The Mixing Time of Glauber Dynamics for Colouring Regular Trees[*] (Revised Version)

Leslie Ann Goldberg[†]   Mark Jerrum[‡]   Marek Karpinski[§]

Abstract

We consider Metropolis Glauber dynamics for sampling proper q-colourings of the n-vertex complete b-ary tree when $3 \le q \le b/(2 \ln b)$. We give both upper and lower bounds on the mixing time. Our upper bound is $n^{O(b/\log b)}$ and our lower bound is $n^{\Omega(b/(q \log b))}$, where the constants implicit in the $O()$ and $\Omega()$ notation do not depend upon n, q or b.

1 Introduction

This paper proves both upper and lower bounds on the mixing time of Glauber dynamics for colourings of regular trees. Specifically, we consider Metropolis Glauber dynamics for sampling proper q-colourings of the n-vertex complete b-ary tree when $3 \le q \le b/(2 \ln b)$. We give both upper and lower bounds on the mixing time, providing reasonably precise information on the dependence of the mixing time on n, b and q. For given q and b, our upper bound is $n^{O(b/\log b)}$ and our lower bound is $n^{\Omega(b/(q \log b))}$, where the constants implicit in the $O()$ and $\Omega()$ notation do not depend upon n, q or b. Since $b/(q \log b)$ and $b/\log b$ are both bounded away from zero, one could equivalently say that the mixing time is sandwiched between $n^{c_1 b/(q \ln b)}$ and $n^{c_2 b/\ln b}$, where $c_1, c_2 > 0$ are constants.

At the time of publication, the results presented here have been sharpened in two stages. Lucier and Molloy, and the same authors with Peres [12], have established the asymptotically correct exponent in the mixing time. (It is our lower bound that is tight.) They also generalised the setting for the upper bounds to non-regular trees. Very recently, Tetali, Vera, Vigoda and Yang [20] have pinned down the constant in the exponent and located the phase transition between "nearly linear" (i.e, $n^{1+o(b)}$) and superlinear mixing time.

[†]Department of Computer Science, University of Liverpool, Liverpool L69 3BX, UK.

[‡]School of Mathematical Sciences, Queen Mary, University of London, Mile End Road, London E1 4NS, UK.

[§]Department of Computer Science and the Hausdorff Center for Mathematics, University of Bonn, 53117 Bonn.

2 Previous work

There has been quite a bit of work on Markov chains for sampling the proper q-colourings of an input graph. Much of this work focusses on *Glauber dynamics*, which is a general term for a Markov chain which updates the colour of one vertex at a time. Proper colourings correspond to configurations in the *zero-temperature Potts model* from statistical physics, and there is a close connection between the mixing time of Glauber dynamics and the qualitative properties of the model. In particular, rapid mixing, specifically $O(n \log n)$ mixing for an n-vertex sub-graph of an infinite graph, often coincides with the case in which the model has a unique infinite-volume Gibbs measure on the infinite graph. See Weitz's PhD thesis [21] and Martinelli's lecture notes [13] for an exposition of this material.

Martinell, Sinclair and Weitz [14] consider Glauber dynamics on the complete n-vertex tree with branching factor b. They show that for $q \geq b+3$, Glauber dynamics for sampling proper q-colourings mixes in $O(n \log n)$ time for arbitrary boundary conditions. This result is optimal in the sense that for $q \leq b+2$ there are boundary conditions for which Glauber dynamics is not even ergodic.

It is also interesting to determine whether Glauber dynamics is rapidly mixing for smaller q in the absence of boundary conditions. Berger, Kenyon, Mossel and Peres [2, Prop. 2.5] prove that the mixing time of Glauber dynamics for quite general spin systems on complete trees is polynomial. Their result includes the case of Glauber dynamics on proper q-colourings, for any $q \geq 3$. Our upper bound applies specifically to q-colourings, but has the advantage of providing an explicit expression for the exponent of the mixing time in terms of b.

Perhaps the main contribution of this paper lies in the lower bound on mixing time. Bhatnagar, Vera and Vigoda [3] consider *reconstruction* on the complete tree with branching factor b. They show that for $C = 2$ and $q > C(b+1)/\ln(b+1)$ non-reconstruction holds, meaning that, over random colourings of the leaves, the expected influence on the root is vanishing. It is known [15] that the expected influence is non-vanishing for a sufficiently large q satisfying $q \leq (1 - \varepsilon)(b + 1)/\ln(b + 1)$ for some $\varepsilon > 0$. This non-vanishing influence implies [2, 3] that the mixing time $\tau(\delta)$ of Glauber dynamics cannot be $O(n(\log n + \log \delta^{-1}))$. Here we provide a much stronger polynomial bound that holds when q is just a modest constant factor smaller than this. Again, the exponent is given explicitly in terms of b and q.

It is worth mentioning that the possibility of polynomial mixing time in the regime $q < b + 1$ (interpreting $b + 1$ as vertex degree) has been studied in situations other than complete trees. Goldberg, Martin and Paterson [9], and Luby, Randall and Sinclair [11] have treated 3-colourings of finite subregions of the 2-dimensional integer lattice. Hayes Vera and Vigoda [10], and Dyer, Flaxman, Frieze and Vigoda [6] have investigated q-colourings in planar graphs and random graphs, respectively.

3 Proof techniques

The upper bound argument is based on canonical paths. The lower bound argument is based on conductance. Essentially, the argument is that it takes a while to move from a colouring in which the colour of the root is forced to be one colour by the induced colouring on the leaves to a colouring in which the colour of the root is forced to be another colour. This is similar to the *recursive majority* idea [2] used to prove a lower bound for the Ising model.

4 The problem

Fix $b \geq 2$ and $q \geq 3$. Let $[q] = \{0, \ldots, q-1\}$. Suppose T is a complete b-ary tree of height H — meaning that there are H edges on a path from the root $r(T)$ to any leaf. Let V be the set of vertices of T and $n = |V|$. Let L be the set of leaves of T. Note that

$$n = \frac{b^{H+1} - 1}{b - 1}$$

so

$$H = \frac{\log((b-1)n + 1)}{\log b} - 1. \tag{1}$$

The *height* $h(v)$ of a vertex $v \in V$ is the number of edges on a path from v down to a leaf. So a vertex $v \in L$ has $h(v) = 0$ and $h(r(T)) = H$.

For any vertex v of T, T_v denotes the subtree of T rooted at v. For any subtree T_v, let $V(T_v)$ be the set of vertices of T_v and let $L(T_v)$ be the set of leaves. A *proper q-colouring* of T_v is a labelling of the vertices with elements of $[q]$ such that neighbouring vertices receive different colours. Let $\Omega(T_v)$ be the set of proper q-colourings of T_v and $\Omega = \Omega(T_{r(T)})$ be the set of proper q-colourings of T.

For a colouring $x \in \Omega$, let $x(T_v)$ denote the restriction of x to the vertices in the subtree T_v. Similarly, for a set $U \subseteq V(T_v)$ and a colouring $x \in \Omega(T_v)$, $x(U)$ denotes the restriction of x to U.

Let $\mathcal{M}$ be the Metropolis Glauber dynamics for sampling from Ω. To move from one colouring to another, this chain selects a vertex v and a colour c uniformly at random. The vertex v is re-coloured with c if and only if this results in a proper colouring. If $q \geq 3$ then the set of proper colourings is connected and $\mathcal{M}$ converges to the uniform distribution on Ω, which we call π. The goal is to study the mixing time of $\mathcal{M}$ as a function of n, b and q. Let P be the transition matrix of $\mathcal{M}$. The *variation distance* between distributions θ_1 and θ_2 on Ω is

$$||\theta_1 - \theta_2|| = \frac{1}{2} \sum_i |\theta_1(i) - \theta_2(i)| = \max_{A \subseteq \Omega} |\theta_1(A) - \theta_2(A)|.$$

For a state $x \in \Omega$, the mixing time of $\mathcal{M}$ from starting state x is

$$\tau_x(\mathcal{M}, \delta) = \min \left\{ t > 0 : ||P^{t'}(x, \cdot) - \pi(\cdot)|| \leq \delta \text{ for all } t' \geq t \right\}.$$

The mixing time of $\mathcal{M}$ is given by

$$\tau(\mathcal{M}, \delta) = \max_x \tau_x(\mathcal{M}, \delta).$$

Our results are as follows, where lg denotes the base-2 logarithm and ln denotes the natural logarithm.

Theorem 1. *Suppose $q \geq 3$. Let $\mathcal{M}$ be the Metropolis Glauber dynamics for sampling proper q-colourings of the n-vertex complete b-ary tree. Then for fixed q and b the mixing time $\tau(\mathcal{M}, 1/(2e))$ is $n^{O(b/\log b)}$ where the constant implicit in the $O()$ notation does not depend upon n, q or b. In particular,*

$$\tau(\mathcal{M}, 1/(2e)) \leq 3bq^2(1 + \lg(n))n^{3+3b/\ln b}.$$

Theorem 2. *Suppose $3 \leq q \leq b/(2\ln b)$. Let $\mathcal{M}$ be the Metropolis Glauber dynamics for sampling proper q-colourings of the n-vertex complete b-ary tree. Then for fixed q and b the mixing time $\tau(\mathcal{M}, 1/(2e))$ is $n^{\Omega(b/(q\log b))}$ where the constant implicit in the $\Omega()$ notation does not depend upon n, q or b. In particular,*

$$\tau(\mathcal{M}, 1/(2e)) \geq \left(\frac{1}{2} - \frac{1}{2e}\right)\frac{2}{9}n^{\frac{b-2}{6(q-1)\ln b}}.$$

5   Bounds on H

The calculations arising in the derivation of Theorems 1 and 2 involve H. It is clear from Equation (1) that $H = \Theta(\log(n)/\log b)$. Since we give explicit bounds in the statement of the theorems, we also require upper and lower bounds on H. We record these here. Note that the bounds can be improved, but we prefer to avoid the complication.

Lemma 3. $H + 1 \leq \lg(n) + 1$ *and* $H \leq \ln(n)/\ln b$. *If* $n \geq b^3$ *then* $H - 1 \geq \ln(n)/(3\ln b)$.

Proof. For the first upper bound, use Equation (1) to see that

$$H + 1 = \log_b((b-1)n + 1) \leq \log_b(bn) = 1 + \log_b n \leq 1 + \lg n,$$

since $n \geq 1$ and $b \geq 2$. For the second upper bound, note that

$$H = \frac{\ln\left(n\left(b - 1 + 1/n\right)\right)}{\ln b} - 1 = \frac{\ln n}{\ln b} - \frac{\ln b - \ln(b - 1 + 1/n)}{\ln b} \leq \frac{\ln n}{\ln b}.$$

Finally, for the lower bound, note that

$$H - 1 = \frac{\ln\left(n\left(b - 1 + 1/n\right)\right)}{\ln b} - 2 = \frac{\ln n}{\ln b} + \frac{\ln\left(b - 1 + 1/n\right)}{\ln b} - 2.$$

Dropping the non-negative middle term, this is at least $\ln(n)/\ln(b) - 2$, which gives the result since $\ln(n)/(3\ln b) \geq 1$. $\square$

6 The upper bound

In this section we prove Theorem 1. We will use the canonical paths method of Jerrum and Sinclair [19]. Let $\mathcal{M}'$ be the trivial Markov chain on Ω that moves from a state x to a new state y by selecting y u.a.r. from Ω. Let P' be the transition matrix of $\mathcal{M}'$. Clearly, for any $\delta' > 0$, $\tau(\mathcal{M}', \delta') = 1$. we will definine canonical paths between pairs of colourings in Ω. These canonical paths will constitute what is called an $(\mathcal{M}, \mathcal{M}')$-flow. Then Theorem 1 follows from the following proposition (which is Observation 13 in the expository paper [8]) taking $A(f)$ to be the *congestion* of the flow and c to be $1/q$. The proof of Proposition 4 combines Diaconis and Saloff Coste's comparison method [4] with upper and lower bounds on mixing time [1, 5, 18] along lines first proposed by Randall and Tetali [17]. See [8] for details.

Proposition 4. *Suppose that $\mathcal{M}$ is a reversible ergodic Markov chain with transition matrix P and stationary distribution π and that $\mathcal{M}'$ is another reversible ergodic Markov chain with the same stationary distribution. Suppose that f is a $(\mathcal{M}, \mathcal{M}')$-flow. Let $c = \min_x P(x, x)$, and assume $c > 0$. Then, for any $0 < \delta' < 1/2$,*

$$\tau_x(\mathcal{M}, \delta) \leq \max\left\{ A(f) \left[\frac{\tau(\mathcal{M}', \delta')}{\ln(1/2\delta')} + 1 \right], \frac{1}{2c} \right\} \ln \frac{1}{\delta\pi(x)}.$$

$\square$

For each pair of distinct colourings $x, y \in \Omega$ we will construct a path $\gamma_{x,y}$ from x to y using transitions of $\mathcal{M}$. This gives an $(\mathcal{M}, \mathcal{M}')$-flow f with congestion

$$\begin{aligned}
A(f) &= \max_{z,w} \frac{1}{\pi(z)P(z,w)} \sum_{x,y:(z,w)\in\gamma_{x,y}} |\gamma_{x,y}| \pi(x) P'(x, y) \\
&= \frac{nq}{|\Omega|} \max_{z,w} \sum_{x,y:(z,w)\in\gamma_{x,y}} |\gamma_{x,y}|,
\end{aligned} \tag{2}$$

where the maximum is over pairs of distinct states z and w in Ω with $P(z, w) > 0$ (hence, $P(z, w) = 1/nq$) and $|\gamma_{x,y}|$ denotes the length of $\gamma_{x,y}$, which is the number of transitions on the path. We will prove the following lemma below.

Lemma 5. *The canonical paths correspond to an $(\mathcal{M}, \mathcal{M}')$-flow f with $A(f) \leq bq(H + 1)n^2 9^{bH}$.*

Theorem 1 follows. Combining Proposition 4 with $\delta' = 1/2e^2$ and Lemma 5, we get

$$\tau_x(\mathcal{M}, \delta) \leq bq(H + 1)n^2 9^{bH} \left(\frac{1}{2} + 1 \right) \ln(|\Omega|/\delta).$$

Since $|\Omega| \leq q^n$,

$$\begin{aligned}
\tau_x(\mathcal{M}, 1/(2e)) &\leq bq(H + 1)n^2 9^{bH} \left(\frac{1}{2} + 1 \right) \ln(2eq^n) \\
&\leq (H + 1)bqn^2 \frac{3}{2}(2 + n\ln(q))9^{bH} \\
&\leq (H + 1)bq^2 n^3 3 e^{3bH}.
\end{aligned}$$

Theorem 1 then follows by applying the two upper bounds in Lemma 3.

Proof of Lemma 5

Defining the canonical paths: a special case

We start by defining paths between colourings x and y for the special case in which, for all $v \in V$, $y(v) = x(v) + 1 \pmod q$. The sequence of colourings on the path is defined to be the sequence of colourings visited by procedure $Cycle^+$ below when it is called with the input T, which is initially coloured x.

Here is the description of procedure $Cycle^+(\widehat{T})$, where $\hat{x}$ is a global variable, representing the current colouring of tree T, and the input parameter $\widehat{T}$ may be any of the subtrees T_v.

1. Let $\widehat{T}_1, \ldots, \widehat{T}_b$ be the subtrees rooted at the children of $r(\widehat{T})$ and let $S = \big\{i : x(r(\widehat{T}_i)) + 1 \neq x(r(\widehat{T})) \pmod q\big\}$,

2. For each $i \in S$ do $Cycle^+(\widehat{T}_i)$.

3. Recolour the root $r(\widehat{T})$ so that $\hat{x}(r(\widehat{T})) = x(r(\widehat{T})) + 1 \pmod q$.

4. For each $i \notin S$ do $Cycle^+(\widehat{T}_i)$.

Since $q \geq 3$, we are guaranteed that $x(r(\widehat{T})) + 1 \neq \hat{x}(r(\widehat{T}_i)) \pmod q$, for all i, after line 2; this ensures that the root can be recoloured in line 3.

The procedure $Cycle^+$ is useful in the construction of more general canonical paths. When recolouring a non-leaf vertex v it will generally be necessary to change the colours of some of the children of v to avoid clashes. Applying the procedure $Cycle^+$ to the subtrees beneath v provides a systematic means to achieve this.

Analysis of the special case

Suppose we observe a transition at some point during the execution of the procedure $Cycle^+(\widehat{T})$, in which the colouring $\hat{x}$ is transformed by adding 1 to the colour of some vertex v (modulo q). How many initial colourings $x(\widehat{T})$ (and hence how many final colourings $y(\widehat{T})$) are consistent with this observed transition?

We will let $s(h)$ denote the maximum number of consistent initial colourings $x(\widehat{T})$, maximised over all trees $\widehat{T}$ of height h and over all possible transitions. We will compute an upper bound on $s(h)$.

Case 1: Suppose that $v = r(\widehat{T})$.

The subtrees $\widehat{T}_i$ with $i \in S$ have already been processed by the time that the transition takes place, so $\hat{x}(T_i) = y(T_i)$ for these trees. The subtrees with $i \notin S$ are yet to be processed, so for these trees we have $\hat{x}(T_i) = x(T_i)$. However, we do not know the set S from observing the transition from $\hat{x}$. Thus, as many as 2^b initial colourings $x(\widehat{T})$ may be consistent with the observed transition from $\hat{x}$.

Case 2: Otherwise, v is in one of the subtrees $\widehat{T}_k$ rooted at one of the children of $r(\widehat{T})$. Then, by the argument of Case 1, there are two choices for the initial colouring $x(T_i)$ of every subtree with $i \neq k$; also there are two possibilities for $x(r(\widehat{T}))$, since we don't know whether line (3) has been executed at the point of the transition. Then $s(h)$ satisfies the recurrence $s(h) \leq \max\{2^b, 2^b s(h-1)\}$ with initial condition $s(0) = 1$. Solving the recurrence, we discover that at most

$$s(h) \leq 2^{bh} \tag{3}$$

initial colourings $x(\widehat{T})$ are consistent with the observed transition, so there are at most $s(H) \leq 2^{bH}$ initial colourings x of T consistent with an observed transition of the procedure $Cycle^+(T)$

Defining the canonical paths: the general case

Let $Cycle^-$ be defined analogously to $Cycle^+$ but implementing the permutation of colours that subtracts 1 (modulo q) from every colour; that is, $y(v) = x(v) - 1 \pmod{q}$ for all $v \in V$.

Let $F \subset [q]$ be a set of "forbidden colours" of size at most two. Given $Cycle^+$ and $Cycle^-$ it is easy to implement a procedure $Cycle(\widehat{T}, F)$ that systematically recolours the tree $\widehat{T}$ so that the new colour assigned to $r(\widehat{T})$ avoids the forbidden colours F: simply apply $Cycle^+$ or $Cycle^-$ or neither in order to bring a colour not in F to the root of $\widehat{T}$. If we observe a transition during the execution of $Cycle(\widehat{T}, F)$ we can tell whether it comes from $Cycle^+$ or from $Cycle^-$.

The recursive procedure $Recolour$, to be described presently, provides a systematic approach to transforming an arbitrary initial colouring x to an arbitrary final colouring y using single-vertex updates. In doing so, it defines canonical paths between arbitrary pairs of proper colourings x and y of T. The sequence of colourings on the path $\gamma_{x,y}$ is defined to be the sequence of colourings visited by procedure $Recolour$ when it is called with the input T (which is initially coloured x) and with colouring y.

Like $Cycle^+$, the procedure $Recolour$ takes an argument $\widehat{T}$, which is the tree which will be recoloured from $x(\widehat{T})$ to $y(\widehat{T})$. It also takes the argument y. As before, $\hat{x}$ is a global variable representing the current colouring of the tree T, which is initially coloured x. Here is the description of procedure $Recolour(\widehat{T}, y)$.

1. Let $\widehat{T}_1, \ldots, \widehat{T}_b$ be the subtrees rooted at the children of $r(\widehat{T})$.

2. For each i, $1 \leq i \leq b$, do $Cycle(\widehat{T}_i, \{x(r(\widehat{T})), y(r(\widehat{T}))\})$. (This step permutes the colours in a subtree, to allow the root to be recoloured in the following step.)

3. Assign the root $r(\widehat{T})$ its final colour $y(r(\widehat{T}))$.

4. For each i, $1 \leq i \leq b$, do $Recolour(\widehat{T}_i, y)$.

Analysis of the canonical paths

Suppose we observe a transition at some point during the execution of a procedure call $Recolour(\widehat{T}, y)$ when $\widehat{T}$ has height h. Let $P(h)$ be an upper bound on the number of pairs $(x(\widehat{T}), y(\widehat{T}))$ consistent with this transition, maximised over all trees $\widehat{T}$ of height h and over all possible transitions. Let

$$C(h) = q(q-1)^{(b^{h+1}-1)/(b-1)-1}$$

be the number of proper colourings of a b-ary tree of height h. Note that $P(H)$ is an upper bound on the number of canonical paths $\gamma_{x,y}$ using a given transition. In order to compute the congestion $A(f)$ using Equation (2), we need to compute an upper bound on $P(H)$. We will compute an upper bound on $P(h)$ by induction on h. The base case is $P(0) = 1$.

Now suppose $h > 0$. Suppose that the transition starts at a colouring $\hat{x}$ and changes the colour of vertex v from $\hat{x}(v)$ to a new colour.

Case 1: First, suppose $v = r(\widehat{T})$. We start by bounding the number of colourings $x(\widehat{T})$ that are consistent with the transition. ¿From the transition, we know the initial colour of the root, $x(r(\widehat{T}))$. For each subtree $\widehat{T}_i$, we know that the initial colouring $x(\widehat{T}_i)$ can be obtained by permuting the colours in $\hat{x}(\widehat{T}_i)$. There are three possible permutations (corresponding to adding $-1, 0$ or 1 modulo q). So the number of possibilities for $x(\widehat{T})$ is at most 3^b. Next we bound the number of consistent colourings $y(\widehat{T})$. The colour $y(r(\widehat{T}))$ is fixed by the transition, but we know nothing about the colourings of the subtrees $\widehat{T}_i$ beyond the fact that they must be consistent with the root being coloured $y(r(\widehat{T}))$. Thus there are at most $((q-1)C(h-1)/q)^b$ possibilities for $y(\widehat{T})$. Overall, we have the upper bound

$$P(h) \le (3(q-1)C(h-1)/q)^b \tag{4}$$

in the case $v = r(T)$.

Case 2: Now suppose v is contained in one of the subtrees $\widehat{T}_k$. It could be that the transition under consideration is employed during Step 2 of $Recolour$ (we call such a transition Type A), or in Step 4 (Type B).

Case 2A: Consider first pairs of Type A. How many pairs $(x(\widehat{T}), y(\widehat{T}))$ of initial and final colourings may use the transition? We'll bound this number by considering separately the pairs $(x(r(\widehat{T})), y(r(\widehat{T})))$ and $(x(\widehat{T}_i), y(\widehat{T}_i))$ and multiplying the results. For the root, $x(r(\widehat{T})) = \hat{x}(r(\widehat{T}))$, while there are q possibilities for $y(r(\widehat{T}))$. For $i < k$, there are at most three possibilities for the colouring $x(\widehat{T}_i)$, and at most $C(h-1)$ for $y(\widehat{T}_i)$. For $i > k$, $x(\widehat{T}_i)$ is fixed by the transition, while there are at most $C(h-1)$ possibilities for $y(\widehat{T}_i)$. Now consider the possibilities for $x(\widehat{T}_k)$ and $y(\widehat{T}_k)$, starting with $x(\widehat{T}_k)$. Given the transition from $\hat{x}(v)$ to its new colour we can tell whether the instance of $Cycle(\widehat{T}_k, \{x(r(\widehat{T})), y(r(\widehat{T}))\})$ is applying $Cycle^+$ to $\widehat{T}_k$ or $Cycle^-$ to $\widehat{T}_k$. In either case, (3) guarantees that the number of initial colourings $x(\widehat{T}_k)$ that are consistent with the transition is at most $2^{b(h-1)}$. Since the number of possibilities for $y(\widehat{T}_k)$ is at most $C(h-1)$, the number for the pair $(x(\widehat{T}_k), y(\widehat{T}_k))$ is bounded by $2^{b(h-1)}C(h-1)$. This gives

an upper bound of $3^b q (2^{h-1} C(h-1))^b$ on the total number of pairs $(x(\widehat{T}), y(\widehat{T}))$ such that the given transition is a Type A transition.

Case 2B: Finally, consider pairs of Type B. For the root, $x(r(\widehat{T}))$ is arbitrary, while $y(r(\widehat{T})) = \hat{x}(r(\widehat{T}))$, so there are q possibilities in all. For $i < k$, there are at most $C(h-1)$ possibilities for the colouring $x(\widehat{T}_i)$, while $y(\widehat{T}_i)$ is fixed. For $i > k$, there are three possibilities for $x(\widehat{T}_i)$, while there are at most $C(h-1)$ possibilities for $y(\widehat{T}_i)$. Inductively, the number of possibilities for the pair $(x(\widehat{T}_k), y(\widehat{T}_k))$ is $P(h-1)$. This gives an upper bound of $3^b q C(h-1)^{b-1} P(h-1)$ on the total number of pairs $(x(\widehat{T}), y(\widehat{T}))$ such that the given transition is a Type B transition.

Completing Case 2: Summing the bounds on the number of pairs $(x(\widehat{T}), y(\widehat{T}))$ such that the given transition is a Type A or Type B transition we obtain an upper bound of

$$P(h) \le 3^b q \, C(h-1)^{b-1} \big[2^{(h-1)b} C(h-1) + P(h-1) \big] \tag{5}$$

on the total number of canonical paths using a given transition in the case $v \ne r(\widehat{T})$. Notice that (5) always dominates (4) since $h \ge 1$. Now let $\chi(h) = P(h)/C(h)$. Since $q^{b-1} C(h) = (q-1)^b C(h-1)^b$, we have the recurrence:

$$\chi(h) \le \left(\frac{3q}{q-1} \right)^b \big[2^{(h-1)b} + \chi(h-1) \big], \tag{6}$$

with initial condition $\chi(0) = q^{-1}$. Now note that the recurrence (6) satisfies $\chi(h) \le 9^{bh}$.

Completing the Analysis: Let $\lambda(h)$ be an upper bound on the number of updates performed by $Recolour(\widehat{T}, y)$ when $\widehat{T}$ has height h. Thus, $\lambda(H)$ is an upper bound on the length of a canonical path $\gamma_{x,y}$.

Now, by Equation (2),

$$A(f) = \frac{nq}{|\Omega|} \max_{z,w} \sum_{x,y:(z,w)\in\gamma_{x,y}} |\gamma_{x,y}| \le \lambda(H) \frac{nq}{|\Omega|} P(H) = \lambda(H) \, n \, q \, \chi(H),$$

so to prove Lemma 5 we need an upper bound on $\chi(h)$.

The subroutine *Cycle* creates paths of length $(b^{h+1} - 1)/(b-1)$. The recurrence governing $\lambda(h)$ is thus $\lambda(h) = (b^{h+1} - 1)/(b-1) + b\lambda(h-1)$, with initial condition $\lambda(0) = 1$. Note that $\lambda(h) \le (h+1)b^{h+1}$. This can be verified by induction on h. For the inductive step,

$$\lambda(h) = \sum_{j=0}^{h} b^j + b\lambda(h-1) \le \sum_{j=0}^{h} b^j + bhb^h,$$

which is at most $(h+1)b^{h+1}$ since $\sum_{j=0}^{h} b^j \le b^{h+1}$ for $b \ge 2$. Thus $\lambda(H) \le (H+1)b^{H+1} \le b(H+1)n$. Putting it all together, the congestion $A(f)$ is bounded above by $qn\chi(H)\lambda(H)$ which proves Lemma 5.

9

7   The lower bound

Suppose
$$2q \leq b/\ln b. \tag{7}$$

The lower bound proof will use the following fact.

Lemma 6. *If $q \geq 3$ and $2q \leq b/\ln b$ then $b - 2 \geq 2(q-1)\ln(q-1)$.*

Proof. By (7), $q - 1 \leq q \leq b/(2\ln b)$ so

$$
\begin{aligned}
2(q-1)\ln(q-1) &\leq \frac{b}{\ln b} \ln\left(\frac{b}{2\ln b}\right) \\
&= \frac{b}{\ln b}\big(\ln(b) - \ln(2\ln b)\big) \\
&= b - \frac{b\ln(2\ln b)}{\ln b} \\
&\leq b - 2,
\end{aligned}
$$

where the final inequality holds since $q \geq 3$ so $b \geq 6$ so $b \geq 2\ln(b)/\ln(2\ln b)$. $\qquad\square$

Given a colouring $x \in \Omega$, define

$$F(x) = \{w \in V \mid \forall y \in \Omega(T_w) \text{ with } y(L(T_w)) = x(L(T_w)) \text{ we have } y(w) = x(w)\}.$$

Informally, $F(v)$ is the set of vertices w of T whose colour is *forced* by $x(L(T_w))$. Our lower bound will be based on a conductance argument which shows that it takes a while to move from a colouring x in which $r(T)$ is forced to be one colour to a colouring y in which $r(T)$ is forced to be another colour. It is useful to note that $F(x)$ can be defined recursively using the structure of T. If w is a child of v we say that w is c-permitting for v in x if either $x(w) \neq c$ or $w \notin F(x)$ (or both).

Observation 7. *If $h(v) = 0$ then $v \in F(x)$. If $h(v) > 0$ then $v \in F(x)$ if and only if, for every colour $c \neq x(v)$, there is a child w of v which is not c-permitting for v in x.*

The recursive definition of $F(x)$ illustrates the connection between our conductance argument and lower-bound arguments based on *recursive majority functions* [2, 16].

Consider a colouring x chosen uniformly at random from Ω. Observe that the events $v \notin F(x)$, with v ranging over all vertices at height h, are independent. To see this, think about constructing the colouring downwards from the root, with each vertex choosing a colour uniformly at random from the colours not used by its parent. Now observe that each event $v \notin F(x)$ is completely determined by decisions made in the subtrees rooted at v. Let u_h be probability of the event $v \notin F(x)$, for any vertex v at height h; note that by symmetry this probability depends only on h.

Lemma 8. $u_h \leq 1/b$.

Proof. The proof is by induction on h. Note that $u_0 = 0$. For the inductive step, let v be a vertex at height $h > 0$. Consider a colouring x chosen uniformly at random from Ω. Fix a colour $c \neq x(v)$ and a child w of v. The probability that $x(w) = c$ is $1/(q-1)$. Also, the probability that $w \in F(x)$ is $1 - u_{h-1}$ and this is independent of the the event $x(w) = c$. (This can be seen by considering the downwards colouring scheme, as before, noting that $F(x)$ is determined by the restriction of x to the subtrees rooted at the children of w.) So the probability that w is c-permitting for v in x is $1 - (1 - u_{h-1})/(q-1)$. These events are independent for different children w of v so the probability that every child w is c-permitting for v in x is

$$\left(1 - \frac{1 - u_{h-1}}{q-1} \right)^b.$$

By Observation 7, the event $v \notin F(v)$ occurs when there exists a colour $c \neq x(v)$ such that every child w if c-permitting for v in x, so by the union bound:

$$
\begin{aligned}
u_h = \Pr(v \notin F(x)) &\leq (q-1) \left(1 - \frac{1 - u_{h-1}}{q-1} \right)^b \\
&\leq (q-1) \exp \left(-\frac{b(1 - u_{h-1})}{q-1} \right) \\
&\leq (q-1) \exp \left(-\frac{b-1}{q-1} \right) \quad (8) \\
&\leq (q-1)b^{-2} \quad (9) \\
&\leq b^{-1},
\end{aligned}
$$

where (8) applies the induction hypothesis and (9) uses assumption (7). $\qquad\square$

Consider a vertex v of T with $h(v) \geq 1$ and a leaf ℓ that is a descendant of v. Consider $x \in \Omega$. Say that v is ℓ-loose in x if there is a $c \neq x(v)$ such that every child w of v, except possibly the one on the path to ℓ, is c-permitting for v in x.

Let $\Psi_{v,\ell}$ be the probability that v is ℓ-loose in x when x is chosen u.a.r. from Ω. Let $\varepsilon = (q-1) \exp \left(-\frac{b-2}{q-1} \right)$.

Lemma 9. *Consider a vertex v of T with $h(v) \geq 1$ and a leaf ℓ that is a descendant of v. Then $\Psi_{v,\ell} \leq \varepsilon$.*

Proof. The calculation is very similar to the calculation in the proof of Lemma 8, with $b - 1$ replacing b. Let $h = h(v)$. Then

$$
\begin{aligned}
\Psi_{v,\ell} &\leq (q-1) \left(1 - \frac{1 - u_{h-1}}{q-1} \right)^{b-1} \\
&\leq (q-1) \exp \left(-\frac{b-2}{q-1} \right),
\end{aligned}
$$

where we have used the fact $u_{h-1} \leq b^{-1}$. $\qquad\square$

We are now ready to give the lower bound argument. The *conductance* of a set $S \subseteq \Omega$ is given by

$$\Phi_S(\mathcal{M}) = \frac{\sum_{x \in S} \sum_{y \in \overline{S}} \pi(x) P(x, y) + \sum_{x \in \overline{S}} \sum_{y \in S} \pi(x) P(x, y)}{2\pi(S)\pi(\overline{S})}.$$

The conductance of $\mathcal{M}$ is $\Phi(\mathcal{M}) = \min_S \Phi_S(\mathcal{M})$, where the min is over all $S \subset \Omega$ with $0 < \pi(S) < 1$. The inverse of the conductance of $\mathcal{M}$ gives a lower bound on the mixing time of $\mathcal{M}$. In particular,

$$\tau(\mathcal{M}, 1/(2e)) \geq (1/2 - 1/(2e))/\Phi(\mathcal{M}). \tag{10}$$

Equation (10) is due to Dyer, Frieze and Jerrum [7]. The formulation used here is Theorem 17 of the expository paper [8].

For $c \in [q]$, let $S_c = \{x \in \Omega \mid (r(T) \in F(x)) \wedge (x(r(T)) = c)\}$. Let $S_q = \{x \in \Omega \mid r(T) \notin F(x)\}$. Clearly, $S_0, \ldots, S_q$ form a partition of Ω. Let $S = S_0 \cup \cdots \cup S_{\lfloor q/2 \rfloor - 1}$. Then $\Phi(\mathcal{M}) \leq \Phi_S(\mathcal{M})$.

Now by Lemma 8 we have $0 \leq \pi(S_q) \leq 1/b$. Also, by symmetry, $\pi(S_c) = \pi(S_{c'})$ for $c, c' \in [q]$. So

$$\left(1 - \frac{1}{b}\right)\frac{\lfloor q/2 \rfloor}{q - 1} \leq \pi(S) \leq \frac{\lfloor q/2 \rfloor}{q - 1}.$$

Since $b \geq 6$ and $q \geq 3$ this gives $\frac{5}{6} \cdot \frac{1}{2} \leq \pi(S) \leq \frac{2}{3}$, so $\pi(S)\pi(\overline{S}) \geq \frac{1}{3} \cdot \frac{2}{3} = \frac{2}{9}$ Thus

$$\Phi_S(\mathcal{M}) \leq \frac{9}{4}\left(\sum_{x \in S} \sum_{y \in \overline{S}} \pi(x) P(x, y) + \sum_{x \in \overline{S}} \sum_{y \in S} \pi(x) P(x, y)\right),$$

and by reversibility

$$\Phi(\mathcal{M}) \leq \frac{9}{2} \sum_{x \in S} \sum_{y \in \overline{S}} \pi(x) P(x, y) \leq \frac{9}{2} \sum_{x,y} \pi(x) P(x, y), \tag{11}$$

where the summation is over x and y for which $r(T) \in F(x)$ and either $r(T) \notin F(y)$ or $x(r(T)) \neq y(r(T))$. Note that if x and y contribute to the summation in (11) then since $P(x, y) > 0$, they differ on a single vertex. Since $r(T) \in F(x)$ we cannot move from x to a proper colouring y by changing the colour of $r(T)$. Thus the only possibility is that $r(T) \notin F(y)$ and x and y differ on a leaf. Also, given the dynamics, we have $P(x, y) = 1/(nq)$.

Lemma 10. $\Phi(\mathcal{M}) \leq \frac{9}{2}\varepsilon^{H-1}$.

Proof. ¿From Equation (11) and the discussion above we have

$$\Phi(\mathcal{M}) \leq \frac{9}{2} \sum_{x,y} \pi(x) P(x, y)$$

12

where the sum is over all colourings x and y for which $r(T) \in F(x)$ and $r(T) \notin F(y)$ and x and y differ on exactly one leaf, ℓ. Letting $c = y(\ell)$, we can write

$$\Phi(\mathcal{M}) \leq \frac{9}{2} \sum_{x \in \Omega} \sum_{\ell \in L} \sum_{c \in [q]} 1_{x,\ell,c} \pi(x) \frac{1}{nq},$$

where $1_{x,\ell,c}$ is the indicator for the event that $r(T) \notin F(y)$ when y denotes the colouring formed from x by recolouring leaf ℓ with colour c. Multiplying by the q possibilities for c and noting that $\pi(X) = 1/|\Omega|$, we get

$$\Phi(\mathcal{M}) \leq \frac{9}{2} \frac{1}{|\Omega|} \frac{1}{nq} q \sum_{x \in \Omega, \ell \in L} 1_{x,\ell},$$

where $1_{x,\ell}$ is the indicator variable for the event that there is a colour c such that, when y is obtained from x by changing the colour of leaf ℓ to c, we have $r(T) \notin F(y)$. This event implies that every vertex v on the path from ℓ to $r(T)$ is ℓ-loose in x. When x is chosen uniformly a random these events are independent and by Lemma 9 they all have probability at most ε. So

$$\Phi(\mathcal{M}) \leq \frac{9}{2} \frac{1}{|\Omega|} \frac{1}{n} b^H |\Omega| \varepsilon^{H-1},$$

where b^H is the number of ℓ in the summation and $|\Omega|$ is the number of x. $\qquad \square$

Theorem 2 follows from Lemma 10 since, by Equation (10), the lemma implies

$$\tau(\mathcal{M}, 1/(2e)) \geq (1/2 - 1/(2e)) \frac{2}{9} \varepsilon^{-(H-1)}.$$

Also

$$\begin{aligned} \varepsilon^{-(H-1)} &= \left(\frac{1}{(q-1)\exp(-(b-2)/(q-1))} \right)^{H-1} \\ &= e^{(H-1)\left(\frac{b-2}{q-1} - \ln(q-1)\right)}. \end{aligned}$$

Using Lemma 6, this is at least

$$e^{(H-1)\left(\frac{b-2}{2(q-1)}\right)}.$$

Using Lemma 3, this is at least

$$e^{\frac{\ln n}{3 \ln b}\left(\frac{b-2}{2(q-1)}\right)} = n^{\frac{b-2}{6(q-1)\ln b}},$$

which gives Theorem 2.

13

References

[1] David J. Aldous. Some inequalities for reversible Markov chains. *J. London Math. Soc. (2)*, 25(3):564–576, 1982.

[2] Noam Berger, Claire Kenyon, Elchanan Mossel, and Yuval Peres. Glauber dynamics on trees and hyperbolic graphs. *Probab. Theory Related Fields*, 131(3):311–340, 2005.

[3] Nayantara Bhatnagar, Juan Vera, and Eric Vigoda. Reconstruction for colorings on trees, arxiv:0711.3664, 2007.

[4] Persi Diaconis and Laurent Saloff-Coste. Comparison theorems for reversible Markov chains. *Ann. Appl. Probab.*, 3(3):696–730, 1993.

[5] Persi Diaconis and Daniel Stroock. Geometric bounds for eigenvalues of Markov chains. *Ann. Appl. Probab.*, 1(1):36–61, 1991.

[6] Martin Dyer, Abraham D. Flaxman, Alan M. Frieze, and Eric Vigoda. Randomly coloring sparse random graphs with fewer colors than the maximum degree. *Random Struct. Algorithms*, 29(4):450–465, 2006.

[7] Martin Dyer, Alan Frieze, and Mark Jerrum. On counting independent sets in sparse graphs. *SIAM J. Comput.*, 31(5):1527–1541, 2002.

[8] Martin Dyer, Leslie Ann Goldberg, Mark Jerrum, and Russell Martin. Markov chain comparison. *Probab. Surv.*, 3:89–111 (electronic), 2006.

[9] Leslie Ann Goldberg, Russell Martin, and Mike Paterson. Random sampling of 3-colorings in $\mathbb{Z}_2$. *Random Struct. Algorithms*, 24(3):279–302, 2004.

[10] Thomas P. Hayes, Juan C. Vera, and Eric Vigoda. Randomly coloring planar graphs with fewer colors than the maximum degree. In *STOC'07—Proceedings of the 39th Annual ACM Symposium on Theory of Computing*, pages 450–458. ACM, New York, 2007.

[11] Michael Luby, Dana Randall, and Alistair Sinclair. Markov chain algorithms for planar lattice structures. *SIAM J. Comput.*, 31(1):167–192, 2001.

[12] Brendan Lucier, Michael Molloy, and Yuval Peres. The Glauber dynamics for colourings of bounded degree trees. In *Approximation, Randomization and Combinatorial Optimization*, volume 5687 of *Lecture Notes in Computer Science*, pages 631–645. Springer-Verlag, Berlin, 2009.

[13] Fabio Martinelli. Lectures on Glauber dynamics for discrete spin models. In *Lectures on probability theory and statistics (Saint-Flour, 1997)*, volume 1717 of *Lecture Notes in Math.*, pages 93–191. Springer, Berlin, 1999.

[14] Fabio Martinelli, Alistair Sinclair, and Dror Weitz. Fast mixing for independent sets, colorings, and other models on trees. *Random Struct. Algorithms*, 31(2):134–172, 2007.

[15] E. Mossel. *Survey: Information flow on trees*, volume 63 of *DIMACS Ser. Discrete Math. Theoret. Comput. Sci.*, pages 155–170. AMS Press, 2004.

[16] Elchanan Mossel. Recursive reconstruction on periodic trees. *Random Struct. Algorithms*, 13(1):81–97, 1998.

[17] Dana Randall and Prasad Tetali. Analyzing Glauber dynamics by comparison of Markov chains. *J. Math. Phys.*, 41(3):1598–1615, 2000. Probabilistic techniques in equilibrium and nonequilibrium statistical physics.

[18] Alistair Sinclair. Improved bounds for mixing rates of Markov chains and multicommodity flow. *Combin. Probab. Comput.*, 1(4):351–370, 1992.

[19] Alistair Sinclair and Mark Jerrum. Approximate counting, uniform generation and rapidly mixing Markov chains. *Inform. and Comput.*, 82(1):93–133, 1989.

[20] Prasad Tetali, Juan C. Vera, Eric Vigoda, and Linji Yang. Phase transition for the mixing time of the Glauber dynamics for coloring regular trees, August 2009. `arXiv:0908.2665v1`.

[21] Dror Weitz. *Mixing in time and space for discrete spin systems.* PhD thesis, U.C. Berkeley, 2004.