



HAL
open science

The Glauber dynamics for edges colourings of trees

Michelle Delcourt, Marc Heinrich, Guillem Perarnau

► **To cite this version:**

Michelle Delcourt, Marc Heinrich, Guillem Perarnau. The Glauber dynamics for edges colourings of trees. 2018. hal-02156373

HAL Id: hal-02156373

<https://hal.science/hal-02156373>

Preprint submitted on 14 Jun 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The Glauber dynamics for edges colourings of trees

Michelle Delcourt¹, Marc Heinrich², and Guillem Perarnau³

¹*Department of Combinatorics and Optimization, University of Waterloo, Canada.*

²*LIRIS, Université Claude Bernard, Lyon, France.*

³*School of Mathematics, University of Birmingham, Birmingham, United Kingdom.*

December 14, 2018

Abstract

Let T be a tree on n vertices and with maximum degree Δ . We show that for $k \geq \Delta + 1$ the Glauber dynamics for k -edge-colourings of T mixes in polynomial time in n . The bound on the number of colours is best possible as the chain is not even ergodic for $k \leq \Delta$. Our proof uses a recursive decomposition of the tree into subtrees; we bound the relaxation time of the original tree in terms of the relaxation time of its subtrees using block dynamics and chain comparison techniques. Of independent interest, we also introduce a monotonicity result for Glauber dynamics that simplifies our proof.

The Glauber dynamics is a Markov chain over the set of configuration of spin systems in the vertices of a graph. In this paper we consider the discrete-time Metropolis Glauber dynamics and study the particular case of sampling k -colourings of graphs. The Glauber dynamics has attracted interest from many different areas. In statistical physics for example, this dynamics gives a single-update sampler for the Gibbs distribution, with the particular case of k -colourings corresponding to the antiferromagnetic Potts model at zero temperature. In computer science, rapid convergence of the Glauber dynamics gives a fully polynomial-time randomised approximation scheme for the number of k -colourings of a graph, a counting problem that is $\sharp P$ -complete. Finally, this dynamics is a popular Monte Carlo method for simulating physical systems as it is both relatively simple and easy to implement. The efficiency of these algorithms depends on the time the chain takes to reach equilibrium, known as the mixing time. This is the main object of study in this paper.

A well-known conjecture in this area states that the Glauber dynamics for k -colourings of graphs with n vertices and maximum degree Δ mixes in polynomial time in n for every $k \geq \Delta + 2$. This bound on the number of colours is best possible as the chain might not be ergodic for $k = \Delta + 1$. Jerrum [16], and independently, Salas and Sokal [25], proved that polynomial mixing happens for $k \geq 2\Delta$. Vigoda [30] improved the bound to $k \geq 11\Delta/6$ by analysing the flip chain. Recently, Chen and Moitra [3], and independently, Delcourt, Perarnau and Postle [4] showed that polynomial mixing holds for $k \geq (11/6 - \epsilon)\Delta$, for some small $\epsilon > 0$. This is the best known bound for general graphs. The conjecture has also been studied for several classes of graphs such as graphs with large girth [8, 14], random graphs [10, 23], planar graphs [13] and graphs with bounded tree-width [29].

One of the central topics in the area is the study of Glauber dynamics on trees. Although it is quite simple to sample a uniformly random k -colouring of a tree, Glauber dynamics on trees often shows an extremal behaviour that helps in the understanding of its performance on general graphs.

In this context, two important thresholds appear: the *uniqueness threshold* at $k = \Delta + 2$ [17] that corresponds to the existence of an infinite-volume Gibbs distribution under any boundary condition and the *reconstruction threshold* at $k = (1 + o(1))\Delta / \log \Delta$ [27], which marks the existence of an infinite Gibbs distribution under free boundary conditions and has connections to algorithmic barriers for local algorithms on trees and sparse random graphs [1].

Most of the existing results study the complete b -ary tree, (where $\Delta = b + 1$). Martinelli, Sinclair and Weitz [22] proved that the Glauber dynamics for k -colourings on b -ary trees with n vertices mixes in time $O(n \log n)$ for every $k \geq \Delta + 2$ and for any boundary condition. Hayes, Vera and Vigoda [13] showed that this dynamic mixes in polynomial time (even for planar graphs) provided that $k \geq C\Delta / \log \Delta$ for some constant C and that it takes superpolynomial time to mix for $k = o(\Delta / \log \Delta)$. Tetali et al. [28] refined this result showing that the threshold for polynomial mixing coincides with the reconstruction threshold. The mixing time has also been studied beyond the reconstruction threshold. Berger et al. [2] proved that the mixing time in a complete tree of maximum degree Δ is polynomial for $k \geq 3$ and constant Δ . Lucier et al. [20, 19] obtained the explicit polynomial bound $n^{O(1 + \Delta/k \log \Delta)}$ for every $k \geq 3$. This result extends to any tree with maximum degree Δ and it is best possible up to the constant in the exponent [11, 19].

In this paper we initiate the study of Glauber dynamics for edge colourings of graphs by bounding the mixing time for the dynamics on edge-colourings of trees with maximum degree Δ . To our knowledge, there are no previous results specific to edge-colourings, although the results for vertex-colourings of general graphs can be transferred as the set of edge-colourings of a graph corresponds to the set of vertex-colourings of its line graph. Note that if a graph has maximum degree Δ , then its line graph has maximum degree at most $2\Delta - 2$. Thus, the results in [3, 4] imply that the Glauber dynamics on edge-colourings on any graph with maximum degree Δ mixes in polynomial time for every $k \geq (11/3 - \epsilon)\Delta$.

Edge-colourings behave quite differently to vertex-colourings. Vizing's theorem states that an edge-colouring exists if $k \geq \Delta + 1$, although the line graph can have maximum degree $2\Delta - 2$. Following the corresponding argument for vertex-colourings, we see that the Glauber dynamics for k -edge-colourings is ergodic for every $k \geq 2\Delta$. However, it is not known what is the smallest $k = k(\Delta)$ such that the chain is ergodic on every graph with maximum degree Δ .

The analysis for trees is simpler. The edges of a tree with maximum degree Δ can be coloured with $k = \Delta$ colours, but no k -colouring exists for $k \leq \Delta - 1$. However, at $k = \Delta$, the Glauber dynamics has multiple stationary distributions, as the colours on the edges incident to a vertex of degree Δ are frozen. It is not difficult to see that the Glauber dynamics on edge-colourings of a tree with maximum degree Δ is ergodic if $k \geq \Delta + 1$. Note that this bound on k is roughly half the bound for ergodicity of vertex-colourings in general graphs, as the maximum degree of the line graph is at most $2\Delta - 2$.

This paper shows that this dynamics has polynomial mixing as soon as it becomes ergodic.

Theorem 1. *Let T be a tree on n vertices and maximum degree $\Delta \geq 3$. For $k \geq \Delta + 1$ the Glauber dynamics for k -edge-colourings on T mixes in time $n^{O(1)}$.*

The constant in the exponent obtained in the proof can be bounded by 60; we make no effort to optimise it. Note that for a complete tree of constant degree Δ rooted at v and $k \in [\Delta + 1, 2\Delta - 1]$ we can find a k -edge-colouring such that the ball of radius $R = \lfloor \log_{\Delta} n \rfloor$ around v is frozen. Using Lemma 4.2 in [12] for the discrete-time dynamics, this shows that the mixing time of the Glauber dynamics for k -edge-colourings is $\Omega(n \log n)$.

Our theorem proves polynomial mixing when $\Delta \geq 3$ and $k \geq \Delta + 1$. If $\Delta = 2$, then the graph is a path, so vertex-colourings and edge-colourings are essentially the same. Dyer, Goldberg and Jerrum [9] proved that Glauber dynamics for $k = 3$ mixes in time $\Theta(n^3 \log n)$.

Part of our motivation for analysing edge-colourings of trees comes from the study of Glauber dynamics in sparse random graphs. Hayes and Vigoda [14] proved that if a graph has girth at least 9, then Glauber dynamics on vertex-colourings using at least $(1 + \epsilon)\Delta$ colours is rapidly mixing, for any $\epsilon > 0$. As sparse random graphs are locally tree-like, this result can be transferred to vertex-colourings of random Δ -regular graphs, giving an asymptotic tight bound. It is an interesting open problem to determine the smallest $k = k(\Delta)$ such that Glauber dynamics mixes in polynomial time for k -edge-colourings of a random Δ -regular graph.

The problem of sampling uniformly random edge-colourings of some particular graph is closely related to some fundamental questions in combinatorics. For instance, $2n$ -edge-colourings of the complete graph K_{2n} correspond to 1-factorisations, and n -edge-colourings of the complete bipartite graph $K_{n,n}$ are in bijection with Latin squares. Markov Chain Monte Carlo methods have been introduced to sample such combinatorial objects (see e.g. [6, 15]), but it is not known if they rapidly mix. A promising approach to the problem is to find the smallest $k = k(n)$ for which Glauber dynamics mixes in polynomial time for k -edge-colourings of K_{2n} or $K_{n,n}$.

Outline of the proof

We now give a brief overview of the strategy used to prove Theorem 1. The bound on the mixing time of the Glauber dynamics will follow from bounding its relaxation time. We first show that it suffices to bound the relaxation time of d -regular trees where $d = k - 1$, which notably simplifies the proof. To this end, we prove a monotonicity result on the relaxation time of the Glauber dynamics that can be of independent interest (see Section 2.3).

The main idea of the proof is to recursively decompose the tree into subtrees, and study the process restricted to each one. The decomposition approach has already been used in the literature to bound the mixing time of the Glauber dynamics for vertex-colourings of trees [2, 19, 20]. To analyse the decomposition procedure we need to study the associated block dynamics, which can be informally described as follows: given a partition of the tree into subtrees, at each step we select one subtree and update its colouring by choosing a uniformly random colouring that is compatible with the boundary condition. We then use a result of Martinelli [21] on block dynamics to upper bound the relaxation time of the whole process in terms of the relaxation time of each block, and the relaxation time of the block dynamics. Our decomposition procedure will satisfy two properties,

- i) the Glauber dynamics on each subtree is ergodic for every possible boundary condition;
- ii) the number of recursive decompositions is small (i.e., logarithmic in the size of the tree).

Condition i) is necessary to apply Martinelli's result on block dynamics. Moreover, if ii) holds, the upper bound we obtain on the relaxation time is polynomial in n , with an exponent that is independent from Δ and k .

The splitting strategy is described precisely in Section 4.2. Informally speaking, to ensure that ii) holds, at every step we split into several trees of at most half the size of the original one. To ensure that i) holds, we make sure that every subtree has at most two edges incident to the boundary, and that these two edges are non-adjacent. The latter condition is not necessary for

$k \geq \Delta + 2$ but it is crucial for the case $k = \Delta + 1$. The strategy is implemented by splitting the tree so that all subtrees pend from a vertex, or an edge.

The second ingredient of the proof is a bound on the relaxation time of the block dynamics, where each block corresponds to a subtree hanging from a root vertex. This dynamics is very similar to the Glauber dynamics on the star spanned by the edges incident to the root, constrained to a boundary condition. In Section 4.1 we reduce the block dynamics to an analogous problem: bounding the relaxation time of list-vertex-colourings of a clique (the line graph of a star), where the lists are given by the boundary constraints and each vertex is updated at a different rate. In Sections 3.1 and 3.2, we use the weighted canonical paths method of Lucier and Molloy [19] and a multi-commodity flow argument to bound the relaxation time of the list-colourings of the clique.

To conclude the proof in the case $k = \Delta + 1$, we need to obtain a bound on the relaxation time of the block dynamics where the blocks correspond to trees hanging from a root edge instead of a vertex. In this case, it suffices to study the list-vertex-colouring dynamics on the graph formed by two cliques intersecting at a vertex, which we do in Section 3.3.

1 Preliminaries

Let $G = (V, E)$ be a finite graph. The *line graph* G^ℓ of G is the graph with vertex set E where two edges are adjacent in G^ℓ if and only if they are incident in G . We define some notions for vertex-colourings of graphs, but these can be naturally transferred to the setting of edge-colourings by considering line graphs.

For $v \in V$, the *neighbourhood of v* , denoted by $N(v)$, is the set of vertices in G adjacent to v . For $k \in \mathbb{N}$, denote $[k] = \{1, \dots, k\}$. A (*proper*) k -*vertex-colouring* is a function $\mu : V \rightarrow [k]$ such that $\mu(v) \neq \mu(u)$ for every $u \in N(v)$. All colourings in this paper are proper. Let Ω_V denote the set of k -vertex-colourings of G . The set of k -edge-colourings (i.e., k -vertex-colourings of G^ℓ) will be denoted by Ω_E . Throughout the paper, we use similar notation to distinguish the vertex and the edge-version of each set or parameter.

For $\mu \in \Omega_V$ and $U \subseteq V$, we denote the restriction of μ to U by $\mu|_U$. We denote by $\mu(U)$ the set of colours in U . We write Ω_U^μ to denote the set of $\sigma \in \Omega_V$ which agree with μ on $V \setminus U$, i.e., with $\sigma|_{V \setminus U} = \mu|_{V \setminus U}$. Informally, we think of Ω_U^μ as colourings of U which are compatible with μ in the boundary of U .

A list assignment of V is a function $L : V \rightarrow 2^{[k]}$. An L -*colouring of G* is a k -colouring $\mu \in \Omega_V$ satisfying $\mu(v) \in L(v)$ for all $v \in V$. We denote by Ω_V^L the set of all L -colourings of G . Note that if $U \subseteq V$ and H is the subgraph of G induced by U , then any $\mu \in \Omega_V$ yields a list assignment for U where $L^\mu(u) = [k] \setminus \mu(N(u) \setminus U)$. This gives a natural bijection between Ω_U^μ and $\Omega_U^{L^\mu}$.

In this paper we will focus on k -edge-colourings of a tree G on n vertices with maximum degree at most Δ . Note that G^ℓ is a union of cliques of size at most Δ such that two cliques intersect in at most one vertex and every cycle is contained in some clique.

The vertices of G with degree 1 are called *leaves* and the vertices with degree at least 2, *internal vertices*. If T is a subtree of G , we define the *exterior and interior (edge) boundary of T* respectively as

$$\begin{aligned} \partial_G T &= \{e \in E \setminus E(T) : N(e) \cap E(T) \neq \emptyset\}, \\ \bar{\partial}_G T &= \{e \in E(T) : N(e) \cap \partial_G T \neq \emptyset\}, \end{aligned}$$

where $N(e)$ is the set of edges in G incident to e . If G is clear from the context, we will denote them by ∂T and $\bar{\partial} T$. We said that T has a *fringe boundary* if all edges in $\bar{\partial} T$ have an endpoint that was a leaf of T . We will use t to denote the size of $\bar{\partial} T$.

In this paper, we will always consider t to be a constant with respect to n, Δ, k . In this sense, for functions f and g we use $f = O_t(g)$ if there exists $c = c(t)$ such that $\limsup f/g \leq c(t)$. We also use $f = \Theta_t(g)$ if $f = O_t(g)$ and $g = O_t(f)$.

1.1 Markov chains and Glauber dynamics

Let X_t be a discrete-time Markov chain with state space Ω . It is often convenient to consider a continuous-time analogue of (X_t) . For that, consider the continuous-time Markov chain with state space Ω and transitions given by non-diagonal entries of the generator matrix \mathcal{L} (the diagonal entries are such that each row of \mathcal{L} adds to 0). In a slight abuse of notation, we denote the chain by \mathcal{L} . Here $\mathcal{L}[x \rightarrow y]$ represents the rate at which x jumps to $y \neq x$. For a dynamics \mathcal{L} we write $(x, y) \in \mathcal{L}$ if and only if $\mathcal{L}[x \rightarrow y] > 0$.

Given probability distributions ν, π on Ω , we define their *total variation distance* as

$$d_{TV}(\nu, \pi) = \frac{1}{2} \sum_{x \in \Omega} |\nu(x) - \pi(x)|.$$

Let \mathcal{L} be an ergodic Markov chain with stationary distribution π . The *mixing time* of \mathcal{L} is defined as

$$t_{\text{mix}}(\mathcal{L}) = \inf \left\{ t : \max_{x \in \Omega} d_{TV}(\nu_x^t, \pi) < 1/4 \right\},$$

where ν_x^t is the probability distribution on Ω obtained by running \mathcal{L} for time t starting at x . The *spectral gap* of \mathcal{L} , denoted by $\text{Gap}(\mathcal{L})$, is the second smallest eigenvalue of $-\mathcal{L}$. The *relaxation time* of \mathcal{L} is defined as

$$\tau(\mathcal{L}) = \frac{1}{\text{Gap}(\mathcal{L})},$$

and it satisfies (see, e.g. Theorem 12.3 in [18]),

$$t_{\text{mix}}(\mathcal{L}) \leq (n \log k) \tau(\mathcal{L}). \tag{1}$$

For a graph $G = (V, E)$ with $V = \{v_1, \dots, v_n\}$ and a positive integer k , the *Glauber dynamics for k -vertex-colourings of G* is a discrete-time Markov chain X_t on Ω_V , where X_{t+1} is obtained from X_t by choosing a $v \in V$ and $c \in [k]$ uniformly at random, and updating v with c if this colour does not appear in $N(v)$.

The *continuous-time Glauber dynamics for k -vertex-colourings of G with parameters (p_1, \dots, p_n)* is a continuous-time Markov chain on Ω_V with generator matrix given by

$$\mathcal{L}_V[\sigma \rightarrow \eta] = \begin{cases} p_i & \text{if } \sigma, \eta \text{ differ at } v_i, \\ 0 & \text{otherwise.} \end{cases}$$

One can imagine this stochastic process as every vertex v_i having an independent rate kp_i Poisson clock. When the clock at v rings, one chooses a colour $c \in [k]$ uniformly at random and recolours v with c if possible. We call \mathcal{L}_V *uniform* if $p_i = 1/k$ for every $i \in [n]$. As \mathcal{L}_V is symmetric for any

set of parameters (i.e. $\mathcal{L}_V[\sigma \rightarrow \eta] = \mathcal{L}_V[\eta \rightarrow \sigma]$), if \mathcal{L}_V is ergodic then its stationary distribution π is uniform on Ω_V .

Note that the continuous version updates vertices faster than the discrete one. In particular, we have $\mathcal{L}_V = n(P - I)$, where P is the transition matrix of X_t and I is the identity matrix. It follows from standard Markov chain comparison results (see e.g. [24]) that

$$t_{\text{mix}}(X_t) = O(n t_{\text{mix}}(\mathcal{L}_V)). \quad (2)$$

If $G = (V, E)$ is a tree with maximum degree Δ its line graph G^ℓ is $(\Delta - 1)$ -degenerate (i.e. every subgraph has a vertex of degree at most $\Delta - 1$). It is known [7] that in this case the Glauber dynamics \mathcal{L}_E for k -edge-colourings of G is ergodic provided $k \geq \Delta + 1$.

For $U \subseteq V$ and $\mu \in \Omega_V$, denote by \mathcal{L}_U^μ the dynamics defined by \mathcal{L}_V restricted to the set Ω_U^μ . Similarly, for a list assignment L of U , let \mathcal{L}_U^L be the dynamics defined on the state space Ω_U^L . All these chains are symmetric but their ergodicity will depend on the size of the boundary of μ in U , and the size of the lists $L(u)$ for $u \in U$. If ergodic, then their stationary distribution will be uniform in the corresponding state space.

2 Comparison of Markov chains

This section introduces the methods in the analysis of the relaxation time of reversible Markov chains that we will use later in the proofs. The first result is a unified framework of two well-known chain comparison techniques. We then define the (reduced) block dynamics and present known results on how this can be used to bound the original chain. Finally, we provide a result on the monotonicity of the Glauber dynamics that will allow us to notably simplify the proof.

2.1 The weighted multi-commodity flow method

In this subsection we present a unified framework for two similar techniques for comparing the relaxation time of two different Markov chains \mathcal{L} and \mathcal{L}' on the same state space. These two techniques are (1) the *fractional paths* (or *multi-commodity flows*) method [26] and (2) the *weighted canonical paths* method [19]. Both methods are generalisations of the *canonical paths method* (see e.g. [5]). The main idea of this method is to simulate the transitions of \mathcal{L}' using transitions of \mathcal{L} in such a way that no transition of \mathcal{L} is used too often, to obtain an upper bound on the ratio between $\tau(\mathcal{L})$ and $\tau(\mathcal{L}')$.

Consider two continuous-time ergodic reversible Markov chains \mathcal{L} and \mathcal{L}' on Ω with stationary distributions π and π' , respectively. In the following, we denote by $\omega: \Omega \times \Omega \rightarrow \mathbb{R}$ a weight function on the transitions of \mathcal{L} . To each $\alpha, \beta \in \Omega$ with $(\alpha, \beta) \in \mathcal{L}'$, we associate a set of paths $\Gamma_{\alpha, \beta}$ where every $\gamma \in \Gamma_{\alpha, \beta}$ is a sequence $\alpha = \xi^0, \dots, \xi^m = \beta$, for some $m \geq 1$ with $(\xi^{i-1}, \xi^i) \in \mathcal{L}$ for every $i \in [m]$. We define a *flow* to be a function $g: \Gamma_{\alpha, \beta} \rightarrow [0, 1]$ such that $\sum_{\gamma \in \Gamma_{\alpha, \beta}} g(\gamma) = 1$. The *weight* of γ with respect to ω is defined as

$$|\gamma|_\omega := \sum_{i=1}^m \omega(\xi^{i-1}, \xi^i).$$

If the weight function ω is a constant equal to 1, then the weight of γ is simply denoted by $|\gamma|$, and corresponds to the length of the transformation sequence.

Let $b := \max_{\alpha \in \Omega} \frac{\pi(\alpha)}{\pi'(\alpha)}$. For every $(\sigma, \eta) \in \mathcal{L}$ we define its *congestion* as

$$\rho_{\sigma, \eta} := \frac{1}{\pi(\sigma) \mathcal{L}[\sigma \rightarrow \eta] \omega(\sigma, \eta)} \sum_{(\alpha, \beta) \in \mathcal{L}'} \sum_{\substack{\gamma \in \Gamma_{\alpha, \beta} \\ \gamma \ni (\sigma, \eta)}} g(\gamma) \pi'(\alpha) \mathcal{L}'[\alpha \rightarrow \beta] |\gamma| \omega.$$

Let $\rho_{\max} = \max\{\rho_{\sigma, \eta} : (\sigma, \eta) \in \mathcal{L}\}$ be the maximum congestion over all transitions of \mathcal{L} . We are now ready to state the main lemma.

Proposition 2 (Weighted multi-commodity flows method). *We have $\tau(\mathcal{L}) \leq b^2 \rho_{\max} \tau(\mathcal{L}')$.*

The proof of this result is based on the same ideas as the proofs of both the weighted canonical paths method and the fractional paths method, providing a common framework for them. It is included in the Appendix for sake of completeness, and follows from standard computations using the variational characterisation of $\text{Gap}(\mathcal{L})$ that involves the variance and the Dirichlet form. We will not need this result in full generality in our proofs, but we believe it is interesting in its own right and will use it in two particular cases.

The first case is when the stationary distribution of the two chains is uniform over Ω , and for every pair $(\alpha, \beta) \in \mathcal{L}'$ the set of paths $\Gamma_{\alpha, \beta}$ consists of a single path $\gamma_{\alpha, \beta}$, with $g(\gamma_{\alpha, \beta}) = 1$. In this case, Proposition 2 implies the result in [19].

Proposition 3 (Weighted canonical paths method). *If π and π' are uniform, then*

$$\tau(\mathcal{L}) \leq \tau(\mathcal{L}') \cdot \max_{(\sigma, \eta) \in \mathcal{L}} \left(\frac{1}{\mathcal{L}[\sigma \rightarrow \eta] \omega(\sigma, \eta)} \sum_{\substack{(\alpha, \beta) \in \mathcal{L}' \\ \gamma_{\alpha, \beta} \ni (\sigma, \eta)}} \mathcal{L}'[\alpha \rightarrow \beta] |\gamma_{\alpha, \beta}| \omega \right).$$

In the second case, all transitions have weight $\omega(\sigma, \eta) = 1$, and again both chains have uniform stationary distributions. In this case, we obtain the following.

Proposition 4 (Fractional paths method). *If π and π' are uniform, then*

$$\tau(\mathcal{L}) \leq \tau(\mathcal{L}') \cdot \max_{(\sigma, \eta) \in \mathcal{L}} \left(\frac{1}{\mathcal{L}[\sigma \rightarrow \eta]} \sum_{(\alpha, \beta) \in \mathcal{L}'} \sum_{\substack{\gamma \in \Gamma_{\alpha, \beta} \\ \gamma \ni (\sigma, \eta)}} g(\gamma) \mathcal{L}'[\alpha \rightarrow \beta] |\gamma| \right).$$

Finally, we state a very simple corollary of Proposition 2.

Corollary 5. *Suppose there exists a constant $c > 1$ such that for every $\alpha, \beta \in \Omega$ it holds that*

- $\frac{1}{c} \pi'(\alpha) \leq \pi(\alpha) \leq c \pi'(\alpha)$;
- $\frac{1}{c} \mathcal{L}'[\alpha \rightarrow \beta] \leq \mathcal{L}[\alpha \rightarrow \beta] \leq c \mathcal{L}'[\alpha \rightarrow \beta]$.

Then, there is a constant $K > 0$ such that

$$\frac{1}{K} \tau(\mathcal{L}') \leq \tau(\mathcal{L}) \leq K \tau(\mathcal{L}').$$

This result simply states that if two Markov chains have the same transitions with similar transition rates, and similar stationary distributions, then their mixing time is also the same up to a constant factor.

2.2 Weighted and Reduced Block Dynamics

This subsection describes block dynamics and its reduced version. Informally speaking, block dynamics is a generalisation of Glauber dynamics where one splits the set of vertices into “blocks” (usually with few intersections/interactions between them), and updates each block at a time, according to some probabilities (see [21]). In this paper we will only consider disjoint blocks that partition the vertex set. Note that if all blocks are singletons, we recover Glauber dynamics. The method is presented for vertex-colourings, but it works the same for edge-colourings by taking an edge partition instead. We will present some known results for bounding the relaxation time of Glauber dynamics in terms of the relaxation time of its (reduced) block dynamics.

Let $G = (V_G, E_G)$ be a graph, $T = (V, E)$ be an induced subgraph and let $\mathcal{V} = \{V_1, \dots, V_r\}$ be a *partition* of V . Let $\mu \in \Omega_{V_G}$. Suppose \mathcal{L}_s^σ is ergodic for every $i \in [r]$ and every $\sigma \in \Omega_{V_i}^\mu$ and let $\pi_{V_i}^\sigma$ denote its stationary distribution, which is also the uniform distribution on $\Omega_{V_i}^\sigma$. The *weighted block dynamics on \mathcal{V} with boundary condition μ* is a continuous-time Markov chain with state space $\Omega_{\mathcal{V}}^\mu$ and generator matrix $\mathcal{B}_{\mathcal{V}}^\mu$ given for any $\sigma \neq \eta$ by

$$\mathcal{B}_{\mathcal{V}}^\mu[\sigma \rightarrow \eta] = \begin{cases} g_i \pi_{V_i}^\sigma(\eta) & \text{if there exists } i \in [r] \text{ such that } \eta \in \Omega_{V_i}^\sigma, \\ 0 & \text{otherwise,} \end{cases}$$

where $g_i = \min_{\sigma \in \Omega_{V_i}^\mu} \text{Gap}(\mathcal{L}_{V_i}^\sigma)$ is the minimum gap for the Glauber dynamics on the block V_i , where the minimum is taken over all possible boundary conditions. Note that $\mathcal{B}_{\mathcal{V}}^\mu$ can be understood as the dynamics where each block V_i updates its entire colouring at times given by an independent Poisson clock of rate g_i . The new colouring of V_i is chosen uniformly among all the possible colourings which are compatible with the current boundary conditions. It is clear that the block dynamics is ergodic if the Glauber dynamics is, since each transition of the Glauber dynamics is a valid transition of the block dynamics. Moreover, both dynamics have the same stationary distribution.

An unweighted version of the block dynamics, corresponding to $g_i = 1$, was used by Martinelli in [21]. Lucier and Molloy generalised this result for weighted block dynamics:

Proposition 6 (Proposition 3.2 in [19]). *For every $\mu \in \Omega_{V_G}$ and partition \mathcal{V} of V , we have*

$$\tau(\mathcal{L}_{\mathcal{V}}^\mu) \leq \tau(\mathcal{B}_{\mathcal{V}}^\mu).$$

Given a block partition \mathcal{V} , let $H = (U, F)$ be the subgraph of T composed of the vertices adjacent to vertices in other blocks. Let $\Omega_{\mathcal{R}}$ be the set of colourings of U induced by the colourings in $\Omega_{\mathcal{V}}^\mu$; we will use $\hat{\sigma}, \hat{\eta}, \dots$ to denote the elements of $\Omega_{\mathcal{R}}$. It is convenient to see $\Omega_{\mathcal{R}}$ as a set of colouring classes of $\Omega_{\mathcal{V}}^\mu$ where two colourings are equivalent if and only if they coincide on U . More precisely, for $\hat{\sigma} \in \Omega_{\mathcal{R}}$, let $\Omega_{\mathcal{V}}^{\hat{\sigma}}$ be the set of colourings $\sigma \in \Omega_{\mathcal{V}}^\mu$ with $\sigma|_U = \hat{\sigma}$. In a slight abuse of notation, we write $\Omega_{V_i}^{\hat{\sigma}}$ to denote the set $\{\eta \in \Omega_{V_i}^\sigma : \sigma \in \Omega_{\mathcal{V}}^{\hat{\sigma}}\}$, and we write $\pi_{V_i}^{\hat{\sigma}}$ for $\pi_{V_i}^\sigma$ where σ is an arbitrary colouring in $\Omega_{\mathcal{V}}^{\hat{\sigma}}$. Note that the projection of $\pi_{V_i}^{\hat{\sigma}}$ onto U is well-defined and independent of the choice of $\sigma \in \Omega_{\mathcal{V}}^{\hat{\sigma}}$.

The *reduced* version of $\mathcal{B}_{\mathcal{V}}^\mu$, is a dynamics with state space $\Omega_{\mathcal{R}}$ and generator matrix $\mathcal{R}_{\mathcal{V}}^\mu$ given for any $\hat{\sigma} \neq \hat{\eta}$ by

$$\mathcal{R}_{\mathcal{V}}^\mu[\hat{\sigma} \rightarrow \hat{\eta}] = \begin{cases} g_i \pi_{\text{proj}}^{i, \hat{\sigma}}(\eta) & \text{if there exists } i \in [r] \text{ such that } \hat{\eta} = \eta|_U \text{ for some } \eta \in \Omega_{V_i}^{\hat{\sigma}}, \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

where $\pi_{\text{proj}}^{i, \hat{\sigma}}$ is the projection of $\pi_{V_i}^{\hat{\sigma}}$ onto U ; that is, for an arbitrary $\sigma \in \Omega_{\mathcal{R}}$, the probability distribution on the colourings of U defined for any $\hat{\eta} \in \Omega_{\mathcal{R}}$ by

$$\pi_{\text{proj}}^{i, \hat{\sigma}}(\hat{\eta}) = \pi_{V_i}^{\hat{\sigma}}(\{\eta \in \Omega_{V_i}^{\hat{\sigma}}, \eta|_U = \hat{\eta}\}). \quad (4)$$

In the particular case where each block contains only one vertex in H , only one vertex of H changes colour during a transition of $\mathcal{R}_{\mathcal{V}}^{\mu}$. In this case, the reduced block dynamic is very similar to the Glauber dynamics on H with some parameters p_i determined by the Glauber dynamics on V_i , only with slightly different transition rates.

We will use another result of Lucier and Molloy that shows that the weighted block dynamics and the reduced block dynamics have the same relaxation time.

Proposition 7 (Proposition 3.3 in [19]). *For every $\mu \in \Omega_{V_G}$ and partition \mathcal{V} of V , we have*

$$\tau(\mathcal{B}_{\mathcal{V}}^{\mu}) = \tau(\mathcal{R}_{\mathcal{V}}^{\mu}).$$

Finally, we will use the following property. If the reduced block dynamics is ergodic, then the projection of $\pi_{\mathcal{V}}^{\mu}$ onto U is its stationary distribution.

Lemma 8. *The reduced block dynamics $\mathcal{R}_{\mathcal{V}}^{\mu}$ is reversible for the projection of $\pi_{\mathcal{V}}^{\mu}$ onto U .*

Proof. Recall that $H = (U, F)$ is the subgraph of T induced by the vertices adjacent to vertices in other blocks. Let π_{proj} be projection of $\pi_{\mathcal{V}}^{\mu}$ onto U ; that is, the probability distribution on the colourings of U defined for any $\hat{\sigma} \in \Omega_{\mathcal{R}}$ by

$$\pi_{\text{proj}}(\hat{\sigma}) = \pi_{\mathcal{V}}^{\mu}(\{\sigma \in \Omega_{\mathcal{V}}^{\mu}, \sigma|_U = \hat{\sigma}\}). \quad (5)$$

First observe that, since H consists of all the vertices in each block which are adjacent to vertices of other blocks, for any $\hat{\sigma} \in \Omega_{\mathcal{R}}$, any extension of $\hat{\sigma}$ to V is obtained by computing an extension on each of the blocks separately. As $\pi_{\mathcal{V}}^{\mu}$ is uniform, we have

$$\pi_{\text{proj}}(\hat{\sigma}) = \frac{|\{\sigma \in \Omega_{\mathcal{V}}^{\mu}, \sigma|_U = \hat{\sigma}\}|}{|\Omega_{\mathcal{V}}^{\mu}|} = \frac{\prod_{j=1}^r |\{\sigma \in \Omega_{V_j}^{\hat{\sigma}}, \sigma|_U = \hat{\sigma}\}|}{|\Omega_{\mathcal{V}}^{\mu}|}.$$

Thus, it follows that, for any two colourings $\hat{\sigma}, \hat{\eta} \in \Omega_{\mathcal{R}}$ which differ only on the block V_i , we have

$$\begin{aligned} \pi_{\text{proj}}(\hat{\sigma}) \mathcal{R}_{\mathcal{V}}^{\mu}[\hat{\sigma} \rightarrow \hat{\eta}] &= \frac{\prod_{j=1}^r |\{\sigma \in \Omega_{V_j}^{\hat{\sigma}}, \sigma|_U = \hat{\sigma}\}|}{|\Omega_{\mathcal{V}}^{\mu}|} \cdot g_i \frac{|\{\eta \in \Omega_{V_i}^{\hat{\eta}}, \eta|_U = \hat{\eta}\}|}{|\Omega_{V_i}^{\hat{\sigma}}|} \\ &= \frac{g_i \cdot |\{\sigma \in \Omega_{V_i}^{\hat{\sigma}}, \sigma|_U = \hat{\sigma}\}| \cdot |\{\eta \in \Omega_{V_i}^{\hat{\eta}}, \eta|_U = \hat{\eta}\}|}{|\Omega_{\mathcal{V}}^{\mu}| |\Omega_{V_i}^{\hat{\sigma}}|} \cdot \prod_{j \neq i} |\{\sigma \in \Omega_{V_j}^{\hat{\sigma}}, \sigma|_U = \hat{\sigma}\}|. \end{aligned}$$

This quantity is symmetric in $\hat{\sigma}$ and $\hat{\eta}$. Indeed, since $\hat{\sigma}$ and $\hat{\eta}$ only differ on V_i , we have $\Omega_{V_i}^{\hat{\sigma}} = \Omega_{V_i}^{\hat{\eta}}$. Moreover, for every $j \neq i$, $\hat{\sigma}$ and $\hat{\eta}$ agree on $V_j \cap U$. As a consequence, we have

$$|\{\sigma \in \Omega_{V_j}^{\hat{\sigma}}, \sigma|_U = \hat{\sigma}\}| = |\{\eta \in \Omega_{V_j}^{\hat{\eta}}, \eta|_U = \hat{\eta}\}|.$$

Thus, it follows that $\pi_{\text{proj}}(\hat{\sigma}) \mathcal{R}_{\mathcal{V}}^{\mu}[\hat{\sigma} \rightarrow \hat{\eta}] = \pi_{\text{proj}}(\hat{\eta}) \mathcal{R}_{\mathcal{V}}^{\mu}[\hat{\eta} \rightarrow \hat{\sigma}]$, and $\mathcal{R}_{\mathcal{V}}^{\mu}$ is reversible for π_{proj} . \square

2.3 Monotonicity of Glauber dynamics

Finally, we introduce a monotonicity statement that will allow us to simplify some of our proofs. The previous subsections gave tools to compare the relaxation time of two Markov chains with the same state space but different transitions. Here we are interested in comparing Markov chains with similar transitions but different state spaces. A natural example is comparing the relaxation time of the Glauber dynamics on G and H , where H is a subgraph of G . In general, it is not clear which of the two relaxation times should be smaller, however, if H and G have a particular structure, we will be able to derive a monotonicity bound.

Proposition 9. *Let $G = (V, E)$ be a graph on n vertices, and k be a positive integer. Let $v \in V$ such that $N(v)$ induces a clique of size at most $k - 2$. For any choice of parameters (p_1, \dots, p_n) , the Glauber dynamics \mathcal{L}_V and $\mathcal{L}_{V \setminus \{v\}}$ for k -colourings of V and $V \setminus \{v\}$ respectively and defined with the same parameters satisfy,*

$$\tau(\mathcal{L}_{V \setminus \{v\}}) \leq \tau(\mathcal{L}_V).$$

The proof follows from standard computations and is included in the Appendix.

3 Glauber dynamics for list-colourings of a clique

In this section we analyse the relaxation time of the Glauber dynamics for list-vertex-colourings of a clique. This will be used later in the proof of our main result for edge-colourings of trees.

Consider the clique with vertex set $U = \{u_1, \dots, u_d\}$ with $d \geq 1$. Throughout this section we fix the number of colours to $k = d + 1$. Recall that given a list assignment of U , we can define Ω_U^L and the dynamics \mathcal{L}_U^L , governed by the parameters (p_1, \dots, p_d) , where p_i is the rate at which vertex u_i changes to $c \in [k]$.

For a positive integer t , a list assignment L is t -feasible if it satisfies

- $|L(u_i)| \geq t + 1$ for every $i \in [t]$;
- $|L(u_i)| = d + 1$ for every $i \in [d] \setminus [t]$.

We say $u \in U$ is *free* if $|L(u)| = d + 1$, and *constrained* otherwise. We should stress here that all the lists are subsets of $[d + 1]$, although the results in this section can be generalised to arbitrary lists of size $k \geq d + 1$ using a variant of Proposition 9 for list colouring. If L is t -feasible, we have

$$\left(\prod_{i=1}^t |L(u_i)| - i + 1 \right) (d + 1 - t)! \leq |\Omega_U^L| \leq \left(\prod_{i=1}^t |L(u_i)| \right) (d + 1 - t)! . \quad (6)$$

The chain \mathcal{L}_U^L is symmetric and it will follow from the results below that, if L is t -feasible, for some t , then the chain is also ergodic. So its stationary distribution is uniform on Ω_U^L . The main goal of this section is to prove the following bound on its relaxation time.

Lemma 10. *If L is t -feasible and $k = d + 1$, then we have*

$$\tau(\mathcal{L}_U^L) = O_t \left(\sum_{i=1}^t \frac{d}{p_i} + \sum_{i=t+1}^d \frac{1}{p_i} \right) .$$

In order to prove the lemma we will use the comparison techniques introduced in Section 2.1. Consider the dynamics $\mathcal{L}_{\text{unif}}^L$ on Ω_U^L with generator matrix given for any $\sigma \neq \eta$ by

$$\mathcal{L}_{\text{unif}}^L[\sigma \rightarrow \eta] = \frac{1}{|\Omega_U^L|}. \quad (7)$$

Clearly $\tau(\mathcal{L}_{\text{unif}}^L) = 1$. The main idea will be to compare $\tau(\mathcal{L}_U^L)$ with $\tau(\mathcal{L}_{\text{unif}}^L)$. For technical reasons, it will be easier to introduce an intermediate chain and compare both to it. Define $\mathcal{L}_{\text{int}}^L$ to be the dynamics on Ω_U^L with generator matrix given for any $\sigma \neq \eta$ by

$$\mathcal{L}_{\text{int}}^L[\sigma \rightarrow \eta] = \begin{cases} \frac{1}{|\Omega_U^L|} & \text{if } (\sigma, \eta) \text{ is a good pair,} \\ 0 & \text{otherwise.} \end{cases} \quad (8)$$

Speaking informally, the dynamics $\mathcal{L}_{\text{int}}^L$ can be seen as $\mathcal{L}_{\text{unif}}^L$ where only moves between “good” pairs of colourings are allowed. We defer the formal definition of a good pair to later in the section.

Lemma 10 follows from the combining these two lemmas (proved in the next two subsections) with the fact that $\tau(\mathcal{L}_{\text{unif}}^L) = 1$.

Lemma 11. *If $k = d + 1$, then we have*

$$\frac{\tau(\mathcal{L}_U^L)}{\tau(\mathcal{L}_{\text{int}}^L)} = O_t \left(\sum_{i=1}^t \frac{d}{p_i} + \sum_{i=t+1}^d \frac{1}{p_i} \right).$$

Lemma 12. *If L is t -feasible and $k = d + 1$, then we have*

$$\frac{\tau(\mathcal{L}_{\text{int}}^L)}{\tau(\mathcal{L}_{\text{unif}}^L)} = O_t(1).$$

3.1 Comparing \mathcal{L}_U^L with $\mathcal{L}_{\text{int}}^L$: the proof of Lemma 11

Let w be an additional vertex with $L(w) = [k]$, so w is free. Consider an order on $U \cup \{w\}$ where w is the smallest vertex. We can extend $\alpha \in \Omega_U^L$ to $U \cup \{w\}$, by letting $\alpha(w)$ be the unique colour not in $\alpha(U)$. To every pair $\alpha, \beta \in \Omega_U^L$, we can assign a permutation $f = f(\alpha, \beta)$ on $U \cup \{w\}$ such that $f(x) = y$ if and only if $\alpha(x) = \beta(y)$. In particular, if $\alpha = \beta$ then f is the identity permutation. One can see the permutation f as a blocking permutation: if $v = f(u) \in U$ for some $u \in U$, then u blocks v from being directly recoloured from $\alpha(v)$ to $\beta(v)$. However, if $v = f(w) \in U$, then v can be directly changed to $\beta(v)$ in the colouring α .

It is useful to think about f using its representation as a union of directed cycles. Throughout this section, by cycle in the permutation we mean a cycle of order at least 2. If a cycle in f contains w , we will be able to recolour every vertex in the cycle by successively recolouring the vertices whose preimage is w . The main difficulty will arise from handling the other cycles that do not contain w , which correspond to circular blockings of vertices in U (e.g. $u, v \in U$, u blocks v and v blocks u). In this case, we will need to insert w into this cycle, and then process it. The merging operation corresponds to recolouring a vertex in α with the only available colour, which must be in its list. This motivates the following classification: a cycle in f is a *1-cycle* if it contains w , a *2-cycle* if it does not contain w but has at least one free vertex and a *3-cycle* if it does not contain w and all its vertices are constrained.

Recolouring a 3-cycle C might be hard because the only colours available in $L(C)$ might already be taken by other vertices in the clique. This motivates the definition of good pairs which governs $\mathcal{L}_{\text{int}}^L$.

Definition 13. *A pair of colourings (α, β) is good if $f(\alpha, \beta)$ contains no 3-cycles.*

Let (α, β) be a good pair and $f = f(\alpha, \beta)$. A v -swap is an operation on f that gives the permutation \hat{f} obtained from f by reassigning $\hat{f}(w) = f(v)$ and $\hat{f}(v) = f(w)$. These operations are in bijection with the valid transitions of \mathcal{L}_U^L .

One can define recolouring sequences between α and β using swaps. The order on $U \cup \{w\}$ gives a canonical way to deal with the cycles in $f(\alpha, \beta)$, processing the cycle with the smallest free vertex, at a time. Precisely, while there is a free vertex in a cycle of length at least 2 of f , let v be the smallest one and

- if v is in a 1-cycle, then $v = w$. While $f(w) \neq w$, update f by performing an $f(w)$ -swap.
- if v is in a 2-cycle, then $v \in U$. Update f by performing a v -swap. Then, while $f(w) \neq w$, update f by performing an $f(w)$ -swap.

As there are no 3-cycles, after termination the procedure produces the identity permutation. As swaps correspond to valid recolouring moves, it gives a recolouring path $\gamma_{\alpha, \beta}$ from α to β that uses transitions from \mathcal{L}_U^L . Note that any vertex in U is recoloured at most twice. In fact, a vertex is only recoloured twice if it is the smallest free vertex in a 2-cycle.

For every transition $(\sigma, \eta) \in \mathcal{L}_U^L$, define

$$\Lambda_{\sigma, \eta} = \{(\alpha, \beta) \in \mathcal{L}_{\text{int}}^L : (\sigma, \eta) \in \gamma_{\alpha, \beta}\}.$$

Our goal is to prove Lemma 11 using the weighted canonical paths method from Section 2.1. To this end, the two following lemmas will help us analyse the congestion resulting from this construction.

Lemma 14. *Given a transition $(\sigma, \eta) \in \mathcal{L}_U^L$ and a permutation f , there are at most two good pairs (α, β) such that $f = f(\alpha, \beta)$ and $(\alpha, \beta) \in \Lambda_{\sigma, \eta}$.*

Proof. Let v be the vertex at which σ and η differ. The order in which we recolour the vertices is fixed by the permutation f . Let v_1, \dots, v_ℓ be the vertices in the order they are recoloured with repetitions, where v_i is the vertex recoloured at the i -th step. Given this sequence, at every step the only valid transition is to perform a v_i -swap. If the recolouring done by the transition (σ, η) corresponds to the i^* -th step in the sequence, then α and β are fully determined. Indeed, β can be recovered from η by sequentially performing v_i -swaps for every $i > i^*$. Symmetrically, α can be recovered from σ by performing v_i -swaps for every $i < i^*$ (recall that the u -swap operation is an involution).

Since every vertex is recoloured at most twice in a recolouring path, v appears at most twice in the sequence v_1, \dots, v_ℓ . So there are at most 2 choices for $i^* \in [\ell]$ with $v_{i^*} = v$, and by the argument above, there are at most two pairs (α, β) , with $f(\alpha, \beta) = f$, and $(\alpha, \beta) \in \Lambda_{\sigma, \eta}$. \square

Using the previous lemma, we can bound the size of $\Lambda_{\sigma, \eta}$

Lemma 15. *Suppose that σ and η differ at $v \in U$. If v is free, then $|\Lambda_{\sigma, \eta}| = O_t(|\Omega_U^L|)$. If v is constrained, then $|\Lambda_{\sigma, \eta}| = O_t(d|\Omega_U^L|)$.*

Proof. By Lemma 14, it suffices to bound the number of permutations f which are compatible with (σ, η) . The key idea is the following claim that not all candidates for f are compatible, as constrained vertices can only block and be blocked by vertices with colours in their lists.

Claim 16. Let $(\alpha, \beta) \in \Lambda_{\sigma, \eta}$ with $f(\alpha, \beta) = f$. For every constrained vertex $u \in U$ with $u \neq v$, one of the following holds:

- i) $\eta(f(u)) = \beta(f(u))$;
- ii) $\sigma(f^{-1}(u)) = \alpha(f^{-1}(u))$;

Moreover, i) holds if and only if u is recoloured in $\gamma_{\alpha, \beta}$ before the transition (σ, η) .

Proof. By construction of the recolouring paths, constrained vertices are only recoloured once. Let C be the cycle containing u . If $v \notin V(C)$, then the vertices in C are recoloured either all before v , or all after v . In particular we have $\sigma|_{V(C)} = \eta|_{V(C)}$, and it is equal to $\beta|_{V(C)}$ if u is recoloured before (σ, η) and to $\alpha|_{V(C)}$ otherwise. Since $f(u), f^{-1}(u) \in V(C)$, either i) or ii) holds.

If $v \in V(C)$, then

- If $f(u)$ is not the smallest free vertex in C , then assume that i) does not hold. Since $f(u)$ is recoloured only once, this means that $\eta(f(u)) = \alpha(f(u))$. Consequently, $f(u)$ is recoloured after the transition (σ, η) , and since $u \neq v$, this is also the case for u . This means that $f^{-1}(u)$ is recoloured either during or after the transition (σ, η) , and in both cases ii) holds.
- Symmetrically, if $f^{-1}(u)$ is not the smallest free vertex in C , then assuming that ii) does not hold we conclude that u has been coloured before the transition (σ, η) and that i) holds.
- Finally, if $f(u) = f^{-1}(u)$ is the smallest free vertex in C , then $v = f(u)$. In this case, either this is the first time v is recoloured, so u is recoloured after (σ, η) and $\sigma(v) = \alpha(v)$, and ii) holds, or it is the second time v is recoloured, so u is recoloured before (σ, η) and $\eta(v) = \beta(v)$ and i) holds.

□

Suppose that v is free. As there are at most t constrained vertices, there are at most 2^t choices to decide which of items i) or ii) holds for them. For $s \in [t]$, let x_1, \dots, x_s be the constraint vertices that satisfy i) and y_1, \dots, y_{t-s} be the ones that satisfy ii). For every $i \in [s]$, by definition of f we have $\eta(f(x_i)) = \beta(f(x_i)) = \alpha(x_i) \in L(x_i)$ and there are at most $|L(x_i)|$ choices $z \in U \cup \{w\}$ for $f(x_i) = z$, namely the ones with $\eta(z) \in L(x_i)$. Analogously, for $i \in [t-s]$, by definition of f we have $\sigma(f^{-1}(y_i)) = \alpha(f^{-1}(y_i)) = \beta(y_i) \in L(y_i)$ and there are at most $|L(y_i)|$ choices $z \in U \cup \{w\}$ for $f(z) = y_i$. Thus, there are at most $\prod_{i=1}^s |L(x_i)| \prod_{i=1}^{t-s} |L(y_i)| = \prod_{i=1}^t |L(u_i)|$ choices for the images of x_i and the preimages of y_i . Note that $f(x_i) \neq y_j$ for every i, j as i) holds if and only if u is recoloured before (σ, η) and v is a free vertex. So these choices fix exactly t images in f . Finally, there are at most $(d+1-t)!$ ways to complete f by choosing successively the image of the remaining elements. Using Lemma 14 and (6), we have

$$|\Lambda_{\sigma, \eta}| \leq 2 \cdot 2^t (d+1-t)! \prod_{i=1}^t (|L(u_i)|) \leq 2^{t+1} \left(\prod_{i=1}^t \frac{|L(u_i)|}{|L(u_i)| - i + 1} \right) |\Omega_U^L| = O_t(|\Omega_U^L|).$$

Assume now that v is constrained. There are at most 2^{t-1} ways to choose a configuration of i) and ii) for the remaining constrained vertices. In comparison to the case where v is free, as neither i) nor ii) holds for $u = v$, there is an extra factor $|L(v)| \leq d + 1$ to choose either the image or the preimage of v in f . Similarly as before, it follows that

$$|\Lambda_{\sigma,\eta}| = O_t(d|\Omega_U^L|).$$

□

We are now in a good situation to prove Lemma 11.

Proof of Lemma 11. In order to apply the weighted canonical paths theorem to \mathcal{L}_U^L and $\mathcal{L}_{\text{int}}^L$, we need to choose a weight function ω for all $(\sigma, \eta) \in \mathcal{L}_U^L$. Let v be the vertex where σ and η differ. We define ω as follows

$$\omega(\sigma, \eta) := \begin{cases} d/p_i & \text{if } v = u_i \text{ for } i \in [t], \\ 1/p_i & \text{if } v = u_i \text{ for } i \in [d] \setminus [t]. \end{cases} \quad (9)$$

As for every recolouring path γ , each element of U is recoloured at most twice, it follows that

$$|\gamma|_\omega \leq 2 \left(\sum_{i=1}^t \frac{d}{p_i} + \sum_{i=t+1}^d \frac{1}{p_i} \right). \quad (10)$$

Both stationary distributions of \mathcal{L}_U^L and $\mathcal{L}_{\text{int}}^L$ are uniform on Ω_U^L . Also recall that $\mathcal{L}_{\text{int}}^L[\alpha \rightarrow \beta] = 1/|\Omega_U^L|$ for any good pair (α, β) .

Using Lemma 15, regardless of whether the vertex where σ and η differ is free or constrained, we can bound the congestion of the transition (σ, η) as follows

$$\begin{aligned} \rho_{\sigma,\eta} &= \frac{1}{\mathcal{L}_U^L[\sigma \rightarrow \eta]\omega(\sigma, \eta)} \sum_{(\alpha, \beta) \in \Lambda_{\sigma,\eta}} \frac{|\gamma_{\alpha,\beta}|_\omega}{|\Omega_U^L|} \\ &\leq \frac{|\Lambda_{\sigma,\eta}|}{\mathcal{L}_U^L[\sigma \rightarrow \eta]\omega(\sigma, \eta)|\Omega_U^L|} \cdot \max_{\alpha, \beta} |\gamma_{\alpha,\beta}|_\omega \\ &= O_t(\max_{\alpha, \beta} |\gamma_{\alpha,\beta}|_\omega) = O_t \left(\sum_{i=1}^t \frac{d}{p_i} + \sum_{i=t+1}^d \frac{1}{p_i} \right). \end{aligned}$$

The desired result follows from Proposition 3. □

3.2 Comparing $\mathcal{L}_{\text{int}}^L$ with $\mathcal{L}_{\text{unif}}^L$: the proof of Lemma 12

The proof of this lemma uses the fractional paths method with uniform weights. To define the paths, we first split the set of constrained vertices into two subsets. Let A be the set of constrained vertices u satisfying $|L(u)| \leq 2(3t + 2)$, and let B be the remaining ones. Before we define the paths, we will need the following result.

Lemma 17. *Let $\alpha, \beta \in \Omega_U^L$ and let ξ_A be an L -colouring of the vertices in A . Assume that $\alpha|_A$ and ξ_A differ on at most one vertex, and similarly for ξ_A and $\beta|_A$. There exists a constant $c(t) > 0$ such that there are at least $c(t)|\Omega_U^L|$ colourings $\xi \in \Omega_U^L$ satisfying that $\xi|_A = \xi_A$, and both (α, ξ) and (ξ, β) are good pairs.*

Proof. Assume that $A = \{u_1, \dots, u_a\}$, so $|B| = t - a$. Construct the colouring ξ by setting $\xi|_A = \xi_A$ and then choosing the colours in $U \setminus A$ one by one, starting from the vertices in B , as follows:

- for each $v \in B$, choose $\xi(v) \notin \xi(A) \cup \alpha(A \cup B) \cup \beta(A \cup B)$ that has not been used already;
- for each free vertex, choose a colour not used by the vertices already coloured in ξ .

When we choose the colour for $v \in B$, there are at most $3t + 2$ forbidden colours. Indeed, there are at most $a + 2(t - a) + 2$ colours in $\xi(A) \cup \alpha(A \cup B) \cup \beta(A \cup B)$, and at most $t - a$ colours used by the previous vertices in B that have already been coloured by ξ . Thus, for $v \in B$ there are at least $|L(v)| - (3t + 2) \geq |L(v)|/2$ choices for $\xi(v)$. Using (6), the total number of colourings extending ξ_A is at least

$$\begin{aligned} \prod_{v \in B} (|L(v)| - (3t + 2)) \cdot (d + 1 - t)! &\geq \frac{(d + 1 - t)!}{2^{t-a}} \prod_{v \in B} |L(v)| \\ &\geq \frac{(d + 1 - t)!}{2^t (3t + 2)^a} \prod_{i=1}^t |L(u_i)| \\ &\geq \frac{|\Omega_U^L|}{(2(3t + 2))^t}. \end{aligned}$$

It suffices to show that for any such extension ξ of ξ_A , (α, ξ) and (ξ, β) are good pairs. We only prove it for (α, ξ) as the other case is symmetric. Assume by contradiction that $f(\alpha, \xi)$ contains a 3-cycle, and let C be this cycle. Then $V(C) \cap B = \emptyset$. Indeed, if $v \in V(C) \cap B$, then there exists a vertex $u \in A \cup B$ with $\alpha(u) = \xi(v)$, but this contradicts the fact that $\xi(v) \notin \alpha(A \cup B)$. Thus, $V(C) \subseteq A$, but this is not possible since $\alpha|_A$ and $\xi|_A$ differ by at most one vertex. \square

We can now compare the relaxation times of $\mathcal{L}_{\text{int}}^L$ and $\mathcal{L}_{\text{unif}}^L$.

Proof of Lemma 12. We use the fractional paths method. Note that both $\mathcal{L}_{\text{int}}^L$ and $\mathcal{L}_{\text{unif}}^L$ are ergodic, reversible and symmetric and that their stationary distributions are uniform on Ω_U^L .

It suffices to define a collection of fractional paths $\Gamma_{\alpha, \beta}$ between any two colourings α and β in Ω_U^L . Since there are at most t constrained vertices and their lists have size at least $t + 1$, we can find a sequence of L -colourings of A , $\alpha|_A = \xi_A^0, \xi_A^1, \dots, \xi_A^m = \beta|_A$, such that any two consecutive colourings differ by one vertex. Let M be an upper bound on the length of these paths for every α, β , which only depends on t .

For $m \in [M]$, let $\Gamma_{\alpha, \beta}^m$ be the collection of all the paths of the form $\alpha = \xi^0, \xi^1, \dots, \xi^m = \beta$ where $\xi^i|_A = \xi_A^i$, and (ξ^{i-1}, ξ^i) is a good pair for $i \in [m]$. By Lemma 17, for each $i \in [m - 1]$ there are at least $c(t)|\Omega_U^L|$ choices for ξ^i , independently of the choices of the other ξ^j for $j \neq i$. Thus, $|\Gamma_{\alpha, \beta}^m| \geq (c(t)|\Omega_U^L|)^{m-1}$, and if g is the uniform flow, then each $\gamma \in \Gamma_{\alpha, \beta}^m$ satisfies $g(\gamma) \leq (c(t)|\Omega_U^L|)^{-(m-1)}$. Let $\Gamma^m = \{\Gamma_{\alpha, \beta}^m : \alpha, \beta \in \Omega_U^L\}$ and $\Gamma = \cup_{m=1}^M \Gamma^m$.

We need to bound the congestion of any good pair (σ, η) . We fix $m \in [M]$ and will bound the contribution of Γ^m to it. Let $\gamma = \xi^0, \xi^1, \dots, \xi^m \in \Gamma^m$ containing (σ, η) . There are at most m choices for $i \in [m]$ such that $\sigma = \xi^{i-1}$ and $\eta = \xi^i$. Then, there are at most $|\Omega_U^L|^{m-1}$ choices for ξ^j with $j \notin \{i - 1, i\}$. Each such path satisfies $g(\gamma) \leq (c(t)|\Omega_U^L|)^{-(m-1)}$. Thus,

$$\sum_{\substack{\gamma \in \Gamma^m \\ \gamma \ni (\sigma, \eta)}} g(\gamma) |\gamma| \leq mc(t)^{-(m-1)}.$$

Hence,

$$\rho_{\sigma,\eta} \leq \sum_{m=1}^M mc(t)^{-(m-1)} = O_t(1),$$

and, by Proposition 4, we obtain the desired result. \square

3.3 Dynamics of two cliques intersecting at a vertex

In this section we study a similar dynamics, that we will also use in the main proof. Let z be a vertex. For $d \geq 1$, let $X = \{z, x_1, \dots, x_{d-1}\}$ and $Y = \{z, y_1, \dots, y_{d-1}\}$ two sets of vertices and consider the graph with vertex set $Z = X \cup Y$ where each set X and Y induces a clique. As before, we fix the number of colours $k = d + 1$. For a list assignment L of Z , recall the definition of Ω_Z^L and the dynamics \mathcal{L}_Z^L , where we denote by $p_z, p_1, \dots, p_{d-1}, q_1, \dots, q_{d-1}$ the parameters for $z, x_1, \dots, x_{d-1}, y_1, \dots, y_{d-1}$, respectively.

Let t, t_X, t_Y be non-negative integers with $t \geq t_X + t_Y$ and $1 \leq t_X, t_Y \leq d - 1$. Without loss of generality, we will assume that d is sufficiently large with respect to t . If this is not the case, then $|\Omega_Z^L|$ is a constant depending on t , and the relaxation time is $O_t(1)$. A list assignment L is (t, t_X, t_Y) -feasible if

- $|L(z)| = d + 1$;
- $|L(x_i)|, |L(y_j)| \geq t$ for every $i \in [t_X]$ and $j \in [t_Y]$;
- $|L(x_i)|, |L(y_j)| = d + 1$ for every $i \in [d - 1] \setminus [t_X]$ and $j \in [d - 1] \setminus [t_Y]$;

We define free and constrained vertices as before, with the exception of z which is considered a constrained vertex. If L is (t, t_X, t_Y) -feasible, then

$$|\Omega_Z^L| \geq \left(\prod_{i=1}^{t_X} |L(x_i)| - i + 1 \right) \left(\prod_{j=1}^{t_Y} |L(y_j)| - j + 1 \right) \cdot (d + 1 - t)(d - t_X)!(d - t_Y)! \quad (11)$$

$$|\Omega_Z^L| \leq \left(\prod_{i=1}^{t_X} |L(x_i)| \right) \left(\prod_{j=1}^{t_Y} |L(y_j)| \right) \cdot d(d - t_X)!(d - t_Y)!. \quad (12)$$

As before, the chain is symmetric and, if L is (t, t_X, t_Y) -feasible for some t, t_X, t_Y satisfying the conditions stated above, it follows from results below that it is ergodic, so its stationary distribution is uniform.

We will prove a bound analogous to the one in Lemma 10 on the relaxation time of \mathcal{L}_Z^L .

Lemma 18. *If L is (t, t_X, t_Y) -feasible and $k = d + 1$, then we have*

$$\tau(\mathcal{L}_Z^L) = O_t \left(\frac{d^2}{p_z} + \sum_{i=1}^{t_X} \frac{d}{p_i} + \sum_{i=t_X+1}^{d-1} \frac{1}{p_i} + \sum_{j=1}^{t_Y} \frac{d}{q_j} + \sum_{j=t_Y+1}^{d-1} \frac{1}{q_j} \right).$$

The proof follows the same lines as the proof of Lemma 10, so we will only sketch it, stressing the parts where the two differ. For $\alpha \in \Omega_Z^L$, denote by α_X and α_Y the restrictions of α onto X and Y , respectively. As we did for the clique, we extend α by adding two artificial vertices: w_X

in X and w_Y in Y , and by assigning to them the only available colour in each set. Define the permutations f_X and f_Y as before.

We say that (α, β) is good if and only if both (α_X, β_X) and (α_Y, β_Y) are good (i.e., there is no 3-cycle in the permutations f_X and f_Y). Remember that z is a constrained vertex, and as a consequence if (α, β) are good, every cycle in the permutations contains a free vertex different from z . Redefine the dynamics $\mathcal{L}_{\text{unif}}^L$ on Ω_Z^L as in (7), and, using the new definition of good pairs, redefine $\mathcal{L}_{\text{int}}^L$ on Ω_Z^L as in (8). We proceed in two steps by bounding the ratios of the relaxation times of $\mathcal{L}_{\text{int}}^L$ and $\mathcal{L}_{\text{unif}}^L$ and of $\mathcal{L}_{\text{int}}^L$ and $\mathcal{L}_{\text{unif}}^L$.

Lemma 19. *If $k = d + 1$, then we have*

$$\frac{\tau(\mathcal{L}_Z^L)}{\tau(\mathcal{L}_{\text{unif}}^L)} = O_t \left(\frac{d^2}{pz} + \sum_{i=0}^{t_X} \frac{d}{p_i} + \sum_{i=t_X+1}^{d-1} \frac{1}{p_i} + \sum_{j=0}^{t_Y} \frac{d}{q_j} + \sum_{j=t_Y+1}^{d-1} \frac{1}{q_j} \right).$$

Sketch of the proof. We reuse the recolouring paths defined in Lemma 11 for the clique. Given two colourings α and β , denote by $\gamma_{\alpha_X, \beta_X}^X$ and $\gamma_{\alpha_Y, \beta_Y}^Y$ the recolouring paths constructed for each of the two sets X and Y independently. Observe that for each path in these sets, each constrained vertex is recoloured at most once. In particular, z changes its colour at most once. Construct the recolouring path $\gamma_{\alpha, \beta}$ in the following way:

- apply the recolourings in $\gamma_{\alpha_X, \beta_X}^X$ until z needs to be recoloured;
- apply the recolourings in $\gamma_{\alpha_Y, \beta_Y}^Y$;
- apply the remaining recolourings in $\gamma_{\alpha_X, \beta_X}^X$.

Note that in the second step, z can be safely recoloured with $\beta(z)$ because its target colour is available in X , since the next move according to $\gamma_{\alpha_X, \beta_X}^X$ would be to recolour z with colour $\beta(z)$. We need to bound the congestion of each transition for this collection of paths.

For a transition (σ, η) , let $\Lambda_{\sigma, \eta}$ be the set of good pairs α, β such that $\gamma_{\alpha, \beta}$ contains (σ, η) . The analogues of Lemmas 14 and 15 hold in this setting. In particular, for every (σ, η) differing at a vertex v , if v is free then $|\Lambda_{\sigma, \eta}| = O_t(|\Omega_Z^L|)$, if $v \neq z$ is constrained then $|\Lambda_{\sigma, \eta}| = O_t(d|\Omega_Z^L|)$, and if $v = z$ then $|\Lambda_{\sigma, \eta}| = O_t(d^2|\Omega_Z^L|)$ as we get an extra factor d for each permutation. This allows us to bound the congestion of a transition as in the previous section, and so Lemma 19 follows. \square

Lemma 20. *If L is (t, t_X, t_Y) -feasible and $k = d + 1$, then we have*

$$\frac{\tau(\mathcal{L}_{\text{int}}^L)}{\tau(\mathcal{L}_{\text{unif}}^L)} = O_t(1).$$

Sketch of the proof. The lemma can be proved using the same steps as in the proof of Lemma 12. Let A_X and A_Y be the set of constrained vertices with lists of size at most $2(3t + 2)$ in X and Y , respectively. Let $A = A_X \cup A_Y$ and let B be the set of constrained vertices that are not in A . As z is a constrained vertex with $|L(z)| = d + 1 > 2(3t + 2)$, we have $z \in B$. The analogous of Lemma 17 still holds in this setting.

For any $\alpha, \beta \in \Omega_Z^L$, we would like to find a sequence $\alpha|_A = \xi_A^0, \dots, \xi_A^m = \beta|_A$ of colourings of A such that each consecutive pair differs only at one vertex. As $z \notin A$ and as all constrained vertices in A have a list of size at least $t + 1$, this sequence can be found by independently recolouring A_X

and A_Y . Arguing as in the proof of Lemma 12, we can construct many recolouring paths between α and β with transitions that correspond to good pairs. Then, Proposition 4 implies the desired result. \square

4 Glauber dynamics on edge-colourings of trees

In this section we prove our main theorem. We follow a similar approach to the one of Lucier and Molloy in [19] for vertex-colourings, by recursively splitting the tree into smaller subtrees using block dynamics. However, there are several points where our strategies differ.

4.1 Relaxation time of block dynamics

In all this section, we will assume that $G = (V_G, E_G)$ is a d -regular tree, that is every internal vertex has degree *exactly* d . We also assume that $k = d + 1$.

Definition 21. *A subtree T of G is splitting if one of the following holds.*

- T is a single edge,
- T has fringe interior boundary $|\bar{\partial}T| \leq 2$. If $\bar{\partial}T = \{e, f\}$, then e and f are not incident.

Fix $\mu \in \Omega_{E_G}$, and fix $T = (V, E)$ a splitting subtree of G . Note that since T has fringe boundary, it is also d -regular. The central point of our proof is to study \mathcal{L}_E^μ by decomposing it into the dynamics of its subtrees using the block dynamics defined in Section 2.2. We will assume that T is rooted in one of the two following ways.

Vertex-rooted trees: The root of T is $r \in V$, an internal vertex of T . Let e_1, \dots, e_d be the edges incident to r . For each $i \in [d]$, we consider the block formed by the edges of the subtree hanging from e_i .

Edge-rooted trees: The root of T is an edge $e = xy$ where x and y are internal vertices of T (so, $e \notin \bar{\partial}T$). Let e_1, \dots, e_{2d-2} be the edges incident to e . We let $\{e\}$ be a block, and for every $i \in [2d-2]$, we consider the block formed by the edges of the subtree hanging from e_i .

In both cases, we denote by $\mathcal{E} = \{E_1, \dots, E_r\}$ the block partition described above. Note that each block in \mathcal{E} contains at most one edge incident to edges in other blocks. Let $H = (U, F)$ be the subgraph induced by these edges. Note that in the case of a vertex rooted tree, H is a star, and in the case of an edge rooted tree, H is a bi-star: two stars joined by an edge. For each $i \in [r]$, let T_i be the subtree with edge set E_i .

Throughout this section we will make the following two assumptions on \mathcal{E} .

- (A1) T_i is splitting for every $i \in [r]$;
- (A2) $\mathcal{L}_{E_i}^\sigma$ is ergodic for every $i \in [r]$ and every $\sigma \in \Omega_E^\mu$.

Thus, we can define the reduced block dynamics on \mathcal{E} with boundary condition μ . Recall from Section 2.2 that its state space is $\Omega_{\mathcal{R}}$, the restriction to H of the colourings in Ω_E^μ . Moreover, its

transition matrix is given for any $\hat{\sigma} \neq \hat{\eta}$:

$$\mathcal{R}_{\mathcal{E}}^{\mu}[\hat{\sigma} \rightarrow \hat{\eta}] = \begin{cases} g_i \pi_{\text{proj}}^{i, \hat{\sigma}}(\hat{\eta}) & \text{if there exists } i \in [r] \text{ such that } \hat{\eta} = \eta|_U \text{ for some } \eta \in \Omega_{E_i}^{\hat{\sigma}}, \\ 0 & \text{otherwise.} \end{cases} \quad (13)$$

where $\pi_{\text{proj}}^{i, \hat{\sigma}}$ is the projection of $\pi_{E_i}^{\hat{\sigma}}$ onto F (see (4)). To bound the relaxation time on T , we will proceed in two steps. First we compare the original dynamics $\mathcal{L}_{\mathcal{E}}^{\mu}$ to the reduced block dynamics on \mathcal{E} using Propositions 6 and 7 from Section 2.2. Then, we bound the relaxation time for the reduced block dynamics using the results from Section 3. The three following lemmas will help us for the second step. They show that the transitions rates and the stationary distribution of the reduced block dynamics are close to uniform. We will first prove bounds on the stationary distribution for the reduced block dynamics, and then proceed to bound the relaxation time for the reduced block dynamics.

Lemma 22. *Assuming (A1)–(A2), the reduced block dynamics $\mathcal{R}_{\mathcal{E}}^{\mu}$ is ergodic and reversible, and its stationary distribution $\pi_{\mathcal{R}}$ is the projection of π_E^{μ} onto F .*

Proof. Let H^{ℓ} be the line graph with vertex set F , the set of edges of H . Consider the Glauber dynamics \mathcal{L}_F^L with the following list constraints on $e \in F$

- $L(e) = [k] \setminus \mu(N(e) \cap \partial T)$ if $e \in \overline{\partial T}$,
- $L(e) = [k]$ for every other edge.

Then, since each block E_i contains only one edge in H , the reduced block dynamics $\mathcal{R}_{\mathcal{E}}^{\mu}$ has exactly the same transitions as \mathcal{L}_F^L , but with possibly different probability transitions. Thus, $\mathcal{R}_{\mathcal{E}}^{\mu}$ is ergodic if and only if \mathcal{L}_F^L is.

If T is vertex-rooted, then H is a star, and H^{ℓ} is a clique. Additionally, since T is splitting, the two edges in $\overline{\partial T}$ are not adjacent, and in particular only one is in H . This edge, if it exists, is assigned a list of length 2, so L is 1-feasible. The ergodicity of \mathcal{L}_F^L follows from Lemma 10 with $t \leq 1$.

If T is edge-rooted, then H is a bi-star, and H^{ℓ} is composed of two cliques intersecting at one vertex. Moreover, the two edges in $\overline{\partial T}$ cannot be in the same side of the bi-clique, and each has a list of length at least 2. So L is (2, 1, 1)-feasible. The ergodicity of \mathcal{L}_F^L follows from Lemma 18 with $t_X, t_Y \leq 1$ and $t \leq 2$.

Lemma 8 implies that the reduced block dynamics is reversible for the projection of π_E^{μ} onto F , concluding the proof. \square

Before giving a bound on the relaxation time of the reduced block dynamics, we will prove some bounds on its stationary distribution $\pi_{\mathcal{R}}$ to show that it deviates from a uniform distribution by at most a constant factor. To this end, the following lemma is a technical tool that we will reuse later. It shows that, given a boundary configuration, under the uniform distribution the probability that an edge e_i is assigned an available colour is close to uniform.

Lemma 23. *For any subtrees T_i with edge-set E_i , any $e_i \in \overline{\partial T_i}$ and any $\sigma \in \Omega_E^{\mu}$, let $C = \sigma(N(e_i) \cap \partial T_i)$. Assuming (A1)–(A2), for every $c \in [k] \setminus C$, we have*

$$\pi_{E_i}^{\sigma}(\{\xi \in \Omega_{E_i}^{\sigma} : \xi(e_i) = c\}) = \frac{1}{2}(1 + O(1/d)).$$

Moreover, if $|\overline{\partial T_i}| = 1$, then

$$\pi_{E_i}^\sigma(\{\xi \in \Omega_{E_i}^\sigma : \xi(e_i) = c\}) = 1/2 .$$

Proof. First assume that $\overline{\partial T_i} = \{e_i, f_i\}$. As G is d -regular, and T_i is splitting, in particular its boundary is fringe. Since $k = d + 1$, this implies that there are exactly two colours available for e_i and two colours for f_i .

We will bound $|\Omega_{E_i}^\sigma|$ and $|\Omega_{E_i}^\sigma(c)|$, the number of colourings in $\Omega_{E_i}^\sigma$ that assign c to e_i . Let $P = (V_P, E_P)$ be the unique path in T_i that connects e_i and f_i and let $s = |V_P|$. As T is splitting and $|\overline{\partial T_i}| = 2$, we have $s \geq 4$. If we fix a colouring ξ_P of E_P , observe that the number of colourings of $\xi \in \Omega_{E_i}^\sigma$, such that $\xi|_P = \xi_P$ is independent of ξ_P . Indeed, if we remove E_P , we obtain a collection of rooted subtrees T'_1, \dots, T'_s with root $v_i \in V_P$. Given ξ_P , there are exactly $(d-1)!$ ways to colour the edges of T'_i incident to v_i , and for each internal vertex, there are exactly $d!$ ways to choose a colouring of the edges hanging from it.

Therefore, in order to bound the ratio $|\Omega_{E_i}^\sigma(c)|/|\Omega_{E_i}^\sigma|$, we only need to compute $|\Omega_{E_P}^\sigma|$, and $|\Omega_{E_P}^\sigma(c)|$, respectively the number of colourings of P compatible with σ , and the number of these colourings ξ_P for which $\xi_P(e_i) = c$. We can obtain a colouring of P by first colouring e_i and f_i , and then choosing the colour of the other edges in P in the order they appear on the path from e_i to f_i . As $s \geq 4$, there is at least one edge in $E_P \setminus \{e_i, f_i\}$. For each of these edges except for the last one, there are d choices of colours. For the last edge there are either d or $d-1$ choices. It follows that

$$\begin{aligned} 4d^{s-1}(d-1) &\leq |\Omega_{E_P}^\sigma| \leq 4d^s , \\ 2d^{s-1}(d-1) &\leq |\Omega_{E_P}^\sigma(c)| \leq 2d^s . \end{aligned}$$

We conclude that

$$\pi_{E_i}^\sigma(\{\xi \in \Omega_{E_i}^\sigma : \xi(e) = c\}) = \frac{|\Omega_{E_i}^\sigma(c)|}{|\Omega_{E_i}^\sigma|} = \frac{|\Omega_{E_P}^\sigma(c)|}{|\Omega_{E_P}^\sigma|} = \frac{1}{2}(1 + O(1/d)) .$$

The second statement follows by a simple symmetry argument. \square

Lemma 24. *Assuming (A1)–(A2), for every $\hat{\sigma} \in \Omega_{\mathcal{R}}$, we have*

$$\pi_{\mathcal{R}}(\hat{\sigma}) = \frac{1 + O(1/d)}{|\Omega_{\mathcal{R}}|} .$$

Proof. Recall that $H = (U, F)$ and that $e_i \in F \cap E_i$ is unique edge of E_i in H . Let $\hat{\sigma} \in \Omega_{\mathcal{R}}$, we will compute $\pi_{\mathcal{R}}(\hat{\sigma})$ by using Lemma 22 and by bounding the number of $\sigma \in \Omega_E^\mu$ such that $\sigma|_F = \hat{\sigma}$. If $t_i = 1$, then e_i is the only boundary edge and, by symmetry, the number of extensions of $\hat{\sigma}$ in E_i does not depend on $\hat{\sigma}(e_i)$. If $t_i = 2$, then there exists $f \in E_i$ with $f \neq e_i$ such that $f \in \overline{\partial T_i}$. As in the proof of Lemma 23, in this case the number of extensions is the same, up to a $1 + O(1/d)$ multiplicative factor. Since T_i is splitting, there are at most two values of $i \in [r]$ with $t_i = 2$. It follows that, up to a $1 + O(1/d)$ multiplicative factor, each $\hat{\sigma}$ has the same number of extensions. This concludes the proof. \square

We will also need the following simple bound on the gap of the dynamics of a single edge.

Lemma 25. *Let $e \in E_G$ be an edge of G . Then,*

$$\min_{\mu \in \Omega_{E_G}} \text{Gap}(\mathcal{L}_{\{e\}}^\mu) \geq \frac{2}{d+1}.$$

Proof. Observe that $\Omega_{\{e\}}^\mu$ is the set of colourings of a single edge with $k_0 \geq 1$ colours, where each transition happens at rate $1/(d+1)$. If $k_0 = 1$, then the relaxation time is 1. If $k_0 \geq 2$, all positive eigenvalues of $-\mathcal{L}_{\{e\}}^\mu$ are equal to $k_0/(d+1)$, so $\text{Gap}(\mathcal{L}_{\{e\}}^\mu) \geq 2/(d+1)$. \square

We can finally obtain a bound on the relaxation time of \mathcal{L}_E^μ .

Lemma 26. *Assuming (A1)–(A2), the following holds,*

$$\tau(\mathcal{L}_E^\mu) = O\left(d^3 + \sum_{i=1}^r \tau_i\right),$$

where $\tau_i := \max_{\sigma \in \Omega_E^\mu} \tau(\mathcal{L}_{E_i}^\sigma)$.

Proof. Using Propositions 6 and 7, we know that the relaxation time of \mathcal{L}_E^μ satisfies

$$\tau(\mathcal{L}_E^\mu) \leq \tau(\mathcal{R}_\mathcal{E}^\mu). \quad (14)$$

Thus, to get the result, we only need to bound the relaxation time of the reduced block dynamics with partition \mathcal{E} . We define an alternative dynamics. Let $\mathcal{R}_{\text{const}}$ be the continuous-time Markov chain with state space $\Omega_\mathcal{R}$ and generator matrix given for any $\hat{\sigma} \neq \hat{\eta}$ by

$$\mathcal{R}_{\text{const}}[\hat{\sigma} \rightarrow \hat{\eta}] = \begin{cases} g_i & \text{if } \hat{\sigma} \text{ and } \hat{\eta} \text{ differ only at } e_i, \\ 0 & \text{otherwise,} \end{cases}$$

where $g_i = 1/\tau_i$. Observe that $\mathcal{R}_\mathcal{E}^\mu$ and $\mathcal{R}_{\text{const}}$ have the same state space and transitions (but different transition probabilities). For every $\hat{\sigma}, \hat{\eta} \in \Omega_\mathcal{R}$, Lemma 23 implies that $\pi_{\text{proj}}^{\hat{\sigma}, i}(\hat{\eta}) = \pi_{E_i}^\sigma(\{\xi \in \Omega_{E_i}^\sigma : \xi(e_i) = \hat{\eta}(e_i)\}) = \Theta(1)$, where σ is an arbitrary colouring in $\Omega_\mathcal{E}^*$. So $\mathcal{R}_\mathcal{E}^\mu[\hat{\sigma} \rightarrow \hat{\eta}] = \Theta(\mathcal{R}_{\text{const}}[\hat{\sigma} \rightarrow \hat{\eta}])$. Moreover, the stationary distribution π_{const} of $\mathcal{R}_{\text{const}}$ is uniform on $\Omega_\mathcal{R}$, and by Lemma 24 we have $\pi_\mathcal{R}(\hat{\sigma}) = \Theta(\pi_{\text{const}}(\hat{\sigma}))$ for every $\hat{\sigma} \in \Omega_\mathcal{R}$. Thus, it follows from Corollary 5 that

$$\tau(\mathcal{R}_\mathcal{E}^\mu) = \Theta(\tau(\mathcal{R}_{\text{const}})) \quad (15)$$

We will bound the relaxation time of $\mathcal{R}_{\text{const}}$ using the results in Section 3, and conclude using (14) and (15).

Suppose first that T is vertex-rooted. This implies that H induces a star with edges e_1, \dots, e_d . Consider the clique with vertex-set $U = \{u_1, \dots, u_d\}$, where u_i is identified with the edge e_i , and the list assignment L of U defined by $L(u_i) = [k] \setminus \mu(N(e_i) \cap \partial T)$. Up to relabelling of the edges, as $k = d+1$, the list assignment L is 1-feasible. Consider the dynamics \mathcal{L}_U^L with probabilities $p_i := g_i$. We can identify $\Omega_\mathcal{R}$ and $\mathcal{R}_{\text{const}}$ with Ω_U^L and \mathcal{L}_U^L . By applying Lemma 10 with $t \leq 1$ we have,

$$\tau(\mathcal{R}_{\text{const}}) = \tau(\mathcal{L}_U^L) = O\left(\sum_{i=1}^t \frac{d}{g_i} + \sum_{i=t+1}^d \frac{1}{g_i}\right) = O\left(d^2 + \sum_{i=1}^d \frac{1}{g_i}\right),$$

where we used Lemma 25 in the last equality.

If T is edge rooted, then H is a bi-star. We can do the same proof replacing Lemma 10 by Lemma 18 obtaining a similar bound with an additive factor of order d^3 . \square

4.2 Proof of Theorem 1

Let $k \geq \Delta + 1$ and $G = (V_G, E_G)$ be a tree on n vertices with maximum degree at most Δ . From (1) and (2) to prove Theorem 1 it suffices to bound $\tau(\mathcal{L}_{E_G})$.

Let $d := k - 1 \geq \Delta$. Construct the d -regularisation $G^d = (V_G^d, E_G^d)$ of G by adding $d - |N(u)|$ leaves adjacent to each internal vertex $u \in V_G$. Note that G^d is d -regular as a tree and it has at most dn vertices. By Proposition 9, we have $\tau(\mathcal{L}_{E_G}) \leq \tau(\mathcal{L}_{E_G^d})$. So, to prove Theorem 1 it suffices to bound the relaxation time of d -regular trees with at most dn vertices, and in the following we assume that G is d -regular.

We will prove the following result by induction on the size of the subtree T :

Claim 27. There is a constant C such that, for every splitting subtree $T = (V, E)$ of G with m edges, and every edge colouring $\mu \in \Omega_{E_G}$, the Glauber dynamics \mathcal{L}_E^μ with parameters $p_i = 1/k$ is ergodic and

$$\tau(\mathcal{L}_E^\mu) \leq d^3 m^C .$$

From this claim, the theorem is obtained immediately by taking $T = G$. If T is composed of a single edge, the Claim 27 follows from Lemma 25. Let m be the number of edges in T and assume $m > 1$. Let $v \in V$ be the vertex such that each subtree T_i hanging from T rooted at v has at most $\lceil \frac{m}{2} \rceil$ edges; this vertex always exists and it is internal. Let e_i be the edge from T_i incident to v . We are going to split T into several subtrees by applying Lemma 26, possibly several times. Note that in order to apply this lemma, we must ensure that each of the subtrees is splitting. This gives constraints on how we can split T . Precisely, while splitting the tree into subtrees, none of the subtrees can have an internal boundary of size at least 3, and for the subtrees with an internal boundary of size 2, the two edges in the boundary must be non-incident.

The splitting procedure is done according to different cases:

1. If all the T_i are splitting, then simply root T at v , and apply Lemma 26. In the following we will assume that not all the T_i are splitting.
2. If $|\bar{\partial}T| = 1$, then there is exactly one subtree, say T_1 , which is not splitting. All other subtrees T_i for $i \neq 1$ are splitting and have fringe boundary of size 1. Root T at e_1 . Now, all subtrees pending from e_1 are splitting and have at most $\lceil \frac{m}{2} \rceil$ edges, and we can apply Lemma 26.
3. If $|\bar{\partial}T| = 2$, with e and f the two edges on the internal boundary, and v is on the path between e and f . Without loss of generality, assume that T_1 and T_2 contain e and f respectively.
 - If exactly one of T_1 or T_2 is not splitting, w.l.o.g. we can assume that it is T_1 . By rooting T at e_1 , all subtrees pending from e_1 are splitting and contain at most $\lceil \frac{m}{2} \rceil$ edges, and we can apply Lemma 26.
 - If both T_1 and T_2 are not splitting, write $e_1 = (v_1, v)$. Since T_1 is not splitting, v_1 must be incident to e . Moreover, since T_2 is not splitting, e_1 and f are not incident, and all the subtrees pending at v_1 are splitting. We apply Lemma 26 a first time by rooting T at v_1 . Let T' be the subtree hanging from v_1 that contains v and root T' at e_2 . Then, all subtrees pending from e_2 are splitting, and we can apply Lemma 26 a second time. The resulting subtrees all have at most $\lceil \frac{m}{2} \rceil$ edges.

4. Finally, if v is not on the path between e and f , let v' be the vertex on this path which is closest to v . We can first split at v' by applying the procedure from the case 3. After this splitting, all the resulting subtrees have at most $\lceil \frac{m}{2} \rceil$ edges except maybe the subtree T' containing v . However T' has a fringe boundary of size 1, and by splitting one more time according to either case 1 or case 2, all resulting subtrees are splitting, and have at most $\lceil \frac{m}{2} \rceil$ edges. In the worst case, we needed to use Lemma 26 three times in this case.

Let T'_1, \dots, T'_s be the subtrees into which T is split by applying the procedure above, and let E'_i be the set of edges of T'_i . Since T'_i is splitting, by the induction hypothesis, $\mathcal{L}_{E'_i}^\sigma$ is ergodic for every $\sigma \in \Omega_E^\mu$, so \mathcal{L}_E^μ is also ergodic. Recall that $\tau_i := \max_{\sigma \in \Omega_E^\mu} \tau(\mathcal{L}_{E'_i}^\sigma)$. Lemma 26 shows that there exists K such that if T is split at a vertex/edge into r subtrees T'_1, \dots, T'_r , then $\tau(\mathcal{L}_E^\mu) \leq K(d^3 + \sum_{i=1}^r \tau_i)$. As we use Lemma 26 at most three times in each step of the splitting procedure described above, we have

$$\tau(\mathcal{L}_E^\mu) \leq K^3 \left(3d^3 + \sum_{i=1}^s \tau_i \right),$$

Using the induction hypothesis on T'_i , if we denote by $m_i \leq \lceil \frac{m}{2} \rceil$ the number of edges in T'_i , we have:

$$\begin{aligned} \tau(\mathcal{L}_E^\mu) &\leq K^3 \left(3d^3 + \sum_{i=1}^s d^3 m_i^C \right) \leq K^3 \left(3d^3 + d^3 \left\lceil \frac{m}{2} \right\rceil^{C-1} \sum_{i=1}^s m_i \right) \\ &\leq K^3 d^3 \left(3 + \frac{m^C}{2^{C-2}} \right) \leq d^3 m^C \cdot \frac{K^3}{2^{C-3}}. \end{aligned}$$

So letting $C \geq 3(1 + \log K)$ gives the desired inequality, and we conclude the proof of Claim 27.

References

- [1] D. Achlioptas and A. Coja-Oghlan. Algorithmic barriers from phase transitions. In *Foundations of Computer Science, 2008. FOCS'08. IEEE 49th Annual IEEE Symposium on*, pages 793–802. IEEE, 2008.
- [2] 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.
- [3] S. Chen and A. Moitra. Linear programming bounds for randomly sampling colorings. *arXiv preprint arXiv:1804.03156*, 2018.
- [4] M. Delcourt, G. Perarnau, and L. Postle. Rapid mixing of Glauber dynamics for colorings below Vigoda’s 11/6 threshold. *arXiv preprint arXiv:1804.04025*, 2018.
- [5] P. Diaconis and L. Saloff-Coste. Comparison theorems for reversible Markov chains. *The Annals of Applied Probability*, pages 696–730, 1993.
- [6] M. Dotan and N. Linial. Efficient generation of random one-factorizations for complete graphs. *arXiv preprint arXiv:1707.00477*, 2017.

- [7] M. Dyer, A. D. Flaxman, A. M. Frieze, and E. Vigoda. Randomly coloring sparse random graphs with fewer colors than the maximum degree. *Random Structures & Algorithms*, 29(4):450–465, 2006.
- [8] M. Dyer, A. Frieze, T. P. Hayes, and E. Vigoda. Randomly coloring constant degree graphs. In *null*, pages 582–589. IEEE, 2004.
- [9] M. Dyer, L. A. Goldberg, and M. Jerrum. Systematic scan for sampling colorings. *The Annals of Applied Probability*, 16(1):185–230, 2006.
- [10] C. Efthymiou, T. P. Hayes, D. Štefankovič, and E. Vigoda. Sampling random colorings of sparse random graphs. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1759–1771. SIAM, 2018.
- [11] L. A. Goldberg, M. Jerrum, and M. Karpinski. The mixing time of Glauber dynamics for coloring regular trees. *Random Structures & Algorithms*, 36(4):464–476, 2010.
- [12] T. P. Hayes and A. Sinclair. A general lower bound for mixing of single-site dynamics on graphs. In *Foundations of Computer Science, 2005. FOCS 2005. 46th Annual IEEE Symposium on*, pages 511–520. IEEE, 2005.
- [13] T. P. Hayes, J. C. Vera, and E. Vigoda. Randomly coloring planar graphs with fewer colors than the maximum degree. *Random Structures & Algorithms*, 47(4):731–759, 2015.
- [14] T. P. Hayes and E. Vigoda. A non-Markovian coupling for randomly sampling colorings. In *Foundations of Computer Science, 2003. Proceedings. 44th Annual IEEE Symposium on*, pages 618–627. IEEE, 2003.
- [15] M. T. Jacobson and P. Matthews. Generating uniformly distributed random Latin squares. *Journal of Combinatorial Designs*, 4(6):405–437, 1996.
- [16] M. Jerrum. A very simple algorithm for estimating the number of k -colorings of a low-degree graph. *Random Structures & Algorithms*, 7(2):157–165, 1995.
- [17] J. Jonasson. Uniqueness of uniform random colorings of regular trees. *Statistics & Probability Letters*, 57(3):243–248, 2002.
- [18] D. A. Levin and Y. Peres. *Markov chains and mixing times*, volume 107. American Mathematical Soc., 2017.
- [19] B. Lucier and M. Molloy. The Glauber dynamics for colorings of bounded degree trees. *SIAM Journal on Discrete Mathematics*, 25(2):827–853, 2011.
- [20] B. Lucier, M. Molloy, and Y. Peres. The Glauber dynamics for colourings of bounded degree trees. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 631–645. Springer, 2009.
- [21] F. Martinelli. Lectures on Glauber dynamics for discrete spin models. In *Lectures on probability theory and statistics*, pages 93–191. Springer, 1999.

- [22] F. Martinelli, A. Sinclair, and D. Weitz. Fast mixing for independent sets, colorings and other model on trees. In *Proceedings of the 15th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 449–458, 2004.
- [23] E. Mossel and A. Sly. Gibbs rapidly samples colorings of $G(n, d/n)$. *Probability theory and related fields*, 148(1-2):37–69, 2010.
- [24] D. Randall and P. Tetali. Analyzing Glauber dynamics by comparison of Markov chains. *Journal of Mathematical Physics*, 41(3):1598–1615, 2000.
- [25] J. Salas and A. D. Sokal. Absence of phase transition for antiferromagnetic Potts models via the Dobrushin uniqueness theorem. *Journal of Statistical Physics*, 86(3-4):551–579, 1997.
- [26] A. Sinclair. Improved bounds for mixing rates of Markov chains and multicommodity flow. *Combinatorics, probability and Computing*, 1(4):351–370, 1992.
- [27] A. Sly. Reconstruction of random colourings. *Communications in Mathematical Physics*, 288(3):943–961, 2009.
- [28] P. Tetali, J. C. Vera, E. Vigoda, and L. Yang. Phase transition for the mixing time of the Glauber dynamics for coloring regular trees. In *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete algorithms*, pages 1646–1656. SIAM, 2010.
- [29] S. Vardi. Randomly coloring graphs of bounded treewidth. *arXiv preprint arXiv:1708.02677*, 2017.
- [30] E. Vigoda. Improved bounds for sampling colorings. *Journal of Mathematical Physics*, 41(3):1555–1569, 2000.

A Missing proofs from Section 2.1

For any function $f : \Omega \rightarrow \mathbb{R}$, the *variance* and the *Dirichlet form* of \mathcal{L} are defined respectively as

$$\begin{aligned}\mathrm{Var}_{\mathcal{L}}(f) &= \frac{1}{2} \sum_{\alpha, \beta \in \Omega} \pi(\alpha)\pi(\beta)(f(\alpha) - f(\beta))^2, \\ \xi_{\mathcal{L}}(f, f) &= \frac{1}{2} \sum_{\alpha, \beta \in \Omega} \pi(\alpha)\mathcal{L}[\alpha \rightarrow \beta](f(\alpha) - f(\beta))^2.\end{aligned}$$

Let $\mathcal{F}_{\mathcal{L}} = \{f : \mathrm{Var}_{\mathcal{L}}(f) > 0\}$ and note that $\mathcal{F}_{\mathcal{L}} = \mathcal{F}_{\mathcal{L}'}$ are the set of non-constant functions, as π and π' are positive on Ω . It is well-known (see e.g. Remark 13.13 in [18]) that the spectral gap of \mathcal{L} satisfies

$$\mathrm{Gap}(\mathcal{L}) = \min_{f \in \mathcal{F}_{\mathcal{L}}} \frac{\xi_{\mathcal{L}}(f, f)}{\mathrm{Var}_{\mathcal{L}}(f)}. \quad (16)$$

A.1 Proof of Propostion 2

We will compare the Dirichlet form and the variance of \mathcal{L} and \mathcal{L}' . We have

$$\begin{aligned}\xi_{\mathcal{L}'}(f, f) &= \frac{1}{2} \sum_{\alpha, \beta \in \Omega} \pi'(\alpha)\mathcal{L}'[\alpha \rightarrow \beta](f(\alpha) - f(\beta))^2 \\ &= \frac{1}{2} \sum_{\alpha, \beta \in \Omega} \sum_{\gamma \in \Gamma_{\alpha, \beta}} g(\gamma)\pi'(\alpha)\mathcal{L}'[\alpha \rightarrow \beta](f(\alpha) - f(\beta))^2 \\ &= \frac{1}{2} \sum_{\alpha, \beta \in \Omega} \sum_{\gamma \in \Gamma_{\alpha, \beta}} g(\gamma)\pi'(\alpha)\mathcal{L}'[\alpha \rightarrow \beta] \left(\sum_{(\sigma, \eta) \in \gamma} \sqrt{\omega(\sigma, \eta)} \frac{f(\sigma) - f(\eta)}{\sqrt{\omega(\sigma, \eta)}} \right)^2 \\ &\leq \frac{1}{2} \sum_{\alpha, \beta \in \Omega} \sum_{\gamma \in \Gamma_{\alpha, \beta}} g(\gamma)\pi'(\alpha)\mathcal{L}'[\alpha \rightarrow \beta] \cdot |\gamma|_{\omega} \sum_{(\sigma, \eta) \in \gamma} \frac{(f(\sigma) - f(\eta))^2}{\omega(\sigma, \eta)} \\ &= \frac{1}{2} \sum_{(\sigma, \eta) \in \mathcal{L}} (f(\sigma) - f(\eta))^2 \cdot \frac{1}{\omega(\sigma, \eta)} \sum_{(\alpha, \beta) \in \mathcal{L}'} \sum_{\substack{\gamma \in \Gamma_{\alpha, \beta} \\ \gamma \ni (\sigma, \eta)}} g(\gamma)|\gamma|_{\omega}\pi'(\alpha)\mathcal{L}'[\alpha \rightarrow \beta] \\ &\leq \rho_{\max} \xi_{\mathcal{L}}(f, f),\end{aligned}$$

where we used the Cauchy-Schwartz inequality in the first inequality. Additionally, we have

$$\begin{aligned}\mathrm{Var}_{\mathcal{L}'}(f) &= \frac{1}{2} \sum_{\alpha, \beta \in \Omega} \pi'(\alpha)\pi'(\beta)(f(\alpha) - f(\beta))^2 \\ &\geq \frac{1}{2b^2} \sum_{\alpha, \beta \in \Omega} \pi(\alpha)\pi(\beta)(f(\alpha) - f(\beta))^2 \\ &= \frac{1}{b^2} \mathrm{Var}_{\mathcal{L}}(f).\end{aligned}$$

Combining the previous two inequalities and using (16), the desired result follows,

$$\mathrm{Gap}(\mathcal{L}') = \min_{f \in \mathcal{F}_{\mathcal{L}'}} \frac{\xi_{\mathcal{L}'}(f, f)}{\mathrm{Var}_{\mathcal{L}'}(f)} \leq b^2 \rho_{\max} \min_{f \in \mathcal{F}_{\mathcal{L}}} \frac{\xi_{\mathcal{L}}(f, f)}{\mathrm{Var}_{\mathcal{L}}(f)} = b^2 \rho_{\max} \mathrm{Gap}(\mathcal{L}).$$

A.2 Proof of Proposition 9

Let $U = \{u_1, \dots, u_m\} = V \setminus \{v\}$. Let p_1, \dots, p_m, p_v be the parameters for u_1, \dots, u_m, v , respectively. Let Ω_V and Ω_U be the set of L -colourings of G and $G[U]$ respectively. Let $d = |N(v)|$ the degree of v . Since $N(v)$ is a clique, we have:

$$|\Omega_V| = |\Omega_U|(k - d).$$

Indeed, for every colouring α_U of U , there are exactly $(k - d)$ possibilities to extend it into a colouring of V . Let π_V and π_U be the stationary distributions of \mathcal{L}_V and \mathcal{L}_U , which are uniform as the transitions are symmetric. Recall that $\mathcal{F}_{\mathcal{L}_V}$ is the set of non-constant functions from Ω_V to \mathbb{R} . Let $\mathcal{F}_{\mathcal{L}_V}^v$ be the subset of these functions which are independent of v , i.e. which satisfy $f(\alpha) = f(\beta)$ whenever α and β agree on U .

Using (16), we have

$$\text{Gap}(\mathcal{L}_V) = \min_{f \in \mathcal{F}_{\mathcal{L}_V}} \frac{\mathcal{E}_{\mathcal{L}_V}(f)}{\text{Var}_{\mathcal{L}_V}(f)} \leq \min_{f \in \mathcal{F}_{\mathcal{L}_V}^v} \frac{\mathcal{E}_{\mathcal{L}_V}(f)}{\text{Var}_{\mathcal{L}_V}(f)}. \quad (17)$$

Let $f \in \mathcal{F}_V$, then we have the following

$$\begin{aligned} \text{Var}_{\mathcal{L}_V}(f) &= \frac{1}{2} \sum_{\alpha, \beta \in \Omega_V} \pi_V(\alpha) \pi_V(\beta) (f(\alpha) - f(\beta))^2 \\ &= \frac{1}{2} \sum_{\alpha, \beta \in \Omega_V} \frac{1}{|\Omega_V|^2} (f(\alpha) - f(\beta))^2 \\ &= \frac{1}{2} \sum_{\alpha_U, \beta_U \in \Omega_U} \sum_{\substack{c \in [k] \\ c \notin \alpha(N(v))}} \sum_{\substack{c' \in [k] \\ c' \notin \beta(N(v))}} \frac{1}{|\Omega_V|^2} (f(\alpha_U) - f(\beta_U))^2 \\ &= \frac{1}{2} \sum_{\alpha_U, \beta_U \in \Omega_U} \frac{(k - d)^2}{|\Omega_V|^2} (f(\alpha_U) - f(\beta_U))^2 \\ &= \text{Var}_{\mathcal{L}_U}(f) \end{aligned} \quad (18)$$

Additionally, if $f \in \mathcal{F}_V^v$, then we have

$$\begin{aligned} \mathcal{E}_{\mathcal{L}_V}(f) &= \frac{1}{2} \sum_{u_i \in U} \sum_{\substack{\alpha, \beta \in \Omega_V \\ \text{differ at } u_i}} \frac{1}{|\Omega_V|} p_i (f(\alpha) - f(\beta))^2 + \sum_{\substack{\alpha, \beta \in \Omega_V \\ \text{differ at } v}} \frac{1}{|\Omega_V|} p_v (f(\alpha) - f(\beta))^2 \\ &= \frac{1}{2} \sum_{u_i \in U} \sum_{\substack{\alpha_U, \beta_U \in \Omega_U \\ \text{differ at } u_i}} \sum_{\substack{c \in [k] \\ c \notin \alpha(N(v)) \\ c \notin \beta(N(v))}} \frac{1}{|\Omega_V|} p_i (f(\alpha_U) - f(\beta_U))^2 \\ &\leq \frac{1}{2} \sum_{u_i \in U} \sum_{\substack{\alpha_U, \beta_U \in \Omega_U \\ \text{differ at } u_i}} \frac{k - d}{|\Omega_V|} p_i (f(\alpha_U) - f(\beta_U))^2 \\ &= \mathcal{E}_{\mathcal{L}_U}(f). \end{aligned} \quad (19)$$

where we used that $N(v)$ is a clique in the inequality. Putting together (17)–(19), gives as required

$$\text{Gap}(\mathcal{L}_V) \leq \text{Gap}(\mathcal{L}_U).$$