The complexity of choosing an H-colouring (nearly) uniformly at random*

Leslie Ann Goldberg and Steven Kelk and Mike Paterson
Department of Computer Science
University of Warwick
Coventry, CV4 7AL, United Kingdom
E-mail: `(leslie|skelk|msp)@dcs.warwick.ac.uk`

August 23, 2001

Abstract

Cooper, Dyer and Frieze studied the problem of sampling H-colourings (nearly) uniformly at random. They considered the family of "cautious" ergodic Markov chains with uniform stationary distribution and showed that, for every fixed connected "nontrivial" graph H, every such chain mixes slowly. In this paper, we give a complexity result for the problem. Namely, we show that for any fixed graph H with no trivial components, there is unlikely to be any *Polynomial Almost Uniform Sampler* (PAUS) for H-colourings. We show that if there were a PAUS for the H-colouring problem, there would also be a PAUS for sampling independent sets in bipartite graphs and, by the self-reducibility of the latter problem, there would be a *Fully-Polynomial Randomised Approximation Scheme* (FPRAS) for #BIS — the problem of counting independent sets in bipartite graphs. Dyer, Goldberg, Greenhill and Jerrum have shown that #BIS is complete in a certain logically-defined complexity class. Thus, a PAUS for sampling H-colourings would give an FPRAS for the entire complexity class. In order to achieve our result we introduce the new notion of *sampling-preserving* reduction which seems to be more useful in certain settings than approximation-preserving reduction.

*Research Report 384, Department of Computer Science, University of Warwick, Coventry CV4 7AL, UK (August 2001). This work was partially supported by EPSRC grants GR/R44560/01 "Analysing Markov-chain based random sampling algorithms" and GR/M96940 "Sharper Analysis of Randomised Algorithms: a Computational Approach" and the IST Programme of the EU under contract numbers IST-1999-14036 (RAND-APX) and IST-1999-14186 (ALCOM-FT).

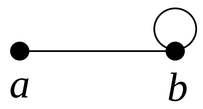

Figure 1: Homomorphisms from G to this graph are independent sets of G.

1 Introduction

Let $H = (V(H), E(H))$ be any fixed graph. An *H-colouring* of a graph $G = (V(G), E(G))$ is just a homomorphism from G to H: The vertices of H correspond to "colours" and the edges of H specify which colours may be adjacent. Thus, an H-colouring of G is a function C from $V(G)$ to $V(H)$ such that for every edge $(u, v) \in E(G)$, the corresponding edge $(C(u), C(v))$ is in $E(H)$. Informally, colours $C(u)$ and $C(v)$ are allowed to be adjacent in the colouring C of G because the edge $(C(u), C(v))$ is an edge of H.

Many combinatorial problems can be viewed as special cases of H-colouring. For example, if H is a k-clique with no self-loops then H-colourings of G correspond to proper k-colourings of G. (In such a colouring, k colours are available for colouring the vertices of G, but no colour may be adjacent to itself.) Here is another example. If H is the graph depicted in Figure 1 then H-colourings of G correspond to independent sets of G — vertices which are coloured "a" are in the independent set, and vertices which are coloured "b" are not. Several models from statistical physics are special cases of H-colouring including the Widom-Rowlinson model, the Beach model, and (for weighted H-colourings) the Potts model and the hard-core lattice gas model. See [2, 9] for details.

The complexity of H-colouring has been well-studied. Many papers considered the following problem: Given a fixed graph H, determine, for an input graph G, whether G has an H-colouring. Hell and Nešetřil [13] completely characterised the set of graphs for which this this problem is NP-complete. They observed that the problem is in P if H has a loop or is bipartite and they showed that it is NP-complete for any other fixed H. See [13] for references to earlier work on this question and [12] for extensions to the case in which the maximum degree of G is bounded. See [4, 5] for extensions to parameterised complexity.

Dyer and Greenhill [9] considered the problem of *counting* H-colourings. Intriguingly, they were able to completely characterise the graphs H for which this problem is #P-complete. A connected component of H is said to be "trivial" if it is a complete graph with all loops present or a complete bipartite graph with no loops present[1]. Dyer and Greenhill showed that counting H-colourings is #P-complete if H has a nontrivial component and that it is in P otherwise. They also extended their result to the case in which the maximum degree of G is bounded.

Other work has focussed on the complexity of *sampling* H-colourings (nearly) uni-

[1]Following the usual notation in the area, we will treat self-loops specially, so it makes sense to refer to bipartite graphs with or without loops. The loop-free single-vertex is viewed as a complete bipartite graph.

formly at random[2]. Positive results for particular graphs H (specifically for the case in which H-colourings are independent sets and for the case in which H-colourings are proper colourings) appear in works such as [8, 14, 17]. A negative result for the independent-set case appears in [6]. The first paper to study the complexity of sampling H-colourings in the general case was Cooper, Dyer and Frieze [3]. They focussed on connected graphs H for which the decision problem "Is there an H-colouring?" is in P, but the counting problem "How many H-colourings are there?" is #P-complete. They showed that for any such H, H-colourings cannot be sampled efficiently using "cautious" Markov chains, which are Markov chains which can change only a constant fraction of the colours of the vertices in a single step. In particular, the mixing time of all such chains is exponential in the number of vertices of G. They also give positive results for certain weighted cases, which are extended in [11]. In particular, [11] shows that for every fixed "dismantleable" H and every degree bound Δ, there are positive vertex-weights which can be assigned to the vertices of H so that weighted H-colourings can be sampled for degree-Δ graphs G. Borgs et al. [1] consider the problem of sampling H-colourings on rectangular subsets of the hypercubic lattice. They show that for every nontrivial connected H there is an assignments of weights to colours for which cautious chains are slowly mixing.

In this work, we study the complexity of sampling H-colourings. We show that if H has no trivial components then the problem of nearly-uniformly sampling H-colourings is intractable in a complexity-theoretic sense. In particular, we show that for any fixed H with no trivial components, there is unlikely to be any *Polynomial Almost Uniform Sampler* (PAUS) for H-colourings. We show that if there were a PAUS for the H-colouring problem, there would also be a PAUS for sampling independent sets in bipartite graphs and, by the self-reducibility of the latter problem, there would be a *Fully-Polynomial Randomised Approximation Scheme* (FPRAS) for #BIS — the problem of counting independent sets in bipartite graphs. Dyer, Goldberg, Greenhill and Jerrum have shown that #BIS is complete in a certain logically-defined subclass of #P which includes problems such as counting downsets in partial orders and counting satisfying assignments in "restricted Horn" CNF Boolean formulas. Thus, a PAUS for sampling H-colourings would give an FPRAS for the entire complexity class. In fact, our result holds even if the input G is restricted to be a connected bipartite graph.

In order to achieve our result we introduce the new notion of *sampling-preserving* reduction. The notion of approximation-preserving reduction (AP-reducibility) from [10] seems to be too demanding. In particular, since AP-reducibility is about *counting* (as opposed to sampling), an AP-reduction is not allowed to inflate the size of the set of structures by a factor which is difficult to compute. Sampling-preserving reductions allow this flexibility while achieving the same final result. The definition of sampling reduction (Section 2) is essentially many-one. Nevertheless the reductions get used in a "Turing reduction" way. In particular, our reduction from SampleBIS to Sample$H$-Col takes an instance of Sam-

[2]Some of this work has been motivated by the well-known connection between almost-uniform sampling and approximate counting [16, 7]. For some graphs H, it can be shown that the problem of approximately counting H-colourings is equivalent to the problem of sampling H-colourings (nearly) uniformly at random. See Section 8.

PLEBIS and constructs many SAMPLEH-COL instances. Since the resulting maps between H-colourings and independent sets are many-one, several reductions can be combined even though they may involve different amounts of inflation of the state space. See Sections 4 and 5 for details.

2 Definitions

The total variation distance between two distributions π and π' on a countable set Ω is given by

$$\mathrm{d_{TV}}(\pi, \pi') = \frac{1}{2} \sum_{\omega \in \Omega} |\pi(\omega) - \pi'(\omega)| = \max_{A \subseteq \omega} |\pi(A) - \pi'(A)|.$$

A sampling problem X maps each instance σ to a set of structures $X(\sigma)$. The goal is to produce a member of $X(\sigma)$ uniformly at random. The size of each structure in $X(\sigma)$ is at most a polynomial in $|\,\sigma\,|$. For a given graph H, the sampling problem SAMPLEH-COL will be defined as follows.

Name. SAMPLEH-COL.

Instance. A loop-free graph G.

Output. An H-colouring of G chosen uniformly at random.

We will be particularly interested in the special case of this problem in which the input graph, G, is connected and bipartite.

Name. SAMPLEBH-COL.

Instance. A loop-free connected bipartite graph G.

Output. An H-colouring of G chosen uniformly at random.

The problem SAMPLEBIS will be defined as follows.

Name. SAMPLEBIS.

Instance. A loop-free connected bipartite graph G.

Output. An independent set of G chosen uniformly at random.

An *almost uniform sampler* [7, 15, 16] for X is a randomised algorithm that takes input σ and accuracy parameter $\epsilon \in (0, 1]$ and gives an output such that the variation distance between the output distribution of the algorithm and the uniform distribution on $X(\sigma)$ is at most ϵ. We will say that algorithm is a *polynomial almost uniform sampler (PAUS)* if its running time is bounded from above by a polynomial in the size of the instance $|\sigma|$ and $1/\epsilon$.

A *sampling-preserving reduction* (SP-reduction) from a sampling problem X to a sampling problem Y (denoted $X \leq_{\mathrm{SP}} Y$) consists of

1. A function f which takes an input (σ, ϵ), in which σ is an instance of X and $\epsilon \in (0,1]$ is an accuracy parameter, and produces an instance $f(\sigma, \epsilon)$ of Y. If $X(\sigma)$ is non-empty then $Y(f(\sigma, \epsilon))$ must be non-empty.

2. A function g which maps each tuple (σ, ϵ, y) with $y \in Y(f(\sigma, \epsilon))$ to a member of $X(\sigma) \cup \{\perp\}$ where "$\perp$" is an error symbol and for every (σ, ϵ) and every $x \in X(\sigma)$,

$$e^{-\epsilon} \frac{|Y(f(\sigma, \epsilon))|}{|X(\sigma)|} \le |\{y \in Y(f(\sigma, \epsilon)) \mid g(\sigma, \epsilon, y) = x\}| \le e^{\epsilon} \frac{|Y(f(\sigma, \epsilon))|}{|X(\sigma)|}. \tag{1}$$

Equation (1) says that for every $x \in X(\sigma)$, the number of $y \in Y(f(\sigma, \epsilon))$ which are mapped to x by g is roughly $\frac{|Y(f(\sigma,\epsilon))|}{|X(\sigma)|}$. Thus, each $x \in X(\sigma)$ is roughly equally represented and the error symbol $\perp$ is represented by only about an ϵ-fraction of $Y(f(\sigma, \epsilon))$.

The functions f and g must be computable in time which is bounded by a polynomial in $|\sigma|$ and $1/\epsilon$.

The problem #BIS is defined as follows.

Name. #BIS.

Instance. A loop-free bipartite graph G.

Output. The number of independent sets of G.

A component of H is *trivial* if it is a complete graph with all loops present or a complete bipartite graph with no loops present. Recall from Dyer and Greenhill [9] that counting H-colourings is in P if H is trivial. The main result of this paper is as follows.

Theorem 1 *Suppose that H is a fixed graph with no trivial components. If* SampleBH-Col *has a PAUS then* SampleBIS *has a PAUS and* #BIS *has an FPRAS. Thus, every problem which is AP-interreducible with* #BIS *(see [10]) has an FPRAS.*

3 Technical lemmas

Let $\nu(a, b)$ denote the number of onto functions from a set of size a to a set of size b. We need to use the following lemma, which is Lemma 18 of [10].

Lemma 2 *(DGGJ) If a and b are positive integers and $a \ge 2b \ln b$ then*

$$b^a (1 - \exp(-a/(2b))) \le \nu(a, b) \le b^a.$$

We also need the following technical lemma.

Lemma 3 *Suppose c_1 and c_2 are fixed positive reals with $c_1 < c_2$. For any $\delta > 0$ and any non-negative integers q and a_0, there are non-negative integers a and b with $a \ge a_0$ which are in $O((a_0 + q)/\delta)$ and satisfy*

$$e^{-\delta} c_2^{a+q} \le c_1^{b+q} \le e^{\delta} c_2^{a+q}.$$

Proof. First, note that it would suffice to find non-negative integers a' and b' which are in $O(q'/\delta)$ and satisfy

$$e^{-\delta} c_2^{a'+q'} \le c_1^{b'+q'} \le e^{\delta} c_2^{a'+q'},$$

where $q' = q + a_0$ because we could simply set $a = a' + a_0$ and $b = b' + a_0$ which would imply $a' + q' = a + q$ and $b' + q' = b + q$.

Taking logarithms, what we need is

$$\left| b' - \frac{a' \log c_2 + q' \log(c_2/c_1)}{\log c_1} \right| \le \frac{\delta}{\log c_1}. \tag{2}$$

Now let ρ be defined by $c_2 = c_1^{1+\rho}$. Then we want

$$|b' - (a'(1+\rho) + q'\rho)| \le \frac{\delta}{\log c_1}. \tag{3}$$

For a positive integer r, we will choose $a' = q'r$, so we want

$$|b' - a' - \rho q'(r+1)| \le \frac{\delta}{\log c_1}. \tag{4}$$

Let $R = \lceil 2 \log c_1/\delta \rceil$. Lemma 19 of [10] says: *For any real $z > 0$ and any positive integer R there is an $x \in [1, \ldots, R]$ such that*

$$\min(zx - \lfloor zx \rfloor, \lceil zx \rceil - zx) \le 1/R.$$

Thus, there is an $x \in [1, \ldots, R]$ such that $\rho q' x$ is within $1/R$ of a non-negative integer. If $x > 1$ we will set $r + 1 = x$. If $x = 1$ then note that $\rho q' 2$ is within $2/R$ of a non-negative integer, so we will set $r = 1$.

Now recall that $a' = q'r$, so $a' \in O(q'/\delta)$ as required. □

4 Demonstrating the existence of SP-reductions: a proof technique

When we introduce an SP-reduction from a sampling problem X to a sampling problem Y, we will need to show that Equation (1) is satisfied. We will typically do this by partitioning $Y(f(\sigma, \epsilon))$ into disjoint sets $Y_0, \ldots, Y_k$. We will show that each of $Y_1, \ldots, Y_k$ is fairly representative of $X(\sigma)$. In particular, for every $x \in X(\sigma)$ and every $i \in [1, k]$,

$$e^{-\epsilon/2} \frac{|Y_i|}{|X(\sigma)|} \le |\{y \in Y_i \mid g(\sigma, \epsilon, y) = x\}| \le e^{\epsilon/2} \frac{|Y_i|}{|X(\sigma)|}. \tag{5}$$

For every $y \in Y_0$, we will have $g(\sigma, \epsilon, y) = \perp$ but we will show that Y_0 is a small part of $Y(f(\sigma, \epsilon))$. In particular,

$$\sum_{i=1}^{k} |Y_i| \ge e^{-\epsilon/2} |Y(f(\sigma, \epsilon))|. \tag{6}$$

Together, (5) and (6) imply (1). Note that (6) follows from

$$|Y_0| \le (\epsilon/4)|Y(f(\sigma,\epsilon))|, \tag{7}$$

since (7) implies $|Y| - |Y_0| \ge (1-\epsilon/4)|Y(f(\sigma,\epsilon))| \ge e^{-\epsilon/2}|Y(f(\sigma,\epsilon))|$.

Quite often the reduction $X \le_{\mathrm{SP}} Y$ will involve several subproblems $Z_1, Z_2, \ldots$ such that, for each of these, an SP-reduction (f_i, g_i) from X to Z_i is already known. The instance $f(\sigma,\epsilon)$ of Y is then formed by "gluing" together instances $f_1(\sigma,\epsilon/2)$ of Z_1, $f_2(\sigma,\epsilon/2)$ of Z_2, and so on. Y_i is (roughly) the portion of $Y(f(\sigma,\epsilon))$ for which, within each $y \in Y_i$, we can "zoom in" on a structure $z \in Z_i(f_i(\sigma,\epsilon/2))$. Each structure in $Z_i(f_i(\sigma,\epsilon/2))$ is represented by an equal number of $y \in Y_i$ so we can get (5) by referring to the SP-reduction from X to Z_i. Establishing (7) is essentially showing that, although $Y(f(\sigma,\epsilon))$ has some structures which don't allow us to "zoom in" on an appropriate sub-problem to find our sample, these aren't so numerous.

Finally, let $Y_i(x) = \{y \in Y_i \mid g(\sigma,\epsilon,y) = x\}$. Suppose that no $y \in Y_i$ has $g(\sigma,\epsilon,y) = \perp$. In this case we can show (5) by showing that for all $x, x' \in X(\sigma)$,

$$|Y_i(x)| \le e^{\epsilon/2}|Y_i(x')|. \tag{8}$$

To see this, note that

$$\frac{|Y_i|}{|X(\sigma)|} = \frac{\sum_{x'\in X(\sigma)}|Y_i(x')|}{|X(\sigma)|} \ge e^{-\epsilon/2}\frac{\sum_{x'\in X(\sigma)}|Y_i(x)|}{|X(\sigma)|} = e^{-\epsilon/2}|Y_i(x)|.$$

5 Sampling *fixed* H-colourings

Suppose that H is connected, loop-free, and bipartite. Denote the vertex partition of H by $(V_L(H), V_R(H))$. We will define the *fixed* H-colouring problem as follows.

Name. SampleFixed$H$-Col

Instance. A loop-free connected bipartite graph G with vertex partition $(V_L(G), V_R(G))$

Output. An H-colouring of G chosen uniformly at random from the set of H-colourings in which vertices of $V_L(G)$ receive colours from $V_L(H)$.

A vertex in $V_L(H)$ is said to be *full* if it is adjacent to every vertex in $V_R(H)$. Similarly, a vertex in $V_R(H)$ is said to be *full* if it is adjacent to every vertex in $V_L(H)$. The graph H is said to be *full* if both $V_L(H)$ and $V_R(H)$ contain at least one full vertex.

Lemma 4 *Suppose that H is a connected nontrivial full loop-free bipartite graph. Then* SampleBIS $\le_{\mathrm{SP}}$ SampleFixed$H$-Col.

Proof. We'll prove the lemma by induction on the number of vertices in H. For the base case, suppose that H has at most 4 vertices. The only connected nontrivial full loop-free bipartite graph H with at most 4 vertices is the path of length 3. There is a one-to-one correspondence between fixed H-colourings of G and independent sets of G: The endpoints of H point out the vertices which are in the independent set (see the proof of Theorem 5 of [10]).

We will now move on to the inductive step. Let F_L be the set of full vertices in $V_L(H)$ and let F_R be the set of full vertices in $V_R(H)$. Let $f_L = |F_L|$ and $f_R = |F_R|$ and $v_L = |V_L(H)|$ and $v_R = |V_R(H)|$. For a subset S of $V_R(H)$, let $N(S)$ be the set of mutual neighbours of S:

$$N(S) = \{v \in V_L(H) \mid \forall u \in S, (u, v) \in E(H)\}.$$

Note that $F_L \subseteq N(S) \subseteq V_L(H)$. S is said to be *left-reducing* if $F_L \subset N(S) \subset V_L(H)$. If S is left-reducing, let H_S be the subgraph of H induced by vertex partition $(N(S), V_R(H))$. Note that H_S has fewer vertices than H. Also, it is connected, full and nontrivial: The set of full vertices in $V_L(H_S)$ is F_L; the set of full vertices in $V_R(H_S)$ includes all of F_R but it does not equal $V_R(H)$ since $N(S) \supset F_L$.

Similarly, a subset S of $V_L(H)$ is *right-reducing* if $F_R \subset N(S) \subset V_R(H)$. If S is right-reducing, let H_S be the subgraph of H induced by vertex partition $(V_L(H), N(S))$. H_S has fewer vertices than H and is connected, full and nontrivial.

Now, let $S_1, \ldots, S_k$ be the left-reducing subsets of $V_R(H)$ and let $S_{k+1}, \ldots, S_{k+j}$ be the right-reducing subsets of $V_L(H)$. (Either or both of k and j could be zero.) For every $i \in \{1, \ldots, k+j\}$, let (f_i, g_i) be an SP-reduction from SampleBIS to SampleFixedH_{S_i}-Col. Take the input (G, ϵ) to SampleBIS and let $G_i = f_i(G, \epsilon/2)$. Let $a_i = |V_L(G_i)|$ and let $b_i = |V_R(G_i)|$. Let $q = \sum_{i=1}^{k+j}(a_i + b_i)$ and let $n = |V_L(G)| + |V_R(G)|$.

Let $f(G, \epsilon)$ be the graph which is constructed as follows, where a and b will be chosen later to satisfy

$$a \geq 2v_L \lceil q \ln(v_R/f_R) + \ln(16n/\epsilon) \rceil, \tag{9}$$

and

$$b \geq 2v_R \lceil q \ln(v_L/f_L) + \ln(16n/\epsilon) \rceil. \tag{10}$$

See Figure 2. For every vertex u of G, put a size-a set $L(u)$ into $V_L(f(G, \epsilon))$ and a size-b set $R(u)$ into $V_R(f(G, \epsilon))$. Also, add edges $L(u) \times R(u)$ to $E(f(G, \epsilon))$. If $u \in V_L(G)$ is connected to $v \in V_R(G)$ by an edge of G then add edges $R(u) \times L(v)$ to $E(f(G, \epsilon))$.

Also, for every vertex u of G and every $i \in [1, \ldots, k+j]$, let $A_{i,u}$ and $B_{i,u}$ be copies of G_i and let $A'_{i,u}$ and $B'_{i,u}$ be copies of G_i in which left-vertices and right-vertices are switched (so the vertices in $V_L(A'_{i,u})$ correspond to the vertices in $V_R(G_i)$ and the vertices in $V_R(A'_{i,u})$ correspond to the vertices in $V_L(G_i)$). Add edges $L(u) \times V_R(A_{i,u})$ and $L(u) \times V_R(A'_{i,u})$ and $R(u) \times V_L(B_{i,u})$ and $R(u) \times V_L(B'_{i,u})$ to $E(f(G, \epsilon))$.

Let

$$V_L(f(G, \epsilon)) = \bigcup_u L(u) \cup \bigcup_{u,i} \{V_L(A_{i,u}) \cup V_L(A'_{i,u}) \cup V_L(B_{i,u}) \cup V_L(B'_{i,u})\}$$

and let Y be the set of fixed H-colourings of $f(G, \epsilon)$. We will partition Y into sets $Y_0, \ldots, Y_{k+j+1}$.

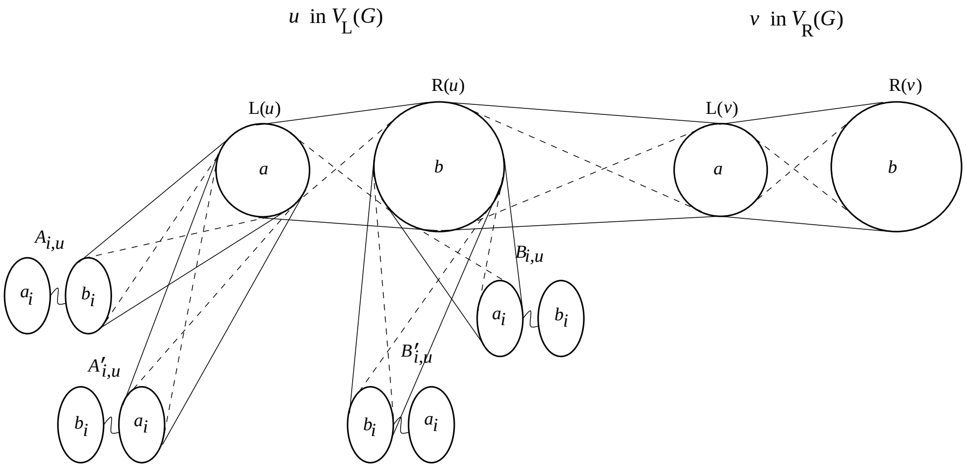

Figure 2: The construction of $f(G, \epsilon)$ in the proof of Lemma 4.

For $i \in [1, \ldots, k]$, Y_i is the set of colourings which are not in $Y_1, \ldots, Y_{i-1}$ but in which some $u \in V_L(G)$ has $R(u)$ coloured with (exactly) the colours in S_i. For $i \in [k+1, \ldots, k+j]$, Y_i is the set of colourings which are not in $Y_1, \ldots, Y_{i-1}$ but in which some $v \in V_R(G)$ has $L(v)$ coloured with S_i.

Now note that for any remaining colourings (any colourings in Y_0 or Y_{k+j+1}), every vertex $u \in V_L(G)$ has $R(u)$ coloured with a set S of colours such that $N(S)$ is either F_L or $V_L(H)$. Similarly, every vertex $v \in V_R(G)$ has $L(v)$ coloured with a set S of colours such that $N(S)$ is either F_R or $V_R(H)$.

Consider a colouring y. Vertex $u \in V_L(G)$ satisfies *Condition (A)* if $R(u)$ is coloured with a set S of colours with $N(S) = F_L$ but $S \subset V_R(H)$. It satisfies *Condition (B)* if $R(u)$ is coloured with a set S of colours with $N(S) = V_L(H)$ but $L(u)$ is coloured with a proper subset of $V_L(H)$. Vertex $v \in V_R(G)$ satisfies *Condition (C)* if $L(v)$ coloured with a set S of colours with $N(S) = F_R$ but $S \subset V_L(H)$. It satisfies *Condition (D)* if $L(v)$ is coloured with a set S of colours with $N(S) = V_R(H)$ but $R(v)$ is coloured with a proper subset of $V_R(H)$.

We now define

$$Y_0 = \{y \in Y - \{Y_1 \cup \cdots \cup Y_{k+j}\} \mid \text{some vertex satisfies Condition A or B or C or D}\}.$$

Now note that colourings in Y_{k+j+1} have the following property. Every vertex u of G either has $R(u)$ coloured $V_R(H)$ or has $L(u)$ coloured $V_L(H)$.

We will first work on establishing Equation (7). Let $Y_{u,A}$ denote the subset of Y in which u satisfies (A). Define $Y_{u,B}$, $Y_{u,C}$ and $Y_{u,D}$ similarly. We will show that the size of each of $Y_{u,A}$, $Y_{u,B}$, $Y_{u,C}$ and $Y_{u,D}$ is at most $(\epsilon/(16n))|Y|$. Equation (7) follows since

$$|Y_0| \leq \sum_{u \in V(G)} |Y_{u,A}| + |Y_{u,B}| + |Y_{u,C}| + |Y_{u,D}|.$$

First, let's show that $|Y_{u,A}| \leq (\epsilon/(16n))|Y|$. Consider the set of colourings in Y in which all neighbours of vertices in $R(u)$ have colours from F_L and let ψ be the number of induced colourings on vertices other than the vertices of $R(u)$. If $\psi = 0$ then $|Y_{u,A}| = 0$, so the claim is trivial. Otherwise, $|Y_{u,A}| \leq \psi(v_R^b - \nu(b, v_R))$ which is at most $\psi v_R^b \exp(-b/(2v_R))$ by Lemma 2. On the other hand, $|Y| \geq \psi v_R^b$, so the claim follows from Equation (10). The proof that $|Y_{u,C}|$ is sufficiently small is similar.

Next, let's show that $|Y_{u,B}| \leq (\epsilon/(16n))|Y|$. Consider the set of colourings in Y in which $R(u)$ is coloured with a subset of F_R and let ψ be the number of induced colourings on all vertices except those in $L(u)$ and $A_{i,u}$ and $A'_{i,u}$ (for $i \in [1, \ldots, j+k]$). If $\psi = 0$ then $|Y_{u,B}| = 0$, so the claim is trivial. Otherwise, $|Y_{u,B}| \leq \psi(v_L^a - \nu(a, v_L))v_R^q v_L^q$ which is at most $\psi v_L^a \exp(-a/(2v_L))v_R^q v_L^q$ by Lemma 2. On the other hand, $|Y| \geq \psi v_L^a f_R^q v_L^q$, so the claim follows from Equation (9). The proof that $|Y_{u,D}|$ is sufficiently small is similar.

We will now work on establishing Equation (5). First consider $i \in [1, \ldots, k]$. Let $Y_{u,i}$ be the set of colourings in Y_i for which $u \in V_L(G)$ is the first vertex in $V_L(G)$ with $R(u)$ coloured S_i. Let Γ be the set of induced colourings on $B_{i,u}$. Note that Γ is the set of fixed H_{S_i}-colourings of $G_i = f_i(G, \epsilon/2)$. Also, each colouring in Γ comes up ψ times in $Y_{u,i}$ for some ψ. (In particular, ψ is the number of colourings of vertices other than $B_{i,u}$ which are induced by colourings in $Y_{u,i}$.) For colouring $y \in Y_{u,i}$ we will let $g(G, \epsilon, y) = g_i(G_i, \epsilon/2, y')$ where y' is the induced colouring on $B_{i,u}$. Then for every independent set x in the set $\mathcal{I}(G)$ of independent sets of G,

$$|\{y \in Y_{u,i} \mid g(G, \epsilon, y) = x\}| = \psi\, |\{y' \in \Gamma \mid g_i(G_i, \epsilon/2, y') = x\}|. \tag{11}$$

Since (f_i, g_i) is an SP-reduction, Equation (1) gives

$$e^{-\epsilon/2}\frac{|\Gamma|}{|\mathcal{I}(G)|} \leq |\{y' \in \Gamma \mid g_i(G_i, \epsilon/2, y') = x\}| \leq e^{\epsilon/2}\frac{|\Gamma|}{|\mathcal{I}(G)|} \tag{12}$$

and Equation (5) follows for $Y_{u,i}$ from Equations (11) and (12) since $|Y_{u,i}| = \psi\,|\Gamma|$. Colourings in $Y_{k+1}, \ldots, Y_{k+j}$ are handled similarly except that we look at induced colourings of $A_{i,u}$ rather than $B_{i,u}$.

It remains to satisfy Equation (5) for $i = k + j + 1$. Note that any colouring y in Y_{k+j+1} points out an independent set of G. A vertex $u \in V_L(G)$ is in the independent set if $R(u)$ is coloured $V_R(H)$. A vertex $v \in V_R(G)$ is in the independent set if $L(u)$ is coloured $V_L(H)$. We will define $g(G, \epsilon, y)$ to be this independent set. Let us focus attention on a given independent set containing w_L vertices in $V_L(G)$ and w_R vertices in $V_R(G)$. We will now calculate how many colourings in Y_{k+j+1} correspond to this independent set.

For any bipartite graph G' with vertex partition $(V_L(G'), V_R(G'))$, let $\phi_H(G')$ denote the number of fixed H-colourings of G'. Then the number of times that this independent set comes up as a colouring in Y_{k+j+1} is the product of the following two quantities.

$$\left(\nu(b, v_R) f_L^a \prod_{i=1}^{k+j} \phi_H(A_{i,u})\phi_H(A'_{i,u}) f_L^{a_i+b_i} v_R^{a_i+b_i}\right)^{w_L+v_R-w_R},$$

$$\left(f_R^b \nu(a, v_L) \prod_{i=1}^{k+j} \phi_H(B_{i,u}) \phi_H(B'_{i,u}) v_L^{a_i+b_i} f_R^{a_i+b_i}\right)^{v_L - w_L + w_R}.$$

Now note that $\phi_H(A_{i,u}) = \phi_H(B_{i,u})$ and $\phi_H(A'_{i,u}) = \phi_H(B'_{i,u})$. So if we let

$$Z = \left(\prod_{i=1}^{k+j} \phi_H(A_{i,u}) \phi_H(A'_{i,u})\right)^{v_L+v_R} (f_L^a \nu(b, v_R) f_L^q v_R^q)^{v_R} \left(f_R^b \nu(a, v_L) v_L^q f_R^q\right)^{v_L},$$

the contribution of the independent set becomes

$$Z(\nu(b, v_R) f_L^a f_L^q v_R^q)^{w_L - w_R} \left(f_R^b \nu(a, v_L) v_L^q f_R^q\right)^{w_R - w_L},$$

which is

$$Z\left(\frac{\nu(b, v_R) v_L^a}{v_R^b \nu(a, v_L)}\right)^{w_L - w_R} \left(\left(\frac{v_R}{f_R}\right)^{b+q} \left(\frac{f_L}{v_L}\right)^{a+q}\right)^{w_L - w_R}.$$

To get Equation (8) we will show that a and b can be chosen so that

$$e^{-\epsilon/(8n)} \leq \left(\frac{\nu(b, v_R) v_L^a}{v_R^b \nu(a, v_L)}\right) \leq e^{\epsilon/(8n)}, \text{ and} \tag{13}$$

and

$$e^{-\epsilon/(8n)} \leq \left(\frac{v_R}{f_R}\right)^{b+q} \left(\frac{f_L}{v_L}\right)^{a+q} \leq e^{\epsilon/(8n)}. \tag{14}$$

This guarantees that the contribution of this independent set is in the range $[e^{-\epsilon/4} Z, e^{\epsilon/4} Z]$, and Equation (8) follows for Y_{k+j+1}. To establish Equation (13), use Lemma 2 to observe that

$$\left(\frac{\nu(b, v_R) v_L^a}{v_R^b \nu(a, v_L)}\right) \leq \frac{1}{1 - \exp(-a/(2v_L))}.$$

Since Equation (9) gives $1 - \exp(-a/(2v_L)) \geq 1 - \epsilon/(16n) \geq e^{-\epsilon/(8n)}$, the right-hand inequality of (13) follows. The left-hand inequality is similar.

We will now show how to choose the values of a and b to satisfy Equation (14). If $v_R/f_R = v_L/f_L$ then simply choose $a = b$ and make them large enough to satisfy Equation (9) and Equation (10). Suppose that $v_R/f_R < v_L/f_L$. Then use Lemma 3 with $c_1 = v_R/f_R$, $c_2 = v_L/f_L$, $\delta = \epsilon/(8n)$, and

$$a_0 = 2v_L \lceil q \ln(v_R/f_R) + \ln(16n/\epsilon) \rceil + 2v_R \lceil q \ln(v_L/f_L) + \ln(16n/\epsilon) \rceil.$$

The lemma gives values of a and b which are in $O((a_0 + q)/\delta)$, which is not too large. Thus, our reduction is sampling-preserving. Note that the reduction can be done in polynomial time — the calculation of a and b does not involve computing Z. The case where $v_L/f_L < v_R/f_R$ is similar. □

6 The proof of Theorem 1

We start with some definitions. First, for every graph H, we will define a loop-free bipartite graph $B[H]$ (this construction was used in [7]). Let the vertices of H be $v_1, \ldots, v_h$. The vertex set of $B[H]$ is $\{x_1, \ldots, x_h\} \cup \{y_1, \ldots, y_h\}$. The edge set of $B[H]$ is

$$\{(x_i, y_j) \mid (v_i, v_j) \in E(H)\}.$$

Thus, a loop (v_i, v_i) in H corresponds to the edge (x_i, y_i) in $B[H]$ and a non-loop (v_i, v_j) in H (for which $i \neq j$) corresponds to two edges (x_i, y_j) and (y_i, x_j) in $B[H]$. For every edge (v_i, v_j) of H, let

$$V_L(H_{i,j}) = \{x_\ell \mid (v_\ell, v_j) \in E(H)\}$$

and

$$V_R(H_{i,j}) = \{y_\ell \mid (v_i, v_\ell) \in E(H)\}$$

and let $H_{i,j}$ be the subgraph of $B[H]$ induced by vertex set $V_L(H_{i,j}) \cup V_R(H_{i,j})$. Note that $x_i \in V_L(H_{i,j})$ and $y_j \in V_R(H_{i,j})$ and x_i is adjacent to all of $V_R(H_{i,j})$ in $H_{i,j}$ and y_j is adjacent to all of $V_L(H_{i,j})$. Thus, $H_{i,j}$ is connected and full. Let $\Delta_1(H)$ be the degree of H. That is,

$$\Delta_1(H) = \max\{\deg(v) \mid v \in V(H)\}.$$

Similarly, let $\Delta_2(H)$ be the maximum degree amongst neighbours of vertices with degree $\Delta_1(H)$:

$$\Delta_2(H) = \max\{\deg(v) \mid \text{ for some } u \in V(H) \text{ with } \deg(u) = \Delta_1(H),\ (u, v) \in E(H)\}.$$

Let

$$R(H) = \{(v_i, v_j) \mid ((v_i, v_j) \in E(H) \text{ and } \deg(v_i) = \Delta_1(H) \text{ and } \deg(v_j) = \Delta_2(H)\}.$$

We will start with the following lemma.

Lemma 5 *Let H be any fixed graph with no trivial components. Then $R(H)$ is non-empty and $\Delta_1(H) > 1$ and $\Delta_2(H) > 1$. Also, for all $(v_i, v_j) \in R(H)$, $H_{i,j}$ is connected, loop-free, bipartite, full and nontrivial.*

Proof. Since H has no non-trivial components, $R(H)$ is non-empty and $\Delta_1(H) > 1$ and $\Delta_2(H) > 1$. Suppose $(v_i, v_j) \in R(H)$. Recall that $H_{i,j}$ is connected, loop-free, bipartite and full. Suppose for contradiction that $H_{i,j}$ is a complete bipartite graph (so vertices in $V_L(H_{i,j})$ have degree $\Delta_1(H)$ in $H_{i,j}$ and vertices in $V_R(H_{i,j})$ have degree $\Delta_2(H)$ in $H_{i,j}$).

This assumption guarantees that $H_{i,j}$ is a connected component of $B[H]$: $B[H]$ cannot have an edge with exactly one endpoint in $V_L(H_{i,j})$ — the endpoint would then have degree exceeding $\Delta_1(H)$ in $B[H]$, which is a contradiction; similarly, $B[H]$ cannot have an edge with exactly one endpoint in $V_R(H_{i,j})$.

Thus, for any $x_\ell \in V_L(H_{i,j})$,

$$\{v_r \mid (v_\ell, v_r) \in E(H)\} = \{v_r \mid y_r \in V_R(H_{i,j})\}. \tag{15}$$

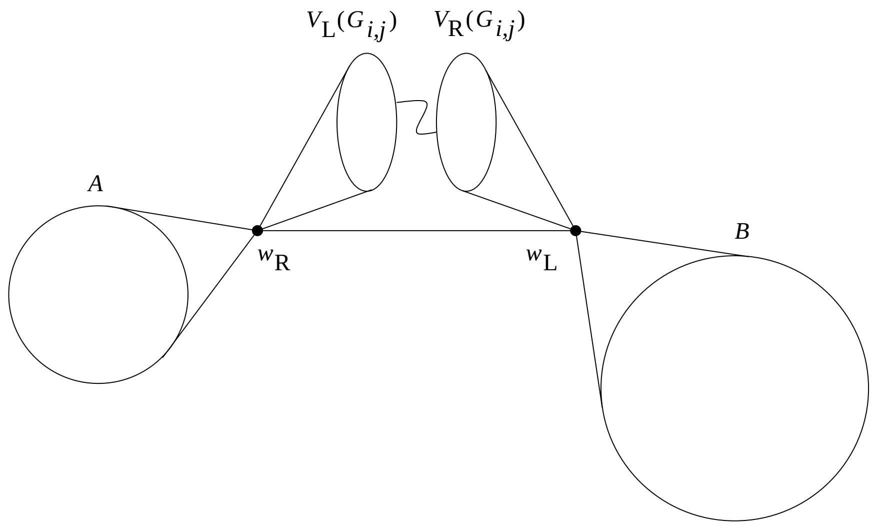

Figure 3: The construction of $f(G, \epsilon)$ in the proof of Lemma 6.

Similarly, for any $y_\ell \in V_R(H_{i,j})$,

$$\{v_r \mid (v_\ell, v_r) \in E(H)\} = \{v_r \mid x_r \in V_L(H_{i,j})\}. \tag{16}$$

Now if H has a vertex v_ℓ such that $(v_i, v_\ell) \in E(H)$ and $(v_j, v_\ell) \in E(H)$ then $x_\ell \in V_L(H_{i,j})$ and $y_\ell \in V_R(H_{i,j})$ so Equations (15) and (16) imply that

$$\{v_r \mid y_r \in V_R(H_{i,j})\} = \{v_r \mid x_r \in V_L(H_{i,j})\}.$$

Thus, $H_{i,j}$ corresponds to a component of H and that component is a looped clique, which contradicts the fact that H has no trivial component.

On the other hand, if there is no v_ℓ with $(v_i, v_\ell) \in E(H)$ and $(v_j, v_\ell) \in E(H)$ then $H_{i,j}$ corresponds to a connected component of H which is a complete bipartite graph, again giving a contradiction. $\square$

We can now prove the main lemma.

Lemma 6 *Suppose that H is a fixed graph with no trivial components. Then* SampleBIS $\leq_{\mathrm{SP}}$ SampleBH-Col.

Proof. Let (G, ϵ) be an input to SampleBIS. For each $(v_i, v_j) \in R(H)$, Lemma 5 and Lemma 4 guarantee that there is a sampling-preserving reduction $(f_{i,j}, g_{i,j})$ from SampleBIS to SampleFixed$H_{i,j}$-Col. Let $G_{i,j} = f_{i,j}(G, \epsilon/2)$. Let $f(G, \epsilon)$ be the graph which is constructed as follows. See Figure 3. Let $q = \sum_{(v_i,v_j)\in R(H)} |V_L(G_{i,j})| + |V_R(G_{i,j})|$. Let

$$V_L(f(G, \epsilon)) = A \cup \{w_L\} \cup \bigcup_{(v_i,v_j)\in R(H)} V_L(G_{i,j}),$$

and

$$V_R(f(G,\epsilon)) = B \cup \{w_R\} \cup \bigcup_{(v_i,v_j)\in R(H)} V_R(G_{i,j}),$$

where A and B are sets of vertices with

$$|A| = \left\lceil \frac{q\ln(|V(H)|) + \ln(8|E(H)|/\epsilon)}{\ln(\Delta_2(H)/(\Delta_2(H)-1))} \right\rceil$$

and

$$|B| = \left\lceil \frac{(q+|A|+1)\ln(|V(H)|) + \ln(8|V(H)|/\epsilon)}{\ln(\Delta_1(H)/(\Delta_1(H)-1))} \right\rceil.$$

Note that there is no division by zero, since $\Delta_1(H)$ and $\Delta_2(H)$ are bigger than one (by Lemma 5). In addition to the edges in the graphs $G_{i,j}$, we add edge (w_L, w_R) and $w_L \times B$ and $w_R \times A$ and, for all $(v_i, v_j) \in R(H)$, we add edges $w_L \times V_R(G_{i,j})$ and $w_R \times V_L(G_{i,j})$.

Let Y be the set of H-colourings of $f(G,\epsilon)$. Y_0 will be the set of colourings in Y in which (w_L, w_R) is not coloured with an edge (v_i, v_j) from $R(H)$. We will now establish Equation (7). For every $v \in V(H)$ with $\deg(v) < \Delta_1(H)$ let $Y_0(v)$ be the set of colourings in Y in which w_L is coloured v_i. Now

$$|Y_0(v)| \leq (\Delta_1(H)-1)^{|B|}|V(H)|^{q+|A|+1}.$$

Now consider any $(v_i, v_j) \in R(H)$. There are at least $\Delta_2(H)^{|A|}\Delta_1(H)^{|B|}$ colourings of $f(G,\epsilon)$ with (w_L, w_R) coloured (v_i, v_j) (for example, the colourings in which all of the vertices of the graphs $G_{i,j}$ are coloured with either v_i or v_j). Thus, $|Y| \geq \Delta_2(H)^{|A|}\Delta_1(H)^{|B|} \geq \Delta_1(H)^{|B|}$. We conclude that

$$|Y_0(v)| \leq (\epsilon/(8|V(H)|))|Y|. \tag{17}$$

Now consider any edge $(v_i, v_k) \in E(H)$ such that $\deg(v_i) = \Delta_1(H)$ but $\deg(v_k) < \Delta_2(H)$. Let $Y_0(v_i, v_k)$ be the set of colourings in Y in which (w_L, w_R) is coloured (v_i, v_k). Now

$$|Y_0(v_i, v_k)| \leq \Delta_1(H)^{|B|}(\Delta_2(H)-1)^{|A|}|V(H)|^q.$$

Also, from before $|Y| \geq \Delta_2(H)^{|A|}\Delta_1(H)^{|B|}$ so

$$|Y_0(v_i, v_k)| \leq (\epsilon/(8|E(H)|))|Y|. \tag{18}$$

Equation (17) and (18) imply Equation (7) since $|Y_0| \leq \sum_{v\in V(H)} Y_0(v) + \sum_{(v_i,v_k)} |Y_0(v_i, v_k)|$.

For an edge $(v_i, v_j) \in R(H)$, let $Y_{i,j}$ be the set of colourings of $f(G,\epsilon)$ with (w_L, w_R) coloured (v_i, v_j). Let Γ be the set of induced colourings on $G_{i,j}$. Note that Γ is the set of fixed $H_{i,j}$-colourings of $G_{i,j}$. Also, each colouring in Γ comes up ψ times in $Y_{i,j}$ where ψ is the number of induced colourings on the vertices other than $G_{i,j}$. For a colouring $y \in Y_{i,j}$ we will set $g(G,\epsilon,y) = g_{i,j}(G_{i,j}, \epsilon/2, y')$ where y' is the induced colouring on $G_{i,j}$. Then Equation (5) follows from the fact that $(f_{i,j}, g_{i,j})$ is an SP-reduction. □

Theorem 1 follows from Lemma 6 and from Lemmas 7 and 8 below.

Lemma 7 *If $X \leq_{\mathrm{SP}} Y$ and Y has a PAUS, then X has a PAUS.*

Proof. Let (f, g) be the reduction from X to Y and let $\mathcal{A}$ be a PAUS for Y. Here is a PAUS for X: On input (σ, ϵ), let y be the output of $\mathcal{A}$ when it is run with inputs $f(\sigma, \epsilon/4)$ and $\epsilon/2$; return $g(\sigma, \epsilon/4, y)$. We must show that the variation distance between the output distribution of this algorithm and the uniform distribution on $X(\sigma)$ is at most ϵ. Let σ be an input with $|X(\sigma)| \geq 1$. Consider any subset A_x of $X(\sigma)$. Let

$$A_y = \{y \in Y(f(\sigma, \epsilon/4)) \mid g(\sigma, \epsilon/4, y) \in A_x\}.$$

Then the probability that $\mathcal{A}$ gives an output in A_y is at most

$$\begin{aligned}
& \frac{|A_y|}{|Y(f(\sigma, \epsilon/4))|} + \frac{\epsilon}{2} \\
\leq\; & \frac{e^{\epsilon/4}|A_x|}{|X(\sigma)|} + \frac{\epsilon}{2} \\
\leq\; & \frac{(1+\epsilon/2)|A_x|}{|X(\sigma)|} + \frac{\epsilon}{2} \\
\leq\; & \frac{|A_x|}{|X(\sigma)|} + \frac{(\epsilon/2)|A_x|}{|X(\sigma)|} + \frac{\epsilon}{2} \\
\leq\; & \frac{|A_x|}{|X(\sigma)|} + \epsilon.
\end{aligned}$$

Also, the probability that $\mathcal{A}$ gives an output in A_y is at least

$$\begin{aligned}
& \frac{|A_y|}{|Y(f(\sigma, \epsilon/4))|} - \frac{\epsilon}{2} \\
\geq\; & \frac{e^{-\epsilon/4}|A_x|}{|X(\sigma)|} - \frac{\epsilon}{2} \\
\geq\; & \frac{(1-\epsilon/2)|A_x|}{|X(\sigma)|} - \frac{\epsilon}{2} \\
\geq\; & \frac{|A_x|}{|X(\sigma)|} - \frac{|A_x|(\epsilon/2)}{|X(\sigma)|} - \frac{\epsilon}{2} \\
\geq\; & \frac{|A_x|}{|X(\sigma)|} - \epsilon.
\end{aligned}$$

□

Recall the following definitions. A *randomised approximation scheme* (RAS) for a counting problem F is a randomised algorithm that takes input σ and accuracy parameter $\epsilon \in (0, 1)$ and produces as output an integer random variable Y satisfying the condition $\Pr(e^{-\epsilon}F(\sigma) \leq Y \leq e^{\epsilon}F(\sigma)) \geq 3/4$. It is a "fully polynomial" randomised approximation scheme (FPRAS) if it runs in time $\mathrm{poly}(|\sigma|, \epsilon^{-1})$. The problem #BIS is "self-reducible" so the following lemma follows from [16].

Lemma 8 *(JVV) If* SAMPLEBIS *has a PAUS then* #BIS *has an FPRAS.*

Proof. The lemma is a special case of Theorem 6.4 of [16]. In order to apply Theorem 6.4 directly we would need to define "self-reducible" formally and to deal with some easy (though annoying) issues:

(i) Inputs to #BIS may be disconnected but inputs to SAMPLEBIS may not.

(ii) In order to apply Theorem 6.4 we technically need a *fully* polynomial almost uniform sampler (FPAUS) for SAMPLEBIS. This can be obtained from a PAUS as [16] explains.

Rather than dealing with these issues, we prefer to simply provide a proof for the lemma. The details given here are from the proof of Proposition 3.4 of [15]. Technically, Jerrum's proof from [15] is about counting *matchings* but the few changes that are needed to yield our lemma are completely routine.

Let (G, ϵ) be an input to #BIS. Suppose that G has components $G_1, \ldots, G_k$. For each i, let the two parts of the vertex set be $V_L(G_i)$ and $V_R(G_i)$ and let the sizes of these parts be ℓ_i and r_i, respectively. Let $N_i = \ell_i r_i$ and let $E(G_i) = \{e_i(1), \ldots, e_i(m_i)\}$. Denote the non-edges of G_i by $\{e_i(m_i+1), \ldots, e_i(N_i)\}$. For $j \in \{1, \ldots, N_i\}$, let $G_i(j)$ be the graph $(V(G_i), \{e_i(1), \ldots, e_i(j)\})$. For any graph G', let $\mathcal{I}(G')$ denote the set of independent sets of G'. Let

$$\rho_i(j) = \frac{|\mathcal{I}(G_i(j+1))|}{|\mathcal{I}(G_i(j))|}.$$

Note that

$$|\mathcal{I}(G_i)| = (\rho_i(m_i)\rho_i(m_i+1)\cdots\rho_i(N_i-1))^{-1}|\mathcal{I}(G_i(N_i))|.$$

Also, the number of independent sets of the complete bipartite graph $G_i(N_i)$ is $2^{\ell_i}+2^{r_i}-1$, so

$$|\mathcal{I}(G_i)| = (2^{\ell_i}+2^{r_i}-1)\prod_{j=m_i}^{N_i-1}\rho_i(j)^{-1}. \tag{19}$$

Furthermore,

$$|\mathcal{I}(G)| = \prod_{i=1}^{k}|\mathcal{I}(G_i)| = \prod_{i=1}^{k}(2^{\ell_i}+2^{r_i}-1)\prod_{j=m_i}^{N_i-1}\rho_i(j)^{-1}. \tag{20}$$

Now let $z = \sum_{i=1}^{k}(N_i - m_i)$. In order to estimate $|\mathcal{I}(G)|$, we need to estimate the z ratios $\rho_i(j)$.

For each ratio $\rho_i(j)$ we can make some observations.

(i) $\rho_i(j) \leq 1$, since $\mathcal{I}(G_i(j+1)) \subseteq \mathcal{I}(G_i(j))$

(ii) $\rho_i(j) \geq 1/2$, since $\mathcal{I}(G_i(j)) \backslash \mathcal{I}(G_i(j+1))$ can be mapped injectively into $\mathcal{I}(G_i(j+1))$ by removing the lexicographically-least endpoint of $e_i(j+1)$.

(iii) Let $\mathcal{A}$ be a PAUS for SAMPLEBIS. For $i \in [1, \ldots, k]$ and $j \in [m_i, \ldots, N_i - 1]$, let $Z_i(j)$ be the indicator variable for the event that, when we run $\mathcal{A}$ with input $G_i(j)$ and accuracy parameter δ, the output is an independent set of $G_i(j+1)$. Note that $\rho_i(j) - \delta \leq E[Z_i(j)] \leq \rho_i(j) + \delta$. This follows immediately from the definition of PAUS, but it is important to note that the input to $\mathcal{A}$, $G_i(j)$, is connected (since all inputs to SAMPLEBIS must be connected).

Let $\overline{Z_i(j)}$ be the result obtained by calling $\mathcal{A}$ $\lceil 74\epsilon^{-2}z \rceil$ times with input $G_i(j)$ and accuracy parameter $\delta = \epsilon/(6z)$ and averaging the value of $Z_i(j)$ which occurs each time. Jerrum shows in his proof that with probability at least 3/4,

$$e^{-\epsilon} \prod_{i=1}^{k} \prod_{j=m_i}^{N_i-1} \rho_i(j) \leq \prod_{i=1}^{k} \prod_{j=m_i}^{N_i-1} \overline{Z_i(j)} \leq e^{\epsilon} \prod_{i=1}^{k} \prod_{j=m_i}^{N_i-1} \rho_i(j).$$

Thus, the quantity

$$\prod_{i=1}^{k} (2^{\ell_i} + 2^{r_i} - 1) \prod_{j=m_i}^{N_i-1} \overline{Z_i(j)}^{-1}$$

is a sufficiently accurate estimate of $|\mathcal{I}(G)|$.

For each of the z pairs (i, j), $O(\epsilon^{-2}z)$ samples were needed, each of which is produced in time poly$(|G|, z/\epsilon)$. Since $z \leq |V(G)|^2$, the total running time is poly$(|G|, \epsilon^{-1})$ and we have an FPRAS. □

7 The presence of trivial components

Theorem 1 shows that sampling H-colourings is difficult if every component of H is nontrivial. Recall from [9] that exactly counting H-colourings is #P-complete if H has even one nontrivial component. Thus, it might appear that Theorem 1 can be improved. In this section, we show that the existence of a single nontrivial component is not enough to make sampling difficult. In particular, we give an example of a graph H with a nontrivial component, for which SAMPLEH-COL has a PAUS. Specifically, let H be the graph depicted in Figure 4.

Observation 9 *Suppose that H is the graph depicted in Figure 4.* SAMPLEH-COL *has a PAUS.*

Proof. Here is a PAUS for SAMPLEH-COL. The input is an instance (G, ϵ) where G has n vertices and, without loss of generality[3], is connected. If $\epsilon < 2^n/(2^n+3^n)$ then the algorithm simply runs for 5^n steps, constructs all of the H-colourings of G (and counts them) and

[3]We can assume that the input is a connected graph without losing generality because we can obtain an H-colouring of a k-component graph G by independently calling our PAUS for each component, specifying accuracy parameter ϵ/k. The final variation distance (between the output distribution and the uniform distribution on H-colourings of G) is at most ϵ.

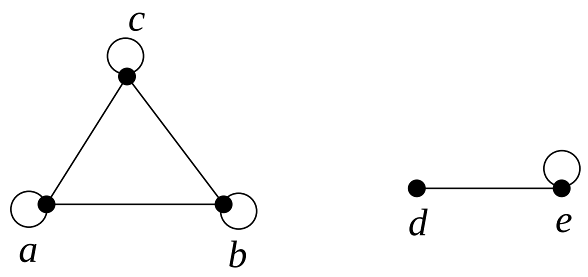

Figure 4: An H with a nontrivial component for which SAMPLEH-COL has a PAUS.

selects one uniformly at random. Note that the running time is at most poly$(1/\epsilon)$ in this case. Otherwise, the algorithm chooses i uniformly at random from $1, \ldots, 3^n + 2^n$. If $i \leq 3^n$, then the algorithm outputs the i'th colouring from the 3^n colourings with colours "a", "b", and "c". Otherwise, let C be the $(i - 3^n)$th of the 2^n (proper and improper) colourings with colours "d" and "e". If C is a legal H-colouring of G, then the algorithm outputs it. Otherwise, it outputs the error symbol $\perp$. The variation distance between the output distribution of the algorithm and the uniform distribution on H-colourings of G is at most the probability that the algorithm outputs $\perp$, which is at most $2^n/(2^n + 3^n) \leq \epsilon$. $\square$

8 Sampling and Counting

Let #BH-COL be defined as follows.

Name. #BH-COL.

Instance. A loop-free connected bipartite graph G.

Output. The number of H-colourings of G.

For certain graphs H, the problem #BH-COL can be expressed as the counting problem associated with a "self-reducible p-relation". For such an H, Theorem 6.3 of Jerrum, Valiant and Vazirani's paper [16] guarantees that if there is an FPRAS for #BH-COL then there is a PAUS for SAMPLEBH-COL. If H has no trivial components, this in turn guarantees (by Theorem 1) an FPRAS for #BIS. Dyer and Greenhill [7] have given a more general framework in which these ideas work: If, for a given graph H, the problem #BH-COL is "self-partitionable" then an FPRAS for #BH-COL can be turned into a PAUS for SAMPLEBH-COL. It is not clear for which graphs H these ideas can be applied, and this is an interesting open question.

9 Acknowledgements

We thank Martin Dyer, Paul Goldberg and Mark Jerrum for useful conversations.

References

[1] C. Borgs, J.T. Chayes, M. Dyer and P. Tetali, On the sampling problem for H-colorings on the hypercubic lattice, pre-print (2000).

[2] G.R. Brightwell and P. Winkler, Graph homomorphisms and phase transitions, *Journal of Combinatorial Theory, Series B,* **77** (1999) 221–262.

[3] C. Cooper, M. Dyer and A. Frieze, On Markov chains for randomly H-colouring a graph, *Journal of Algorithms,* **39(1)** (2001) 117–134.

[4] J. Diaz, M. Serna and D.M. Thilikos, (H, C, K)-coloring: fast, easy, and hard cases, To appear in MFCS (2001).

[5] J. Diaz, M. Serna and D.M. Thilikos, The complexity of parameterized H-colorings: a survey, pre-print (2000).

[6] M.E. Dyer, A.M. Frieze and M.R. Jerrum, On counting independent sets in sparse graphs, FOCS 1997, 223–231.

[7] M. Dyer and C. Greenhill, Random walks on combinatorial objects. In J.D. Lamb and D.A. Preece, editors, *Surveys in Combinatorics*, volume 267 of *London Mathematical Society Lecture Note Series*, pages 101–136. Cambridge University Press, 1999.

[8] M. Dyer and C. Greenhill, On Markov chains for independent sets, *Journal of Algorithms,* **35** (2000) 17–49.

[9] M.Dyer and C. Greenhill, The complexity of counting graph homomorphisms. *Random Structures and Algorithms,* **17** (2000) 260–289.

[10] M. Dyer, C. Greenhill, L.A. Goldberg, and M. Jerrum, On the relative complexity of approximate counting problems. In *Proceedings of APPROX*, volume 1913 of *Springer Lecture Notes in Computer Science*, pages 108–119, 2000.

[11] M. Dyer, M. Jerrum and E. Vigoda, Rapidly mixing Markov chains for dismantleable constraint graphs, Pre-print 2001.

[12] A. Galluccio, P. Hell and J. Nešetřil, The complexity of H-coloring of bounded degree graphs, *Discrete Mathematics,* **222** (2000) 101–109.

[13] P. Hell and J. Nešetřil, On the complexity of H-coloring, *Journal of Combinatorial Theory, Series B,* **48** (1990) 92–110.

[14] M. Jerrum, A very simple algorithm for estimating the number of k-colorings of a low-degree graph, *Random Structures and Algorithms,* **7** (1995) 157–165.

[15] M. Jerrum, Sampling and Counting, Chapter 3 of *Lecture Notes from a recent Nachdiplomvorlesung at ETH-Zürich "Counting, sampling and integrating: algorithms and complexity"*. Available at `http://www.dcs.ed.ac.uk/home/mrj/pubs.html` .

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

[17] E. Vigoda, Improved bounds for sampling colorings, *Journal of Mathematical Physics,* **41** (2000) 1555–1569.