Algorithms to Approximately Count and Sample Conforming Colorings of Graphs

Sarah Miracle [1] and Dana Randall [2]

College of Computing
Georgia Institute of Technology
Atlanta, GA 30332-0765 USA

Abstract

Conforming colorings naturally generalize many graph theory structures, including independent sets, vertex colorings, list colorings, H-colorings and adapted colorings. Given a multigraph G and a function F that assigns a forbidden ordered pair of colors to each edge e, we say a coloring C of the vertices is *conforming to* F if, for all $e = (u, v)$, $F(e) \neq (C(u), C(v))$. We consider Markov chains on the set of conforming colorings and provide some general conditions for when they can be used to construct efficient Monte Carlo algorithms for sampling and counting.

1 Introduction

Adapted colorings [3] have been introduced recently as a generalization of many well-studied discrete models, including independent sets, vertex colorings, list colorings, and H-colorings. We consider "conforming colorings" as a further generalization. Let $G = (V, E)$ be a (multi)graph and for $k \in \mathbb{Z}^+$, let $[k] = \{1, \dots, k\}$ be a set of colors. We are given a set of edge constraints $F : E \to [k] \times [k]$ describing forbidden ordered pairs of colors on the endpoints of each edge, and we are interested in the set of vertex colorings satisfying these constraints. We say that a coloring $C : V \to [k]$ is a *conforming coloring* if, for each edge $e = (u, v)$, we have $F(e) \neq (C(u), C(v))$. Let $\Omega = \Omega(G, F, k)$ be the set of all conforming colorings of G with forbidden pairs F and k colors. A simple application is resource allocation, where vertices represent jobs and edge constraints capture incompatible scheduling assignments. We focus here on approximately counting and sampling conforming and adapted colorings.

[1] Supported by a DOE Office of Science Graduate Fellowship and NSF CCF-1219020.
[2] Supported in part by NSF CCF-1219020.

The connections between conforming colorings and many standard graph theoretic objects is straight-forward. For example, when $k = 2$ and $F(e) = (1, 1)$ for all edges $e \in E$, each allowable vertex coloring is an independent set in G. Likewise, given a graph G, form a multigraph G' where each edge is replaced with k parallel edges, each labeled with distinct (i, i) for $1 \leq i \leq k$, then the conforming colorings of G' are the proper colorings of G. We can formulate weighted versions of these models similarly, including the class of H-colorings, or homomorphisms from a graph G to H preserving adjacencies. *Adapted colorings* are a special case of conforming colorings that has garnered a lot of interest recently where we are given an edge coloring $C : E \rightarrow [k]$ and are looking for a vertex coloring $C' : V \rightarrow [k]$ such that there is no edge $e = (u, v)$ with $C(e) = C'(u) = C'(v)$ [3].

There has been extensive work trying to approximately count graph structures using Monte Carlo approaches. The main ingredient is designing a Markov chain for sampling configurations that is rapidly mixing. For example, for k-colorings, local chains that modify a small number of vertices in each move converge quickly to equilibrium if there are enough colors compared to the maximum degree of the graph [2], whereas even finding a single k-coloring is NP-complete for small degree. When local algorithms are slow, nonlocal variants can be more effective, but they are typically more challenging to analyze. One such example is the Wang-Swendsen-Kotecký chain for proper colorings that uses moves based on Kempe chains [8].

We show that there is a polynomial time algorithm for finding a conforming coloring when the number of colors is at least $\max(3, \Delta)$, where Δ is the maximum degree in G, including multi-edges and self-loops. Moreover, we show that the local Markov chain $\mathcal{M}_L$ that recolors one vertex at a time connects the state space of conforming colorings (Ω), is rapidly mixing and there is a FPRAS (fully polynomial randomized approximation scheme) for approximately counting. For adapted colorings, we prove a stronger result requiring only that $k \geq \max(\Delta_m, 3)$, where Δ_m is related to Δ except at most two parallel edges between any two vertices are counted toward the degree.

When we have 2 colors, the local chain does not always connect the state space so we introduce a new component chain $\mathcal{M}_C$ that reverses colors on (possibly) large components that are predetermined based on the structure of G and F. We provide conditions under which we can find conforming 2-colorings efficiently and show $\mathcal{M}_C$ connects the state space, is rapidly mixing, and there is an FPRAS for counting. Last, we provide examples for which both chains can be slow.

The mixing results build on ideas used to show fast and slow mixing in the context of colorings and independent sets; however, the proofs require careful fine-tuning to fit the more general setting of conforming colorings and to prove the bounds we achieve here. Moreover, unlike sampling colorings and independent sets where we typically restrict to graphs on which connecting the state space is trivial, in this more general setting establishing ergodicity for the two chains has proven considerably more challenging.

2 The Local Markov Chain $\mathcal{M}_L$

We begin exploring how to find a conforming coloring, if one exists. First, we prove that if $k \geq \max(\Delta, 3)$ and Ω is not degenerate, a conforming coloring exists. We define a *degenerate state space* recursively as follows. Let a *flower* be a vertex v with k self-loops, each with a distinct color; note there is no conforming coloring of a flower. A vertex is *color-fixed* if it has exactly $k-1$ self-loops each with a distinct color or k self-loops with exactly one color repeated. A color-fixed vertex has exactly one valid color in every conforming coloring. For each color-fixed vertex v, color v with it's only valid color c and remove v from G. Next, handle each edge constraint $e = (u, v)$, $u \neq v$ as follows. If $c \neq F_v(e)$ remove e. Otherwise, if $c = F_v(e)$ add a self-loop colored $F_u(e)$ to u. Continue this process until either a flower is found at which point G is degenerate or there are no color-fixed vertices and G is not degenerate. The following result is proved constructively using graph theoretic techniques by giving an algorithm that iteratively colors vertices.

Theorem 2.1 *Given a graph G with n vertices, $k \geq \max(\Delta, 3)$ and edge k-coloring F such that Ω is not degenerate, there exists a conforming k-coloring of G and we give an $O(\Delta n^2)$ algorithm for finding one.*

We now consider the local Markov chain $\mathcal{M}_L$ that, at each step, selects a vertex v and a color c uniformly at random and colors vertex v with color c if this results in a valid conforming coloring. First, we show that when $k \geq \max(\Delta, 3)$, the Markov chain $\mathcal{M}_L$ is ergodic (i.e., irreducible and aperiodic). We prove that for each $\sigma, \alpha \in \Omega$, there is a path from σ to α using only transitions of $\mathcal{M}_L$, thus implying the connectivity of $\mathcal{M}_L$. The primary challenge is that there might not be a path between σ and α that only modifies vertices in the symmetric difference. We prove the following theorem.

Theorem 2.2 *For any graph G, $k \geq \max(\Delta, 3)$ and edge k-coloring F, $\mathcal{M}_L$ connects $\Omega(G, F, k)$.*

Next, we study whether $\mathcal{M}_L$ generates random conforming colorings efficiently. Let $(\mathcal{P}, \Omega, \pi)$ be a Markov chain with transition matrix $\mathcal{P}$, stationary distribution π and state space Ω. For all $\epsilon > 0$, the *mixing time* $\tau(\epsilon)$ of $\mathcal{M}$ is defined as $\tau(\epsilon) = \min\{t : \max_{x\in\Omega} \frac{1}{2}\sum_{y\in\Omega} |\mathcal{P}^t(x, y) - \pi(y)| \leq \epsilon, \forall t' \geq t\}$, where

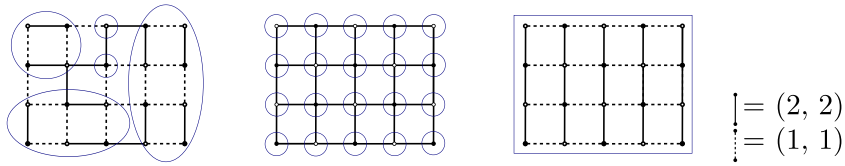

Fig. 1. The color-implied components associated with each edge coloring are circled.

$\mathcal{P}^t(x, y)$ is the t-step transition probability. We say that a Markov chain is *rapidly mixing* if the mixing time is bounded above by a polynomial in n and $\log(\epsilon^{-1})$. Our fast mixing proof uses a detailed application of the path coupling technique due to Dyer and Greenhill [1]. In order to prove a stronger result in the adapted coloring setting we use a more sophisticated coupling introduced by Jerrum [4].

Finally, we can use $\mathcal{M}_L$ to approximately count conforming colorings. To do this we design a FPRAS which, in our context, is a randomized algorithm that given a graph G with n vertices, edge coloring F and error parameter $0 < \epsilon \leq 1$, produces a number N such that $\mathcal{P}[(1+\epsilon)N \leq A(G, F) \leq (1+\epsilon)N] \geq \frac{3}{4}$, where $A(G, F)$ is the number of colorings of G conforming to F and runs in time polynomial in n and ϵ^{-1}. Our proof uses similar techniques to [5].

Theorem 2.3 *Given a graph G, $k \geq \max(\Delta, 3)$ and edge k-coloring F the mixing time of $\mathcal{M}_L$ on $\Omega(G, F, k)$ satisfies $\tau(\epsilon) \leq \lceil ekn^5 \rceil \lceil \ln \epsilon^{-1} \rceil$ and there exists a FPRAS for counting the number of k-colorings conforming to F.*

3 The Chain $\mathcal{M}_C$ and Conforming 2-Colorings

The case when $k = 2$ is interesting because it generalizes independent sets and $\mathcal{M}_L$ typically is not ergodic. Consider, for example, the Cartesian lattice where horizontal edges are colored $(1, 1)$ and vertical edges are colored $(2, 2)$; there are two conforming colorings corresponding to the two proper 2-colorings. To handle the case $k = 2$, we introduce a non-local chain $\mathcal{M}_C$ based on "color-implied" components which we show is ergodic. Some moves of $\mathcal{M}_C$ are Kempe chain moves analogous to those in [8], but in general they are more complicated. Additionally, under conditions based on the degrees of the color-implied components we can find a conforming coloring, $\mathcal{M}_C$ is rapidly mixing and we have a FPRAS. On the other hand, we show that there are settings when $\mathcal{M}_C$ requires exponential time with $\Delta = 4$.

A color-implied component is a connected set of vertices where coloring any vertex in the component implies a unique coloring of the remaining vertices in the component; thus, each component has at most two valid conforming colorings. For $b \in \{1, 2\}$, we define a path $P = v_1, v_2, \ldots v_x$ to be a *b-alternating* *path* from v_1 to v_x if the following two conditions hold: $F_{v_1}(v_1, v_2) = b$ and for all $1 \leq i \leq x-2$, $F_{v_{i+1}}(v_i, v_{i+1}) \neq F_{v_{i+1}}(v_{i+1}, v_{i+2})$. Define two vertices u and v to be *color-implied* if there is a 1-alternating path and a 2-alternating path from u to v or $u = v$. We show that color-implied is an equivalence relation, thus determining a partition of the vertices of G into connected components $C_1, C_2, \ldots C_s$ (e.g. Fig. 1). Using a modified version of DFS we can find this partition in polynomial time. Let G' be the graph whose vertices are the components $C_1, C_2, \ldots, C_s$ and there is an edge $(C_i, C_j) \in G'$ if there exists $(v_i, v_j) \in G : v_i \in C_i, v_j \in C_j$. For each component C_i, let $\rho(C_i)$ and $\overline{\rho}(C_i)$ be the two conforming colorings of C_i. Next, consider any two components C_i and C_j. Let $\rho(u)$ be the color of vertex u in $\rho(C_i)$. We show that for all the edges $(u, v) : u \in C_i, v \in C_j$ either $F_u(u, v) = \rho(u)$ or $F_u(u, v) = \overline{\rho}(u)$. If there exists an edge $e = (C_i, C_j) \in G'$ we will define $F_{C_i}(e) = \rho(C_i)$ if for all edges $(u, v) : u \in C_i, v \in C_j$, $F_u(u, v) = \rho(u)$ and otherwise $F_{C_i}(e) = \overline{\rho}(C_i)$. We now define the Markov chain $\mathcal{M}_C$ which connects the state space $\Omega(G, F, 2)$. Starting at any initial conforming coloring, at each step select an integer $i \in 1, 2, \ldots, s$ uniformly at random. With probability $1/2$ color C_i, $\rho(C_i)$ if this is valid, with probability $1/2$ color C_i, $\overline{\rho}(C_i)$ if this is valid, otherwise do nothing.

Theorem 3.1 *For any graph G and edge 2-coloring F, the Markov chain $\mathcal{M}_C$ connects $\Omega(G, F, 2)$.*

We show that $\mathcal{M}_C$ is rapidly mixing when every vertex v of the auxiliary graph $G' = G'(G, F)$ has $d(v) \leq 2$ or $d(v) \leq 4$ and v is *monochromatic*. We say a vertex $v \in G'$ is *monochromatic* if for any two edges e_1, e_2 adjacent to v, $F_v(e_1) = F_v(e_2)$. In the adapted setting, if C_i is a single vertex then this corresponds to having all adjacent edges colored the same. We use path coupling to prove that a related chain $\mathcal{M}_E$ mixes rapidly and then use the comparison technique (see, [7]) to relate the mixing time of $\mathcal{M}_E$ to the mixing time of $\mathcal{M}_C$. The chain $\mathcal{M}_E$ is a generalization of the edge chain introduced by Luby and Vigoda [6]. Under these same conditions, we give a polynomial time algorithm for finding a conforming coloring and a FPRAS for approximately counting the number of conforming colorings by showing the model is self-reducible and appealing to [5].

Theorem 3.2 *Given a graph G with n vertices and edge 2-coloring F with auxiliary graph $G' = G'(G, F)$ such that for every vertex $v \in G'$ either $d(v) \leq 2$ or $d(v) \leq 4$ and v is monochromatic, we give an $O(n^3)$ algorithm for finding a conforming 2-coloring (if one exists), the mixing time of $\mathcal{M}_C$ on $\Omega(G, F, 2)$ satisfies $\tau(\epsilon) = O(n^4)$ and there exists a FPRAS for counting $|\Omega(G, F, 2)|$.*

Finally, we show there is a graph (see Fig. 2) on which $\mathcal{M}_C$ requires exponential time by demonstrating a bottleneck in the state space.

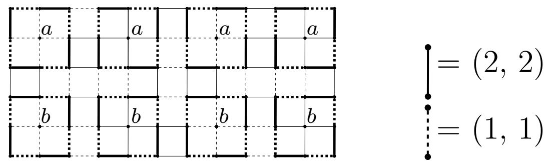

Fig. 2. An edge coloring for which the Markov chain $\mathcal{M}_C$ mixes slowly.

Theorem 3.3 *There exists a graph G with n vertices and $\Delta = 4$, a constant $c > 1$, and an edge 2-coloring F for which the mixing time of $\mathcal{M}_C$ on $\Omega(G, F, 2)$ satisfies $\tau(\epsilon) = \Omega(c^n)$.*

Proof Sketch. The auxiliary graph G' for Fig. 2, has a large component C containing all vertices except for those labelled a and b, each of these forms a single vertex component. For each of the two colorings of C either the a vertices or the b vertices are free to change colors giving an exponential number of configurations. However there is only a single coloring of the a and b vertices that allows the color of C to change, creating a bottleneck. □

References

[1] M. Dyer and C. Greenhill. A more rapidly mixing Markov chain for graph colorings. *Random Structures & Algorithms,* **13:** 285–317, 1998.

[2] A. Frieze and E. Vigoda. A survey on the use of Markov chains to randomly sample colorings. In *Combinatorics, Complexity and Chance: A tribute to Dominic Welsh* (G. Grimmett, C. McDiarmid Eds.), 53–71, 2007.

[3] P. Hell and X. Zhu. Adaptable chromatic number of graphs. *European Journal of Combinatorics,* **29:** 912–921, 2008.

[4] M. Jerrum, A very simple algorithm for estimating the number of k-colourings of a low-degree graph. *Random Structures & Algorithms,* **7:** 157–165, 1995.

[5] M. Jerrum, L. Valiant and V. Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theoretical Computer Science,* **43:** 169–188, 1986.

[6] M. Luby and E. Vigoda. Fast convergence of the Glauber dynamics for sampling independent sets. *Random Structures & Algorithms,* **15:** 229–241, 1999.

[7] D. Randall and P. Tetali. Analyzing Glauber dynamics by comparison of Markov chains. *Journal of Mathematical Physics,* **41:** 1598–1615, 2000.

[8] J. Wang, R. Swendsen and R. Kotecký. Three-state antiferromagnetic Potts models: A Monte Carlo study. *Physical Review B,* **42:** 2465–2474, 1990.