Counting and sampling H-colourings *

Martin Dyer †
School of Computing
University of Leeds

Leslie Ann Goldberg‡
Dep't of Computer Science
University of Warwick

Mark Jerrum§
Division of Informatics
University of Edinburgh

September 1, 2003

Abstract

For counting problems in #P which are "essentially self-reducible", it is known that sampling and approximate counting are equivalent. However, many problems of interest do not have such a structure and there is already some evidence that this equivalence does not hold for the whole of #P. An intriguing example is the class of H-colouring problems, which have recently been the subject of much study, and their natural generalisation to vertex- and edge-weighted versions. Particular cases of the counting-to-sampling reduction have been observed, but it has been an open question as to how far these reductions might extend to any H and a general graph G. Here we give the first completely general counting-to-sampling reduction. For every fixed H, we show that the problem of approximately determining the partition function of weighted H-colourings can be reduced to the problem of sampling these colourings from an approximately correct distribution. In particular, any rapidly-mixing Markov chain for sampling H-colourings can be turned into an FPRAS for counting H-colourings.

1 Introduction

Jerrum, Valiant and Vazirani [14] showed that for *self-reducible* problems in #P, approximate counting and approximate sampling are of similar computational complexity. In particular, a problem has a *fully polynomial randomised approximation scheme* (FPRAS) if and only if it has a *fully polynomial approximate sampler* (FPAS). The techniques of [14] have been applied even to problems that do not seem to be self-reducible, and a generalization of [14] was given by Dyer and Greenhill [5]. In general, however, the situation seems more complicated, as exemplified by the following observation of Brightwell and Goldberg [1].

Observation 1 *There exists a problem in* #P *which has an FPRAS but no FPAS, unless there is a polynomial time algorithm for computing the discrete logarithm.*

*This work was partially supported by the EPSRC grant "Sharper Analysis of Randomised Algorithms: a Computational Approach", the EPSRC grant GR/R44560/01 "Analysing Markov-chain based random sampling algorithms" and the IST Programme of the EU under contract numbers IST-1999-14186 (ALCOM-FT) and IST-1999-14036 (RAND-APX). A preliminary version appeared in the proceedings of the 6th International Workshop on Randomization and Approximation Techniques in Computer Science.

†`http://www.comp.leeds.ac.uk/~dyer/`, School of Computing, University of Leeds, Leeds LS2 9JT, United Kingdom.

‡`http://www.dcs.warwick.ac.uk/~leslie/`, Department of Computer Science, University of Warwick, Coventry, CV4 7AL, United Kingdom.

§`http://www.dcs.ed.ac.uk/~mrj/`, School of Informatics, University of Edinburgh, JCMB, The King's Buildings, Edinburgh EH9 3JZ, United Kingdom.

Proof. Consider the problem with instances $(p, r, C(p, r), y)$, where $C(p, r)$ is a certificate that p is a prime with primitive root r, and $y \in \{1, \ldots, p-1\}$. The input can be verified in polynomial time. (See Section 10.2 and Example 12.2 of [15].) The solution set is defined to be $\{x \mid 0 \leq x \leq p-2, r^x = y \pmod{p}\}$. This problem trivially has an FPRAS, since the solution set is always of size 1 exactly. Furthermore, the problem is in #P, since $r^x \bmod p$ can be computed in polynomial time. However, an FPAS would clearly give a polynomial-time solution to the discrete logarithm problem. □

In fact, the proof of Observation 1 does not rely on the details of the discrete logarithm problem. Any *one-way permutation*[1] could be used to construct a #P-problem with an FPRAS but no FPAS. Thus it seems likely that there exist problems in #P which have an FPRAS but no FPAS. On the other hand, it is an open question as to whether, under any reasonable complexity assumption, there exist problems in #P which possess an FPAS but no FPRAS. A candidate problem might be the *orbit-counting problem* [8]. If a sampling algorithm were discovered which did not essentially implement Burnside's lemma, it would be unclear how to use it for approximate counting.

Despite these issues, it is widely believed that approximate counting and approximate sampling are inter-reducible in polynomial time for most, or even all, "reasonable" problems in #P. H-colouring[2] provides a convenient setting for investigating this issue. It is not known whether the H-colouring problem is self-reducible. Indeed, this is given as an open problem by Diaz [4]. However, we would like to understand the relationship between approximate counting and approximate sampling for this problem. On the one hand, reductions between approximate counting and sampling are known for several of the best-known instances of H-colouring. These include the (usual) vertex-colouring problem [12] (see also section 3 below) and the independent set problem or, more generally, its vertex-weighted version the *hard core lattice gas* model. (See, for instance, Examples 3.3 and 3.4 in [5].) On the other hand, straightforward attempts to apply the method of [14] to H-colouring seem to fail.

Dyer, Jerrum, and Vigoda [7] have shown how to extend the counting-to-sampling reduction from the vertex-colouring setting to the H-colouring setting, but their proof works only if H is dismantleable (which is quite a strong restriction, see [2]) and the input graph, G, has bounded degree. This paper extends their result to any H and to general graphs G. We show that, for every fixed H, the problem of approximately-counting H-colourings can be reduced to the problem of sampling H-colourings from an approximately-correct distribution. Thus, the MCMC method is applicable to H-colouring. In particular, any rapidly-mixing Markov chain for sampling H-colourings can be turned into an FPRAS for counting H-colourings. In fact, we express our results in the more general setting from Section 1.1 of [6] in which vertices and edges of H may have weights. Thus, we show that an algorithm for sampling from the Gibbs distribution leads to an FPRAS for the partition function.

The other direction is still open. The natural reduction from sampling H-colourings to counting H-colourings suffers from the defect that the resultant counting sub-problems correspond to *list colouring* problems rather than to unrestricted colouring problems. Thus, sampling may be reduced to the problem of (approximately) counting *list* H-colourings, but possibly not to the (presumably easier) problem of counting H-colourings. Thus, it is not clear for which graphs H negative sampling results such as [3, 10] yield negative results for approximability. Approximate counting could be easier than approximate sampling for H-colouring. Note that for almost every H, it is #P-hard to *exactly* count H-colourings (see [6]).

[1] The definition of a "one-way permutation" is beyond our scope — think of a one-way function which, for each n, is a permutation on inputs of size n. Details can be found in [11].

[2] See Section 2 for a precise definition of this and other basic notions mentioned in this introduction.

2 Definitions and Statement of Theorem 2

Our definitions are from Section 1.1 of [6]. Let $H = (V(H), E(H))$ be a fixed graph. We will allow H to have self-loops, but not multiple edges between a pair of vertices. Let $V(H) = \{c_1, \ldots, c_h\}$. We refer to the vertices of $V(H)$ as "colours". Every colour c_j has a weight $\lambda_{c_j} > 0$. If an unordered pair of colours (c_i, c_j) is in $E(H)$ then it has a weight $\lambda_{c_i,c_j} > 0$. Otherwise, it has zero weight, i.e. $\lambda_{c_i,c_j} = 0$. Let $\lambda_{\max}$ be the maximum of all vertex and edge weights in H.

Suppose that σ is a function from $V(G)$ to $V(H)$, where G is a simple graph, without multiple edges or self-loops. We assign the weight $w_\sigma(G)$ to σ, where $w_\sigma(G)$ is given by

$$w_\sigma(G) = \prod_{v \in V(G)} \lambda_{\sigma(v)} \prod_{(u,v) \in E(G)} \lambda_{\sigma(u),\sigma(v)}.$$

Note that $w_\sigma(G) > 0$ if and only if σ is a *homomorphism* from G to H. (A homomorphism from G to H is just a function σ from $V(G)$ to $V(H)$ which has the property that for every edge (u, v) of G, $(\sigma(u), \sigma(v))$ is an edge of H. A homomorphism from G to H is also known as an "H-colouring of G".) Let $\Omega_H(G)$ be the set of H-colourings of G. That is,

$$\Omega_H(G) = \{\sigma : V(G) \to V(H) \mid w_\sigma(G) > 0\}.$$

The partition function $Z_H(G)$ is given by

$$Z_H(G) = \sum_{\sigma \in \Omega_H(G)} w_\sigma(G). \tag{1}$$

The Gibbs distribution on H-colourings of G is the distribution in which each colouring σ has probability

$$\pi_{H,G}(\sigma) = \frac{w_\sigma(G)}{Z_H(G)}.$$

If u is a vertex of G and c_i is a colour in $V(H)$, we use the notation $Z_H(G)\{u \to c_i\}$ to denote $\sum_{\sigma \in \Omega_H(G), \sigma(u) = c_i} w_\sigma(G)$. We will use similar notation when we want to restrict more vertices of G to have particular colours.

As a technical matter, we can assume without loss of generality that there are not distinct colours $c_\alpha \in V(H)$ and $c_\beta \in V(H)$ with identical edge weights. That is, we do not have c_α and c_β such that, for all i, $\lambda_{c_\alpha,c_i} = \lambda_{c_\beta,c_i}$. It is straightforward to see that any two such colours can be treated as a single colour with effective vertex weight $\lambda_{c_\alpha} + \lambda_{c_\beta}$.

Since we are interested in computation (which is inherently discrete), we will assume that all of the weights λ_{c_j} and λ_{c_i,c_j} are rational. Now suppose that K is the least common multiple of the denominators of all of the positive weights. Consider what happens when replace the weights with $\hat{\lambda}_{c_j} = K\lambda_{c_j}$ and $\hat{\lambda}_{c_i,c_j} = K\lambda_{c_i,c_j}$. The weight of a colouring is then $\hat{w}_\sigma(G) = K^{n+m} w_\sigma(G)$, where $n = |V(G)|$ and $m = |E(G)|$. Similarly, $\hat{Z}_H(G) = K^{n+m} Z_H(G)$ and $\hat{\pi}_{H,G}(\sigma) = \pi_{H,G}(\sigma)$. Thus, we can assume without loss of generality that all weights λ_{c_j} and λ_{c_i,c_j} are natural numbers. We will make this assumption in the rest of this paper. See [5] and [9] for a further discussion of this issue.

We will consider the complexity of the following problems.

Name. H-Partition.

Instance. A graph G.

Output. The value of the partition function $Z_H(G)$.

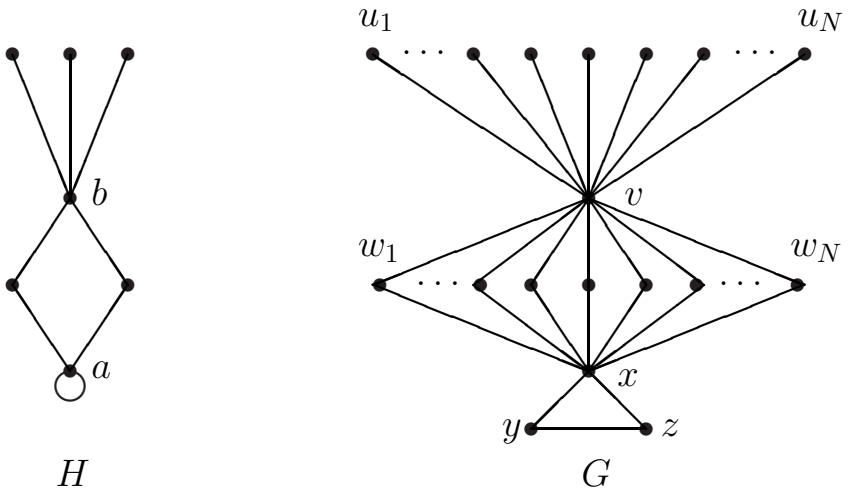

Figure 1: An H-colouring problem.

Name. H-GibbsSample.

Instance. A graph G.

Output. An H-colouring σ of G chosen from distribution $\pi_{H,G}$.

Note that if all vertex and edge weights of H are set to 1 then H-Partition is simply the problem of counting H-colourings of G and H-GibbsSample is the problem of sampling an H-colouring of G uniformly at random.

Figure 1 gives an example. The triangle xyz forces x to be coloured with a or one of its neighbours. It follows that, for large N, there are $\Theta(10^N)$ colourings where x is coloured a and v is coloured b, and $\Theta(9^N)$ other colourings. Thus "almost all" colourings are of the former type.

A *randomised approximation scheme* for H-Partition is a randomised algorithm that takes as input a graph G and an error tolerance $\varepsilon > 0$, and outputs a number $\widehat{Z} \in \mathbb{N}$ (a random variable of the "coin tosses" made by the algorithm) such that

$$\Pr\left[e^{-\varepsilon} Z_H(G) \leq \widehat{Z} \leq e^{\varepsilon} Z_H(G)\right] \geq \frac{3}{4}. \tag{2}$$

The algorithm is a *fully polynomial randomised approximation scheme*, or *FPRAS*, if it runs in time bounded by a polynomial in $|V(G)|$ and ε^{-1}.

In this paper we will simplify the presentation of our FPRAS by presenting it in a slightly different form. Our randomised algorithm will take as input an n-vertex graph G and an error tolerance $\varepsilon > 0$. With probability at least $1 - 2^{-n^5}$, it will succeed. In this case, the running time will be bounded from above by a polynomial in n and ε^{-1}. Also, it will output a number $\widehat{Z} \in \mathbb{N}$ such that

$$\Pr\left[e^{-\varepsilon} Z_H(G) \leq \widehat{Z} \leq e^{\varepsilon} Z_H(G)\right] \geq \frac{7}{8}. \tag{3}$$

If the algorithm fails, the running time might be as large as $\mathrm{poly}(n, \varepsilon^{-1})\, 2^{\binom{n'}{2}} |V(H)|^{n'}$, where $n' \in O(n^2)$. Note that the expected running time of our algorithm is at most a polynomial in n and ε^{-1}. Furthermore, our algorithm can be converted into a standard FPRAS by truncating long runs after polynomially many steps (and outputting an arbitrary answer after truncation).

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

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

An *approximate sampler* [5, 13, 14] for H-GibbsSample is a randomised algorithm that takes as input a graph G and an accuracy parameter $\varepsilon \in (0, 1]$ and gives an output (a random variable) such that the variation distance between the output distribution of the algorithm and the Gibbs distribution $\pi_{H,G}$ is at most ε. The algorithm is a *fully polynomial approximate sampler (FPAS)* if its running time is bounded from above by a polynomial in $|V(G)|$ and $\log(\varepsilon^{-1})$.

Theorem 2 *If there is an FPAS for H-GibbsSample then there is an FPRAS for H-Partition.*

3 An easy reduction

Our general strategy will be to reduce G to a tree by removing edges one by one, but unfortunately the reduction is not straightforward. We will need to attach "gadgets" to the vertices of G in order to exclude some undesirable colourings. These are discussed in section 4 below. But first, to illustrate some of the difficulties, we will sketch a simpler reduction, which suffices for two special cases of counting *unweighted* H-colourings. These are problems in which either *every* or *no* vertex of H has a loop. The usual vertex colouring problem provides an example.

Recall that $h = |V(H)|$. If G has h or fewer vertices, we will count its H-colourings by exhaustive enumeration. Otherwise, by applying the pigeonhole principle to any subset of $V(G)$ of size $(h+1)$, there must exist two vertices $u, v \in V(G)$ such that

$$\Pr\left(\sigma(u) = \sigma(v)\right) = \sum_{\sigma : \sigma(u) = \sigma(v)} \pi_{H,G}(\sigma) \geq \tbinom{h+1}{2}^{-1}.$$

Take sufficiently many samples to locate any pair u, v with $\Pr(\sigma(u) = \sigma(v)) \geq 1/h^2$. Now let G_{uv} be the graph obtained from G by identifying u and v as a single vertex uv. Parallel edges may be removed from G_{uv} since all edge weights of H are 1. However, there may be a loop on the vertex uv, which means it must be coloured with a looped vertex of H. In the case where H has no looped vertices, the situation does not arise (u and v will not be adjacent in G). In the case where H has all looped vertices, the uv loop is no restriction and we may remove the loop. By sampling colourings of G, we can estimate the ratio $\tau_{uv} = |\Omega_H(G_{uv})|/|\Omega_H(G)| \geq 1/h^2$. Now we estimate $|\Omega_H(G_{uv})|$ recursively, and hence estimate $|\Omega_H(G)|$ as $|\Omega_H(G_{uv})|/\tau_{uv}$.

This reduction is clearly invalid if H has vertex weights, since the vertex uv must receive a squared weighting in G_{uv}. Thus G itself becomes vertex-weighted. To proceed further, we must assume that we can sample H-colourings when G is a vertex-weighted graph. Similarly, if H has edge weights, the parallel edges in G_{uv} are significant, and we are soon obliged to deal with edge-weighted G. Even in the case where all weights are 1, but H has both looped and unlooped vertices, the reduction may be invalid, as illustrated by Figure 1. Here $y, z \in V(G)$ are both coloured a with frequency almost one-fifth. But, if we identify y and z, the vertex yz has a loop, signifying that it can only be coloured with the looped vertex $a \in V(H)$. If we ignore the loop on yz (in order to make G_{yz} a simple graph), the number of H-colourings of G_{yz} explodes to $\Theta(25^N)$, by colouring both x and v with b. The ratio τ_{yz} is now an exponentially large quantity rather than a fraction.

In general, this reduction is valid if we assume that the class of graphs from which G can be chosen includes the class from which H can be chosen. But this assumption is not true of the

H-colouring problem as usually stated, particularly in its weighted variants. Therefore we need to proceed more carefully to obtain our reduction.

4 Gadgets

Let $t = 2|V(H)|$. Let P be a path of $2t$ edges from some vertex A to some vertex B. For any colour $c_i \in V(H)$ and any colour $c_j \in V(H)$, recall that $Z_H(P)\{A \to c_i, B \to c_j\}$ denotes $\sum_{\sigma \in \Omega_H(P), \sigma(A)=c_i, \sigma(B)=c_j} w_\sigma(P)$. Let $\delta(c_i, c_j)$ be the quantity

$$\delta(c_i, c_j) = \frac{Z_H(P)\{A \to c_i, B \to c_j\}}{\lambda_{c_i}\lambda_{c_j}}.$$

The quantity $\delta(c_i, c_j)$ is the total weight of all H-colourings of P which start at colour c_i and end at colour c_j except that we exclude the weight of the colours at the two endpoints. For any colour c_i, let $\delta(c_i) = \delta(c_i, c_i)$.

We will be assuming that H is connected and that it has more than one vertex. Thus, every colour $c_i \in V(H)$ has at least one neighbour so $\delta(c_i) > 0$. If all vertices $c_i \in V(H)$ and $c_j \in V(H)$ have $\delta(c_i) = \delta(c_j)$ we will define $\delta^*(H) = 1$. Otherwise, we define $\delta^*(H)$ to be the following positive quantity.

$$\delta^*(H) = \min\left\{\log_2\left(\frac{\delta(c_i)}{\delta(c_j)}\right) \;\middle|\; c_i \in V(H), c_j \in V(H), \delta(c_i) > \delta(c_j)\right\}.$$

We will use the following technical lemma (cf. [6]).

Lemma 3 *If $\delta(c_i) \geq \delta(c_j)$ and $j \neq i$ then $\delta(c_i, c_j) < \delta(c_i)$.*

Proof. If c_i dominates c_j in the sense that $\lambda_{c_i c_\alpha} > \lambda_{c_j c_\alpha}$ for all α then the lemma follows from the definition of δ. Suppose that c_i does not dominate c_j. Let W be a symmetric $h \times h$ matrix in which the entry in row i and column j is λ_{c_i,c_j}. Let Λ be the diagonal matrix in which the entry in row i and column i is λ_{c_i}. Let Ψ be the positive diagonal matrix such that $\Psi^2 = \Lambda$. Let $[\,\cdot\,]_i$ denote the ith column of a matrix and $[\,\cdot\,]_{ij}$ its (i,j)th element. Note that

$$\delta(c_i, c_j) = [(W\Lambda)^{2t-1}W]_{ij} = [\Psi^{-1}(\Psi W \Psi)^{2t}\Psi^{-1}]_{ij}.$$

Since $\Psi W \Psi$ is symmetric, it can be written as $U^T L U$ where L is diagonal and U is orthonormal (i.e., $U^T U = I$).

Now

$$\begin{aligned}\delta(c_i, c_j) &= [\Psi^{-1}(\Psi W \Psi)^{2t}\Psi^{-1}]_{ij} = [\Psi^{-1}U^T L^{2t} U \Psi^{-1}]_{ij} = [L^t U \Psi^{-1}]_i^T [L^t U \Psi^{-1}]_j \\ &\leq \sqrt{[\Psi^{-1}U^T L^{2t} U \Psi^{-1}]_{ii}[\Psi^{-1}U^T L^{2t} U \Psi^{-1}]_{jj}} = \sqrt{\delta(c_i)\delta(c_j)} \leq \delta(c_i)\end{aligned}$$

using Cauchy-Schwartz, with strict inequality unless $[L^t U \Psi^{-1}]_i$ is a multiple of $[L^t U \Psi^{-1}]_j$. But this condition is true if and only if $[LU\Psi^{-1}]_i$ is a multiple of $[LU\Psi^{-1}]_j$, which is true if and only if $[\Psi^{-1}U^T L U \Psi^{-1}]_i$ is a multiple of $[\Psi^{-1}U^T L U \Psi^{-1}]_j$, i.e. $[W]_i$ is a multiple of $[W]_j$. This is impossible since c_i does not dominate c_j and there are not distinct colours with identical edge weights (see Section 2). □

We will let $\delta'(H)$ be the following positive quantity.

$$\delta'(H) = \min\left\{\log_2\left(\frac{\delta(c_i)}{\delta(c_i, c_j)}\right) \;\middle|\; i \neq j, \delta(c_i) \geq \delta(c_j)\right\}.$$

Finally, we let $\delta^\dagger(H) = \min(\delta^*(H), \delta'(H))$.

Let S be a subset of $V(H)$ such that for every colour $c_i \in S$ and every colour $c_j \in S$, $\delta(c_i) = \delta(c_j)$. Let $\delta(S)$ denote $\delta(c_i)$ for $c_i \in S$.

A graph H' with a designated vertex u' is said to be "good" for S if it satisfies the following properties.

i. For every $c \in S$, $Z_H(H')\{u' \to c\} > 0$, and

ii. for every colour $c \in V(H)$ with $\delta(c) > \delta(S)$, $Z_H(H')\{u' \to c\} = 0$.

Informally, (H', u') is good for S if every colour $c \in S$ can be applied to u' in a valid colouring but no vertex of higher δ-value can be applied to u'.

The set S is said to be "good" if there exists an (H', u') which is good for S. If S is good then κ_S is then defined to be the minimum number of vertices in a graph H' such that some pair (H', u') is good for S. κ is defined to be the maximum of κ_S over all good S. We will not assume that κ is known in our algorithm, but we will refer to it in our analysis.

Suppose we have a (fixed-size) graph H' with a designated vertex u' and we want to check whether the pair (H', u') is good for S. We do this by examining each of the (at most $|V(H)|^{|V(H')|}$) colourings in $\Omega_H(H')$. Thus, we can check every graph H' of size at most n' and every possible designated vertex u' by examining at most $(n')^2 2^{\binom{n'}{2}} |V(H)|^{n'}$ colourings.

As we observed in section 3, the function of these gadgets (H', u') is to exclude unwanted colourings. The triangle $H' = xyz$ in Figure 1, with distinguished vertex $u' = x$, illustrates the phenomenon. It has colourings in which x can be coloured a, but none in which it can receive the colour b of larger H-degree. Its attachment to G at x therefore excludes the (otherwise more numerous) colourings of G in which x would be coloured b.

5 Proof of Theorem 2

Suppose that G has multiple connected components, say G_A, G_B and G_C. It is immediate from Equation (1) that $Z_H(G) = Z_H(G_A)Z_H(G_B)Z_H(G_C)$. Thus, we may assume without loss of generality that G is connected. We will do so for the rest of the paper. For connected G, suppose that H has multiple connected components, say H_A, H_B and H_C. Inspection of Equation (1) reveals that $Z_H(G) = Z_{H_A}(G) + Z_{H_B}(G) + Z_{H_C}(G)$. Thus, we may assume without loss of generality that H is also connected. We will do so for the rest of the paper.

Let n denote $|V(G)|$. As in Section 3, we avoid trivialities by assuming $n \geq h$. In the reduction, we will construct a sequence G_0, G_1, ..., G_p of connected graphs. As long as there is no failure in the (randomised) reduction, the following properties will hold.

(i) $G_0 = G$,

(ii) $Z_H(G_p)$ can be calculated in polynomial time (polynomial in n), and

(iii) the construction of $G_0, \ldots, G_p$ will take polynomial time.

Let

$$\varrho_i = \frac{Z_H(G_i)}{Z_H(G_{i+1})}.$$

Then

$$Z_H(G) = \varrho_0 \varrho_1 \cdots \varrho_{p-1} Z_H(G_p).$$

We will estimate $Z_H(G)$ using the method of Jerrum, Valiant and Vazirani [14]. In particular, we will define a quantity s_i for each i such that

(iv) either ϱ_i or ϱ_i^{-1} is an easily-computable multiple of s_i (so an approximation to s_i gives an approximation to ϱ_i), and

(v) there is an experiment which can be performed using a perfect sampler for H-GibbsSample with input G_i or G_{i+1} for which the output is a 0/1 random variable with mean s_i, and

(vi) there is a polynomial q in n and ε^{-1} such that ${s_i}^{-1} \leq q(n, \varepsilon^{-1})$.

It follows (see the proof of Proposition 3.4 of [13]) that $O(\, q(n, \varepsilon^{-1})\, p\, \varepsilon^{-2}\,)$ samples taken from an approximate sampler for H-GibbsSample with accuracy parameter

$$O\left(\frac{\varepsilon}{q(n, \varepsilon^{-1})p}\right)$$

give a sufficiently accurate approximation to s_i, and hence to ϱ_i. That is, if we use the sampler with this many samples, and multiply our estimates of the ϱ_is, the resulting estimate of $|\Omega_H(G)|$ is within the required accuracy with probability at least 7/8.

As mentioned earlier, the overall probability of failure in our reduction (which could cause one or more of (i)–(vi) to fail) will be at most $m2^{-n^6} \leq 2^{-n^5}$, where $m = |E(G)|$. Suppose that $E(G) = \{e_1, \ldots, e_m\}$, and that the edges are ordered in such a way that the graph $(V(G), \{e_1, \ldots, e_{n-1}\})$ is a tree. We will use the notation $u(e_i)$ and $v(e_i)$ to denote the endpoints of the edge e_i. We will now describe the construction of the sequence $G_0, \ldots, G_p$.

In order to shorten our description of the reduction, we will break the sequence $G_0, \ldots, G_p$ into a number of subsequences. The rough intuition is that each subsequence does the job of "removing" one edge of G (with the eventual goal of producing a tree, whose colourings we can count directly). Removing the edge directly could cause the number of colourings to explode as discussed in Section 3 so we use the subsequence to remove the edge in a more controlled manner. In the following pages, we will show how to construct a subsequence

$$\Gamma_0 = G_y, \Gamma_1 = G_{y+1}, \ldots, \Gamma_{2r+2} = G_{y+2r+2},$$

where y is a multiple of $2r+2$ for some number r to be chosen later. The sequence $G_0, \ldots, G_p$ will then be the concatenation of the subsequences. We will rely on some properties which will always be true for the graph Γ_0, which is the first graph in each subsequence. First, it will be the case that for some $j \in \{n, \ldots, m\}$, the graph $\Gamma_0 = G_y$ is identical to $(V(G), \{e_1, \ldots, e_j\})$ except that every vertex $u \in V(G)$ may have one or more gadgets attached to u in Γ_0. Each gadget is simply a graph H' of size at most $\kappa + 1$. Note that κ is $O(1)$ as a function of n and m since H' does not depend upon G. If a gadget H' is attached to vertex u then one of the vertices of H' is identified with u. There are $2(m-j)$ gadgets in all, and these are distributed over the n vertices of G (so some vertices may have multiple gadgets). Since $m \leq \binom{n}{2}$, Γ_0 has $O(n^2)$ vertices. Another property that will always be true is that Γ_0 will be connected. As an invariant in the construction, we will also guarantee that each graph Γ_i has at least one H-colouring. That is, we will guarantee that $Z_H(\Gamma_i) > 0$.

The first sequence of graphs that we will construct will start with $\Gamma_0 = G_0$ and $j = m$, so all of the invariants will be true initially.

The rest of this section has the following structure. Parts 1A and 1B show how to construct the subsequence $\Gamma_1, \ldots, \Gamma_{2r+2}$ given the starting graph Γ_0. Part 1A shows how to do sampling in order to build gadgets that will be used in the subsequence and Part 1B shows how to build the subsequence itself. The sequence $G_0, \ldots, G_p$ is simply the concatentation of the subsequences constructed in Parts 1A and 1B. Part 2 shows how to finish the proof once $G_0, \ldots, G_p$ is constructed. In particular, it shows how $|\Omega_H(G_p)|$ can be computed.

Recall that Γ_0 looks like the graph $(V(G), \{e_1, \ldots, e_j\})$ except for possibly some small gadgets. Our goal in the subsequence $\Gamma_0, \ldots, \Gamma_{2r+2}$ will be to remove the edge e_j.

Part 1A: Learning about the graph Γ_0 and constructing $H(S)$

Before we can remove e_j it will help us to know which colours in $V(H)$ can be used to colour $u(e_j)$ in Γ_0. More particularly, we would like to know which colours in $V(H)$ are good choices for $u(e_j)$ when we modify Γ_0 by attaching a certain structure to $u(e_j)$.

Thus, we will first define a graph Γ_0' which is the same as Γ_0 except that a certain structure (which we will call F) will be attached to $u(e_j)$. We will then use our sampling oracle to study the colourings of Γ_0', paying particular attention to which colours are applied to vertex $u(e_j)$. Once we know the colours, we will use this information in the construction of $\Gamma_1, \Gamma_2, \ldots$ We start with some definitions. Let M be a straightforward upper bound for $Z_H(\Gamma_0)$. In particular, we can take

$$M = |V(H)|^{|V(\Gamma_0)|} \lambda_{\max}^{|V(\Gamma_0)|+|E(\Gamma_0)|}. \tag{4}$$

Recall that $t = 2|V(H)|$ and that $\delta^\dagger(H)$ is the minimum of the quantities $\delta^*(H)$ and $\delta'(H)$ from Section 4. Let r be defined by the following equation.

$$r = \left\lceil \frac{n^7 + \log_2(M)}{\delta^\dagger(H)} \right\rceil.$$

For now, the reader should just think of r as being a sufficiently large polynomial in n. Let F be the graph with the vertex set

$$V(F) = \{f_0\} \cup \bigcup_{p\in[1,\ldots,r], q\in[1,\ldots,2t-1]} \{f_{p,q}\}$$

and the edge set $E(F)$ which is defined to be

$$\bigcup_{p\in[1,\ldots,r]} \{(f_0, f_{p,1})\} \cup \bigcup_{p\in[1,\ldots,r], q\in[1,\ldots,2t-2]} \{(f_{p,q}, f_{p,q+1})\} \cup \bigcup_{p\in[1,\ldots,r]} \{(f_{p,2t-1}, f_0)\}.$$

F looks like a "flower" with vertex f_0 at the centre and r petals. Each petal is a cycle of length $2t$ which starts and ends at f_0.

Let Γ_0' be a graph constructed from Γ_0 by attaching F. Vertex f_0 of F should be identified with vertex $u(e_j)$.

We now need some notation to describe the colourings of Γ_0' and of Γ_0. For any $d \in \mathbb{N}$, let

$$Z_H(\Gamma_0)\{u(e_j) \to [d]\} = \sum_{c\in V(H), \delta(c)=d} Z_H(\Gamma_0)\{u(e_j) \to c\}.$$

Informally, $Z_H(\Gamma_0)\{u(e_j) \to [d]\}$ is the collective weight of all colourings in which $u(e_j)$ is coloured with a colour with δ-value d.

Define δ to be the quantity such that $Z_H(\Gamma_0)\{u(e_j) \to [\delta]\} > 0$ but, for all $d > \delta$, $Z_H(\Gamma_0)\{u(e_j) \to [d]\} = 0$. Informally, δ is the largest δ-value which can be applied to $u(e_j)$.

Let S^+ be the set of all colours c with $\delta(c) = \delta$ and $Z_H(\Gamma_0)\{u(e_j) \to c\} > 0$. Let S^- be $\big\{c \in S^+ \;\big|\; Z_H(\Gamma_0)\{u(e_j) \to c\} \geq (1/n) Z_H(\Gamma_0)\{u(e_j) \to [\delta]\}\big\}$. Thus, S^+ is the set of all value-δ colours which may be applied to $u(e_j)$ and S^- is the set of "frequently used" ones. Note that S^- is non-empty since there are fewer than n colours.

We will now describe an experiment which can be performed on Γ_0' to determine the "likely" colours that colour vertex $u(e_j)$. In the reduction, we will perform the experiment to learn about

these colours. This knowledge will be used in the construction of Γ_1. Suppose that we run H-GibbsSample with input Γ_0' and accuracy parameter $\gamma = 2^{-n^7}$ to collect $s = 2n^8$ samples from $\Omega_H(\Gamma_0')$. Let S be the collection of colours that are assigned to $u(e_j)$ in these samples.

We claim that, except with failure probability at most 2^{-n^6}, we have $S^- \subseteq S \subseteq S^+$. To see that the failure probability is this small first observe that the probability that a colour c with $\delta(c) < \delta$ is in S is at most

$$s\left(\frac{Z_H(\Gamma_0')\{u(e_j) \to c\}}{Z_H(\Gamma_0')} + \gamma\right). \tag{5}$$

Since $\delta(c)$ is the total weight of all colourings of a "petal" of F in which $u(e_j)$ is coloured c, the quantity (5) is at most

$$s\left(\frac{Z_H(\Gamma_0)\{u(e_j) \to c\}\delta(c)^r}{\delta^r} + \gamma\right).$$

Since $\delta^\dagger(H) \leq \delta^*(H)$ (see the definition of $\delta^*(H)$ in Section 4), the definition of r guarantees that the term $\frac{Z_H(\Gamma_0)\{u(e_j)\to c\}\delta(c)^r}{\delta^r} \leq \gamma$. Thus the probability that there exists a colour c with $\delta(c) < \delta$ in S is at most $s\,|V(H)|\,2\gamma$.

Also, the probability that a colour $c \in S^-$ is left out of S is at most

$$\left(1 - \frac{1}{n} + \sum_{c:\delta(c)<\delta} \frac{Z_H(\Gamma_0')\{u(e_j) \to c\}}{Z_H(\Gamma_0')} + \gamma\right)^s \leq \left(1 - \frac{1}{2n}\right)^s \leq \exp(-n^7),$$

so the probability that such a colour exists is at most $|V(H)|\exp(-n^7)$ and the sum of the failure probabilities is $s\,|V(H)|\,2\gamma + |V(H)|\exp(-n^7) \leq 2^{-n^6}$.

We have shown that, except with failure probability at most 2^{-n^6}, we have $S^- \subseteq S \subseteq S^+$. The reduction now begins searching for a graph $H(S)$ with a designated vertex $u'(S)$ which is good for S. (See Section 4.) If we do not have failure, then the pair $(H(S), u'(S))$ exists and $|V(H(S))| \leq \kappa$. Recall that κ is a constant depending only on H, and not on our input Γ_0. If there is no failure, then our input Γ_0 does provide an upper bound for $|V(H(S))|$ since $|V(H(S))| \leq |V(\Gamma_0)|$. The latter follows from the fact that $(\Gamma_0, u(e_j))$ is good for S. Thus we restrict the search to graphs with at most $|V(\Gamma_0)|$ vertices and the expected time of the search is at most a polynomial in n.

Part 1B: Constructing the sequence $\Gamma_1, \ldots, \Gamma_{2r+2}$ from Γ_0

In this part we will show how to construct $\Gamma_1, \ldots, \Gamma_{2r+2}$ assuming that we did not have failure in Part 1A. Recall that Γ_0 looks like the graph $(V(G), \{e_1, \ldots, e_j\})$ except for possibly some small gadgets. Our goal in constructing $\Gamma_1, \ldots, \Gamma_{2r+2}$ is to remove the edge e_j. Removing the edge directly could cause the number of colourings to explode as in Section 3. Instead, we gradually build up some "scaffolding" gadgetry, which will prevent the number of colourings from exploding when edge e_j is removed. After removing e_j, we have to take away the scaffolding, again working gradually to keep the number of colourings under control. Ideally, we would like Γ_{2r+2} to look exactly like Γ_0 except for the removal of e_j. What happens in the construction is that Γ_{2r+2} looks like Γ_0 except for the removal of e_j and the addition of two $O(1)$-sized gadgets. We now describe the construction in detail.

First, the graphs $\Gamma_1, \ldots, \Gamma_r$ are constructed. For $i \in [0, \ldots, r-1]$, Γ_{i+1} is constructed from Γ_i by adding a length-$2t$ cycle $\{u(e_j), f_{i+1,1}, \ldots, f_{i+1,2t-1}, u(e_j)\}$ where $f_{i+1,1}, \ldots, f_{i+1,2t-1}$ are new vertices.

For every $\sigma \in \Omega_H(\Gamma_i)$, let $\mathrm{ext}(\sigma)$ be the non-empty set

$$\mathrm{ext}(\sigma) = \{\sigma' \in \Omega_H(\Gamma_{i+1}) \mid \forall\, v \in V(\Gamma_i), \sigma(v) = \sigma'(v)\}.$$

($\mathrm{ext}(\sigma)$ is non-empty because every colour in H has at least one neighbour.) For every $\sigma' \in \mathrm{ext}(\sigma)$, let $\hat{w}(\sigma,\sigma') = w_{\sigma'}(\Gamma_{i+1})/w_\sigma(\Gamma_i)$. Note that $\hat{w}(\sigma,\sigma') \geq 1$ since all vertex and edge weights are positive integers. Let $s_i = \varrho_i$. Note that (iv) is satisfied for the graph Γ_i since ϱ_i *is* s_i, so it is clearly "an easily-computable multiple of s_i". We now wish to establish (v) for the graph Γ_i. We wish to exhibit an experiment which can be performed using a perfect sampler for H-GibbsSample with input Γ_i or Γ_{i+1} for which the output is a 0/1 random variable with mean s_i. Here is the experiment: Choose σ' from $\pi_{H,\Gamma_{i+1}}$. Let σ be the restriction of σ' to $V(\Gamma_i)$. Output 1 with probability $\left(\hat{w}(\sigma,\sigma')\,|\mathrm{ext}(\sigma)|\right)^{-1}$ and 0 otherwise. The probability that a 1 is output is

$$\begin{aligned}
\frac{1}{Z_H(\Gamma_{i+1})} & \sum_{\sigma' \in \Omega_H(\Gamma_{i+1})} w_{\sigma'}(\Gamma_{i+1}) \frac{1}{\hat{w}(\sigma,\sigma')\,|\mathrm{ext}(\sigma)|} \\
&= \frac{1}{Z_H(\Gamma_{i+1})} \sum_{\sigma \in \Omega_H(\Gamma_i)} \sum_{\sigma' \in \mathrm{ext}(\sigma)} w_\sigma(\Gamma_i) \frac{1}{|\mathrm{ext}(\sigma)|} \\
&= \frac{Z_H(\Gamma_i)}{Z_H(\Gamma_{i+1})} = s_i.
\end{aligned}$$

Thus, (v) is satisfied. Finally, we must satisfy (vi). That is, we must show that there is a polynomial q in n and ε^{-1} such that ${s_i}^{-1} \leq q(n,\varepsilon^{-1})$. Since $Z_H(\Gamma_{i+1}) \leq Z_H(P)Z_H(\Gamma_i)$ where P is a length-$2t$ path, we have $s_i^{-1} \leq Z_H(P)$, so (vi) holds. We have now completed the construction of the graphs $\Gamma_1,\ldots,\Gamma_r$ and the argument that these graphs satisfy our requirements. Note that the graph Γ_r is the same as the graph Γ'_0 which we considered in Part 1A.

Next, the graph Γ_{r+1} is constructed from Γ_r by attaching $H(S)$ to $u(e_j)$, identifying the vertex $u(e_j)$ of Γ_r with the vertex $u'(S)$ in the gadget $H(S)$. That is $V(\Gamma_{r+1}) = V(\Gamma_r) \cup V(H(S))$, but $|V(\Gamma_{r+1})| = |V(\Gamma_r)| + |V(H(S))| - 1$ since the vertex $u(e_j)$ of Γ_r is identified with the vertex $u'(S)$ of $H(S)$. Also, $E(\Gamma_{r+1}) = E(\Gamma_r) \cup E(H(S))$. Since $|V(H(S))| \leq \kappa$, the construction of Γ_{r+1} is fast.

We will now show that (iv), (v) and (vi) hold for Γ_{r+1} (i.e., for $i = r$). Let $s_r = \varrho_r^{-1} Z_H(H(S))^{-1}$. Consider the following experiment. Choose σ from π_{H,Γ_r}. Output 1 with probability

$$\sum_{\sigma' \in \mathrm{ext}(\sigma)} \frac{\hat{w}(\sigma,\sigma')}{Z_H(H(S))},$$

where

$$\mathrm{ext}(\sigma) = \{\sigma' \in \Omega_H(\Gamma_{r+1}) \mid \forall\, v \in V(\Gamma_r), \sigma(v) = \sigma'(v)\}$$

as above and $\hat{w}(\sigma,\sigma') = w_{\sigma'}(\Gamma_{r+1})/w_\sigma(\Gamma_r)$. Output 0 otherwise. The probability that 1 is output is

$$\frac{1}{Z_H(\Gamma_r)} \sum_{\sigma \in \Omega_H(\Gamma_r)} w_\sigma(\Gamma_r) \sum_{\sigma' \in \mathrm{ext}(\sigma)} \frac{\hat{w}(\sigma,\sigma')}{Z_H(H(S))} \;=\; \frac{1}{Z_H(H(S))} \frac{1}{Z_H(\Gamma_r)} Z_H(\Gamma_{r+1}) = s_r.$$

We must now establish (vi).

$$s_r^{-1} = \frac{Z_H(H(S))Z_H(\Gamma_r)}{Z_H(\Gamma_{r+1})} \leq \frac{Z_H(H(S))Z_H(\Gamma_r)}{\sum_{c \in S} Z_H(\Gamma_r)\{u(e_j) \to c\}}, \tag{6}$$

where the inequality follows from the fact that $(H(S), u(e_j))$ is good for S. Now from our analysis

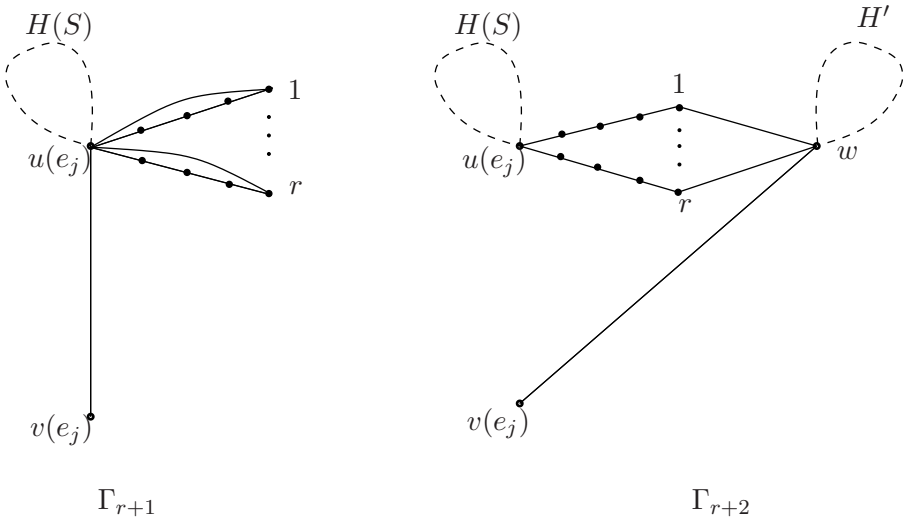

Figure 2: The construction of Γ_{r+2}

in Part 1A we have $Z_H(\Gamma_r) = Z_H(\Gamma'_0) \leq 2\delta^r Z_H(\Gamma_0)\{u(e_j) \to [\delta]\}$. Also, since $S^- \subseteq S$,

$$\begin{aligned}\sum_{c\in S} Z_H(\Gamma_r)\{u(e_j)\to c\} &\geq \delta^r \sum_{c\in S^-} Z_H(\Gamma_0)\{u(e_j)\to c\}\\ &\geq \delta^r \sum_{c\in S^-} (1/n) Z_H(\Gamma_0)\{u(e_j)\to[\delta]\},\end{aligned} \tag{7}$$

where the final inequality follows from the definition of S^-. Thus,

$$s_r^{-1} \leq \frac{Z_H(H(S))Z_H(\Gamma_r)}{\sum_{c\in S} Z_H(\Gamma_r)\{u(e_j)\to c\}} \leq Z_H(H(S))2n,$$

which gives us (vi).

The graph Γ_{r+2} is constructed from Γ_{r+1} as follows. Let H' be a new copy of the gadget $H(S)$. Let w be the designated vertex of H' so that (H', w) is good for S. To form Γ_{r+2}, we join together Γ_{r+1} and H'. Thus, $V(\Gamma_{r+2}) = V(\Gamma_{r+1}) \cup V(H')$. We do the "joining" by deleting the edges $(f_{i,2t-1}, u(e_j))$ (for $i \in [1,\ldots,r]$) and adding in edges $(f_{i,2t-1}, w)$ for each such i. Also, we delete the edge $(u(e_j), v(e_j))$ and add in edge $(w, v(e_j))$. See Figure 2.

Now let $s_{r+1} = \varrho_{r+1}$. Consider the following experiment. Choose σ' from the distribution $\pi_{H,\Gamma_{r+2}}$. If $\sigma'(w) = \sigma'(u(e_j))$ then output 1 with probability $(Z_H(H')\{w \to \sigma'(w)\})^{-1}$. Otherwise output a 0. The probability that 1 is output is

$$\begin{aligned}\frac{1}{Z_H(\Gamma_{r+2})} \sum_{\sigma'\in\Omega_H(\Gamma_{r+2}),\sigma'(w)=\sigma'(u(e_j))} w_{\sigma'}(\Gamma_{r+2}) \frac{1}{Z_H(H')\{w\to\sigma'(w)\}} &= \varrho_{r+1}\\ &= s_{r+1}.\end{aligned}$$

We must now establish (vi).

Now

$$Z_H(\Gamma_{r+2}) = \sum_{c_\alpha, c_\beta} Z_H(\Gamma_{r+2})\{u(e_j)\to c_\alpha, w\to c_\beta\}. \tag{8}$$

Also,

$$Z_H(\Gamma_{r+2})\{u(e_j) \to c_\alpha, w \to c_\beta\} \leq M Z_H(H(S)) Z_H(H') \delta(c_\alpha, c_\beta)^r,$$

where M is our upper bound for $Z_H(\Gamma_0)$ from Equation (4). On the other hand, we have just shown in our proof of (6) and (7) that

$$Z_H(\Gamma_{r+1}) \geq Z_H(\Gamma_{r+1})\{u(e_j) \to [\delta]\} \geq \delta^r (1/n).$$

Thus if $c_\alpha \neq c_\beta$

$$\begin{aligned} \frac{Z_H(\Gamma_{r+2})\{u(e_j) \to c_\alpha, w \to c_\beta\}}{Z_H(\Gamma_{r+1})} &\leq \frac{nM Z_H(H(S)) Z_H(H') \delta(c_\alpha, c_\beta)^r}{\delta^r} \\ &\leq n Z_H(H(S)) Z_H(H') \gamma, \end{aligned} \tag{9}$$

by the definition of r since $\delta^\dagger(H) \leq \delta'(H)$ (see the definition of $\delta'(H)$ in Section 4).

Finally,

$$\begin{aligned} Z_H(\Gamma_{r+2})\{u(e_j) \to c_\alpha, w \to c_\alpha\} &\leq Z_H(\Gamma_{r+1})\{u(e_j) \to c_\alpha\} Z_H(H') \\ &\leq Z_H(\Gamma_{r+1}) Z_H(H'). \end{aligned} \tag{10}$$

Putting together (8) and (9) and (10) we get

$$\begin{aligned} s_{r+1}^{-1} &= \varrho_{r+1}^{-1} \\ &= \frac{Z_H(\Gamma_{r+2})}{Z_H(\Gamma_{r+1})} \\ &\leq \sum_{c_\alpha \neq c_\beta} \left(n Z_H(H(S)) Z_H(H') \gamma \right) + \sum_{c_\alpha} Z_H(H'), \end{aligned}$$

which gives us (vi).

For $i \in \{0, \ldots, r-2\}$, graph $\Gamma_{r+2+i+1}$ is constructed from graph Γ_{r+2+i} by deleting vertices $f_{i+1,1}, \ldots, f_{i+1,2t-1}$ (and the edges adjacent to these vertices).

To establish Property (v) we define a notion which is analogous to $\text{ext}(\sigma)$. In particular, for every $\sigma \in \Omega_H(\Gamma_{r+2+i+1})$ let $\text{bext}(\sigma)$ be the non-empty set

$$\text{bext}(\sigma) = \{\sigma' \in \Omega_H(\Gamma_{r+2+i}) \mid \forall\, v \in V(\Gamma_{r+2+i+1}), \sigma(v) = \sigma'(v)\}.$$

For every $\sigma' \in \text{bext}(\sigma)$, let $\hat{w}(\sigma, \sigma') = w_{\sigma'}(\Gamma_{r+2+i}) / w_\sigma(\Gamma_{r+2+i+1}) \geq 1$. The following experiment has mean $s_{r+2+i} = \varrho_{r+2+i}^{-1}$. Choose σ' from $\pi_{H,\Gamma_{r+2+i}}$. Let σ be the restriction of σ' to $V(\Gamma_{r+2+i+1})$. With probability $(\hat{w}(\sigma, \sigma') \, |\text{bext}(\sigma)|)^{-1}$, output 1. Otherwise, output 0. Since $Z_H(\Gamma_{r+2+i}) \leq Z_H(P) Z_H(\Gamma_{r+2+i+1})$, we have $\varrho_{r+2+i} \leq Z_H(P)$, so (vi) holds.

Note that the graph Γ_{2r+1} is the same as the graph Γ_0 except that the gadget $H(S)$ has been attached to $u(e_j)$ and the edge $(u(e_j), v(e_j))$ has been replaced with the path $u(e_j), f_{r,1}, \ldots, f_{r,2t-1}, w, v(e_j)$ and the gadget H' has been attached to w.

Finally, the graph Γ_{2r+2} is constructed from Γ_{2r+1} by deleting vertices $f_{r,1}, \ldots, f_{r,2t-1}$ (and the edges adjacent to theses vertices).

The proof that Property (v) and Property (vi) hold is similar to what we have just done with $s_{2r+1} = \varrho_{2r+1}^{-1}$. The new difficulty is showing that for every $\sigma \in \Omega_H(\Gamma_{2r+2})$, the set

$$\text{bext}(\sigma) = \{\sigma' \in \Omega_H(\Gamma_{2r+1}) \mid \forall\, v \in V(\Gamma_{2r+2}), \sigma(v) = \sigma'(v)\}$$

is non-empty.

Suppose that σ is a colouring $\in \Omega_H(\Gamma_{2r+2})$ in which $u(e_j)$ is coloured with colour a and w is coloured with colour b. We must show that there is a colouring of the path $u(e_j), f_{r,1}, \ldots, f_{r,2t-1}, w$ in which $u(e_j)$ is coloured a and w is coloured b. We will do this by looking at two cases.

Case 1: H is a loopless bipartite graph. Recall that (by construction) Γ_{2r+1} has at least one H-colouring. This means that Γ_{2r+1} is bipartite. Also, $u(e_j)$ and w are in the same part of the vertex partition of Γ_{2r+1}. The graph Γ_{2r+2} is still connected (by construction) with $u(e_j)$ and w in the same part. This means that a and b are from the same side of H's vertex partition. Since H is connected, there is an even-length path from a to b of length at most $|V_H| - 1$. Thus, there is a walk of length $2t$ from a to b. (Take the path above and go back and forth on the last edge.)

Case 2: H is not a loopless bipartite graph, so it has an odd-length cycle of length at most $|V(H)|$. In this case, let c be some node on the cycle. We will construct an even-length path from a to b of length less than $2t$: First go from a to c using at most $|V(H)|-1$ edges. Then go from c to b using at most $|V(H)| - 1$ edges. Finally, if the constructed path has odd length, then go around the odd-length cycle in the middle. The total number of edges is at most $3|V(H)| - 2 < 2t$. Once again, we can find a walk of length $2t$ from a to b by going back and forth on the last edge.

This completes the argument that Γ_{2r+2} is properly constructed and it completes the construction of $\Gamma_1, \ldots, \Gamma_{2r+2}$. Thus, we have constructed the sequence

$$\Gamma_0 = G_y, \Gamma_1 = G_{y+1}, \ldots, \Gamma_{2r+2} = G_{y+2r+2}$$

as required. Note that the graph Γ_{2r+2} is identical to $(V(G), \{e_1, \ldots, e_{j-1}\})$ except that every vertex $u \in V(G)$ may have some gadgets attached to u in Γ_{2r+2}. The gadgets that are present in Γ_{2r+2} which were not present in Γ_0 are the new gadget $H(S)$ (of size at most κ) which is attached to $u(e_j)$ and the new gadget consisting of vertex w and the graph H' (of total size at most $\kappa + 1$) which is attached to $v(e_j)$. If $j = n$ then we are finished and $y + 2r + 2 = p$. Otherwise, we start Part 1A again with $\Gamma_0 = G_{y'} = G_{y+2r+2}$.

Part 2: Computing $|\Omega_H(G_p)|$

We have now shown how to construct $G_0, \ldots, G_p$. We have shown that our construction satisfies (i), (iii), (iv), (v) and (vi). It remains to show that property (ii) is satisfied – namely, that we can compute $Z_H(G_p)$ in polynomial time (polynomial in n).

By construction, G_p is identical to the tree $T = (V(G), \{e_1, \ldots, e_{n-1}\})$ except that every vertex $u \in V(G)$ may have some gadgets attached to u in G_p. Each gadget is a graph H' of size at most $\kappa + 1$. One of the vertices of H' is identified with u. There are $2(m - n + 1)$ gadgets in all.

We can compute $Z_H(G_p)$ by dynamic programming. For each gadget (H', u') and each colour c, we first compute $Z_H(H')\{u' \to c\}$.

Now consider a rooted version of T. For each vertex $v \in V(G)$, let $G_p(v)$ denote the portion of G_p corresponding to the sub-tree rooted at v in T (including attached gadgets). We can calculate $Z_H(G_p(v))\{v \to c\}$ using the values of $Z_H(G_p(v'))\{v \to c'\}$ for all children v' of v in T and all colours $c' \in V(H)$ and all quantities $Z_H(H')\{u' \to c''\}$.

References

[1] G.R. Brightwell and L.A. Goldberg, personal communication.

[2] G.R. Brightwell and P. Winkler, Gibbs measures and dismantlable graphs, *J. Combin. Theory Ser. B* 78(1) 141–166 (2000)

[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. Díaz, H-colorings of Graphs, The Algorithmics Column, Bulletin of the EATCS **72** (October, 2001) 82–92.

[5] 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.

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

[7] M. Dyer, M. Jerrum and E. Vigoda, Rapidly mixing Markov chains for dismantleable constraint graphs. In J. Nesetril and P. Winkler, editors, *Proceedings of a DIMACS/DIMATIA Workshop on Graphs, Morphisms and Statistical Physics*, March 2001, to appear.

[8] L.A. Goldberg, Computation in permutation groups: counting and randomly sampling orbits. In J.W.P. Hirschfeld, editor, *Surveys in Combinatorics*, volume 288 of *London Mathematical Society Lecture Note Series*, pages 109–143. Cambridge University press, 2001.

[9] L.A. Goldberg, M. Jerrum and M. Paterson, The computational complexity of two-state spin systems, Pre-print (2001).

[10] L.A. Goldberg, S. Kelk and M. Paterson, The complexity of choosing an H-colouring (nearly) uniformly at random, To appear in STOC 2002.

[11] O. Goldreich, *The Foundations of Cryptography - Volume 1,* (Cambridge University Press, 2001)

[12] 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.

[13] M. Jerrum, Sampling and Counting. Chapter 3 of *Counting, Sampling and Integrating: Algorithms and Complexity*, Birkhäuser, Basel. (In preparation.)

[14] 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.

[15] C.H. Papadimitriou, *Computational Complexity,* (Addison-Wesley, 1994)