach for upper bounding the mixing time and relaxation time. Let $G^* = (V, E)$ be the star graph on $b + 1$ vertices, i.e., the complete tree T_1 of height 1 with b leaves, and H be the height of the complete tree T_H, i.e., $H = \lfloor \log_b n \rfloor$. Let τ^* be the relaxation time of the Glauber dynamics on the star graph G^* using k colors.

We use the following decomposition result of Lucier and Molloy [21], which is an application of the block dynamics technique (see, Proposition 3.4 in [23]) to the Glauber dynamics on the complete trees combined with Lemma 2 in Mossel and Sly [27].

THEOREM 2.1. *The relaxation time T_{relax} of the Glauber dynamics on the complete tree of height H with branching factor b satisfies*

$$T_{\text{relax}} \leq (\max\{b, \tau^*\})^H.$$

Therefore, proving the upper bounds in Theorem 1.1 reduces to the problem of getting tight upper bounds of the relaxation time τ^* of the Glauber dynamics on G^*. In [21], the authors used a canonical path argument to bound $\tau^* = O(b^{2+1/C} k)$ for any $C > 0$. Instead, here we use two different coupling arguments to show the following two theorems for τ^*.

THEOREM 2.2. *For any $C < 1$, there exists $b_0 > 0$ such that, for any $b > b_0$, the mixing and relaxation times of the Glauber dynamics on G^* using $k = Cb/\ln b$ colors are $O(b^{1/C} \ln^2 b)$. When $C = 1$, the mixing and relaxation times are $O(b \ln^4 b)$.*

THEOREM 2.3. *For any $C > 1$, there exists $b_0 > 0$ such that, for any $b > b_0$, the mixing and relaxation times of the Glauber dynamics on G^* using $k \geq Cb/\ln b$ colors are $O(b \ln b)$.*

REMARK 2.1. *It can be shown that the relaxation time is actually $O(b)$ when $C > 1$, from our analysis. However, unless we can also eliminate the constant factors and thereby show a very sharp bound of at most b, the extra $\ln b$ factor makes little difference to the relaxation time of the dynamics on the whole tree.*

The most difficult (and also interesting) case turns out to be when $C \leq 1$. We will prove Theorem 2.2 in Section 4 and the proof of Theorem 2.3 is in the full version of this paper [30]. We sketch the high-level idea of the proof of Theorem 2.2 in Section 4.1. Having Theorems 2.2 and 2.3 in hand, we can then apply Theorem 2.1 to get the upper bounds on the relaxation time as stated in Theorem 1.1. We get

$$T_{\text{relax}} = \begin{cases} O(b \ln b)^H = O(n^{1 + \frac{\ln \ln b + O(1)}{\ln b}}), & \text{if } C > 1; \\ O(b \ln^4 b)^H = O(n^{1 + \frac{4 \ln \ln b + O(1)}{\ln b}}), & \text{if } C = 1; \\ O(b^{\frac{1}{C}} \ln^2 b)^H = O(n^{\frac{1}{C} + \frac{2 \ln \ln b + O(1)}{\ln b}}), & \text{if } C < 1. \end{cases}$$

To then get the desired upper bounds on the mixing time of the whole tree we need a slightly more advanced tool, the log-Sobolev constant of the Markov chain. By adapting Theorem 5.7 in Martinelli, Sinclair and Weitz [24] to our setting of colorings, we are able to establish (the proof is omitted here) the following relationship between the inverse of the log-Sobolev constant c_{sob}^{-1} and the relaxation time T_{relax} of the Glauber dynamics on trees.

THEOREM 2.4.

$$c_{sob}^{-1} \leq T_{\text{relax}}(2 \log_b(n) \ln(k)) \leq T_{\text{relax}}(2 \ln(n)).$$

Since the inverse of the log-Sobolev constant gives a relatively tight upper bound on the mixing time (see Inequality (3.2) in Section 3), using Theorem 2.4 we are able to complete the proofs of the upper bounds in Theorem 1.1.

2.2 Lower bounds. Our proof of the lower bound in Theorem 1.1 when $C < 1$ builds upon the approach used in [12]. They lower bounded the relaxation time by upper bounding the conductance of the Glauber

dynamics on the subset $S \subseteq \Omega$ where the root is *frozen* (meaning that the configuration at the leaves uniquely determines the color of the root) to some color in $\{1, 2, ..., \lfloor k/2 \rfloor\}$. They showed the conductance of S satisfies $\Phi_S = O(n^{-1/6C})$ when $0 < C < 1/2$, which implies (by (3.1) and (3.3) in Section 3) that $T_{\text{mix}} \geq T_{\text{relax}} - 1 = \Omega(n^{1/6C})$.

We improve their bound on the conductance of S by analyzing the probability that for a given leaf z, in a random coloring σ of the complete tree, the root is frozen and changing the color of z in σ to some other color unfreezes the root. We prove that the number of such leaves in most colorings that freeze the root is $O(n^{-1/C+1+o_b(1)})$. Since the probability of recoloring a specific leaf is $1/n$, then intuitively we have $\Phi_S = O(n^{-1/C+o_b(1)})$, and hence $T_{\text{mix}} \geq T_{\text{relax}} - 1 = \Omega(n^{1/C-o_b(1)})$. A complete analysis of the lower bound is in Section 5, and in the analysis we will see that the $o_b(1)$ error term is $b^{1-1/C}/C$ when $1/2 < C < 1$ and zero when $C \leq 1/2$.

Finally, we will show in Section 6 how all of the proofs generalize for $k = o(b/\ln b)$, and thus prove Theorem 1.2.

3 Technical Preliminaries.

Let $P(\cdot, \cdot)$ denote the transition matrix of the Glauber dynamics, and $P^t(\cdot, \cdot)$ denote the t-step transition probability. The total variation distance at time t from initial state σ is defined as

$$\|P^t(\sigma, \cdot) - \pi\|_{TV} := \frac{1}{2} \sum_\eta |P^t(\sigma, \eta) - \pi(\eta)|.$$

The mixing time T_{mix} for a Markov chain is then defined as

$$T_{\text{mix}} = \min_t \{\max_\sigma \{\|P^t(\sigma, \cdot) - \pi\|_{TV}\} \leq 1/2e\}.$$

Given two copies, (X_t) and (Y_t), of the Markov chain at time $t > 0$, recall that a (one-step) coupling of (X_t) and (Y_t), is a joint distribution whose left and right marginals are identical to the (one-step) evolution of (X_t) and (Y_t), respectively. The Coupling Lemma [2] (c.f., Theorem 5.2 in [20]) guarantees that if, there is a coupling and time $t > 0$, so that for every pair (X_0, Y_0) of initial states, $\Pr[X_t \neq Y_t] \leq 1/2e$ under the coupling, then $T_{\text{mix}} \leq t$.

Let $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_{|\Omega|}$ be the eigenvalues of the transition matrix P. The spectral gap c_{gap} is defined as $1 - \lambda$ where $\lambda = \max\{\lambda_2, |\lambda_{|\Omega|}|\}$ denotes the second largest eigenvalue in absolute value. The relaxation time T_{relax} of the Markov chain is then defined as c_{gap}^{-1}, the inverse of the spectral gap. It is an elementary fact that the mixing time gives a good upper bound on the

relaxation time (see, e.g., [10] for the following bound), which we will use in our analysis:

$$(3.1) \qquad T_{\text{relax}} \leq T_{\text{mix}} + 1.$$

For the upper bounds on the mixing time of the dynamics on the whole tree, we also use the following well-known relationship between the mixing time and the inverse of the log-Sobolev constant (see e.g. [6]):

$$(3.2) \quad T_{\text{mix}} = O\left(c_{sob}^{-1} \ln \ln \frac{1}{\min_{\sigma \in \Omega}\{\pi(\sigma)\}}\right).$$

Readers can refer to [6] for definitions and more details about the log-Sobolev constant.

To lower bound the mixing and relaxation times we analyze the conductance. The conductance of the Markov chain on Ω with transition matrix P is given by $\Phi = \min_{S \subseteq \Omega}\{\Phi_S\}$, where Φ_S is the conductance of a specific set $S \subseteq V$ defined as

$$\Phi_S = \frac{\sum_{\sigma \in S} \sum_{\eta \in \bar{S}} \pi(\sigma) P(\sigma, \eta)}{\pi(S)\pi(\bar{S})}.$$

Thus, a general way to find a good upper bound on the conductance is to find a set S such that the probability of escaping from S is relatively small. The well-known relationship between the relaxation time and the conductance is established in [18] and [28] and we will use the form

$$(3.3) \qquad T_{\text{relax}} = \Omega(1/\Phi),$$

for proving the lower bounds.

Finally, in much of our analysis below, we use the following version of a Chernoff-type bound; see, e.g., Theorem 4.4 and 4.5 in [26].

PROPOSITION 3.1. (CHERNOFF BOUND) *Let random variables $X_1, \ldots, X_n$ correspond to n independent Bernoulli trials with $\Pr[X_i = 1] = p_i$ respectively. Then if $X = \sum X_i$ and $\mu = E[X]$, for any $\delta < 2e - 1$, we have*

$$\Pr[X > (1 + \delta)\mu] \leq \exp(-\delta^2 \mu/4).$$

4 Upper Bound on Mixing Time for $C \leq 1$: Proof of Theorem 2.2.

In this section, we upper bound the mixing time of the Glauber dynamics on the star graph $G^* = (V, E)$ when $k = Cb/\ln b$ for any $C \leq 1$. To be more precise, let $V = \{r, \ell_1, ..., \ell_b\}$, where r refers to the root and $\ell_1, ..., \ell_b$ are the b leaves and $E = \{(r, \ell_1), ..., (r, \ell_b)\}$. For convenience, here we let

$$\epsilon := 1/C - 1,$$

and hence $k = b/((1 + \epsilon) \ln b)$.

We use the maximal one-step coupling, originally studied for colorings by Jerrum [16] to upper bound the mixing time of the Glauber dynamics on general graphs. For a coloring $X \in \Omega$, let $A_X(v)$ denote the set of available colors of v in the coloring X, i.e., $A_\sigma(v) = \{c \in \mathcal{C} : \forall u \in N(v), \sigma(u) \neq c\}$. The coupling (X_t, Y_t) of the two chains is done by choosing the same random vertex v_t for recoloring at step t and maximizing the probability of the two chains choosing the same update for the color of v_t. Thus, for each color $c \in A_{X_t}(v) \cap A_{Y_t}(v)$, with probability $1/\max\{|A_{X_t}(v)|, |A_{Y_t}(v)|\}$ we set $X_{t+1}(v) = Y_{t+1}(v) = c$. With the remaining probability, the color choices for $X_{t+1}(v)$ and $Y_{t+1}(v)$ are coupled arbitrarily.

We prove the theorem by analyzing the coupling in rounds, where each round consists of $T := 20b \ln b$ steps. Our main result is the following lemma which says that in each round we have a good probability of coalescing (i.e., achieving $X_t = Y_t$).

LEMMA 4.1. *For all $\epsilon \geq 0$, there exists b_0 such that for all $b > b_0$ if $k = b/((1 + \epsilon) \ln b)$ and $T = 20b \ln b$ for all $(x_0, y_0) \in \Omega \times \Omega$, the following holds:*

$$\Pr\left[X_T = Y_T \mid X_0 = x_0, Y_0 = y_0\right]$$
$$\geq \begin{cases} (20(1 + \epsilon)b^\epsilon \ln b)^{-1}, & \text{if } \epsilon > 0; \\ (20 \ln^3 b)^{-1}, & \text{if } \epsilon = 0. \end{cases}$$

It is then straightforward to prove Theorem 2.2.

Proof of Theorem 2.2. For $\epsilon > 0$, let $p_T := (20(1 + \epsilon)b^\epsilon \ln b)^{-1}$; and for $\epsilon = 0$ let $p_T := (20 \ln^3 b)^{-1}$. By repeatedly applying Lemma 4.1 we have, for all (x_0, y_0),

$$\Pr\left[X_{2iT} \neq Y_{2iT} \mid X_0 = x_0, Y_0 = y_0\right]$$
$$\leq (1 - p_T)^{2i} \leq 1/2e$$

when $i = 1/p_T$. Therefore, by applying the Coupling Lemma, mentioned in Section 3, the mixing time is $O((1 + \epsilon)b^{1+\epsilon} \ln^2 b)$ for $\epsilon > 0$ and $O(b \ln^4 b)$ for $\epsilon = 0$. □

4.1 Overview of the Coupling Argument. Before formally proving Lemma 4.1 we give a high-level overview of its proof. We will analyze the maximal one-step coupling on the star graph G^*. We say a vertex v "disagrees" at time t if $X_t(v) \neq Y_t(v)$, otherwise we say the vertex v "agrees". We denote the set of disagreeing vertices at time t of our coupled chains by

$$D_t = \{v \in V : X_t(v) \neq Y_t(v)\},$$

and we use $D_t^L = D_t \setminus \{r\}$ to represent the set of disagreeing leaves. When we use the term "with

high probability" in this section, it means that the probability goes to 1 as b goes to infinity.

In the maximal one-step coupling of the Glauber dynamics for the star graph, at each step we select a random vertex to recolor in both chains, and then we use the best way to couple the colors of that vertex. Therefore, if the coupling selects a leaf ℓ to recolor at time t, then the probability that ℓ will be disagree in X_t and Y_t is at most $1/(k-1)$, and with probability at least $(k-2)/(k-1)$ the leaf will use the same color which is chosen uniformly random from $\mathcal{C} \setminus \{X_t(r), Y_t(r)\}$. We also know that if we simply assign a random color from $\mathcal{C}$ to each leaf, with probability at least $\Omega(1/(b^\epsilon \ln b))$ there is a color in $\mathcal{C}$ which is unused in any leaf. This last point hints at the success probability in the statement of Lemma 4.1.

We analyze the T-step epoch in four stages. The warm-up round is of length $T_0 := 4(b + 1) \ln b$ steps. In the warm-up round we just want to make sure we recolor each leaf at least once and we recolor the root at most $20 \ln b$ times. This is straightforward to prove via Chernoff bounds. Proposition 4.2 is the formal statement. We then run for a further $4(b+1) \ln b$ steps, during which we prove (in Lemma 4.2 for $\epsilon > 0$ and Lemma 4.3 for $\epsilon = 0$) that with high probability for $\epsilon > 0$ and with probability at least $1/(2 \ln^2 b)$ for $\epsilon = 0$, the root does not change colors in either chain, and each leaf is recolored at least once. Consequently at the end of these $T_w := 8(b+1) \ln b$ steps, with a good probability, all of the leaf disagreements will be of the same form in the sense that they will have the same pair of colors.

The next stage is of a random length T_1, which is defined as the first time (after T_w) where we are recoloring the root and the root has a common available color in (X_t) and (Y_t). We prove in Lemma 4.4, that with probability $\Omega(1/b^\epsilon \ln b)$, $T_1 < 4(b+1) \ln b$. We then have probability at least $1/2$ of the root agreeing after the update, and then after at most $T_2 := 4(b + 1) \ln b$ further steps we are likely to coalesce since we just need to recolor each leaf at least once before the root changes back to a disagreement.

4.2 Coupling Argument: Proof of Lemma 4.1. We begin with a basic observation about the maximal one-step coupling.

PROPOSITION 4.1. *Let $\mathcal{C}(D_t^L) := \bigcup_{\ell \in D_t^L} \{X_t(\ell), Y_t(\ell)\}$ denote the set of colors that appear in the disagreeing leaves at time t. Then, $A_{X_t}(r) \oplus A_{Y_t}(r) \subseteq \mathcal{C}(D_t^L)$.*

This is simply because those colors that appear on the leaves with agreements are both unavailable in X_t and Y_t for the root.

We now analyze the first stage of the T-step epoch.

PROPOSITION 4.2. *The probability that in $T_0 = 4(b+1)\ln b$ steps, the coupling (X_t, Y_t) or the Glauber dynamics (X_t) will recolor the root at most $20\ln b$ times and recolor every leaf at least once is at least $1 - 2b^{-3}$.*

Proof. It is a simple fact which follows from the Chernoff bound and the coupon collector problem. Using the union bound the probability that there is a leaf which is not recolored in T_0 steps is at most

$$b\left(1 - \frac{1}{b+1}\right)^{4(b+1)\ln b} \le b^{-3}.$$

Now, let N be the number of times the root is recolored in T_0 steps. The expectation $\mathrm{E}[N]$ is simply $4\ln b$. Then, by the Chernoff bound

$$\Pr[N \ge 20\ln b] \le \Pr[N \ge (1+4)\mathrm{E}[N]] \le b^{-3}.$$

Therefore the lemma holds by the union bound. □

Then we will prove that after $T_\mathrm{w} = T_0 + 4(b+1)\ln b$ steps, with high probability all of the leaf disagreements are of the same type when $\epsilon > 0$.

LEMMA 4.2. *When $\epsilon > 0$, starting from any pair of initial states (x_0, y_0), after T_w steps, with high probability, for all $\ell \in D_{T_w}^L$, $X_{T_w}(\ell) = Y_{T_w}(r)$ and $Y_{T_w}(\ell) = X_{T_w}(r)$.*

Proof. The idea is that if we just look at one chain, say (X_t), then after T_0 steps, with high probability the root is frozen. Moreover, the root is likely to continue to be frozen for the remainder of the T_w steps since we recolor the root at most $O(\ln b)$ times. In the worst case the root is frozen to a disagreement, say $X_t(r) = c_1$, $Y_t(r) = c_2$ and $c_1 \ne c_2$. Then after recoloring a leaf ℓ at time t' where $t < t' < T_\mathrm{w}$, the only possible disagreement is $X_{t'}(\ell) = c_2, Y_{t'}(\ell) = c_1$. Hence, it suffices to recolor each leaf at least once.

Let $\mathcal{E}$ be the event that in the first T_0 steps, every leaf is recolored at least once and in another $4(b+1)\ln b$ steps, every leaf is recolored again at least once and the root is recolored at most $20\ln b$ times. We are first going to bound that for $t > T_0$,

$$(4.4) \qquad \Pr[|A_{X_t}(r)| > 1 \mid \mathcal{E}] \le \frac{1}{(1+\epsilon)b^\epsilon \ln b} := p_0,$$

and the same thing happens for Y_t.

Let G_W be the graph with b isolated vertices $\{v_1, ..., v_b\}$, corresponding to the leaves $\{\ell_1, ..., \ell_b\}$. Let (W_t) be a Glauber process on G_W using $k-1$ colors from another color set $\mathcal{C}_W$. We are going to define W_0 and couple (W_t) with (X_t) such that $|A_{X_t}(r)| = |A_{W_t}|+1$ at any time t, where $A_{W_t} := \{c \in \mathcal{C}_W : \forall v_i, W_t(v_i) \ne c\}$.

To do this, for every t we are going to define a bijection $f_t : \mathcal{C} \setminus \{X_t(r)\} \to \mathcal{C}_W$ such that $f_t(X_t(\ell_i)) = W_t(v_i)$ for all i. Notice that if such a bijection exists then $|A_{X_t}(r)| = |A_{W_t}| + 1$.

At time $t = 0$, pick any bijection f_0 from $\mathcal{C}_W$ to $\mathcal{C} \setminus \{X_0(r)\}$. Define W_0 by $W_0(v_i) = f(X_0(\ell_i))$ for all i. We will update f_t only when we choose the root to recolor at time t in the coupling of (W_t) and (X_t). To do the coupling at time $t+1$, we first choose a vertex v in G^* to recolor:

- if $v = \ell_i$, then we choose a random color c different from $X_t(r)$ to recolor v. Correspondingly, we choose the vertex v_i in G_W to recolor using color $f_t(c)$.

- if $v = r$, then we choose a random color c from $A_{X_t}(r)$ to recolor the root in G^*. Correspondingly, we update the mapping f_t in the following natural way: $f_t(X_{t-1}(r)) = f_{t-1}(c)$, (and $f_t(c)$ is undefined).

Since (W_t) itself is a Glauber process that recolors the vertices of G_W uniformly at random from $\mathcal{C}_W$, conditioning on $\mathcal{E}$, simple calculations yield that for any $t > T_0$,

$$\Pr[|A_{W_t}| \ge 1 \mid \mathcal{E}] \le \frac{1}{(1+\epsilon)b^\epsilon \ln b}.$$

Then (4.4) follows by a simple coupling argument.

Since the same thing happens for (Y_t) and the root is recolored at most $20\ln b$ times, then by the union bound, conditioning on $\mathcal{E}$, the probability that at each time we try to recolor the root after T_0 steps, the root is always frozen in both copies is at least $1 - (40\ln b)(p_0) = 1 - 40/((1+\epsilon)b^\epsilon)$. Finally, by Proposition 4.2, $\mathcal{E}$ happens with high probability, and hence the lemma is proven. □

For the threshold case $\epsilon = 0$, we use a slightly weaker lemma for the warm-up stage, in the sense that the successful probability will only be at least $\Omega(1/\ln^2 b)$. The proof is deferred to the full version of the paper [30].

LEMMA 4.3. *When $\epsilon = 0$, starting from any pair of initial states (x_0, y_0), after $T_w' = T_0 + 2b\ln\ln b$ steps, with probability at least $1/(2\ln^2 b)$, for all $\ell \in D_{T_w'}^L$, $X_{T_w'}(\ell) = Y_{T_w'}(r)$ and $Y_{T_w'}(\ell) = X_{T_w'}(r)$.*

After we succeed in the warm-up stage meaning that all of the leaf disagreements are of the same type, we enter the root-coupling stage, where we try to couple the root. Let T_1 be the first time that there is a common

available color in the root and the coupling chain select the root to recolor, that is

$$T_1 := T_1^{XY} = \arg \min_t \{ A_{X_t}(r) \bigcap A_{Y_t}(r) \neq \emptyset$$

$$\text{and the root } r \text{ is selected at step } t \}.$$

Clearly T_1 is a stopping time.

LEMMA 4.4. *For $\epsilon \geq 0$, for any pair of initial states (x_0, y_0) where all of the leaf disagreements are of the same type (i.e., there is a pair of colors c_1, c_2 such that for all $\ell \in D_0^L$, we have $x_0(\ell) = c_1$ and $y_0(\ell) = c_2$), we have,*

$$\Pr\left[T_1^{XY} < 4(b+1)\ln b \mid (X_0, Y_0) = (x_0, y_0) \right]$$
$$> \frac{1}{4(1+\epsilon)b^\epsilon \ln b}.$$

Proof. First of all, by Proposition 4.1, $|A_{X_0}(r) \oplus A_{Y_0}(r)| \leq 2$. We are interested in the time t when there is a common color available for the root in (X_t, Y_t).

Let (Z_t) be a Glauber process on the graph G_Z of $b+1$ isolated vertices $\{v_0, v_1, v_2, ..., v_b\}$ in which v_0 corresponds to the root and v_i corresponds to the leaves ℓ_i for any $i > 0$. The color set used in the process (Z_t) is $\mathcal{C}_Z = [k] \setminus \{c_1, c_2\}$. Each step, (Z_t) chooses a random vertex and recolors it with a random color from the set $\mathcal{C}_Z$. Let T_Z be the stopping time on Z satisfying:

$$T_1^Z = \arg \min_t \{ |A_{Z_t}| \geq 1 \text{ and } v_0 \text{ is selected at step } t \},$$

where t is greater or equal to $2(b+1)\ln b$ and $A_{Z_t} = \{ c \in \mathcal{C}_Z : \forall i \in [1, .., b], Z_t(v_i) \neq c \}$ is the set of unused colors in the vertices $\{v_1, v_2, ... v_b\}$. We want to couple (Z_t) with (X_t, Y_t) in such a way that $T_1^Z \geq T_1^{XY}$ for all the runs, and then if we show that for any initial state z_0, we have
(4.5)
$$\Pr\left[T_1^Z < 4(b+1)\ln b \mid Z_0 = z_0 \right] > \frac{1}{4(1+\epsilon)b^\epsilon \ln b}.$$

Then by the coupling, we know that the lemma is also true.

Now we are going to construct the coupling between (Z_t) and (X_t, Y_t) for $t \leq T_1^{XY}$. Let z_0 be the initial state satisfying that for any $i \in [1, .., b]$, if $x_0(\ell_i) = y_0(\ell_i) \in \mathcal{C}_Z$ then $z_0(v_i) = x_0(\ell_i)$, otherwise we give a arbitrary color to the vertex v_i. On each step t, we first randomly select a vertex in G^* to update in (X_t, Y_t) and accordingly we select the corresponding vertex in G_Z to update in Z_t:

- If the vertex is a leaf ℓ_i:

  (X_t, Y_t) selects a random color c or a disagreement to update. If $c \in \mathcal{C}_Z$ then we give the same color to v_i in Z_t, otherwise we give a random color to v_i.

- If the vertex is the root r:

  Recolor the root on (X_t, Y_t) according to the maximal one-step coupling and pick a random color in $\mathcal{C}_Z$ to recolor v_0 in Z.

Observe that, $A_{Z_t} \subseteq A_{X_t}(r) \bigcap A_{Y_t}(r)$ for any $0 \leq t \leq T_1^{XY}$, which implies that $T_1^Z \geq T_1^{XY}$ holds with probability 1. Now we will show that Inequality (4.5) holds.

Let t_z be the first time when we hit the root after $2(b+1)\ln b$ steps. Since Z is a purely random process, we know that with probability at least $1/(4(1+\epsilon)b^\epsilon \ln b)$, $T_1^Z = t_z$. This is because, similar to Proposition 4.2, by the coupon collector problem, we can easily prove that:

FACT 4.1. *With high probability, every vertex in the graph G_Z will be recolored at least once after $2(b+1)\ln b$ steps and they are recolored to a random color in $\mathcal{C}_Z$.*

Therefore, by the fact that for each color c, the indicator random variable of whether c is used by some leaves or not is negatively associated to each other (c.f., Theorem 14 in [8]), it follows by some elementary calculations that for any fixed t_z and large enough b,

$$\Pr\left[A_{Z_\tau} \neq \emptyset \mid t_z = \tau \right] \geq 0.9\left(1 - (1 - (1 - \frac{1}{|\mathcal{C}_Z|})^b)^{|\mathcal{C}_Z|}\right)$$

$$(4.6) \qquad\qquad \geq \frac{1}{3(1+\epsilon)b^\epsilon \ln b}$$

where the constant .9 in the penultimate inequality follows from Fact 4.1 for b sufficiently large. Moreover, for t_z, by simple calculations, we know that:

FACT 4.2. *The time t_z is with high probability less than $4(b+1)\ln b$.*

Thus, by applying (4.6), we have

$$\Pr\left[T_1^Z < 4(b+1)\ln b \mid Z_0 = z_0 \right]$$

$$\geq \sum_{\tau=2(b+1)\ln b}^{4(b+1)\ln b} \Pr\left[A_{Z_\tau} \neq \emptyset \mid t_z = \tau \right] \cdot \Pr\left[t_z = \tau \right]$$

$$\geq \frac{1}{3(1+\epsilon)b^\epsilon \ln b} \cdot \Pr\left[t_z \in [2(b+1)\ln b, 4(b+1)\ln b) \right]$$

$$\geq \frac{1}{4(1+\epsilon)b^\epsilon \ln b},$$

where the last inequality follows in this case from Fact 4.2 for b sufficiently large.

This completes the proof of Lemma 4.4. $\qquad\square$

We also know that when the root is recolored, if $|A_X(r) \oplus A_Y(r)| \leq 2$ and $|A_X(r) \bigcap A_Y(r)| \geq 1$ holds, then the probability that the root will be recolored

to the same color in both X and Y is at least $1/2$. Hence, at time $T_1 = T_1^{XY}$, with probability at least $1/2$ the root will become an agreement. Combining with Lemma 4.2, we proved that with probability at least $1/O((1+\epsilon)b^\epsilon \ln b)$ when $\epsilon > 0$, starting from arbitrary initial states (x_0, y_0), the root will couple in at most $12(b+1)\ln b$ steps and by that time all the disagreements (if there is any) in the leaves are of the same type. When $\epsilon = 0$, combining with Lemma 4.3, we get that the probability of the same event happening is at least $1/O(\ln^3 b)$.

The last step is to let all of the disagreements in the leaves go away without changing the root to a disagreement, again with constant probability, after $T_2 = 4(b+1)\ln b$ more steps. Here is the precise statement of the lemma, the proof of which is deferred to the full version of this paper [30].

LEMMA 4.5. *For $\epsilon \geq 0$, consider a pair of initial states (x_0, y_0) where the root r agrees (i.e., $x_0(r) = y_0(r)$) and all of the leaf disagreements are of the same type. Then, with probability at least $1/2$ after $T_2 = 4(b+1)\ln b$ steps, we have $X_{T_2} = Y_{T_2}$.*

Finally, by combining Lemmas 4.2, 4.4 and 4.5 together, we can conclude that: when $\epsilon > 0$, with probability at least $1/(20(1+\epsilon)b^\epsilon \ln b)$ after $t = T_w + T_1 + T_2 < T$ steps of the coupling, we have $X_t = Y_t$; when $\epsilon = 0$, from Lemmas 4.3, 4.4 and 4.5, we have that with probability at least $1/(20\ln^3 b)$ after $t = T'_w + T_1 + T_2 < T$ steps of the coupling, we have $X_t = Y_t$, which proves Lemma 4.1.

5 Proof of the Lower Bounds Below the Threshold in Theorem 1.1.

In this section we prove that when $C < 1$:

$$T_{\text{relax}} = \Omega(n^{1/C - o(1)}).$$

In the remainder of this section, $T = T_H$ denotes a complete tree of height $H = \lfloor \log_b n \rfloor$ where the root is denoted by r. Let $L(T)$ or simply L denote the leaves of T. For a vertex v of T, let T_v denote the subtree of T rooted at v and T_v^* denote $T_v \backslash \{v\}$. For the convenience, in this section, let

$$\epsilon := 1/C - 1,$$

and hence $k = b/(1+\epsilon)\ln b$.

In coloring $\sigma \in \Omega(T)$, we say a vertex v is *frozen* in σ if in the subtree T_v the coloring $\sigma(L(T_v))$ of the leaves of T_v forces the color for v. In other words, v is frozen in σ if: for all $\eta \in \Omega$ where $\eta(L(T_v)) = \sigma(L(T_v))$, we have $\eta(v) = \sigma(v)$. Note, by definition, the leaves are always frozen. Observe that for a vertex to be frozen its frozen children must "block" all other color choices. This is formalized in the following observation as in [12].

FACT 5.1. *A vertex v where $h(v) > 0$ is frozen in coloring σ if and only if, for every color $c \neq \sigma(v)$, there is a child w of v where $\sigma(w) = c$ and w is frozen.*

Using this inductional way of defining a vertex being "frozen" in a coloring, we can further show the following lemma, which is a generalization of Lemma 8 in [12] to the case when $0 < \epsilon < 1$, i.e., $1 > C > 1/2$.

LEMMA 5.1. *In a random coloring of tree T, the probability that a vertex of T is not frozen is at most $b^{-\epsilon}$. For the leaves in T, by definition, they are always frozen.*

5.1 Upper Bound on the Conductance.
Let $S_c = S_c(T)$ denote those colorings in $\Omega(T)$ where the root of T is frozen to color c. Let $S = \cup_{1 \leq c \leq k/2} S_c$. We will analyze the conductance of S to lower bound the mixing time.

To upper bound the conductance of S we need to bound the number of colorings $\sigma \in S$ which can leave S with one transition, and in that case how many transitions leave S. To unfreeze the root, we need to recolor a leaf. Thus, we need to bound the number of colorings frozen at the root which can become unfrozen by one recoloring, and in that case, we need to bound the number of leaves which can be recolored to unfreeze the root. For a coloring σ, vertex v and color c, let $\sigma^{v \to c}$ denote the coloring obtained by recoloring v to c.

We capture the colorings on the "frontier" of S as follows. For tree T, coloring $\sigma \in \Omega(T)$, a vertex v and a leaf z of T_v, let $\mathcal{E}_{v,z}^\sigma$ denote the event that the coloring σ is frozen at the vertex v of T and there exists a color c where the coloring $\sigma^{z \to c}$ is not frozen at the vertex v. By definition, this event only depends on the configurations at the leaves of the subtree T_v. In particular, for the root of the tree, let $\mathcal{E}(\sigma, z) := \mathcal{E}_{r,z}^\sigma$ and $\mathbf{1}_{\sigma,z}$ be the indicator of it.

We can convert the above intuition into the following upper bound on conductance of S (similar to Lemma 10 in [12]):

LEMMA 5.2.

$$\Phi_S \;\leq\; \frac{6}{n} \sum_{z \in L(T)} \sum_{\sigma \in \Omega(T)} \frac{\mathbf{1}_{\sigma,z}}{|\Omega(T)|}$$

(5.7)
$$= \frac{6}{n} \sum_{z \in L(T)} \sum_{\sigma \in \Omega(T)} \Pr_{\sigma \in \Omega}\left[\mathcal{E}(\sigma, z)\right]$$

Now if we can prove that

(5.8)
$$\Pr_{\sigma \in \Omega}\left[\mathcal{E}(\sigma, z)\right] \leq b^{-(1+\epsilon - o(1))H},$$

where $o(1)$ is an inverse polynomial of b when $\epsilon < 1$ and equals to zero when $\epsilon \geq 1$. This will be clarified later in Lemma 5.3. Then by plugging this back into the upper bound (5.7) we get

$$\Phi_S \leq \frac{6}{n} \cdot b^H \cdot b^{-(1+\epsilon-o(1))H} \leq 20n^{-1-\epsilon+o(1)}.$$

Therefore, we can conclude that the conductance of this Glauber dynamics is $O(n^{-1-\epsilon+o(1)})$, and hence by (3.1) and (3.3), the mixing time and the relaxation time is $\Omega(n^{1/C-o(1)})$.

In the following section we prove Inequality (5.8) which is the heart of the proof of the lower bound part in Theorem 1.1. The proofs of Lemmas 5.1 and 5.2 are in the full version of this paper [30].

5.2 Proof of Inequality (5.8). Fix the color of the root to be color $c^* \in \mathcal{C}$. Let $\Omega^* = \{\sigma \in \Omega : \sigma(r) = c^*\}$ be the set of colors where the root is colored c^*. The conditional probability $\Pr_{\sigma \in \Omega^*}[\mathcal{E}(\sigma, z)] := \Pr_{\sigma \in \Omega}[\mathcal{E}(\sigma, z) \mid \sigma(w_0) = c^*]$ will be the same for all c^*. Hence,

$$\Pr_{\sigma \in \Omega}[\mathcal{E}(\sigma, z)] = \Pr_{\sigma \in \Omega^*}[\mathcal{E}(\sigma, z)].$$

For the remainder of the proof we condition on the root being colored c^*.

For the event $\mathcal{E}(\sigma, z)$ to occur we need that along the path from the leaf z to the root r, unfreezing each of these vertices will "free" a color for their parent. More precisely, let $w_0, \ldots, w_H$ where $w_0 = r$ and $w_H = z$ denote the path in T from the root r down to the leaf z. For σ to be in $\mathcal{E}(\sigma, z)$, w_1 has to be frozen because the color of z only affects the root through w_1, and if w_1 is not frozen then it can not affect the root becoming unfrozen. Moreover, in order for the root to become unfrozen by changing the color of the leaf z, it must also occur that w_1 becomes unfrozen at the same time, hence $\sigma \in \mathcal{E}_{w_1,z}^\sigma$. Applying this argument in a similar manner down to the leaf z, we can observe that

$$\Pr_{\sigma \in \Omega^*}\left[\mathcal{E}_{r,z}^\sigma\right] \leq \Pr_{\sigma \in \Omega^*}\left[\mathcal{E}_{w_1,z}^\sigma\right]$$
$$\leq \cdots \leq \Pr_{\sigma \in \Omega^*}\left[\mathcal{E}_{w_{H-1},z}^\sigma\right].$$

This suggests an inductive proof to bound $\Pr_{\sigma \in \Omega^*}\left[\mathcal{E}_{r,z}^\sigma\right]$. Actually, we need a much stronger result:

$$\Pr_{\sigma \in \Omega^*}\left[\mathcal{E}_{r,z}^\sigma\right] \leq b^{-(1+\epsilon-o(1))}\Pr_{\sigma \in \Omega^*}\left[\mathcal{E}_{w_1,z}^\sigma\right]$$
$$\leq b^{-2(1+\epsilon-o(1))}\Pr_{\sigma \in \Omega^*}\left[\mathcal{E}_{w_2,z}^\sigma\right]$$
$$\leq \cdots \leq b^{-H(1+\epsilon-o(1))}.$$

Intuitively, the event $\mathcal{E}_{w_i,z}^\sigma$ implies the fact that w_{i+1} is the only child that causes w_i simultaneously being

frozen and being blocked from using color $\sigma(w_{i+1})$, which further implies that both $\mathcal{E}_{w_{i+1},z}^\sigma$ and all the siblings of w_{i+1} should be either using colors other than $\sigma(w_{i+1})$ or not frozen. Here, the sibling's of w_i mean the children of w_{i-1} except w_i. And $b^{-(1+\epsilon-o(1))}$ comes from the probability of the last event above concerning the siblings of w_{i+1}.

Now, we are going to analyze $\Pr_{\sigma \in \Omega^*}\left[\mathcal{E}_{r,z}^\sigma\right]$ more carefully and formally. For each $1 \leq i \leq H$, let $\mathcal{A}_{i,z}^\sigma$ denote the event that no sibling y of w_i satisfies: both $\sigma(y) = \sigma(w_i)$ and σ is frozen at y. Notice that, if $\sigma \notin \mathcal{A}_{1,z}^\sigma$, meaning that there is a sibling of w_1 being frozen to the color of w_1 in σ, then changing the colors in the leaves of T_{w_1} to make w_1 unfrozen will not be sufficient to make all the children of the root colored with $\sigma(w_1)$ become unfrozen. Hence, $\sigma \notin \mathcal{E}_{r,z}^\sigma$, and we have $\mathcal{E}_{r,z}^\sigma$ implies $\mathcal{A}_{1,z}^\sigma$. As noted earlier, we also know that $\mathcal{E}_{r,z}^\sigma$ implies $\mathcal{E}_{w_1,z}^\sigma$. Therefore,

$$(5.9) \quad \Pr_{\sigma \in \Omega^*}\left[\mathcal{E}_{r,z}^\sigma\right] \leq \Pr_{\sigma \in \Omega^*}\left[\mathcal{E}_{w_1,z}^\sigma \bigcap \mathcal{A}_{1,z}^\sigma\right]$$
$$= \sum_{c_1 \in \mathcal{C}^*} \Pr_{\sigma \in \Omega^*}\left[(\sigma(w_1) = c_1)\bigcap \mathcal{E}_{w_1,z}^\sigma \bigcap \mathcal{A}_{1,z}^\sigma\right],$$

where $\mathcal{C}^* = \mathcal{C} - c^*$.

We will bound the terms in the last equation separately for each c_1. We have that for each $c_1 \in \mathcal{C}^*$,

$$(5.10) \quad \Pr_{\sigma \in \Omega^*}\left[(\sigma(w_1) = c_1)\bigcap \mathcal{E}_{w_1,z}^\sigma \bigcap \mathcal{A}_{1,z}^\sigma\right] =$$
$$\Pr_{\sigma \in \Omega^*}\left[\mathcal{E}_{w_1,z}^\sigma \bigcap \mathcal{A}_{1,z}^\sigma \mid \sigma(w_1) = c_1\right]\Pr_{\sigma \in \Omega^*}\left[\sigma(w_1) = c_1\right].$$

Consider the following method to generate a random coloring: First we choose a random color c_0 for the root, then we choose a random color for each child v_i of the root from $\mathcal{C}^*$, and then we do the same things for each subtree T_{v_i} recursively, where $1 \leq i \leq b$. Hence, we can first generate the configurations $\eta \in \Omega(T \setminus T_{w_1}^*)$ for those vertices not inside the subtree rooted at w_1 and then we generate the configurations $\tau \in \Omega(T_{w_1}^*)$ inside the subtree T_{w_1}. From this perspective, it is clear that the events $\mathcal{A}_{1,z}^\sigma$ and $\mathcal{E}_{w_1,z}^\sigma$ are independent, conditioned on the fixed colors of the root and w_1. Therefore, we have

$$(5.11) \quad \Pr_{\sigma \in \Omega^*}\left[\mathcal{E}_{w_1,z}^\sigma \bigcap \mathcal{A}_{1,z}^\sigma \mid \sigma(w_1) = c_1\right] =$$
$$\Pr_{\sigma \in \Omega^*}\left[\mathcal{E}_{w_1,z}^\sigma \mid \sigma(w_1) = c_1\right]\Pr_{\sigma \in \Omega^*}\left[\mathcal{A}_{1,z}^\sigma \mid \sigma(w_1) = c_1\right].$$

Observe that, the first term on the right hand side is actually the same as the probability of the event $\mathcal{E}(\sigma, z)$ for a tree of height $H - 1$ with the color of the root being fixed to c_1, i.e. $\Pr_{\eta \in \Omega(T_{w_1})}\left[\mathcal{E}_{w_1,z}^\eta \mid \eta(r) = c_1\right]$. As we discussed before, this probability is the same for all

$c_1 \in \mathcal{C}^*$ because of symmetry, and hence we denote it as $\Pr_{\eta \in \Omega^*(T_{w_1})}\left[\mathcal{E}^\eta_{w_1,z}\right]$. Putting Equations (5.9), (5.10) and (5.11) together, we have

$$(5.12) \quad \Pr_{\sigma \in \Omega^*}\left[\mathcal{E}^\sigma_{w_1,z} \bigcap \mathcal{A}^\sigma_{1,z}\right] \leq \Pr_{\eta \in \Omega^*(T_{w_1})}\left[\mathcal{E}^\eta_{w_1,z}\right]$$
$$\times \sum_{c_1 \in \mathcal{C}^*}\left(\Pr_{\sigma \in \Omega^*}\left[\mathcal{A}^\sigma_{1,z} \mid \sigma(w_1) = c_1\right]\Pr_{\sigma \in \Omega^*}\left[\sigma(w_1) = c_1\right]\right).$$

Finally, we can calculate a good upper bound on $\Pr_{\sigma \in \Omega^*}\left[\mathcal{A}^\sigma_{1,z} \mid \sigma(w_1) = c_1\right]$ as stated in the following lemma. And by the symmetry, the probabilities are the same for different colors $c_1 \in \mathcal{C}^*$ we are fixing for the vertex w_1.

LEMMA 5.3.

$$\Pr_{\sigma \in \Omega^*}\left[\mathcal{A}^\sigma_{1,z} \mid \sigma(w_1) = c_1)\right] \leq b^{-(1+\epsilon-o(1))},$$

where $o(1)$ is the function $(1+\epsilon)/b^\epsilon$ when $\epsilon < 1$ and equals to zero when $\epsilon \geq 1$.

Plugging Lemma 5.3 into Inequality (5.12), we get:

$$\begin{aligned}
\Pr_{\sigma \in \Omega^*}\left[\mathcal{E}^\sigma_{r,z}\right] &\leq \Pr_{\sigma \in \Omega^*}\left[\mathcal{E}^\sigma_{w_1,z} \bigcap \mathcal{A}^\sigma_{1,z}\right] \\
&\leq \Pr_{\eta \in \Omega^*(T_{w_1})}\left[\mathcal{E}^\eta_{w_1,z}\right] \cdot b^{-(1+\epsilon-o(1))}.
\end{aligned}$$

By induction, applied on $\Pr_{\eta \in \Omega^*(T_{w_1})}\left[\mathcal{E}^\eta_{w_1,z}\right]$, we have that:

$$\Pr_{\sigma \in \Omega}\left[\mathcal{E}^\sigma_{r,z}\right] \leq b^{-(1+\epsilon-o(1))H},$$

which completes the proof of (5.8). The proof of Lemma 5.3 is in the full version of this paper [30].

6 A Simple Generalization to $k = o(b/\ln b)$: Proof of Theorem 1.2.

In all of the previous sections, we assumed $k = Cb/\ln b$ where C is constant. But we are also interested in the case when k is constant, say a hundred colors, and what the mixing time of the Glauber dynamics will be in this case. Let $\alpha = \alpha(k,b) := b/(k \ln b)$. We would like also to see how to generalize the upper bound and lower bound analysis assuming α is any function growing with b, that is when k is $o(b/\ln b)$. Intuitively, it is not hard to see the analysis will go through in the same way since the hardest case is when α is around the non-reconstruction threshold. Actually, all of our proofs will be the same and we just need to modify slightly the statements.

For the upper bound, we change Lemma 4.1 and Lemma 4.4 into the following ones.

LEMMA 6.1. *Let $T = 20b \ln b$. There exists b_0, for all $(x_0, y_0) \in \Omega \times \Omega$, all $\alpha(k,b) \geq 2$, and all $b > b_0$ the following holds:*

$$\Pr\left[X_T = Y_T \mid X_0 = x_0, Y_0 = y_0\right]$$
$$\geq 1/(20\alpha(k,b)b^{\alpha(k,b)} \ln b).$$

LEMMA 6.2. *For any pair of initial states (x_0, y_0) where all of the leaf disagreements are of the same type, then*

$$\Pr\left[T_1^{XY} < 4b \ln b \mid (X_0, Y_0) = (x_0, y_0)\right]$$
$$\geq 1/(4\alpha(k,b)b^{\alpha(k,b)-1} \ln b).$$

Then by the same argument as in Section 4, we are able to show that the relaxation time of the Glauber dynamics on G^* is upper bounded by $O(\alpha b^\alpha \ln b)$. Thus, the mixing time of the Glauber dynamics on the complete tree is bounded by

$$T_{\text{mix}} = O\left(n^{\alpha + \frac{\ln \alpha + 2 \ln \ln b + 20}{\ln b}} \ln^2 n\right),$$

and the relaxation time is bounded by

$$T_{\text{relax}} = O\left(n^{\alpha + \frac{\ln \alpha + 2 \ln \ln b + 20}{\ln b}}\right).$$

For the lower bound, we change Lemma 5.1 and Lemma 5.3 into the following lemmas.

LEMMA 6.3. *In a random coloring of the tree T, the probability that a vertex of T is not frozen is at most b^{-1}.*

LEMMA 6.4.

$$\Pr_{\sigma \in \Omega^*}\left[\mathcal{A}^\sigma_{1,z} \mid \sigma(w_1) = c_1)\right] \leq b^{-\alpha(k,b)}.$$

Then, by exactly the same way as in Section 5, we can show that the mixing time and the relaxation time of the Glauber dynamics on the complete tree T when $\alpha \geq 2$ is lower bounded by $\Omega(n^\alpha) = \Omega(n^{b/(k \ln b)})$.

7 Conclusions.

In the context of spin systems on sparse graphs and more generally for random instances of constraint satisfaction problems (CSPs), an informal conjecture of Andrea Montanari [5] asserts that in the nonreconstruction regime for such models, the Glauber dynamics is always fast (as in $O(n \log n)$) in converging to stationarity on almost all (as in $(1 - o(1))$) of the state space. While the community is far from establishing such a precise connection between the reconstruction threshold and a (dynamical) transition in mixing time, the present contribution provides further evidence towards such a conjecture, by establishing tight estimates on the mixing time of Glauber dynamics on colorings at and near the reconstruction threshold. Results of similar flavor for other instances of CSPs are natural open problems of interest.

References

[1] D. Achlioptas and A. Coja-Oghlan. Algorithmic barriers from phase transitions. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science* (FOCS), 793-802, 2008.

[2] D. Aldous. Random walks on finite groups and rapidly mixing Markov chains. *Séminaire de Probabilités XVII,* Springer Lecture Notes in Mathematics 986, 243-297, 1983.

[3] N. Berger, C. Kenyon, E. Mossel, and Y. Peres. Glauber dynamics on trees and hyperbolic graphs. *Probability Theory and Related Fields*, 131(3):311–340, 2005.

[4] N. Bhatnagar, J. Vera, E. Vigoda, and D. Weitz. Reconstruction for colorings on tree. Preprint, 2008. Available from arXiv at: http://arxiv.org/abs/0711.3664

[5] J. Chayes, F. Martinelli, M. Molloy, and P. Tetali. Workshop report: Phase transitions, hard combinatorial optimization problems and message passing algorithms. Available at: http://www.birs.ca/workshops/2008/08w5109/report08w5109.pdf

[6] P. Diaconis and L. Saloff-Coste. Logarithmic Sobolev inequalities for Markov chains. *Annals of Applied Probability*, 6(3):695–750, 1996.

[7] J. Ding, E. Lubetzky, and Y. Peres. Mixing time of critical Ising model on trees is polynomial in the height. Preprint, 2009. Available from arXiv at: http://arxiv.org/abs/0901.4152

[8] D. Dubhashi and D. Ranjan. Balls and bins: A study in negative dependence. *Random Struct. Algorithms*, 13(2):99–124, 1998.

[9] M. Dyer and A. Frieze. Randomly coloring graphs with lower bounds on girth and maximum degree. *Random Struct. Algorithms*, 23(2):167–179, 2003.

[10] M. Dyer, L. Goldberg, M. Jerrum, and R. Martin. Markov chain comparison. *Probability Surveys*, 3:89–111, 2006.

[11] A. Frieze and E. Vigoda. A survey on the use of Markov chains to randomly sample colourings. *Combinatorics, Complexity and Chance,* Oxford University Press, 2007.

[12] L. Goldberg, M. Jerrum, and M. Karpinski. The mixing time of Glauber dynamics for colouring regular trees. Preprint, 2008. Available from arXiv at: http://arxiv.org/abs/0806.0921

[13] T. Hayes and A. Sinclair. A general lower bound for mixing of single-site dynamics on graphs. *Annals of Applied Probability*, 17(3):931–952, 2007.

[14] T. Hayes, J. Vera, and E. Vigoda. Randomly coloring planar graphs with fewer colors than the maximum degree. In *Proceeding of the 39th ACM Symposium on Theory of Computing* (STOC), 450–458, 2007.

[15] T. Hayes and E. Vigoda. Coupling with the stationary distribution and improved sampling for colorings and independent sets. *Annals of Applied Probability*, 6(4):1297–1318, 2006.

[16] M. Jerrum. A very simple algorithm for estimating the number of k-colorings of a low-degree graph. *Random Struct. Algorithms*, 7(2):157–166, 1995.

[17] J. Jonasson. Uniqueness of uniform random colorings of regular trees. *Statistics and Probability Letters*, 57:243–248, 2002.

[18] G. Lawler and A. Sokal. Bounds on the L^2 spectrum for Markov chains and Markov processes: a generalization of Cheeger's inequality. *Transactions of the American Mathematical Society*, 309:557–580, 1988.

[19] M. Ledoux. *The Concentration of Measure Phenomenon.* Mathematical Surveys and Monographs (Vol. 89), American Mathematical Society, 2001.

[20] D.A. Levin, Y. Peres, and E.L. Wilmer. *Markov chains and mixing times.* American Mathematical Society, 2009.

[21] B. Lucier and M. Molloy. The Glauber dynamics for colourings of bounded degree trees. Preprint, 2008. Available at: http://www.cs.toronto.edu/~molloy/webpapers/wb-jour.pdf

[22] B. Lucier, M. Molloy, and Y. Peres. The Glauber dynamics for colourings of bounded degree trees. To appear in *Proceedings of the 13th Intl. Workshop on Randomization and Computation* (RANDOM), 2009.

[23] F. Martinelli. *Lectures on Glauber dynamics for discrete spin models.* Springer Lecture Notes in Mathematics (Vol. 1717), 2000.

[24] F. Martinelli, A. Sinclair, and D. Weitz. The Ising model on trees: Boundary conditions and mixing time. *Communications in Mathematical Physics*, 250:301–334, 2004.

[25] F. Martinelli, A. Sinclair, and D. Weitz. Fast mixing for independent sets, colorings, and other models on trees. *Random Struct. Algorithms*, 31(2):134–172, 2007.

[26] M. Mitzenmacher and E. Upfal. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis.* Cambridge University Press, 2005.

[27] E. Mossel and A. Sly. Gibbs rapidly samples colorings of $G(n, d/n)$. To appear in *Probability Theory and Related Fields.*

[28] A. Sinclair and M. Jerrum. Approximate counting, uniform generation and rapidly mixing Markov chains. *Information and Computation*, 82:93–133, 1989.

[29] A. Sly. Reconstruction of random colourings. *Communications in Mathematical Physics*, 288(3):943–961, 2009.

[30] P. Tetali, J. C. Vera, E. Vigoda, and L. Yang. Phase transition for the mixing time of the Glauber dynamics for coloring regular trees. Preprint, 2009. Available from arXiv at: http://arxiv.org/abs/0908.2665

[31] E. Vigoda. Improved bounds for sampling colorings. *Journal of Mathematical Physics*, 41:155–1569, 2000.