Discrete Applied Mathematics 210 (2016) 133–149

Contents lists available at ScienceDirect

Discrete Applied Mathematics

journal homepage: www.elsevier.com/locate/dam

Algorithms to approximately count and sample conforming colorings of graphs

Sarah Miracle [a,*], Dana Randall [b]

[a] *Computer and Information Sciences, College of Arts and Sciences, University of St. Thomas, 2115 Summit Avenue Saint Paul, Minnesota 55105, USA*

[b] *2140 Klaus Advanced Computing Building, College of Computing, Georgia Institute of Technology, Atlanta, GA 30332-0765, USA*

ARTICLE INFO

Article history:
Received 3 January 2014
Received in revised form 26 November 2014
Accepted 2 May 2015
Available online 1 June 2015

Keywords:
Adapted colorings
Markov chains
Independent sets
H-colorings

ABSTRACT

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 $e = (u, v)$, $(C(u), C(v)) \neq F(e)$. Conforming colorings generalize many natural graph theoretic concepts, including independent sets, vertex colorings, list colorings, H-colorings and adapted colorings and consequently there are known complexity barriers to sampling and counting. We introduce natural Markov chains on the set of conforming colorings and provide general conditions for when they can be used to design efficient Monte Carlo algorithms for sampling and approximate counting.

1. Introduction

Adapted colorings [13,14] have been studied as a natural generalization of many well-studied discrete models, including independent sets, colorings, list colorings, and H-colorings [9,12]. Here we introduce "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 vertex 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. Conforming colorings formalize constraints in many applications including resource allocation, where vertices represent jobs and edge constraints capture incompatible scheduling assignments. We focus on approximately counting and sampling conforming and adapted colorings.

The connection 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 $e \in E$, then in each conforming coloring the vertices colored 1 form an independent set. 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 $1 \le i \le k$, then the conforming colorings of G' are exactly the proper k-colorings of G. We also can formulate weighted versions of these standard models. For example in the hard-core lattice gas model, we are given an activity λ which represents the fugacity of the gas and are interested in sampling independent sets from a weighted distribution where each independent set I occurs with probability $\lambda^{|I|}/Z$, where $|I|$ is the size of the independent set and $Z = \sum_J \lambda^{|J|}$ is the normalizing constant known as the partition function. If $\lambda > 1$, then dense independent sets are more likely while if $\lambda < 1$, sparse independent sets are favored. To formulate this model using conforming colorings, we let $k > 2$

* Corresponding author.
E-mail addresses: sarah.miracle@stthomas.edu (S. Miracle), randall@cc.gatech.edu (D. Randall).

http://dx.doi.org/10.1016/j.dam.2015.05.003

and $F(e) = (1, 1)$ for all edges. Again, in each conforming coloring the vertices colored 1 form an independent set. However, now for each independent set I there are $(n - |I|)^{k-1}$ ways to color the remaining vertices. When sampling uniformly from the set of conforming colorings, each independent set I thus occurs with probability $\lambda^{|I|}/Z$, where $\lambda = 1/(k - 1)$. Valid configurations are weighted independent sets with low fugacity (favoring sparse independent sets). Likewise, we can construct instances of conforming colorings that correspond to independent sets with high fugacity (favoring dense sets). Given a graph G, we construct a new graph G_r with $r|V|$ vertices and $r^2|E|$ edges as follows. We replace each vertex with r copies and replace each edge (u, v) in G with the complete bipartite graph $K_{r,r}$ connecting each copy of u in G_r with each copy of v. We then set $k = 2$ with colors 1 and 2 and let $F(e) = (1, 1)$ for all edges e in G_r. Notice that if $(u, v) \in E$ and any of the copies of u are colored 1 in G_r, then all of the copies of v must be colored 2. Thus, the 1 vertices in G_r correspond to an independent set in G (where we include a vertex in the independent set if at least one copy in G_r is colored 1) and each independent set has weight $\lambda^{|I|}/Z$ where $\lambda = 2^r - 1$ and $Z = \sum_J \lambda^{|J|}$ is the normalizing constant.

The same type of construction based on parallel edges allows us to capture the class of H-colorings, or homomorphisms from a graph G to H that preserve adjacency. H-colorings themselves have been studied as a natural generalization of many discrete models including colorings, independent sets and the Widom–Rowlinson model from statistical physics (see, i.e., [2,7,4,12]). Conforming colorings can model these problems by replacing each edge in G with parallel edges representing all of the edges of H that are *not* present, including self-loops. Conforming colorings are more general than the H-coloring problem because labeling the edges of a graph with forbidden colorings allows non-homogeneity in the coloring restrictions on neighboring vertices. Hell and Nešetřil [12] showed that if H does not have a loop, then deciding whether there exists an H-coloring is NP-complete, and it is in P otherwise. Dyer and Greenhill [9] proved that counting H-colorings exactly is ♯-P complete unless each component of H is trivial (i.e., a complete graph with loops or a complete bipartite graph), in which case the counting problem is also in P.

A special case of conforming colorings that has garnered interest recently is known as *adapted (or adaptable) colorings*. Given an edge coloring $C : E \to [k]$, a vertex coloring $C' : V \to [k]$ is adapted to C if there is no edge $e = (u, v)$ with $C(e) = C'(u) = C'(v)$. Hell and Zhu [14] introduced the *adaptable chromatic number* in 2008. Subsequently there have been a flurry of papers deriving bounds on the adaptable chromatic number in graphs and hypergraphs, the adaptable list chromatic number, and determining when a graph G is adaptably k-choosable, where each of these is a natural generalization of the standard graph theoretic notions (see, e.g., [10,13,14,17,19]). Recently, Cygan et al. [5] gave a polynomial time algorithm for finding an adapted 3-coloring given a fixed edge 3-coloring of a complete graph, resolving the so-called "stubborn problem" in the classification of constraint satisfaction problems [3].

In this paper we focus on the problems of approximately counting and randomly sampling conforming and adapted colorings of graphs. Previous research on approximation algorithms in the context of H-colorings yielded both positive and negative results [4,7], which is not surprising since they include independent sets and colorings as special cases. Our motivation for studying approximation algorithms in the more general class of conforming colorings is similar to that for H-colorings. Not only does the model capture many fundamental problems that are interesting in their own right, but such a study allows us to examine which approaches to randomized approximate counting can be extended to this more general class of problems.

1.1. Previous work

There has been extensive work trying to approximately count various graph structures using Monte Carlo approaches. The main ingredient is designing a Markov chain for sampling configurations that is rapidly mixing. For example, for independent sets we are given a fugacity λ and are interested in sampling independent sets I from the Gibbs distribution $\pi(I) = \lambda^{|I|}/Z$, where Z is the normalizing constant. Local chains that modify a small number of vertices in each move are known to be efficient on $\mathbb{Z}^2$ at fugacity $\lambda < 2.48$ [23] and inefficient when $\lambda > 5.3646$ [1]. Similarly, local chains on the space of k-colorings of graphs are efficient if there are enough colors compared to the maximum degree of the graph [11], whereas even finding a single k-coloring is NP-complete for small degree.

Dyer and Greenhill [9] proved that counting H-colorings exactly is ♯-P complete unless each component of H is trivial (i.e., a complete graph with loops or a complete bipartite graph), in which case the counting problem is also in P. Previous research on approximation algorithms in the context of H-colorings yielded both positive and negative results [4,7], which is not surprising since they include independent sets and colorings as special cases. Cooper, Dyer and Frieze [4] considered sampling algorithms for H-colorings and showed that for a large class of Markov chains convergence will be slow, particularly on graphs with high degree, although they give an efficient algorithm for sampling H-colorings when H is a tree with self-loops everywhere. Likewise, Borgs et al. [2] showed that for any finite, connected, non-trivial H, there are weights on the edges such that all quasi-local ergodic Markov chains will be slowly mixing on finite regions of Allison Martin-Attix $\mathbb{Z}^d$. Dyer, Goldberg and Jerrum [7] explored the connection between approximate counting and random sampling in the context of H-colorings. Sampling and counting are known to be equivalent for problems that are "self-reducible" [16], but H-colorings do not, in general, have this nice property. Dyer et al. succeed in one direction for H-colorings, namely showing that an efficient algorithm for random sampling can be used to design an FPRAS for approximate counting. The other direction remains open, as does the reduction for the general class of conforming colorings.

When local algorithms are slow, nonlocal variants can be more effective, but they are typically more challenging to analyze. Examples include the Swendsen Wang algorithm for the Ising and Potts models [22] and the Wang–Swendsen–Kotecký