

Improved Distributed Algorithms for Random Colorings

Charlie Carlson

Department of Computer Science, University of California, Santa Barbara, CA, USA

Daniel Frishberg

Department of Computer Science and Software Engineering, California Polytechnic State University, San Luis Obispo, CA, USA

Eric Vigoda

Department of Computer Science, University of California, Santa Barbara, CA, USA

Abstract

Markov Chain Monte Carlo (MCMC) algorithms are a widely-used algorithmic tool for sampling from high-dimensional distributions, a notable example is the equilibrium distribution of graphical models. The Glauber dynamics, also known as the Gibbs sampler, is the simplest example of an MCMC algorithm; the transitions of the chain update the configuration at a randomly chosen coordinate at each step. Several works have studied distributed versions of the Glauber dynamics and we extend these efforts to a more general family of Markov chains. An important combinatorial problem in the study of MCMC algorithms is random colorings. Given a graph G of maximum degree Δ and an integer $k \geq \Delta + 1$, the goal is to generate a random proper vertex k -coloring of G .

Jerrum (1995) proved that the Glauber dynamics has $O(n \log n)$ mixing time when $k > 2\Delta$. Fischer and Ghaffari (2018), and independently Feng, Hayes, and Yin (2018), presented a parallel and distributed version of the Glauber dynamics which converges in $O(\log n)$ rounds for $k > (2 + \varepsilon)\Delta$ for any $\varepsilon > 0$. We improve this result to $k > (11/6)\Delta$ for a fixed $\varepsilon > 0$. This matches the state of the art for randomly sampling colorings of general graphs in the sequential setting. Whereas previous works focused on distributed variants of the Glauber dynamics, our work presents a parallel and distributed version of the more general flip dynamics presented by Vigoda (2000) (and refined by Chen, Delcourt, Moitra, Perarnau, and Postle (2019)), which recolors local maximal two-colored components in each step.

2012 ACM Subject Classification Theory of computation → Distributed algorithms; Theory of computation → Graph algorithms analysis; Theory of computation → Random walks and Markov chains

Keywords and phrases Distributed Graph Algorithms, Local Algorithms, Coloring, Glauber Dynamics, Sampling, Markov Chains

Digital Object Identifier 10.4230/LIPIcs.OPODIS.2023.13

Related Version *Full Version:* <https://doi.org/10.48550/arXiv.2309.07859>

Funding Eric Vigoda: Research supported in part by NSF grant CCF-2147094.

1 Introduction

This paper presents parallel and distributed algorithms for sampling from high-dimensional distributions. An important application is sampling from the equilibrium distribution of a graphical model. The equilibrium distribution is often known as the Gibbs or Boltzmann distribution, and efficient sampling from the Gibbs/Boltzmann distribution is a key step for Bayesian inference [24, 28].



Charlie Carlson, Daniel Frishberg, and Eric Vigoda;

licensed under Creative Commons License CC-BY 4.0

27th International Conference on Principles of Distributed Systems (OPODIS 2023).

Editors: Alysson Bessani, Xavier Défago, Junya Nakamura, Koichi Wada, and Yukiko Yamauchi; Article No. 13; pp. 13:1–13:18



Leibniz International Proceedings in Informatics

LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

13.2 Improved Distributed Algorithms for Random Colorings

Our focus is algorithms in the LOCAL model for the k -colorings problem. The k -colorings problem is a graphical model of particular combinatorial interest and has played an important role in the development of algorithmic sampling techniques with provable guarantees. The LOCAL model is a standard model of distributed computation due to Linial [26].

In the LOCAL model, the input to a problem is generally a graph $G = (V, E)$. Each vertex is identified with a processor and is assigned a unique identifier. In each *round* of an algorithm, each vertex is allowed to send an unbounded amount of information (a *message*) to each of its neighbors, and may perform an unbounded amount of computation locally.

For an input graph $G = (V, E)$ and integer $k \geq 2$, let Ω denote the proper (vertex) k -colorings of G , namely $\Omega = \{\sigma : V \rightarrow \{1, \dots, k\} : \text{for all } (v, w) \in E, \sigma(v) \neq \sigma(w)\}$ is the collection of assignments of k colors to the vertices so that neighboring vertices receive different colors. The associated Gibbs distribution μ is the uniform distribution over Ω , the space of proper k -colorings.

Under mild conditions on G (e.g., triangle-free [3]), the number of k -colorings is exponentially large, i.e., $|\Omega| = \exp(\Omega(n))$. Nevertheless, our goal is to sample from μ , the uniform distribution over this exponentially large set, in time $\text{poly}(n)$, and ideally in time $O(n \log n)$. Furthermore, in the distributed setting our goal is to generate samples ideally in time $O(\log n)$.

A common technique for sampling from the Gibbs distribution in a wide range of scientific fields is the *Markov Chain Monte Carlo (MCMC)* method. The simplest example of an MCMC algorithm is the Glauber dynamics, also known as the Gibbs sampler.

Consider an input graph $G = (V, E)$ with maximum degree Δ , and $k \geq \Delta + 2$. The Glauber dynamics updates the color of a randomly chosen vertex in each step. In particular, from a coloring $X_t \in \Omega$ at time t , the transitions $X_t \rightarrow X_{t+1}$ of the Glauber dynamics work as follows. We choose a random vertex v uniformly at random from V , and a color c uniformly at random from the set of colors $\{1, \dots, k\}$. If no neighbor of v has color c in the current coloring X_t , i.e., $c \notin X_t(N(v))$ where $N(v)$ are the neighbors of vertex v , then we recolor v as $X_{t+1}(v) = c$ and otherwise we set $X_{t+1}(v) = X_t(v)$. For all other vertices $w \neq v$ we set $X_{t+1}(w) = X_t(w)$. This corresponds to the Metropolis version of the Glauber dynamics. Alternatively one can choose the color c uniformly from $\{1, \dots, k\} \setminus X_t(N(v))$, which is the set of colors that do not appear in the neighborhood of v in X_t ; this is the heat-bath version of the Glauber dynamics.

When $k \geq \Delta + 2$ then the Glauber dynamics is ergodic and the unique stationary distribution is uniform over Ω . The *mixing time* is the number of steps, from the worst initial state X_0 , so that the chain is within total variation distance $\leq 1/4$ of the stationary distribution (see Section 2.2 for a more formal definition).

There are various attempts at running asynchronous versions of the Glauber dynamics in the distributed setting, namely HOGWILD! [35, 40], but there are few theoretical results and the resulting process is not guaranteed to have the correct asymptotic distribution [8, 7, 36]. There is also considerable work in constructing distributed sampling algorithms, including distributed versions of the Glauber dynamics [14, 13, 22, 12, 11, 27]; we discuss below the relevant results in our setting of the colorings problem. An important caveat about previous results is that they require a strong form of decay of correlations, such as the Dobrushin uniqueness condition, and our results hold in regions where Dobrushin's uniqueness condition does not hold.

In the sequential setting, a seminal work of Jerrum [21] proved $O(n \log n)$ mixing time of the Glauber dynamics whenever $k > 2^\Delta$ where Δ is the maximum degree. Vigoda [37] presented an alternative dynamics which we will refer to as the flip dynamics and proved

$O(n \log n)$ mixing time of the flip dynamics when $k > \frac{11}{6}$. The flip dynamics is a generalization of the Glauber dynamics which “flips” maximal 2-colored components (clusters) in each step by interchanging the pair of colors on the chosen cluster; Vigoda’s analysis chooses particular flip probabilities which depend on the size of the chosen cluster and do not flip any cluster larger than size six.

Vigoda’s result was recently improved to $k > (\frac{11}{6} - \varepsilon_0)$ for some fixed $\varepsilon_0 \approx 10^{-5}$ by Chen, Delcourt, Moitra, Perarnau, and Postle [5]. This later result of $k > (\frac{11}{6} - \varepsilon_0)$ is the best known result for general graphs. There are various improvements (e.g., [9, 6]), however they all require particular girth or maximum degree assumptions; the girth is the length of the shortest cycle.

In the distributed setting, Feng, Sun and Yin [12] achieved $O(\log n)$ rounds in LOCAL model when $k > (2 + \varepsilon)$ and $O(\log n)$ rounds when $k > (2 + \sqrt{2})$. Fischer and Ghaffari [14], and independently, Feng, Hayes and Yin [11], presented a distributed algorithm which converges in $O(\log n)$ rounds for k -colorings on any graph of maximum degree when $k > (2 + \varepsilon)$ for any $\varepsilon > 0$. These results match Jerrum’s result (in the sequential setting) for general graphs. We improve upon these works to match the current state of the art results in the sequential setting for general graphs for $k > (11/6)$.

We present the following improved result:

Theorem 1. *For all $\varepsilon > 0$, all ≥ 2 , all $\delta > 0$, and any $k > (11/6 + \varepsilon)$, for any graph $G = (V, E)$ of maximum degree Δ , a random k -coloring within total variation distance $\leq \delta$ from uniform can be generated in $O(\log(n/\delta))$ rounds, where $n = |V|$.*

The above result is optimal as there is a matching $\Omega(\log(n/\delta))$ lower bound due to Feng, Sun, and Yin [12]. Moreover, combining our analysis with the refined analysis of Chen et al. [5] we obtain the following result.

Theorem 2. *There exists $\varepsilon^* > 0$, for all ≥ 2 , all $\delta > 0$, and any $k > (11/6 - \varepsilon^*)$, for any graph of maximum degree Δ , a random k -coloring within total variation distance $\leq \delta$ from uniform can be generated in $O(\log(n/\delta))$ rounds.*

The Dobrushin uniqueness condition, which is a sufficient condition in several previous distributed sampling works, holds for colorings on general graphs of maximum degree iff $k > 2$ [33]. Thus, our results hold beyond the Dobrushin uniqueness threshold, and thereby resolves an open problem of [14] who asked whether efficient distributed algorithms intrinsically need to be stuck at Dobrushin’s condition.”

Our proof of fast convergence of our new distributed flip dynamics utilizes the path coupling framework of Bubley and Dyer [4], which is an important tool in the analysis of the mixing time for sequential Markov chains. In a coupling analysis path coupling allows one to only consider “neighboring pairs”. In the special case of the Glauber dynamics, path coupling is related to Dobrushin’s uniqueness condition but path coupling is a weaker condition (namely, Dobrushin’s uniqueness condition implies path coupling). We believe our work raises the following intriguing open question. For any spin system, or equivalently any undirected graphical model, does the path coupling condition for a local (sequential) Markov chain imply the existence of an efficient distributed algorithm which converges in $O(\log n)$ steps?

1.1 Motivation

Designing a distributed algorithm for constructing a coloring is a seminal problem in the study of distributed algorithms [26, 29]. It is an important problem in the study of symmetry breaking and is useful in the design of networking algorithms [2, 34, 25, 26]. One of the

fundamental problems in this context that has received significant attention is minimizing the number of rounds required to construct a $(\Delta + 1)$ -coloring in the LOCAL model; see Barenboim, Elkin, and Goldenberg [1] for a recent breakthrough, and see [15, 16] for more recent follow-up works.

Our focus is on generating a random coloring, in other words to generate a sample from the uniform distribution over all colorings, or more precisely, from a distribution that is arbitrarily close (in total variation distance) to the uniform distribution. More generally, our goal is to sample from the equilibrium distribution of a graphical model.

Graphical models are a fundamental tool in machine learning [28], and the associated sampling problem is important for associated learning, inference, and testing problems. A noteworthy example in the history of graphical models and in the importance of the associated sampling problem is the work on Restricted Boltzmann Machines (RBMs) of Hinton [18]. An RBM is an instance of the Ising model on a bipartite graph. The Ising model is a simpler variant of the random colorings problem in which we are sampling labelings of the vertices of a bipartite graph with only 2 colors where the labelings are weighted exponentially by the number of monochromatic edges; the generalization to $k > 2$ colors is the Potts model, and the zero-temperature (antiferromagnetic) Potts model is the random colorings problem that we study. The design of fast learning algorithms for RBMs was fundamental in the development of deep learning algorithms [19, 20, 30, 31, 32].

Given the proliferation of machine learning tasks on high-dimensional data, there is a clear need for distributed sampling algorithms for graphical models. For example, speeding up inference in *latent Dirichlet allocation* models via parallel and distributed Gibbs sampling [38, 23] and via the stochastic gradient sampler [39] has received attention in the machine learning community, as has the distributed problem of finding a k -coloring as a subroutine for Gibbs sampling [17].

Sampling colorings is a natural combinatorial problem to address particularly because of its importance in the study of sequential sampling algorithms. Jerrum's sampling algorithm [21] for $k > 2$ colors was a seminal work as it pioneered the use of the coupling method for sampling problems on graphical models. As mentioned earlier, Vigoda [37] improved Jerrum's result to $k > 11/\epsilon/6$ and this was the state of the art until the recent improvement to $k > (11/\epsilon)^{1/\epsilon}$ [5]. One of the major open problems in the area of sequential sampling is to obtain an efficient sampling scheme when $k > \Delta + 1$, see [6] for the most recent progress.

Our general question is whether efficient sequential sampling schemes yield efficient distributed sampling algorithms, by which we mean an $O(\log n)$ round algorithm in the LOCAL model. A distributed version of the Metropolis version of the Glauber dynamics for colorings was introduced in [14, 11] and was proved to be an efficient distributed sampling scheme when $k > (2 + \epsilon)\Delta$ for all $\epsilon > 0$. Our work goes beyond the single-site Glauber dynamics to designing efficient distributed sampling schemes for more general dynamics.

1.2 Technical Contribution

Recall that the Glauber dynamics updates a single vertex in each step. Several recent works present and analyze distributed versions of the Glauber dynamics (specifically, the Metropolis version) in various contexts [14, 11, 27, 12]. For more general MCMC algorithms which update larger regions than a vertex in each step, do efficient convergence results in the sequential setting for such Markov chains yield efficient distributed sampling algorithms?

A prime example to consider for this more general question is Vigoda's flip dynamics [37]. Attaining a distributed version of the flip dynamics is more challenging as we need to simultaneously recolor clusters of up to 6 vertices; here a cluster refers to a maximal 2-colored

component and the recoloring acts by interchanging the respective pair of colors on each cluster. Our first contribution is presenting a distributed version of Vigoda's flip dynamics. The challenge is to make a distributed version which is efficient but simple enough that we can still analyze it.

To parallelize the cluster recolorings, we need to ensure that no two overlapping clusters are simultaneously active, and that no two neighboring clusters that share colors are both active. On the other hand, we need to "activate" each cluster for potential recoloring with a sufficiently large probability to obtain a mixing time that is independent of the maximum degree, namely $O(\log n)$.

Our analysis of our distributed version of Vigoda's flip dynamics follows the high-level coupling presented in Vigoda's original work [37]. A coupling analysis of a Markov chain, considers two copies of the Markov chain (in this case the distributed flip dynamics), each with arbitrary starting states. Our aim is that there are "coupled transitions" for the two chains so that after $O(\log n)$ steps the two chains have coalesced in the same state with sufficiently large probability; by coupled transition we mean that the two chains can couple their transitions as long as when viewed in isolation, each is a faithful copy of the original Markov chain. The idea is that if we consider one of the chains to be in the stationary distribution, then we showed that after $O(\log n)$ steps our algorithm has likely reached the stationary distribution and hence the mixing time is $O(\log n)$.

There are several important technical challenges that arise when doing a coupling analysis in the distributed setting for the flip dynamics. First, we need to ensure that the clusters we flip (which means swap the pair of colors in a maximal 2-colored component) do not interfere with any other clusters we might flip by either overlapping, or by neighboring and containing a common color. Subsequently when we do try to couple a pair of flips in the two coupled chains, we need to consider the case that one of these two clusters is not flippable in only one chain due to one of these aforementioned conflicts (such as an overlapping cluster in only one of the chains).

Finally, similar to the original analysis of Vigoda [37], we use the path coupling framework [4] from which we only need to design and analyze a coupling for pairs of chains that differ at a single vertex, which we call v^* ; in contrast, without path coupling we need to analyze pairs of chains that differ on an arbitrary number of vertices. However, the coupling analysis for a pair of chains X_t, Y_t that differ at this single vertex v^* is more complicated than in Vigoda's sequential setting. In Vigoda's original analysis, the only pertinent cluster flips in X_t or Y_t are those clusters that either include v^* or include a neighbor of v^* . In our analysis in the distributed setting, we also need to consider the effect from clusters that are distance exactly 2 away from v^* , where distance is measured by cluster adjacencies. These distance-2 clusters are identical sets of vertices in both chains X_t and Y_t but they may be flippable in only one of the chains (due to differing distance-1 clusters).

Our work suggests that a more general phenomenon is at play. We conjecture that, for any graphical model, a path coupling analysis for any local Markov chain in the sequential setting yields an efficient distributed sampling scheme. We believe our work will be an important step towards proving this general conjecture.

1.3 Paper Overview

In Section 3, we present a parallel and distributed version of Vigoda's flip dynamics. We analyze the mixing time of our distributed flip dynamics when $k > (11/6 + \varepsilon)$ for any $\varepsilon > 0$, thereby proving Theorem 1, in Sections 4 and 5. We use a coupling argument that builds upon the analysis in Vigoda [37]. Our analysis is more complicated than the original

argument of Vigoda due to clusters which appear in both coupled chains, but possibly being “flippable” in one chain but not the other chain due to differing conflicts with neighboring clusters, see Section 1.2 for a very high-level overview. In the appendix of the full version of this paper, we further utilize the linear programming (LP) framework and the refined metric on colorings presented in Chen et al. [5] to achieve the further improved result as stated in Theorem 2. This proof combines our proof approach for Theorem 1 with the more technical analysis of Chen et al. [5].

2 Preliminaries

Let $[k] = \{1, \dots, k\}$. For a graph $G = (V, E)$ let $i \sim j$ denote $(i, j) \in E$, and for $v \in V$, let $N(v) = \{w \in V : v \sim w\}$ denote the neighbors of a vertex v . For integer $k \geq 2$, let $\Omega^* = [k]^V$ denote the set of k -labelings and $\Omega = \{\sigma \in [k]^V : \text{for all } i \sim j, \sigma(i) \neq \sigma(j)\}$ denote the set of k -colorings of G . Throughout this paper, a coloring (or k -coloring) refers to a proper vertex k -coloring.

2.1 Clusters

For a coloring σ , a cluster S in σ is a maximal 2-colored component of *size at most 6*; this is formally defined in the following definition.

Definition 3. Let $G = (V, E)$ be a graph and $\sigma \in \Omega^*$. For a vertex $v \in V$ and color $c \in [k]$ let $S_\sigma(v, c)$ denote the set of vertices reachable from v by a $(\sigma(v), c)$ alternating path, i.e., a path of vertices $v = v_1, v_2, \dots, v_\ell \in V$ for some $\ell \geq 1$ that alternate between colors $\sigma(v)$ and c . When $|S_\sigma(v, c)| \leq 6$ then we refer to $S = S_\sigma(v, c)$ as a cluster. Let

$$\mathcal{S}_\sigma = \bigcup_{v \in V, c \in [k]} \{S_\sigma(v, c) : |S_\sigma(v, c)| \leq 6\},$$

denote the collection of all clusters in σ of size at most 6, where the size of a cluster refers to the number of vertices in the cluster. The restriction to size at most 6 is due to the Markov chain used as in previous works [37, 5].

The key operation of our Markov chain is “flipping” clusters which we define now.

Definition 4. For a labeling $\sigma \in \Omega^*$, vertex $v \in V$, and color $c \in [k]$, the flip of cluster $S_\sigma(v, c)$ interchanges colors $\sigma(v)$ and c on the set $S_\sigma(v, c)$.

Let σ' denote the resulting coloring after this flip of cluster $S_\sigma(v, c)$. Notice that if $\sigma \in \Omega$ then $\sigma' \in \Omega$, i.e., if it is a proper coloring before the flip, then after the flip it remains a proper coloring since the clusters are maximal 2-colored components. Hence for Vigoda’s flip dynamics, if we start the flip dynamics at a proper coloring, i.e., $X_0 \in \Omega$, then we are guaranteed to stay at proper colorings, i.e., $X_t \in \Omega$ for all $t \geq 0$.

The subsequent flip dynamics defined in Section 3 is defined on the set Ω , which is the set of proper k -colorings, and for algorithmic purposes one only needs to consider proper colorings. The extension of the state space to Ω^* , which is the set of all labelings, is only needed in the proof for technical reasons pertaining to the use of the path coupling method [4], which we present in Section 2.3. The introduction of improper colorings in the coupling analysis arises in all related path coupling proofs for colorings [4, 37, 5], see Section 4.2 for further discussion of this technicality of introducing improper colorings in the proof.

Consider a coloring $\sigma \in \Omega$ and a vertex $v \in V$. For every color c which does not appear in the neighborhood of v , i.e., $c \notin \sigma(N(v))$ then the corresponding cluster is of size 1, i.e., $|S_\sigma(v, c)| = 1$ since $S_\sigma(v, c) = \{v\}$. Flips of these singleton clusters are exactly the transitions

of the Glauber dynamics. The flip dynamics of Vigoda [37] is a generalization of the Glauber dynamics in which clusters of size at most 6 are flipped with positive probability (depending on the size of the cluster). Note, for $c = \sigma(v)$ then we get a singleton cluster and the flip does not change the coloring, hence the flip dynamics has a non-zero self-loop probability and thus is aperiodic.

For clusters $S, T \in \mathcal{S}_\sigma$, we say S and T are neighboring clusters, which we denote as $S \sim T$, if there exists $v \in S$ and $w \in T$ where $v \sim w$.

2.2 Markov Chains

Consider a Markov chain (X_t) with state space Ω and transition matrix P and unique stationary distribution π . We say that the chain is *aperiodic* if $\gcd\{t: P^t(x, x) > 0\} = 1$ for all $x \in \Omega$ and *irreducible* if for all $x, y \in \Omega$, there exists a t such that $P^t(x, y) > 0$. Recall that if the chain is both aperiodic and irreducible, then it is *ergodic* and the chain has a unique *stationary distribution* π where: for all $x, y \in \Omega$, $\lim_{t \rightarrow \infty} P^t(x, y) = \pi(y)$. If P is symmetric, then π is the uniform distribution over Ω .

The *mixing time* is the number of steps, from the worst initial state X_0 , until the chain is within total variation distance $\leq 1/4$ of the stationary distribution:

$$T_{\text{mix}} := \max_{x \in \Omega} \min\{t \mid \|P^t(\sigma, \cdot) - \pi\|_{\text{TV}} \leq 1/4\},$$

where d_{TV} is the *total variation distance*, $\|\mu - \omega\|_{\text{TV}} := \frac{1}{2} \sum_{x \in \Omega} |\mu(x) - \omega(x)|$. The choice of constant $1/4$ is somewhat arbitrary since, for any $\varepsilon > 0$, we can obtain total variation distance $\leq \varepsilon$ after $\leq \log(1/\varepsilon)T_{\text{mix}}$ steps.

2.3 Path Coupling

Consider an ergodic Markov chain \mathcal{MC} with state space Ω and transition matrix P . A *coupling* for \mathcal{MC} defines, for all pairs $X_t, Y_t \in \Omega$, a joint transition $(X_t, Y_t) \rightarrow (X_{t+1}, Y_{t+1})$ such that the individual transitions $(X_t \rightarrow X_{t+1})$ and $(Y_t \rightarrow Y_{t+1})$, when viewed in isolation from each other, act according to the transition matrix P . The goal is to find a coupling that minimizes the coupling time: $T_{\text{couple}} := \min \{t \mid \text{for all } X_0, Y_0 \in \Omega, \Pr(X_t \neq Y_t \mid X_0, Y_0) \leq 1/4\}$. This implies that $T_{\text{mix}} \leq T_{\text{couple}}$.

To bound the coupling time and hence the mixing time, we use the *path coupling* method of Bubley and Dyer [4] which allows us to only consider a small subset of pairs of states. We will analyze the coupling with respect to the *Hamming distance* $H(\sigma, \tau) := \sum_{v \in V} (\sigma(v) \neq \tau(v))$. We present the more general form of path coupling in the appendix of the full version of this paper which allows more general metrics.

Theorem 5 [4, 10]). Consider an ergodic Markov chain on $\Omega^* = [k]^V$. Let $\beta > 0$. If for all pairs of states $X_t, Y_t \in \Omega^*$ where $H(X_t, Y_t) = 1$, there exists a coupling such that:

$$\mathbf{E}(H(X_{t+1}, Y_{t+1}) \mid X_t, Y_t) \leq (1 - \beta),$$

then the mixing time is bounded by $T_{\text{mix}} \leq O\left(\frac{\log(|V|)}{\beta}\right)$. Moreover, the mixing time within total variation distance $\leq \delta$, for any $\delta > 0$, in time $O(\log(|V|)/(\beta\delta))$.

3 Algorithm Description: Distributed Flip Dynamics

We begin by defining a sequential process and then show that this process can be implemented efficiently in a distributed manner.

We have the following parameters in our algorithm. Let $\gamma = \varepsilon/(5000k)$ where $k \geq (11/6 + \varepsilon)$ for some $\varepsilon > 0$. The parameter γ will be used for the activation probability of a cluster. In the full version of this paper when we strengthen the main result for $k < (11/6)$ we redefine γ so that it depends on the distance of k below $(11/6)$.

Let $1 \geq f_i \geq 0$ for all $i \geq 1$ be a sequence of “flip” probabilities that contain the following key properties: $f_1 = 1$, $f_i \geq f_{i+1}$ for all i , and $f_i = 0$ for all $i \geq 7$. The following process is well-defined for any choice of flip probabilities with these properties. To prove Theorems 1 and 2 we will choose slightly different flip probabilities. In particular, to prove the slightly weaker result (Theorem 1) in Section 4 we will choose flip probabilities as in [37], and then to get the refined result (Theorem 2) in the appendix of the full version of this paper we will use the setting in [5].

We now define the Markov chain $\mathcal{MC}_{\text{flip}}$ with state space Ω . For a coloring $X_t \in \Omega$, the transitions $X_t \rightarrow X_{t+1}$ of $\mathcal{MC}_{\text{flip}}$ are defined as follows:

1. Independently for each $S \in \mathcal{S}_\sigma$, cluster S is active with probability γ .
2. A cluster $S = S_{X_t}(v, c)$ is flippable if the following hold:
 - (a) S is active;
 - (b) *Overlapping clusters*: There is no active $S' \neq S$ where $S \cap S' \neq \emptyset$;
 - (c) *Conflicting neighboring clusters*: For all active clusters $T = T_{X_t}(w, c')$ where $S \sim T$, $\{X_t(v), c\} \cap \{X_t(w), c'\} = \emptyset$.
3. Independently for each flippable cluster S , flip S with probability f_i where $i = |S|$.
4. Let X_{t+1} denote the resulting coloring.

Notice that step 2c is saying that for a pair of active and neighboring clusters S and T , the pair of colors defining cluster S are disjoint from the pair of colors defining cluster T .

Lemma 6. *The Markov chain $\mathcal{MC}_{\text{flip}}$ is ergodic and symmetric and hence the unique stationary distribution is the uniform distribution over Ω .*

Proof. Observe that with positive probability, no cluster is active and $P(\sigma, \sigma) > 0$ for all $\sigma \in \Omega$. Thus, the Markov chain is aperiodic. For irreducibility, since $f_1 > 0$, the irreducibility of $\mathcal{MC}_{\text{flip}}$ follows from irreducibility of the Glauber dynamics which holds whenever $k \geq \gamma + 2$ (see, e.g., Jerrum [21]). Hence, the chain is ergodic. Moreover, the chain is symmetric, for $\sigma \in \Omega$, let σ' be the coloring obtained from σ after flipping clusters $S_\sigma(v_1, c_1), \dots, S_\sigma(v_\ell, c_\ell)$ in one step of $\mathcal{MC}_{\text{flip}}$. Then, starting from σ' and flipping clusters $S_{\sigma'}(v_1, \sigma(v_1)), \dots, S_{\sigma'}(v_\ell, \sigma(v_\ell))$ recovers σ . Since $\mathcal{MC}_{\text{flip}}$ is ergodic and symmetric then the uniform distribution is the unique stationary distribution. \blacktriangleleft

Lemma 7. *Each step of the Markov chain $\mathcal{MC}_{\text{flip}}$ can be implemented in the LOCAL model in $O(1)$ rounds.*

Proof. We describe the steps of the algorithm and how to implement them in the LOCAL model. At a given time step t , denote the current coloring as $\sigma = X_t$.

1. For each vertex $v \in V$ and for each color $c \in [k]$, identify the cluster $S_\sigma(v, c)$. We accomplish this step by (i) sending a message indicating the index of v to each neighboring vertex w with $\sigma(w) = c$, (ii) passing this message, along with the index of w , to each neighbor x of w with $\sigma(x) = \sigma(v)$, and (iii) repeating this process for up to six rounds. After the six rounds, each vertex has received the identities of all other vertices in its six-hop neighborhood with which it might share a cluster, and thus can determine the clusters (and their sizes) to which it belongs. Moreover, any 2-colored components of size > 6 will be identified and discarded.

2. Fix an arbitrary ordering of the vertex set of G and for each cluster S , identify $\text{pres}(S)$, the lowest-index vertex $v \in S$. We can accomplish this step by letting each vertex $u \in S$ compare its own index to each of the indices of other vertices in S , which have been passed during step 1.
3. For each cluster S , activate S with probability \dots . More precisely, for each $v \in V$ and for each $c \in [k]$, if $v = \text{pres}(S_\sigma(v, c))$, activate $S_\sigma(v, c)$ by sending a message to every $u \in S_\sigma(v, c)$.
4. Detect conflicts:
 - (a) *Overlapping clusters*: For all $v \in V$, if $S_\sigma(v, c), S_\sigma(v, c')$ are both active for some $c \neq c'$, send messages to $\text{pres}(S_\sigma(v, c)), \text{pres}(S_\sigma(v, c'))$ to “deactivate” $S_\sigma(v, c), S_\sigma(v, c')$.
 - (b) *Conflicting neighboring clusters*: For all $v \in V$, for every neighbor w of v , if there exist clusters $S_\sigma(v, c) \neq S_\sigma(w, c')$ such that $\{\sigma(v), c\} \cap \{\sigma(w), c'\} \neq \emptyset$ and if $S_\sigma(v, c)$ and $S_\sigma(w, c')$ are both active, deactivate $S_\sigma(v, c)$ and $S_\sigma(w, c')$ (by sending messages to $\text{pres}(S_\sigma(v, c))$ and $\text{pres}(S_\sigma(w, c'))$).
5. For all $v \in V$, for all $c \in [k]$, if $S_\sigma(v, c)$ is still active and $v = \text{pres}(S_\sigma(v, c))$, flip $S_\sigma(v, c)$ with probability f_i , where $i = |S_\sigma(v, c)|$ (by sending a message to each $w \in S_\sigma(v, c)$ to change its color from c to $\sigma(v)$ or vice versa).

Since, in step 5, only $\text{pres}(S)$ is responsible for flipping S , the probability of a given cluster S being flipped, conditioned on S being active and having no active neighboring or overlapping cluster, is $f_{|S|}$.

Each of the above steps requires a constant number of rounds, proving the claim. Furthermore, the amount of computation performed locally at each vertex depends only (and polynomially) on the maximum degree of the graph and the number of colors. That is, not only is the number of rounds in the LOCAL model $O(1)$, but also the algorithm is efficient with respect to the local computation performed in each round. \blacktriangleleft

4 Analysis of Distributed Flip Dynamics

Here we prove our main result Theorem 1, namely fast convergence of the distributed flip dynamics when $k > (11/6 + \varepsilon)$ for any $\varepsilon > 0$. Hence, fix $\varepsilon > 0$ and $k > (11/6 + \varepsilon)$. Our specific choice of flip probabilities for this section and for Section 5 are the following:

$$f_1 = 1, f_2 = 13/42, f_3 = 1/6, f_4 = 2/21, f_5 = 1/21, f_6 = 1/84. \quad (1)$$

These parameters match the original paper of Vigoda [37]; there are other parameter choices for which the analysis works, e.g., see [5], in fact, we will utilize these alternative parameters in the appendix of the full version of this paper.

4.1 Overview

We will analyze the mixing time of the chain $\mathcal{MC}_{\text{flip}}$ using path coupling. Consider a pair of colorings X_t, Y_t which differ at exactly one vertex and let v^* denote the disagreement, i.e., $X_t(v^*) \neq Y_t(v^*)$ and for all $w \neq v^*$, $X_t(w) = Y_t(w)$. Our coupling is the identity coupling for all clusters that are the same in both chains, i.e., for all clusters S where $S = S_{X_t}(w, c) = S_{Y_t}(w, c)$ for some $w \in V, c \in [k]$, we use the identity coupling for the activation probability. By the identity coupling for the activation probability we mean that with probability \dots the cluster S is active in both chains, and with probability $1 - \dots$ it is inactive in both chains. Moreover, if the cluster S is flippable in both chains then we also use

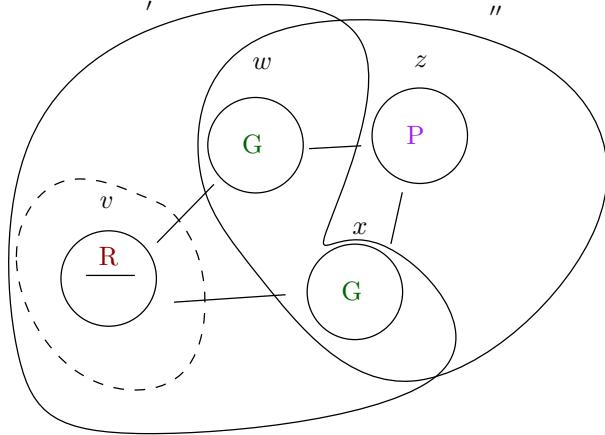


Figure 1 The vertex v has $X_t(v) = R$ and $Y_t(v) = B$. Here, $T = S(v, R) = S_\tau(v, B)$, where $\sigma = X_t$ and $\tau = Y_t$. By Definition 8, $\text{dist}(v, T) = 0$, $\text{dist}(v, T') = \text{dist}(v, T'') = 1$.

the identity coupling for the flip probability, which means that if both clusters are flippable then with probability $\frac{1}{2}$ we flip the cluster in both chains and with probability $\frac{1}{2}$ we flip the cluster in neither of the chains.

We will define the distance of cluster T from the disagree vertex v^* based on the shortest path via neighboring clusters.

Definition 8. For a coloring $\sigma \in \{X_t, Y_t\}$, and a cluster $T \in \mathcal{S}_\sigma$, we define $\text{dist}(v^*, T)$ inductively as follows. If $T = \{v^*\}$ then let $\text{dist}_\sigma(v^*, T) = 0$. In general, let

$$\text{dist}_\sigma(v^*, T) = \min\{i : \text{there exists } S \in \mathcal{S}_\sigma \text{ where } S \sim T, \text{dist}_\sigma(v^*, S) = i - 1\}.$$

Remark 9. Note, this notion of distance is equivalent to the shortest path distance from the singleton cluster $\{v^*\}$ in the *cluster graph*; the cluster graph is the graph on all clusters in coloring σ where clusters S and T are adjacent if $S \sim T$. Distance 0 clusters are the singleton sets $\{v^*\}$ for every color which does not appear in the neighborhood of v^* . Distance 1 clusters are those that contain a neighbor of v^* (regardless of whether they also contain v^*).

Any clusters T where no vertex in T is adjacent to v^* are identical in the two chains, and thus, for every $i \geq 2$:

$$T \in \mathcal{S}_{X_t}, \text{dist}_{X_t}(v^*, T) = i \iff T \in \mathcal{S}_{Y_t}, \text{dist}_{Y_t}(v^*, T) = i.$$

Similarly, the only clusters T which “disagree” in the sense that they appear in only one chain then T is at distance 1 from v^* ; more formally, if $T \in \mathcal{S}_{X_t} \setminus \mathcal{S}_{Y_t}$, then $\text{dist}_{X_t}(v^*, T) = 1$, and if $T \in \mathcal{S}_{Y_t} \setminus \mathcal{S}_{X_t}$, then $\text{dist}_{Y_t}(v^*, T) = 1$. We use $\text{dist}(v^*, T)$ when the distances are equal, i.e., $\text{dist}_{X_t}(v^*, T) = \text{dist}_{Y_t}(v^*, T)$.

For such clusters T where $\text{dist}(v^*, T) \geq 2$ we use the identity coupling for the activation probability in X_t and Y_t , and thus the cluster T is active in both chains or in neither chain. It follows that for clusters T' with $\text{dist}(v^*, T') \geq 3$ then the cluster is flippable in both chains or in neither chain, as their neighboring active clusters are identical in the two chains. Therefore, we can use the identity coupling for the flip probability of this cluster T if the cluster is flippable, and such clusters are flipped in both chains or neither chain; this leads to the following observation.

Observation 10. For any cluster T' where $\text{dist}(v^*, T') \geq 3$, $X_{t+1}(T') = Y_{t+1}(T')$.

For clusters T where $\text{dist}_G(T, v^*) = 2$, it can occur that T is flippable in only one of the chains (due to a neighboring cluster at distance 1 that occurs in only one of the chains). Hence, there is a probability that such clusters can be a new disagreement. The upcoming Lemma 11 proves that this occurs with an arbitrarily small constant probability.

The following lemma bounds the expected increase in Hamming distance from flips on clusters at distance exactly 2 from v^* .

Lemma 11.

$$\sum_{T: \text{dist}(v^*, T)=2} |T| \mathbf{Pr}(X_{t+1}(T) \neq Y_{t+1}(T)) \leq \varepsilon / 10,$$

where $k \geq (\frac{11}{6} + \varepsilon)$.

We will account for these potential disagreements at distance 2 via the clusters at distance 1. For a cluster T at distance 1 to occur in only one of the chains, the pair of colors defining T must include color $X_t(v^*)$ or color $Y_t(v^*)$.

Proof of Lemma 11. Let $\mathcal{S}_{X_t} \oplus \mathcal{S}_{Y_t} := (\mathcal{S}_{X_t} \setminus \mathcal{S}_{Y_t}) \cup (\mathcal{S}_{Y_t} \setminus \mathcal{S}_{X_t})$ denote the set of clusters that appear in one chain but not in the other chain. Consider a cluster $S \in \mathcal{S}_{X_t} \oplus \mathcal{S}_{Y_t}$. Note, all such S are at $\text{dist}(v^*, S) = 1$.

Let $c_X = X_t(v^*)$ and $c_Y = Y_t(v^*)$. These clusters $S \in \mathcal{S}_{X_t} \oplus \mathcal{S}_{Y_t}$ are either:

$$S_{X_t}(w, c_X), S_{X_t}(w, c_Y), S_{Y_t}(w, c_X), \text{ or } S_{Y_t}(w, c_Y),$$

for some neighbor $w \in N(v^*)$. Hence, there are ≤ 4 such clusters $S \in \mathcal{S}_{X_t} \oplus \mathcal{S}_{Y_t}$.

Each such cluster S has size ≤ 6 and hence it has $\leq 6 \cdot 2$ neighboring clusters T that share a color with S . These clusters T are at distance $= 2$ from v^* . Note that if S and T are both active then T is not flippable in one of the chains, but it may be flippable in the other chain where S does not appear; hence, the chains X_{t+1} and Y_{t+1} potentially differ at T . This yields the following:

$$\sum_{T: \text{dist}(v^*, T)=2} |T| \mathbf{Pr}(X_{t+1}(T) \neq Y_{t+1}(T)) \leq 6 \times (4)(12) \leq 288 \leq \varepsilon / 10. \quad \blacktriangleleft$$

We now account for the “good moves” where the disagreement at v^* is removed. This occurs by Glauber updates at v^* where we update v^* to an available color, which is a color that does not appear in its neighborhood.

Definition 12. Denote the set of available colors for v^* in X_t as:

$$A(v^*) = A_{X_t}(v^*) := \{c : c \notin X_t(N(v^*))\}.$$

Note, the sets $A_{X_t}(v^*) = A_{Y_t}(v^*)$ since v^* is the only disagreement at time t . Consider a color $c \in A(v^*)$. The clusters involving c to which v^* belongs satisfy $S_{X_t}(v^*, c) = S_{Y_t}(v^*, c) = \{v^*\}$ and hence the identity coupling is used for this cluster. Therefore, with probability the cluster is active in both chains and if no active clusters overlap and no neighboring clusters have a common color then v^* is recolored to c .

We can now bound the probability of v^* agreeing at time $t+1$ in terms of the number of available colors for v^* .

Lemma 13. $\mathbf{Pr}(X_{t+1}(v^*) = Y_{t+1}(v^*)) \geq |A(v^*)| (1 - \varepsilon/500).$

13:12 Improved Distributed Algorithms for Random Colorings

Proof. For each color $c \in A(v^*)$ note $S_{X_t}(v^*, c) = S_{Y_t}(v^*, c) = \{v^*\}$. Hence, for $c \in A(v^*)$, let $S_c = S_{X_t}(v^*, c) = S_{Y_t}(v^*, c)$ denote this cluster of size 1 which appears in both chains. Since S_c appears in both chains we use the identity coupling for being active so that with probability $\frac{1}{2}$ the cluster S_c is active in both chains, and with probability $\frac{1}{2}$ the cluster S_c is inactive in both chains. The cluster S_c may have different neighboring clusters in the two chains (which affects whether it is flippable) but if it is flippable in both chains then with probability $f_1 = 1$ we flip the cluster in both chains.

There are at most $2k-1$ neighboring clusters in each chain that share a color with one of the respective cluster, and there are $k-1$ clusters (namely those at v^*) that overlap with these clusters. If none of the $2k-1$ neighboring clusters is active, and none of the $2(k-1)$ overlapping clusters is active in either chain, then we can flip $\{v^*\}$ in both chains. After this flip, v^* agrees in both chains, and hence we obtain:

$$\begin{aligned} \Pr(X_{t+1}(v^*) = Y_{t+1}(v^*)) &\geq |A(v^*)| \cdot (1 - \frac{1}{2})^{4k-2} \\ &\geq |A(v^*)| \cdot \exp(-\varepsilon(4k-2)/2500k) \\ &\geq |A(v^*)| \cdot \exp(-\varepsilon/500) \\ &\geq |A(v^*)| \cdot (1 - \varepsilon/500), \end{aligned}$$

where the second inequality uses the fact that $1 - x \geq \exp(-2x)$ for $x \leq 1/2$. \blacktriangleleft

The upcoming lemma captures the potential disagreements that arise from flipping clusters at distance one. The coupling on clusters containing v^* or neighboring v^* in at least one chain will be coupled based on the new color c .

Definition 14. For a color $c \in [k]$, let $N_c(v^*) = \{w \in N(v) : X_t(w) = c\} = \{w \in N(v) : Y_t(w) = c\}$ denote the neighbors of v^* with color c , and let $d_c(v^*) = |N_c(v^*)|$ denote the number of neighbors of v^* with color c at time t .

Let $\mathcal{S}_{X_t}(c)$ denote the collection of clusters at distance 1 in X_t that involve color c :

$$\mathcal{S}_{X_t}(c) := \{S_{X_t}(w, X_t(v^*)) \mid w \in N_c(v^*)\} \cup \{S_{X_t}(w, Y_t(v^*)) \mid w \in N_c(v^*)\},$$

and similarly let $\mathcal{S}_{Y_t}(c)$ denote the corresponding collection for the coloring Y_t .

The sets $\mathcal{S}_{X_t}(c)$ and $\mathcal{S}_{Y_t}(c)$ are coupled with each other. We will specify the detailed coupling later, for now all that is needed is that these sets $\mathcal{S}_{X_t}(c)$ and $\mathcal{S}_{Y_t}(c)$ are coupled with each other. We can now state the key lemma bounding the increase in Hamming distance when we do a coupled update on these sets $\mathcal{S}_{X_t}(c), \mathcal{S}_{Y_t}(c)$.

In the following statement, recall, that for $\sigma, \tau \in \Omega$, $H(\sigma, \tau) = |\{v \in V : \sigma(v) \neq \tau(v)\}|$ is the Hamming distance.

Lemma 15. Let $c \in [k]$ where $d_c(v^*) > 0$. Recall that the flips of clusters in $\mathcal{S}_{X_t}(c)$ for $X_t \rightarrow X_{t+1}$ are coupled with clusters in $\mathcal{S}_{Y_t}(c)$ for $Y_t \rightarrow Y_{t+1}$. Let \mathcal{F}_c denote the event that one of these coupled flips occurred in at least one of the chains. Then,

$$\sum_{c \in [k]} \mathbf{E}(H(X_{t+1}, Y_{t+1}) \mid \mathcal{F}_c) \leq \sum_{c \in [k]} \left[1 + (1 + \varepsilon/5) \left(\frac{11}{6} d_c(v^*) - 1 \right) \right] + \varepsilon / 10.$$

The proof of Lemma 15 is deferred to Section 5. Note the error term $\varepsilon / 10$ is coming from Lemma 11. Combining the above lemmas we can prove the main result (this is the slightly weaker version for $k \geq (11/6 + \varepsilon)$ for any $\varepsilon > 0$).

4.2 Proof of Theorem 1

We can extend the definition of our Markov chain (see Section 3) to be over all labelings $\Omega^* = [k]^V$ instead of just proper colorings Ω . This is necessary to apply path coupling Theorem 5. An identical approach is used in both [37] and [5]. The reason for this extension of the state space is the following. In the path coupling analysis we start with a pair of chains X_t, Y_t that differ at a single vertex v^* . In the coupling analysis we may introduce a new disagreement at time $t + 1$ at a neighbor $w \in N(v^*)$ where the pair of disagreements at time $t + 1$ are colored as $X_{t+1}(w) = Y_{t+1}(v^*)$, $Y_{t+1}(w) = X_{t+1}(v^*)$ and $X_{t+1}(v^*) \neq Y_{t+1}(v^*)$. Hence, the Hamming distance between X_{t+1} and Y_{t+1} is two but the number of Glauber dynamics steps (or cluster flips) to go from X_{t+1} to Y_{t+1} is three, and therefore in the path coupling analysis this new disagreement at w increases the distance by two (even though the Hamming distance increases by just one). This issue is resolved by extending the space to labelings Ω , subsequently there is an intermediate improper coloring so that X_{t+1} and Y_{t+1} are distance two apart.

The definition of the Markov chain described in Section 3 is identical, we simply extend the state space. A set $S_\sigma(v, c)$ is still defined as the set of vertices reachable from v by a $(\sigma(v), c)$ alternating path. And hence the notion of a cluster is still the same as before. Note, that while the chain restricted to proper colorings is symmetric, this is not necessarily true for improper colorings. All of the bounds stated in Section 4 hold for possibly improper colorings $X_t, Y_t \in \Omega^*$.

Consider a labeling $X_0 \in \Omega^* \setminus \Omega$; note, X_0 is not a proper coloring since $X_0 \notin \Omega$. For $k \geq + 2$, there is a sequence of transitions with non-zero probability (e.g., a sequence of Glauber moves as in the proof of irreducibility) so that it reaches a proper coloring, i.e., $X_t \in \Omega$ for some $t \geq 0$. Moreover, for any proper coloring $X_t \in \Omega$ then it stays on proper colorings, i.e., $X_s \in \Omega$ for all $s \geq t$, as the process does not introduce improper colorings. Therefore, states in Ω are the only ones which have positive probability in the stationary distribution, and hence the stationary distribution of the chain is uniform over the set of proper colorings Ω , even though the state space is all labelings Ω^* .

If the initial state is restricted to Ω , i.e., X_0 is a proper coloring, then the chain is identical to the process defined in Section 2.2. Furthermore, since the mixing time is defined from the worst initial state then a mixing time upper bound for the chain defined on Ω^* implies the same bound on the mixing time for the chain from Section 2.2 defined only on Ω .

We now have all the tools necessary to prove Theorem 1.

Proof of Theorem 1. First consider the available colors for v^* . Note that, since there is an extra available color for every time a color repeats in $N(v^*)$, we have

$$|A(v^*)| \geq k - d(v^*) + \sum_{c \in [k]: d(v^*) \geq 2} (d_c(v^*) - 1), \quad (2)$$

where $d(v^*) = \sum_{c \in [k]} d_c(v^*) = |N(v^*)|$ is the degree of v^* .

Now by combining Lemmas 11, 13, and 15 we can complete the proof of the theorem:

$$\begin{aligned} & \mathbf{E}(H(X_{t+1}, Y_{t+1}) \mid X_t, Y_t) \\ & \leq 1 - \mathbf{Pr}(X_{t+1}(v^*) = Y_{t+1}(v^*)) + \sum_{c: d(v^*) > 0} (\mathbf{E}(H(X_{t+1}, Y_{t+1}) \mid \mathcal{F}_c) - 1) \\ & \quad + \sum_{T: \text{dist}(v^*, T) = 2} |T| \mathbf{Pr}(X_{t+1}(T) \neq Y_{t+1}(T)) \\ & \leq 1 - |A(v^*)|(1 - \varepsilon/5) + (1 + \varepsilon/5) \sum_{c \in [k]: d(v^*) > 0} \left(\frac{11}{6} d_c(v^*) - 1 \right) + \varepsilon/5 \quad (3) \end{aligned}$$

13:14 Improved Distributed Algorithms for Random Colorings

$$\begin{aligned}
&= 1 - |A(v^*)|(1 - \varepsilon/5) + (1 - \varepsilon/5) \sum_{c \in [k]: d_c(v^*) > 0} \left(\frac{11}{6} d_c(v^*) - 1 \right) \\
&\quad + (2\varepsilon/5) \sum_{c \in [k]: d_c(v^*) > 0} \left(\frac{11}{6} d_c(v^*) - 1 \right) + \varepsilon/5 \\
&\leq 1 - |A(v^*)|(1 - \varepsilon/5) + (1 - \varepsilon/5) \sum_{c \in [k]: d_c(v^*) > 0} \left(\frac{11}{6} d_c(v^*) - 1 \right) \\
&\quad + \varepsilon (22/30) + \varepsilon/5
\end{aligned}$$

where Equation (3) follows from Lemmas 11, 13, and 15. Then using Equation (2) we get

$$\mathbf{E}(H(X_{t+1}, Y_{t+1}) | X_t, Y_t) \leq 1 - (1 - \varepsilon/5) \left[k - \frac{11}{6} d(v^*) \right] + \varepsilon (28/30) \quad (4)$$

$$\leq 1 - (1 - \varepsilon/5)\varepsilon + \varepsilon (28/30) \quad (5)$$

$$\leq 1 - \varepsilon (29/30) + \varepsilon (28/30) \quad (6)$$

$$\leq 1 - \varepsilon^2/60000,$$

where Equation (5) uses that $k \geq (1 + \varepsilon)\frac{11}{6}$ and Equation (6) uses that $(1 - \varepsilon/5) \geq 29/30$ when $\varepsilon \leq 1/6$. Note, the case when $\varepsilon > 1/6$ and $k > 2$ is handled by [14, 11] or can be handled in our analysis by setting ε in terms of $1/k$ instead of $1/k$. Finally, applying the path coupling Theorem 5 we obtain mixing time $O(\log n)$. Moreover, we obtain mixing time within total variation distance $\leq \delta$, for any $\delta > 0$, in time $O(\log(n/\delta))$. \blacktriangleleft

5 Coupling Analysis for Neighboring Clusters

We now prove Lemma 15. Before delving into the proof we state several key properties of the settings for the flip probabilities in Equation (1):

1. For all integer $i, j \geq 1$, $i(f_i - f_{i+1}) + (j-1)(f_j - f_{j+1}) \leq 5/6$.
2. For all integer $i \geq 1$, $2(i-1)f_i + f_{2i+1} \leq 2/3$.

Fix a color $c \in [k]$ where $d_c(v^*) > 0$; we will consider two cases: $d_c(v^*) = 1$ or $d_c(v^*) \geq 2$.

5.1 Flippable Difference

Lemma 16. *For any cluster C , $\Pr(C \text{ is active and not flippable}) \leq \varepsilon/250$.*

Proof. The cluster C is active with probability ε . Assuming C is active, there are two ways that C is not flippable, either (i) an overlapping cluster, or (ii) a neighboring cluster that shares a color with C . For case (i), since $|C| \leq 6$ and each vertex is in k clusters, then the probability of a cluster that overlaps C also being active is $\leq 6k\varepsilon$. For case (ii), there are ≤ 6 neighboring vertices, each has ≤ 2 clusters that share a color, and hence the probability of case (ii) is $\leq 12\varepsilon$. Combining the above calculations we have the following:

$$\Pr(C \text{ is active and not flippable}) \leq (6k\varepsilon + 12\varepsilon) \leq \frac{18k\varepsilon}{5000k} < \varepsilon/250. \quad \blacktriangleleft$$

5.2 Color Appears Once

Suppose $d_c(v^*) = 1$. Let $w \in N(v^*)$ be the unique neighbor where $X_t(w) = Y_t(w) = c$, and let $R := X_t(v^*)$ and $B := Y_t(v^*)$. We are coupling the clusters in the set $\mathcal{S}_{X_t}(c)$ with $\mathcal{S}_{Y_t}(c)$, and since $d_c(v^*) = 1$ these sets are the following:

$$\mathcal{S}_{X_t}(c) = \{S_{X_t}(w, R), S_{X_t}(w, B)\} \text{ and } \mathcal{S}_{Y_t}(c) = \{S_{Y_t}(w, R), S_{Y_t}(w, B)\}.$$

Observe $S_{X_t}(w, R) = S_{Y_t}(w, R) \cup \{v^*\}$ and $S_{Y_t}(w, B) = S_{X_t}(w, B) \cup \{v^*\}$. Let $i := |S_{Y_t}(w, R)|$ (hence, $|S_{X_t}(w, B)| = i + 1$), and let $j := |S_{X_t}(w, B)|$ ($|S_{Y_t}(w, B)| = j + 1$). Note, $i, j \geq 1$.

We couple the clusters in the following manner. With probability $\frac{1}{2}$, cluster $S_{X_t}(w, R)$ is active in X_t and $S_{Y_t}(w, R)$ is active in Y_t , while with probability $\frac{1}{2}$ both of these clusters are inactive. Similarly, with probability $\frac{1}{2}$ then both: cluster $S_{X_t}(w, B)$ is active in X_t and $S_{Y_t}(w, B)$ is active in Y_t .

Suppose that $S_{X_t}(w, R)$ and $S_{Y_t}(w, R)$ are both flippable. In this case we maximize the probability that we flip both clusters. Since $f_i \geq f_{i+1}$ then with probability f_{i+1} we flip both clusters $S_{X_t}(w, R)$ and $S_{Y_t}(w, R)$, assuming they were both flippable. Similarly, with probability f_{j+1} we flip both clusters $S_{X_t}(w, B)$ and $S_{Y_t}(w, B)$, assuming they were both flippable. Note in both of these cases where we flip both $S_{X_t}(w, R)$ and $S_{Y_t}(w, R)$ or we flip both $S_{X_t}(w, B)$ and $S_{Y_t}(w, B)$, then the Hamming distance does not change as the chains only differ at v^* after the coupled update.

Suppose that all 4 clusters were flippable. (Recall an active cluster S is flippable if there is no overlapping active cluster and no neighboring active cluster which shares one of the two colors with S .) Then after the above coupling of $S_{X_t}(w, R)$ with $S_{Y_t}(w, R)$, and $S_{X_t}(w, B)$ with $S_{Y_t}(w, B)$, there remains probability $f_j - f_{j+1}$ to flip $S_{X_t}(w, B)$, and probability $f_i - f_{i+1}$ to flip $S_{Y_t}(w, R)$. We maximally couple these remaining flips and hence with probability $\min\{f_i - f_{i+1}, f_j - f_{j+1}\}$ we couple the flips of clusters $S_{X_t}(w, B)$ and $S_{Y_t}(w, R)$. Note in this case where we flip both $S_{X_t}(w, B)$ and $S_{Y_t}(w, R)$ then the Hamming distance increases by $\leq (i + j - 1)$ since $S_{X_t}(w, B) \cap S_{Y_t}(w, R) \supseteq \{w\}$.

In the above coupling, we considered 3 coupled flips of cluster pairs $S_{X_t}(w, R), S_{Y_t}(w, R)$; $S_{X_t}(w, B), S_{Y_t}(w, B)$; and $S_{X_t}(w, B), S_{Y_t}(w, R)$. For each pair, it may occur that one of these clusters is flippable and the other is not flippable (due to a neighboring or overlapping cluster also being active). In that case we flip the flippable cluster by itself. In which case, the Hamming distance increases by at most 6 since the cluster is of size at most 6. By Lemma 16 the probability of this occurring for a specific cluster is at most $\varepsilon/250$, and since there are 3 pairs we have the effect is at most $36 \cdot \varepsilon/250 = \varepsilon/5$.

Let us assume without loss of generality that $i \leq j$ and hence $f_i - f_{i+1} \geq f_j - f_{j+1}$. Now we can simplify and summarize the effect of the above coupled flips that change the Hamming distance. Since the clusters are active with probability $\frac{1}{2}$, with probability $\leq (f_j - f_{j+1})$ we flip $S_{X_t}(w, B)$ and $S_{Y_t}(w, R)$ and then the Hamming distance increases by $\leq (i + j - 1)$. Moreover, with probability $f_i - f_{i+1} - (f_j - f_{j+1})$ we flip $S_{X_t}(w, B)$ by itself and the Hamming distance increases by i . Therefore, we have the following:

$$\begin{aligned} \mathbf{E}(H(X_{t+1}, Y_{t+1}) \mid \mathcal{F}_c) &\leq 1 + (i + j - 1)(f_j - f_{j+1}) + i((f_i - f_{i+1}) - (f_j - f_{j+1})) + \varepsilon/5 \\ &= 1 + i(f_i - f_{i+1}) + (j - 1)(f_j - f_{j+1}) + \varepsilon/5 \\ &\leq 1 + \left(\frac{5}{6} + \varepsilon/5\right) \end{aligned}$$

where the last inequality follows by Property 1.

5.3 Color Appears More Than Once

The analysis of the case when the color appears more than once, i.e., $d_c(v^*) > 1$, follows the same general approach as in Section 5.2 for the case $d_c(v^*) = 1$. In particular, we use the same high-level coupling as used by Vigoda [37] but in addition we use Lemma 16 to bound the probability that a cluster is flippable in one chain and the coupled cluster is not flippable in the other chain. We refer the reader to the appendix of the full version of this paper for details.

6 Proof of Theorem 2: Mixing below 11/6

Sections 4 and 5 present the proof of Theorem 1 which establishes $O(\log n)$ mixing time of the distributed flip dynamics when $k > (11/6 + \varepsilon)$ for all $\varepsilon > 0$. The improved result for $k > (11/6 - \varepsilon^*)$ for a fixed $\varepsilon^* > 0$ as stated in Theorem 2 is proved in the appendix of the full version of this paper.

The proof of Theorem 2 uses the new metric introduced in [5], which is a weighted Hamming distance. In particular, in [5] they identify the configurations on the local neighborhood of the disagree vertex v^* for which the coupling analysis is tight, these are referred to as extremal configurations. Hence, for a pair of configurations X_t, Y_t which differ at a single vertex v^* , let γ denote the fraction of neighbors of v^* in non-extremal configurations. Then, [5] defines a new weighted Hamming distance as $\mathcal{H}(X_t, Y_t) = 1 - \gamma\eta$ for an appropriately defined small constant $\eta > 0$.

Using this new weighting, [5] proves rapid mixing of the flip dynamics in the sequential setting for $k > (11/6 - \varepsilon^*)$. The challenge in their analysis is that one has to consider the effect of coupled flips which do not change the Hamming distance but simply change whether some neighbors of v^* are in extremal configurations.

To obtain Theorem 2 we combine the approaches of [5] with our analysis in Sections 4 and 5 of the effect of the distributed synchronization. However the analysis becomes considerably more complicated than in [5] because multiple clusters in the neighborhood of v^* can flip in a single step, this leads to many new cases where the new weighted Hamming distance can change. The detailed analysis is contained in the appendix of the full version of this paper.

References

- 1 Leonid Barenboim, Michael Elkin, and Uri Goldenberg. Locally-iterative distributed (± 1)-coloring and applications. *Journal of the ACM*, 69(1):5:1–5:26, 2021. doi:10.1145/3486625.
- 2 Leonid Barenboim, Michael Elkin, Seth Pettie, and Johannes Schneider. The locality of distributed symmetry breaking. *Journal of the ACM*, 63(3), 2016. doi:10.1145/2903137.
- 3 Anton Bernshteyn, Tyler Brazelton, Ruijia Cao, and Akum Kang. Counting colorings of triangle-free graphs. *Journal of Combinatorial Theory, Series B*, 161:86–108, 2023. doi:10.1016/J.JCTB.2023.02.004.
- 4 Russ Bubley and Martin E. Dyer. Path coupling: a technique for proving rapid mixing in Markov chains. In *Proceedings of the 38th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 223–231, 1997. doi:10.1109/SFCS.1997.646111.
- 5 Sitan Chen, Michelle Delcourt, Ankur Moitra, Guillem Perarnau, and Luke Postle. Improved bounds for randomly sampling colorings via linear programming. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2216–2234, 2019. doi:10.1137/1.9781611975482.134.
- 6 Zongchen Chen, Kuikui Liu, Nitya Mani, and Ankur Moitra. Strong spatial mixing for colorings on trees and its algorithmic applications. In *Proceedings of the 64th IEEE Symposium on Foundations of Computer Science (FOCS)*, 2023.

- 7 Constantinos Daskalakis, Nishanth Dikkala, and Siddhartha Jayanti. Hogwild!-Gibbs can be panaccurate. In *Proceedings of the 31st Advances in Neural Information Processing Systems (NeurIPS)*, pages 32–41, 2018. URL: <https://proceedings.neurips.cc/paper/2018/hash/a5bfc9e07964f8dddeb95fc584cd965d-Abstract.html>.
- 8 Christopher De Sa, Kunle Olukotun, and Christopher Ré. Ensuring rapid mixing and low bias for asynchronous Gibbs sampling. In *Proceedings of the 33rd International Conference on Machine Learning (ICML)*, pages 1567–1576, 2016. URL: <http://proceedings.mlr.press/v48/sa16.html>.
- 9 Martin Dyer, Alan Frieze, Thomas P. Hayes, and Eric Vigoda. Randomly coloring constant degree graphs. *Random Structures & Algorithms*, 43(2):181–200, 2013. doi:10.1002/RSA.20451.
- 10 Martin Dyer and Catherine Greenhill. Random walks on combinatorial objects. *Surveys in Combinatorics*, pages 101–136, 1999.
- 11 Weiming Feng, Thomas P. Hayes, and Yitong Yin. Distributed symmetry breaking in sampling (optimal distributed randomly coloring with fewer colors). *arXiv preprint*, 2018. arXiv:1802.06953.
- 12 Weiming Feng, Yuxin Sun, and Yitong Yin. What can be sampled locally? In *Proceedings of the International Symposium on Principles of Distributed Computing (PODC)*, pages 121–130, 2017. doi:10.1145/3087801.3087815.
- 13 Weiming Feng and Yitong Yin. On local distributed sampling and counting. In *Proceedings of the 37th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 189–198, 2018. doi:10.1145/3212734.3212757.
- 14 Manuela Fischer and Mohsen Ghaffari. A simple parallel and distributed sampling technique: Local Glauber dynamics. In *32nd International Symposium on Distributed Computing (DISC)*, volume 121, pages 26:1–26:11, 2018. doi:10.4230/LIPICS.DISC.2018.26.
- 15 Xinyu Fu, Yitong Yin, and Chaodong Zheng. Locally-iterative ($\epsilon + 1$)-coloring in sublinear (in ϵ) rounds. *arXiv preprint*, 2023. arXiv:2207.14458.
- 16 Marc Fuchs and Fabian Kuhn. Brief announcement: List defective colorings: Distributed algorithms and applications. In *Proceedings of the 35th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, 2023. doi:10.1145/3558481.3591319.
- 17 Joseph Gonzales, Yucheng Low, Arthur Gretton, and Carlos Guestrin. Parallel Gibbs sampling: From colored fields to thin junction trees. In *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 15, pages 324–332, 2011. URL: <http://proceedings.mlr.press/v15/gonzalez11a/gonzalez11a.pdf>.
- 18 G.E. Hinton. Training products of experts by minimizing contrastive divergence. *Neural computation*, 14(8):1771–1800, 2002. doi:10.1162/089976602760128018.
- 19 G.E. Hinton, S. Osindero, and Y-W. Teh. A fast learning algorithm for deep belief nets. *Neural computation*, 18(7):1527–1554, 2006. doi:10.1162/NECO.2006.18.7.1527.
- 20 G.E. Hinton and R.R. Salakhutdinov. Replicated softmax: an undirected topic model. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 1607–1614, 2009. URL: <https://proceedings.neurips.cc/paper/2009/hash/31839b036f63806cba3f47b93af8ccb5-Abstract.html>.
- 21 Mark 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. doi:10.1002/RSA.3240070205.
- 22 Michael I. Jordan, Jason D. Lee, and Yun Yang. Communication-efficient distributed statistical inference. *Journal of the American Statistical Association*, 2018. doi:10.1080/01621459.2018.1429274.
- 23 Christos Karras, Aristeidis Karras, Dimitrios Tsolis, Konstantinos C. Giotopoulos, and Spyros Sioutas. Distributed gibbs sampling and lda modelling for large scale big data management on pyspark. In *2022 7th South-East Europe Design Automation, Computer Engineering, Computer Networks and Social Media Conference (SEEDA-CECNSM)*, pages 1–8, 2022. doi:10.1109/SEEDA-CECNSM57760.2022.9932990.

24 Daphne Koller and Nir Friedman. *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, 2009. URL: <http://mitpress.mit.edu/catalog/item/default.asp?ttype=2&tid=11886>.

25 Fabian Kuhn. Weak graph colorings: Distributed algorithms and applications. In *Proceedings of the 21st ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 138–144, 2009. doi:10.1145/1583991.1584032.

26 Nathan Linial. Locality in distributed graph algorithms. *SIAM Journal on Computing*, 21(1):193–201, 1992. doi:10.1137/0221015.

27 Hongyang Liu and Yitong Yin. Simple parallel algorithms for single-site dynamics. In *54th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 1431–1444, 2022. doi:10.1145/3519935.3519999.

28 Kevin P. Murphy. *Machine learning: a probabilistic perspective*. MIT Press, 2012.

29 Moni Naor and Larry Stockmeyer. What can be computed locally? In *Proceedings of the 25th Annual ACM Symposium on Theory of Computing (STOC)*, pages 184–193, 1993. doi:10.1145/167088.167149.

30 S. Osindero and G.E. Hinton. Modeling image patches with a directed hierarchy of Markov random fields. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 1121–1128, 2008.

31 R. Salakhutdinov and G.E. Hinton. Deep Boltzmann machines. In *Proceedings of the 12th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 448–455, 2009.

32 R. Salakhutdinov, A. Mnih, and G.E. Hinton. Restricted Boltzmann machines for collaborative filtering. In *Proceedings of the 24th International Conference on Machine Learning (ICML)*, pages 791–798, 2007.

33 Jesús Salas and Alan 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. doi:10.1007/BF02199113.

34 Johannes Schneider and Roger Wattenhofer. A new technique for distributed symmetry breaking. In *Proceedings of the 29th ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing (PODC)*, pages 257–266, 2010. doi:10.1145/1835698.1835760.

35 Alexander Smola and Shravan Narayananmurthy. An architecture for parallel topic models. In *Proceedings of the VLDB Endowment (PVLDB)*, 2010. doi:10.14778/1920841.1920931.

36 Alexander Terenin, Daniel Simpson, and David Draper. Asynchronous Gibbs sampling. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 144–154, 2020. URL: <http://proceedings.mlr.press/v108/terenin20a.html>.

37 Eric Vigoda. Improved bounds for sampling colorings. *Journal of Mathematical Physics*, 41(3):1555–1569, 2000. doi:10.1063/1.533196.

38 Feng Yan, Ningyi Xu, and Yuan Qi. Parallel inference for latent dirichlet allocation on graphics processing units. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 22, 2009. URL: <https://proceedings.neurips.cc/paper/2009/hash/ed265bc903a5a097f61d3ec064d96d2e-Abstract.html>.

39 Yuan Yang, Jianfei Chen, and Jun Zhu. Distributing the stochastic gradient sampler for large-scale LDA. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 1975–1984, 2016. doi:10.1145/2939672.2939821.

40 Ce Zhang and Christopher Ré. Dimmwitted: A study of main-memory statistical analytics. In *Proceedings of the VLDB Endowment (PVLDB)*, 2014. doi:10.14778/2732977.2733001.