Combinatorics, Probability and Computing (2007) 16, 43–61. © 2006 Cambridge University Press
doi:10.1017/S096354830600767X First published online 14 August 2006 Printed in the United Kingdom

The Complexity of Ferromagnetic Ising with Local Fields

LESLIE ANN GOLDBERG[1†] and MARK JERRUM[2†]

[1]Department of Computer Science, University of Warwick, Coventry, CV4 7AL, UK
[2]Division of Informatics, University of Edinburgh,
JCMB, The King's Buildings, Edinburgh EH9 3JZ, UK

Received 17 August 2004; revised 26 November 2005

We consider the complexity of approximating the partition function of the ferromagnetic Ising model with varying interaction energies and local external magnetic fields. Jerrum and Sinclair provided a fully polynomial randomized approximation scheme for the case in which the system is consistent in the sense that the local external fields all favour the same spin. We characterize the complexity of the general problem by showing that it is equivalent in complexity to the problem of approximately counting independent sets in bipartite graphs, thus it is complete in a logically defined subclass of #P previously studied by Dyer, Goldberg, Greenhill and Jerrum. By contrast, we show that the corresponding computational task for the q-state Potts model with local external magnetic fields and $q > 2$ is complete for all of #P with respect to approximation-preserving reductions.

1. Introduction

1.1. The Ising model and the Potts model

An Ising system is defined by a graph $G = (V, E)$. Each edge $(i, j) \in E$ has an associated *interaction strength* $J_{i,j}$ (a real number). Each vertex $v \in V$ has an associated local *external magnetic field* which corresponds to the parameter ℓ_v (a real number). A *configuration* of the system is an assignment $\sigma : V \to \{-1, +1\}$ of 'spins' to the vertices of G. We associate each configuration σ with an energy

$$H(\sigma) = -\sum_{(i,j)\in E} J_{i,j}\sigma(i)\sigma(j) - \sum_{v\in V} \ell_v \sigma(v). \tag{1.1}$$

† This work was partially supported by the EPSRC grant 'Discontinuous Behaviour in the Complexity of Randomized Algorithms'.

https://doi.org/10.1017/S096354830600767X Published online by Cambridge University Press

The partition function corresponding to 'inverse temperature' β (a positive real number) is

$$Z(G,\beta,J_{i,j},\ell_v) = \sum_{\sigma:V\to\{-1,+1\}} \exp(-\beta H(\sigma))$$
$$= \sum_{\sigma:V\to\{-1,+1\}} \prod_{(i,j)\in E} e^{\beta J_{i,j}\sigma(i)\sigma(j)} \prod_{v\in V} e^{\beta\ell_v\sigma(v)}.$$

The system is *ferromagnetic* if every interaction energy $J_{i,j}$ is nonnegative.

The most commonly studied version of the Ising model (see, for example, [2]) is the version in which the interaction energies $J_{i,j}$ are uniform over the edges $(i,j)\in E$ and the local magnetic fields ℓ_v are uniform over vertices $v\in V$. The general model is also studied, particularly in the situation in which the interaction energies $J_{i,j}$ and the local fields ℓ_v are random variables (see [1, 6, 14]).

To avoid the exponentials in our notation, we define edge weights $\lambda_{i,j}=\exp(2\beta J_{i,j})$ and vertex weights $\mu_v=\exp(2\beta\ell_v)$ so that we can rewrite the partition function as follows:

$$Z(G,\lambda_{i,j},\mu_v) = \sum_{\sigma:V\to\{-1,+1\}} \prod_{(i,j)\in E} \lambda_{i,j}^{\frac{1}{2}\sigma(i)\sigma(j)} \prod_{v\in V} \mu_v^{\frac{1}{2}\sigma(v)} \tag{1.2}$$
$$= \prod_{(i,j)\in E} \lambda_{i,j}^{-\frac{1}{2}} \prod_{v\in V} \mu_v^{-\frac{1}{2}} \sum_{\sigma:V\to\{-1,+1\}} \prod_{(i,j)\in E:\sigma(i)=\sigma(j)} \lambda_{i,j} \prod_{v\in V:\sigma(v)=+1} \mu_v. \tag{1.3}$$

Note that the system is ferromagnetic if and only if every edge weight $\lambda_{i,j}$ is at least 1. We say that the system is *consistent* if the external field favours the same spin at all vertices. That is, the system is consistent if μ_v is either uniformly at least 1 or uniformly at most 1.

Jerrum and Sinclair [13] have shown that the problem of exactly computing the partition function of an Ising system is #P-complete, even in the ferromagnetic case. They show that, in the *antiferromagnetic* case, the partition function cannot even be computed *approximately* in polynomial time unless NP = RP. However, they present an efficient algorithm for the ferromagnetic case, which approximates the partition function, *provided that the system is consistent*. Theorem 1.1, below, characterizes the complexity of approximately computing the partition function of a ferromagnetic Ising system that is not consistent. In particular, we show that this problem is complete with respect to approximation-preserving reductions in a logically defined subclass of #P from [7]. The subclass may be of intermediate complexity between those problems that have a fully polynomial randomized approximation scheme and those problems that are complete for #P with respect to approximation-preserving reductions. By contrast, the corresponding problem for the Potts model with $q>2$ spins is complete for #P in this sense – this is Theorem 1.2 below. The theorem implies that there is no fully polynomial randomized approximation scheme (for the ferromagnetic Potts model) unless NP = RP. Background on the complexity of approximate counting is provided in Section 1.2.

First we give the definitions for the Potts model. In this model there are q spins and a *configuration* is an assignment $\sigma : V\to\{1,\ldots,q\}$ of spins to the vertices of the underlying

https://doi.org/10.1017/S096354830600767X Published online by Cambridge University Press

graph $G = (V, E)$. In the most general version of the model, we have a distinct interaction strength $J_{i,j}$ for each edge $(i, j) \in E$ and a distinct external field corresponding to a real number $h_{v,c}$ associated to each vertex v and each possible spin c. Thus, the energy of a configuration σ (see equation (1.1)) is

$$H(\sigma) = -\sum_{(i,j)\in E} J_{i,j}\,\chi(\sigma(i), \sigma(j)) - \sum_{v\in V} h_{v,\sigma(v)}, \tag{1.4}$$

where

$$\chi(s, s') = \begin{cases} +1, & \text{if } s = s', \\ -1, & \text{otherwise.} \end{cases}$$

Letting $\lambda_{i,j} = \exp(2\beta J_{i,j})$ as before and letting $\mu_{v,c} = \exp(\beta h_{v,c})$, we get

$$Z(G, \lambda_{i,j}, \mu_{v,c}) = \prod_{(i,j)\in E} \lambda_{i,j}^{-\frac{1}{2}} \sum_{\sigma: V\to\{1,\ldots,q\}} \prod_{(i,j)\in E:\sigma(i)=\sigma(j)} \lambda_{i,j} \prod_{v\in V} \mu_{v,\sigma(v)}. \tag{1.5}$$

A version of this model in which the interaction energy is uniform over edges $(i, j) \in E$ and the external field $\mu_{v,c}$ is uniform over vertices v (but varies with c) is studied in [3]. Note that the $q = 2$ case of the partition function (1.5) is essentially the same as the partition function of the Ising model (1.3). Furthermore, Jerrum and Sinclair's inapproximability result for the *antiferromagnetic* Ising model extends to the Potts model [15, p. 138]. We will focus on the *ferromagnetic* case in this paper.

1.2. The complexity of approximate counting

A *randomized approximation scheme* is an algorithm for approximately computing the value of a function f. The approximation scheme has a parameter $\varepsilon > 0$ which specifies the error tolerance. For concreteness, suppose that f is a function from Σ^* to $\mathbb{R}$. For example, f might map an encoding of a graph G to the number of independent sets of G. A *randomized approximation scheme* for f is a randomized algorithm that takes as input an instance $x \in \Sigma^*$ (*e.g.*, an encoding of a graph G) and an error tolerance $\varepsilon > 0$, and outputs a number $z \in \mathbb{Q}$ (a random variable of the 'coin tosses' made by the algorithm) such that, for every instance x,

$$\Pr\left[e^{-\varepsilon} f(x) \leqslant z \leqslant e^{\varepsilon} f(x)\right] \geqslant \frac{3}{4}. \tag{1.6}$$

The randomized approximation scheme is said to be a *fully polynomial randomized approximation scheme*, or *FPRAS*, if it runs in time bounded by a polynomial in $|x|$ and ε^{-1}. Note that the quantity $3/4$ in equation (1.6) could be changed to any value in the open interval $(\frac{1}{2}, 1)$ without changing the set of problems that have randomized approximation schemes.

Dyer, Goldberg, Greenhill and Jerrum [7] studied the complexity of approximate counting. They identified three classes of counting problems that are interreducible under approximation-preserving reductions. These are (i) the problems that admit an FPRAS, (ii) the problems that are complete for #P with respect to approximation-preserving reducibility, and (iii) a third class, of intermediate complexity, that can be characterized as the hardest problems in a logically defined subclass of #P. We will use the notion of

https://doi.org/10.1017/S096354830600767X Published online by Cambridge University Press

approximation-preserving reduction from [7]. Suppose that f and g are functions from Σ^* to $\mathbb{R}$. An 'approximation-preserving reduction' from f to g gives a way to turn an FPRAS for g into an FPRAS for f. An *approximation-preserving reduction* from f to g is a randomized algorithm $\mathscr{A}$ for computing f using an 'oracle' for g (which we can think of as an unwritten sub-routine for g). The algorithm $\mathscr{A}$ takes as input a pair $(x, \varepsilon) \in \Sigma^* \times (0,1)$, and satisfies the following three conditions: (i) every oracle call made by $\mathscr{A}$ is of the form (w, δ), where $w \in \Sigma^*$ is an instance of g, and $0 < \delta < 1$ is an error bound satisfying $\delta^{-1} \leqslant \text{poly}(|x|, \varepsilon^{-1})$; (ii) the algorithm $\mathscr{A}$ meets the specification for being a randomized approximation scheme for f (as described above) whenever the oracle meets the specification for being a randomized approximation scheme for g; and (iii) the run-time of $\mathscr{A}$ is polynomial in $|x|$ and ε^{-1}.

If an approximation-preserving reduction from f to g exists we write $f \leqslant_{\text{AP}} g$, and say that f *is AP-reducible to* g. If $f \leqslant_{\text{AP}} g$ and $g \leqslant_{\text{AP}} f$ then we say that f *and* g *are AP-interreducible*, and write $f \equiv_{\text{AP}} g$.

We can now say a little bit more about the three classes identified by Dyer, Goldberg, Greenhill and Jerrum [7]. The first class, containing the problems that admit an FPRAS, are trivially AP-interreducible since all the work can be embedded into the reduction (which declines to use the oracle). The second class is the set of problems that are AP-interreducible with #SAT, which is defined as follows.

Name. #SAT.
Instance. A Boolean formula φ in conjunctive normal form.
Output. The number of satisfying assignments to φ.

All problems in #P are AP-reducible to #SAT. Zuckerman [16] has shown that #SAT cannot have an FPRAS unless NP = RP. The same is obviously true of any problem in #P to which #SAT is AP-reducible. See [7] for details. The third class is the set of problems that are AP-interreducible with #BIS, which is defined as follows.

Name. #BIS.
Instance. A bipartite graph B.
Output. The number of independent sets in B, which we denote #IS(B). (An *independent set* is a set of vertices that does not contain both end-points of any edge.)

Dyer, Goldberg, Greenhill and Jerrum [7] have shown that this class includes a number of natural counting problems such as counting downsets in a partial order, counting configurations in the 2-particle Widom–Rowlinson model and counting configurations in the Beach model. Furthermore, these problems are complete for the logically defined complexity class #RHΠ_1 with respect to AP-reducibility. No function AP-interreducible with #BIS is known to admit an FPRAS, or to be AP-interreducible with #SAT. Thus, it is possible that the complexity of this class of problems in some sense lies strictly between the class of problems admitting an FPRAS and #SAT.

We will study the following computational problems.

https://doi.org/10.1017/S096354830600767X Published online by Cambridge University Press

Name. FERROMAGNETIC ISING.
Instance. A ferromagnetic Ising system consisting of a graph G, edge weights $\lambda_{i,j} \geqslant 1$ and vertex weights μ_v.
Output. The partition function $Z(G, \lambda_{i,j}, \mu_v)$.

Name. FERROMAGNETIC POTTS(q).
Instance. A ferromagnetic Potts system consisting of a graph G, edge weights $\lambda_{i,j} \geqslant 1$ and vertex weights $\mu_{v,c}$.
Output. The partition function $Z(G, \lambda_{i,j}, \mu_{v,c})$.

Jerrum and Sinclair [13] gave an FPRAS for the physically realistic special case of FERROMAGNETIC ISING in which the system is consistent. We can now state our result, which is that the general problem, in which the system may or not be consistent, is AP-interreducible with #BIS.

Theorem 1.1. FERROMAGNETIC ISING $\equiv_{\mathrm{AP}}$ #BIS.

Note that, for ease of presentation, Jerrum and Sinclair described their FPRAS for the special case in which all of the vertex weights μ_v are identical. Their result does not require the vertex weights to be identical, provided they are consistent; for a proof of this assertion, see the proof of Theorem 1 of [10].

It is an open question [11] whether there is an FPRAS for the ferromagnetic Potts model for any fixed $q > 2$ even in the zero-field case (in which every $h_{v,c}$ is equal to zero). We show that, for fixed $q > 2$, there is unlikely to be an FPRAS for FERROMAGNETIC POTTS(q). In particular, we have the following.

Theorem 1.2. *Suppose* $q > 2$. *Then* FERROMAGNETIC POTTS(q) $\equiv_{\mathrm{AP}}$ #SAT.

Corollary 1.3. *Suppose* $q > 2$. *Then there is no FPRAS for* FERROMAGNETIC POTTS(q) *unless* NP = RP.

1.3. Specifying the input parameters

In order to make our definitions of the problems FERROMAGNETIC ISING and FERROMAGNETIC POTTS(q) precise, we need to specify how the input parameters $\lambda_{i,j}$, μ_v and $\mu_{v,c}$ should be encoded. Since we are interested in randomized approximation schemes (rather than in exact counting algorithms), we lose no generality by insisting that these parameters be rational (see Observation 1.4 below). Thus, we adopt this convention.

Observation 1.4. *Consider the ferromagnetic Ising system given by* G, $\lambda_{i,j}$ *and* μ_v. *Let* ε *be a positive constant. Suppose that* G *has* n *vertices and* m *edges. Let* k *be any integer that exceeds* $\lceil \log_{10}(n/\varepsilon) \rceil$ *and* $\lceil \log_{10}(m/\varepsilon) \rceil$. *Let* $\hat{\lambda}_{i,j}$ *be the rational number derived from* $\lambda_{i,j}$ *by retaining the first* $k+1$ *digits in the decimal expansion of* λ (*and setting the other digits to zero*). *Define* $\hat{\mu}_v$ *similarly. Then*

$$e^{-\varepsilon} Z(G, \lambda_{i,j}, \mu_v) \leqslant Z(G, \hat{\lambda}_{i,j}, \hat{\mu}_v) \leqslant e^{\varepsilon} Z(G, \lambda_{i,j}, \mu_v).$$

https://doi.org/10.1017/S096354830600767X Published online by Cambridge University Press

Proof. Consider the decimal expansion of $\lambda_{i,j} \geqslant 1$, namely $\lambda_{i,j} = \sum_{r=-\infty}^{\infty} b_r 10^r$, where $b_r \in \{0, \ldots, 9\}$. Let s be the largest r such that $b_r > 0$. Let

$$\hat{\lambda}_{i,j} = \sum_{r=s-k}^{s} b_r 10^r.$$

Then

$$\begin{aligned}\hat{\lambda}_{i,j} &= \lambda_{i,j} - \sum_{r<s-k} b_r 10^r \geqslant \lambda_{i,j} - 10^{s-k} \\ &\geqslant (1 - 10^{-k})\lambda_{i,j} \geqslant \exp\left(-\frac{1}{10^{k-1}}\right)\lambda_{i,j} \geqslant \exp\left(-\frac{\varepsilon}{m}\right)\lambda_{i,j}.\end{aligned}$$

Also

$$\hat{\lambda}_{i,j} \leqslant \lambda_{i,j} \leqslant \exp\left(\frac{\varepsilon}{m}\right)\lambda_{i,j}.$$

Thus, the contribution of all edges to the partition function only differs by at most an $\exp(\varepsilon/2)$ factor if we use the approximate edge weights $\hat{\lambda}_{i,j}$ instead of the actual ones. The approximation of vertex weights is similar: if μ_v is less then 1 then we can keep $k+1$ digits of the decimal expansion of $1/\mu_v$. □

1.4. Outline of the paper

Theorem 1.1 follows from Lemma 2.1, which is proved in Section 2, and from Lemma 3.1, which is proved in Section 3. Theorem 1.2 is proved in Section 4.

2. Reduction from #BIS to Ferromagnetic Ising

The result of this section is the following.

Lemma 2.1. #BIS $\leqslant_{\mathrm{AP}}$ Ferromagnetic Ising.

The proof of Lemma 2.1 consists of two reductions which are given in Lemmas 2.2 and 2.3 below. The first of these reduces #BIS to a 'permissive' version of #BIS, which we refer to as #PermissiveBIS(γ), where $\gamma > 0$ is a parameter to be explained presently. The second of these reduces #PermissiveBIS(γ) to Ferromagnetic Ising. The instance of Ferromagnetic Ising that is produced by the combined reduction has an underlying graph which is bipartite.

We start by defining #PermissiveBIS(γ). The usual version of #BIS is 'hard core' in the sense that a configuration is an independent set which is not allowed to contain adjacent (IN, IN) pairs. In the permissive version, all assignments from vertices to {IN, OUT} are configurations. However, configurations are given weights which discourage adjacent (IN, IN) pairs. In particular, each edge between two vertices that are both assigned IN is weighted by a factor of $\gamma^2 < 1$.

Name. #PermissiveBIS(γ).
Instance. A bipartite graph B with vertex set V.

https://doi.org/10.1017/S096354830600767X Published online by Cambridge University Press

Output. The quantity $Z_\gamma(B) = \sum_{\tau:V\to\{\mathrm{IN},\mathrm{OUT}\}} \gamma^{2b(\tau)}$, where $b(\tau)$ denotes the number of $(\mathrm{IN},\mathrm{IN})$ edges in τ.

2.1. Reducing #BIS to #PermissiveBIS(1/4)

In this section we prove the following lemma, which shows that #BIS can be reduced to the problem #PermissiveBIS(1/4). The constant 1/4 was chosen for convenience; other constants would also work. Formally, our reduction applies to any graph (whether or not it is bipartite) but we are interested in the application to bipartite graphs.

Lemma 2.2. $\#\mathrm{BIS} \leqslant_{\mathrm{AP}} \#\mathrm{PermissiveBIS}(1/4)$.

Let $G = (V, E)$ be an instance of #BIS. Let $n = |V|$. For $i \in V$, let S_i and T_i be disjoint sets of size $r = 6n$. Construct an instance $\widehat{G} = (\widehat{V}, \widehat{E})$ of #PermissiveBIS(1/4) as follows:

$$\widehat{V} = \bigcup_{i\in V} S_i \cup T_i$$

and

$$\widehat{E} = \bigcup_{(i,j)\in E} S_i \times S_j \cup \bigcup_{i\in V} S_i \times T_i.$$

Let Ω be the set of configurations corresponding to $\widehat{G}$. That is, Ω is the set of assignments τ from the vertex set $\widehat{V}$ to $\{\mathrm{IN},\mathrm{OUT}\}$. We will define some subsets of Ω. For $i \in V$, let $\Omega_i^{\mathrm{small}}$ be the set of configurations in which few vertices from $S_i \cup T_i$ are mapped to IN. That is,

$$\Omega_i^{\mathrm{small}} = \{\tau \in \Omega : |\{v \in S_i \cup T_i : \tau(v) = \mathrm{IN}\}| < r/10\}, \text{ and } \Omega^{\mathrm{small}} = \bigcup_{i\in V} \Omega_i^{\mathrm{small}}.$$

Let $\Omega_i^{\mathrm{split}}$ be the set of configurations in which both S_i and T_i have vertices which are mapped to IN. That is,

$$\Omega_i^{\mathrm{split}} = \{\tau \in \Omega : \mathrm{IN} \in \tau(S_i) \text{ and } \mathrm{IN} \in \tau(T_i)\}, \text{ and } \Omega^{\mathrm{split}} = \bigcup_{i\in V} \Omega_i^{\mathrm{split}}.$$

Let Ω^{bad} be the set of remaining configurations with $(\mathrm{IN},\mathrm{IN})$ edges. That is,

$$\Omega^{\mathrm{bad}} = \{\tau \in \Omega - (\Omega^{\mathrm{small}} \cup \Omega^{\mathrm{split}}) : \tau \text{ has an } (\mathrm{IN},\mathrm{IN}) \text{ edge}\}.$$

Let Ω^{good} be the rest of the configurations: $\Omega^{\mathrm{good}} = \Omega - (\Omega^{\mathrm{small}} \cup \Omega^{\mathrm{split}} \cup \Omega^{\mathrm{bad}})$. For any subset Ψ of Ω, let $Y(\Psi)$ denote the contribution to the partition function $Z_\gamma(\widehat{G})$ from configurations in Ψ. That is, $Y(\Psi) = \sum_{\tau\in\Psi} \gamma^{2b(\tau)}$, where $\gamma = 1/4$. Let Y denote $Y(\Omega)$.

We start by deriving upper bounds for $Y(\Omega^{\mathrm{small}})$, $Y(\Omega^{\mathrm{split}})$ and $Y(\Omega^{\mathrm{bad}})$. Let i be any vertex in V. First note that since $\gamma^2 < 1$, we have

$$Y(\Omega_i^{\mathrm{small}}) \leqslant \sum_{k=0}^{r/10} \binom{2r}{k} \sum_{\tau':\widehat{V}-S_i\cup T_i\to\{\mathrm{IN},\mathrm{OUT}\}} \gamma^{2b(\tau')}.$$

https://doi.org/10.1017/S096354830600767X Published online by Cambridge University Press

Also,

$$Y \geqslant \sum_{\tau\in\Omega:\tau(S_i)=\{\mathrm{OUT}\}} \gamma^{2b(\tau)} = 2^r \sum_{\tau':\widehat{V}-S_i\cup T_i\to\{\mathrm{IN},\mathrm{OUT}\}} \gamma^{2b(\tau')},$$

so[1]

$$Y(\Omega_i^{\mathrm{small}}) \leqslant 2^{-r}\sum_{k=0}^{r/10}\binom{2r}{k} Y \leqslant 2^{-r}\left(\frac{2re}{r/10}\right)^{r/10} Y \leqslant 4^{-n}Y. \tag{2.1}$$

Similarly,

$$\begin{aligned} Y(\Omega_i^{\mathrm{split}}) &\leqslant \sum_{k,\ell=1}^{r}\binom{r}{k}\binom{r}{\ell}\gamma^{2k\ell} \sum_{\tau':\widehat{V}-S_i\cup T_i\to\{\mathrm{IN},\mathrm{OUT}\}} \gamma^{2b(\tau')} \\ &\leqslant 2^{-r}\sum_{k,\ell=1}^{r}\binom{r}{k}\binom{r}{\ell}\gamma^{2k\ell} Y \\ &\leqslant 2^{-r}\sum_{k,\ell=1}^{r}\binom{r}{k}\binom{r}{\ell}\gamma^{k+\ell} Y \\ &\leqslant 2^{-r}(1+\gamma)^r(1+\gamma)^r Y \\ &\leqslant 4^{-n}Y. \end{aligned} \tag{2.2}$$

Every configuration $\tau \in \Omega^{\mathrm{bad}}$ has an $(\mathrm{IN},\mathrm{IN})$ edge between some sets S_i and S_j and each of S_i and S_j have at least $r/10$ vertices mapped to IN by τ. Thus, every such τ has at least $(r/10)^2$ $(\mathrm{IN},\mathrm{IN})$ edges. This means that

$$Y(\Omega^{\mathrm{bad}}) \leqslant \gamma^{2(r/10)^2} 2^n 2^{rn} \leqslant 2^{-n}2^{rn}, \tag{2.3}$$

where the final inequality assumes n sufficiently large (in fact $n \geqslant 2$ suffices). Let $t = \frac{2n4^{-n}}{1-2n4^{-n}}$. From (2.1) and (2.2),

$$Y(\Omega^{\mathrm{good}} \cup \Omega^{\mathrm{bad}}) = Y(\Omega - (\Omega^{\mathrm{small}} \cup \Omega^{\mathrm{split}})) \geqslant (1-2n4^{-n})Y = Y/(1+t),$$

so, using (2.3),

$$\begin{aligned} Y &\leqslant (1+t)Y(\Omega^{\mathrm{good}} \cup \Omega^{\mathrm{bad}}) \\ &\leqslant (1+t)(Y(\Omega^{\mathrm{good}}) + 2^{-n}2^{rn}) \\ &\leqslant Y(\Omega^{\mathrm{good}}) + tY(\Omega^{\mathrm{good}}) + (1+t)2^{-n}2^{rn}. \end{aligned} \tag{2.4}$$

Consider the configurations $\tau \in \Omega^{\mathrm{good}}$. For every $i \in V$, there are two possibilities:

- $\tau(S_i) = \{\mathrm{OUT}\}$ and $|\{v \in T_i : \tau(v) = \mathrm{IN}\}| \geqslant r/10$, or
- $\tau(T_i) = \{\mathrm{OUT}\}$ and $|\{v \in S_i : \tau(v) = \mathrm{IN}\}| \geqslant r/10$.

Furthermore, $(i,j) \in E$ implies that $\tau(S_i) = \{\mathrm{OUT}\}$ or $\tau(S_j) = \{\mathrm{OUT}\}$ (or both). From this, we see that each $\tau \in \Omega^{\mathrm{good}}$ points out an independent set of G – a vertex $i \in V$ is

[1] To see that the second inequality is correct, note that $\sum_{i=0}^{t}\binom{n}{i} \leqslant \left(\frac{en}{t}\right)^t$. This upper bound on the sum of binomial coefficients is well known. For example, a proof is in [9] or [5].

https://doi.org/10.1017/S096354830600767X Published online by Cambridge University Press

in the independent set if and only if $\tau(T_i) = \{\text{OUT}\}$. Also, each independent set of G corresponds to exactly

$$\ell := \left(\sum_{k=r/10}^{r} \binom{r}{k} \right)^n$$

configurations $\tau \in \Omega^{\text{good}}$. Using the bound from our earlier calculation from the derivation of (2.1),

$$(1-4^{-n})2^{rn} \leqslant (1-5^{-n})^n 2^{rn} \leqslant \ell \leqslant 2^{rn},$$

for n sufficiently large, so

$$2^{rn}\#\text{IS}(G) - 4^{-n}2^{rn}2^n \leqslant (1-4^{-n})2^{rn}\#\text{IS}(G) \leqslant \ell\#\text{IS}(G) = Y(\Omega^{\text{good}}) \leqslant Y.$$

Also, using (2.4),

$$Y \leqslant Y(\Omega^{\text{good}}) + tY(\Omega^{\text{good}}) + (1+t)2^{-n}2^{rn} \leqslant 2^{rn}\#\text{IS}(G) + t2^{rn}2^n + (1+t)2^{-n}2^{rn}.$$

Combining these two equations, dividing by 2^{rn}, and re-arranging, we get

$$\frac{Y}{2^{rn}} - (t2^n + (1+t)2^{-n}) \leqslant \#\text{IS}(G) \leqslant \frac{Y}{2^{rn}} + 4^{-n}2^n.$$

For $n \geqslant 6$, this gives us

$$\frac{Y}{2^{rn}} - \frac{1}{4} \leqslant \#\text{IS}(G) \leqslant \frac{Y}{2^{rn}} + \frac{1}{4}.$$

This equation gives us a simple AP-reduction from #BIS to #PermissiveBIS(1/4): given an accuracy parameter ε and an instance G of #BIS, obtain an approximation $\widehat{Y}$ to Y satisfying

$$e^{-\varepsilon/21}Y \leqslant \widehat{Y} \leqslant e^{\varepsilon/21}Y,$$

and round $\widehat{Y}/2^{rn}$ to the nearest integer. See the proof of Theorem 3 of [7] for details showing that this provides a sufficiently accurate approximation to #IS(G). This concludes the proof of Lemma 2.2. □

2.2. Reduction from #PermissiveBIS(1/4) to Ferromagnetic Ising

In this section we prove the following lemma.

Lemma 2.3. #PermissiveBIS(1/4) $\leqslant_{\text{AP}}$ Ferromagnetic Ising.

Let $B = (V(B), E(B))$ be an instance of #PermissiveBIS(1/4), and denote by L and R the bipartition of $V(B)$. Let $m = |E(B)|$. For every vertex $v \in V(B)$, let $d(v)$ denote the degree of v in B.

Construct an instance G of Ferromagnetic Ising as follows. For every vertex $v \in V(B)$, let W_v be a set of $2d(v)$ distinct vertices. The vertex set of G is

$$V(G) = V(B) \cup \bigcup_{v \in V(B)} W_v.$$

https://doi.org/10.1017/S096354830600767X Published online by Cambridge University Press

Table 1.

	$\tau(i)$	$\sigma(i)$
$i \in L$	IN	$+1$
$i \in L$	OUT	-1
$i \in R$	IN	-1
$i \in R$	OUT	$+1$

The edge set of G is

$$E(G) = E(B) \cup \bigcup_{v \in V(B)} v \times W_v.$$

For every edge $(i,j) \in E(G)$, let $\lambda_{i,j} = \lambda = 4$. For every vertex $v \in V(B)$, let $\mu_v = 1$. For every vertex $v \in L$ and every vertex $w \in W_v$, let $\mu_w = \frac{2}{7}$. Finally, for every vertex $v \in R$ and every vertex $w \in W_v$, let $\mu_w = \frac{7}{2}$.

Consider a function $\tau : V(B) \to \{\text{IN}, \text{OUT}\}$. The function τ induces an assignment of spins to vertices in $V(B)$ as shown in Table 1.

Now let

$$Z'(G, \lambda, \mu_i) = \sum_{\sigma : V(G) \to \{-1,+1\}} \prod_{(i,j) \in E(G) : \sigma(i) = \sigma(j)} \lambda_{i,j} \prod_{i \in V(G) : \sigma(i) = +1} \mu_i$$

be the partition function of G, omitting some easily computed scaling factors (see (1.3)). Consider the contribution to $Z'(G, \lambda, \mu_i)$ from configurations σ induced by a particular map τ. Note that τ fixes the value of $\sigma(i)$ for every $i \in L \cup R$.

First, consider a vertex $i \in L$. If $\sigma(i) = +1$ then the contribution from the vertices in W_i and the edges connecting them to i is a factor of $(1 + \frac{2}{7}\lambda)^{2d(i)}$. If $\sigma(i) = -1$, the contribution is $(\frac{2}{7} + \lambda)^{2d(i)}$. Thus, the contribution of vertex i may be summarized as

$$\left[\left(1 + \tfrac{2}{7}\lambda\right)\left(\tfrac{2}{7} + \lambda\right)\right]^{d(i)} \left(\frac{1 + \frac{2}{7}\lambda}{\frac{2}{7} + \lambda}\right)^{d(i)\sigma(i)}.$$

Similarly, the contribution for a vertex $i \in R$ is

$$\left[\left(1 + \tfrac{7}{2}\lambda\right)\left(\tfrac{7}{2} + \lambda\right)\right]^{d(i)} \left(\frac{1 + \frac{7}{2}\lambda}{\frac{7}{2} + \lambda}\right)^{d(i)\sigma(i)}.$$

Let

$$A = (1 + \tfrac{2}{7}\lambda)(\tfrac{2}{7} + \lambda)(1 + \tfrac{7}{2}\lambda)(\tfrac{7}{2} + \lambda)\lambda.$$

Putting our observations together,

$$Z'(G, \lambda, \mu_i) = A^m \sum_{\tau : V(B) \to \{\text{IN},\text{OUT}\}} \left[\prod_{i \in L} \left(\frac{1 + \frac{2}{7}\lambda}{\frac{2}{7} + \lambda} \right)^{d(i)\sigma(i)} \times \prod_{j \in R} \left(\frac{1 + \frac{7}{2}\lambda}{\frac{7}{2} + \lambda} \right)^{d(j)\sigma(j)} \prod_{(i,j) \in E(B)} \lambda^{(\sigma(i)\sigma(j)-1)/2} \right],$$

https://doi.org/10.1017/S096354830600767X Published online by Cambridge University Press

where the values $\sigma(i)$ and $\sigma(j)$ are induced by τ. Recalling that $\lambda = 4$, this simplifies to

$$Z'(G, \lambda, \mu_i) = A^m \sum_{\tau} \prod_{i \in L} 2^{-d(i)\sigma(i)} \prod_{j \in R} 2^{d(j)\sigma(j)} \prod_{(i,j) \in E(B)} 2^{\sigma(i)\sigma(j)-1}$$
$$= A^m \sum_{\tau} \prod_{(i,j) \in E(B)} 2^{\sigma(i)\sigma(j)-1+\sigma(j)-\sigma(i)}.$$

Now observe

$$\sigma(i)\sigma(j) - 1 + \sigma(j) - \sigma(i) = \begin{cases} -4 & \text{if } \tau(i) = \text{IN and } \tau(j) = \text{IN}, \\ 0 & \text{otherwise}, \end{cases}$$

so $Z'(G, \lambda, \mu_i) = A^m Z_{1/4}(B)$.

Lemma 2.3 follows from the fact that $Z'(G, \lambda, \mu_i)$ is an easily computed multiple of $Z(G, \lambda, \mu_i)$. □

3. Reduction from Ferromagnetic Ising to #BIS

In this section we prove the following result.

Lemma 3.1. Ferromagnetic Ising $\leqslant_{\text{AP}}$ #BIS.

We start by defining a restricted version of Ferromagnetic Ising.

Name. Restricted Ferromagnetic Ising.
Instance. A ferromagnetic Ising system consisting of a graph $G = (V, E)$ in which every edge (i, j) has edge weight $\lambda_{i,j} = \frac{4}{3}$ and every vertex $v \in V$ has $\mu_v \in \{\frac{1}{2}, 1, 2\}$.
Output. The partition function $Z(G, \lambda_{i,j}, \mu_i)$.

The reduction corresponding to Lemma 3.1 is in two parts. First, in Section 3.1, we reduce Restricted Ferromagnetic Ising to #BIS. Then, in Section 3.2, we reduce Ferromagnetic Ising to Restricted Ferromagnetic Ising. The choice of the constant 4/3 in the definition of Restricted Ferromagnetic Ising is not critical: other constants would also work (though 4/3 is convenient in the proof of Lemma 3.2). It is important that $\lambda_{i,j}$ be greater than 1 since this is the ferromagnetic case.

3.1. Reduction from Restricted Ferromagnetic Ising to #BIS

This section proves the following lemma.

Lemma 3.2. Restricted Ferromagnetic Ising $\leqslant_{\text{AP}}$ #BIS.

Suppose that we are given an accuracy parameter ε and an instance $G = (V, E)$ of Restricted Ferromagnetic Ising with $|V| = n$ and $|E| = m$. Let V_+ be the set of vertices $v \in V$ with $\mu_v = 2$ and let V_- be the set of vertices $v \in V$ with $\mu_v = \frac{1}{2}$. Let $r = \lceil 4m + 2n + \log_2 \varepsilon^{-1} \rceil$. For every $i \in V$, let S_i and T_i be disjoint sets of size r. For every

https://doi.org/10.1017/S096354830600767X Published online by Cambridge University Press

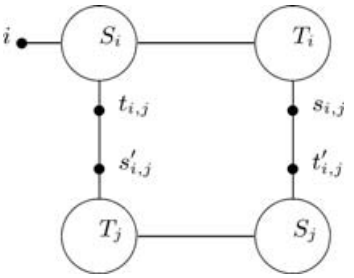


Figure 1. Some of the edges of the instance $\widehat{G}$ of #BIS, where $i \in V_-$, $j \notin V_+ \cup V_-$, and $(i,j) \in E$.

$(i,j) \in E$, let $s_{i,j}$, $t_{i,j}$, $s'_{i,j}$ and $t'_{i,j}$ be vertices. Construct an instance $\widehat{G} = (\widehat{V}, \widehat{E})$ of #BIS as follows:

$$\widehat{V} = \bigcup_{i \in V} (S_i \cup T_i) \cup \bigcup_{(i,j) \in E} \{s_{i,j}, t_{i,j}, s'_{i,j}, t'_{i,j}\} \cup V_+ \cup V_-,$$

$$\widehat{E} = \bigcup_{i \in V} (S_i \times T_i) \cup \bigcup_{i \in V_-} (S_i \times \{i\}) \cup \bigcup_{i \in V_+} (T_i \times \{i\})$$
$$\cup \bigcup_{(i,j) \in E} \big(S_i \times \{t_{i,j}\} \cup S_j \times \{t'_{i,j}\} \cup T_i \times \{s_{i,j}\} \cup T_j \times \{s'_{i,j}\} \cup \{(s_{i,j}, t'_{i,j}), (s'_{i,j}, t_{i,j})\}\big).$$

See Figure 1.

Say an independent set $\tau : \widehat{V} \to \{\mathrm{IN}, \mathrm{OUT}\}$ is *complete* if $\tau(S_i \cup T_i) = \{\mathrm{IN}, \mathrm{OUT}\}$ for all $i \in V$; otherwise τ is *incomplete*. Denote by $Y = \#\mathrm{BIS}(\widehat{G})$ the total number of independent sets in $\widehat{G}$, and by Y' the number of complete independent sets.

Let $u = n2^{n-1}2^{r(n-1)}2^{4m+n}$. A crude upper bound on incomplete independent sets shows $Y - Y' \leqslant u$. The first factor of n in u represents the choice of a vertex i with $\tau(S_i \cup T_i) = \{\mathrm{OUT}\}$. The $2^{n-1}2^{r(n-1)}$ represents the number of configurations on the other $n-1$ sets $S_j \cup T_j$: one of $\tau(S_j)$ and $\tau(T_j)$ is $\{\mathrm{OUT}\}$ and the other can be assigned arbitrarily. The remaining factor corresponds to all assignments of the remaining vertices. Similarly, let $\ell = 2^n(2^r - 1)^n$. A lower bound on the number of complete independent sets gives $Y' \geqslant \ell$.

Note that for sufficiently large n (in fact, for $n > 2$) we have $4u/\ell \leqslant \varepsilon$, so

$$Y - Y' \leqslant u \leqslant \frac{\varepsilon}{4}\ell \leqslant \frac{\varepsilon}{4}Y.$$

Therefore

$$Ye^{-\varepsilon/2} \leqslant Y\left(1 - \frac{\varepsilon}{4}\right) \leqslant Y' \leqslant Y,$$

where the first inequality uses $\varepsilon \leqslant 1$. Thus, a sufficiently accurate estimate for $Y = \#\mathrm{BIS}(\widehat{G})$ (say with error bound $\delta = \varepsilon/2$) is also a sufficiently accurate estimate for Y' (*i.e.*, within $e^{\pm\varepsilon}$). It only remains to show that Y' is directly related to the partition function of the instance (G, V_+, V_-) of Restricted Ferromagnetic Ising.

https://doi.org/10.1017/S096354830600767X Published online by Cambridge University Press

Table 2.

Case	Number of ways to extend τ to $\{s_{i,j}, s'_{i,j}, t_{i,j}, t'_{i,j}\}$	Case	Number of ways to extend τ to i
$\sigma(i) = \sigma(j)$	4	$i \in V_+, \sigma(i) = +1$	2
$\sigma(i) \neq \sigma(j)$	3	$i \in V_+, \sigma(i) = -1$	1
		$i \in V_-, \sigma(i) = -1$	2
		$i \in V_-, \sigma(i) = +1$	1

To each complete independent set $\tau : \widehat{V} \to \{\mathrm{IN}, \mathrm{OUT}\}$ in the #BIS instance $\widehat{G}$, there naturally corresponds an Ising configuration

$$\sigma(i) = \begin{cases} +1, & \text{if } \mathrm{IN} \in \tau(S_i), \\ -1, & \text{if } \mathrm{IN} \in \tau(T_i). \end{cases}$$

Since τ is complete, these cases are exhaustive. Now fix $\sigma : V \to \{+1, -1\}$, and consider the number of independent sets τ in $\widehat{G}$ associated with σ under the above correspondence. There are $(2^r - 1)^n$ ways to choose the restriction of τ to the set $\bigcup_{i \in V} S_i \cup T_i$. The number of ways to extend τ to the other vertices is as shown in Table 2.

Collecting these observations,

$$Y' = (2^r - 1)^n (4 \times 3)^{m/2} 2^{(|V_+| + |V_-|)/2} \times \sum_{\sigma : V \to \{-1,+1\}} \prod_{(i,j) \in E} \left(\frac{4}{3}\right)^{\sigma(i)\sigma(j)/2} \prod_{i \in V_+} 2^{\sigma(i)/2} \prod_{i \in V_-} 2^{-\sigma(i)/2}.$$

Comparing with (1.2), it can be seen that Y', up to an easily computable factor, is exactly the partition function $Z(G, \lambda_{i,j}, \mu_v)$. This completes the proof of Lemma 3.2. □

3.2. Reduction from Ferromagnetic Ising to Restricted Ferromagnetic Ising

This section contains the proof of the following lemma.

Lemma 3.3. Ferromagnetic Ising $\leqslant_{\mathrm{AP}}$ Restricted Ferromagnetic Ising.

Let $G = (V, E)$ be an instance of Ferromagnetic Ising with edge weights $\lambda_{i,j}$ and vertex weights μ_v. Let $n = |V|$ and $m = |E|$. We will construct an instance $\widehat{G} = (\widehat{V}, \widehat{E})$ of Restricted Ferromagnetic Ising with edge weights $\hat{\lambda}_{i,j} = \frac{4}{3}$ and vertex weights $\hat{\mu}_i \in \{\frac{1}{2}, 1, 2\}$, such that

$$\exp(-\varepsilon) Z(\widehat{G}, \hat{\lambda}_{i,j}, \hat{\mu}_i) \leqslant Z(G, \lambda_{i,j}, \mu_i) \leqslant \exp(\varepsilon) Z(\widehat{G}, \hat{\lambda}_{i,j}, \hat{\mu}_i).$$

The general strategy is to replace each edge (i, j) of G by a gadget $G_{i,j}$ in $\widehat{G}$, all of whose edges have the standard weight $\lambda = \frac{4}{3}$. The gadget will have effective weight λ_{eff} close to $\lambda_{i,j} > 1$. In the terminology of Jaeger, Vertigan and Welsh [12], we first t-thicken the edge (i, j), and then ℓ-stretch all t edges so formed. (A formal description of the construction

https://doi.org/10.1017/S096354830600767X Published online by Cambridge University Press

Table 3.

$\sigma(i)$	$\sigma(j)$	Contribution to $Z(G_{i,j}, \frac{4}{3}, 1)$
$+1$	$+1$	$a_\ell^t \lambda_{\text{eff}}$
$+1$	-1	a_ℓ^t
-1	$+1$	a_ℓ^t
-1	-1	$a_\ell^t \lambda_{\text{eff}}$

will be given presently.) We shall see that, by a suitable choice of t and ℓ, the ratio $\lambda_{\text{eff}}/\lambda_{i,j}$ may be made close to 1; furthermore, the construction is reasonably efficient, in the sense that a close approximation may be achieved using relatively small values of t and ℓ.

More formally, for $\ell \geqslant 1$, let P_ℓ be an ℓ-edge path in which all edges have weight $\lambda = \frac{4}{3}$ and all vertices have weight 1. Let f_ℓ denote the contribution to the partition function $Z(P_\ell, \lambda, 1)$ from the assignment $(+1, +1)$ to the end-points (by symmetry, this is the same as the contribution from the assignment $(-1, -1)$) and let a_ℓ denote the contribution from the assignment $(+1, -1)$. Observe from (1.2) that these satisfy the recurrences

$$f_\ell = \lambda^{1/2} f_{\ell-1} + \lambda^{-1/2} a_{\ell-1},$$

and

$$a_\ell = \lambda^{-1/2} f_{\ell-1} + \lambda^{1/2} a_{\ell-1}$$

with $f_1 = \lambda^{1/2}$ and $a_1 = \lambda^{-1/2}$. Thus, the solution is

$$f_\ell = \frac{\lambda^{-\ell/2}}{2}\left((\lambda+1)^\ell + (\lambda-1)^\ell\right),$$

and

$$a_\ell = \frac{\lambda^{-\ell/2}}{2}\left((\lambda+1)^\ell - (\lambda-1)^\ell\right).$$

For a suitable choice of $t \geqslant 0$ and $\ell > 1$ (chosen below), let $G_{i,j}$ be the graph with vertex set

$$V_{i,j} = \{i, j\} \cup \bigcup_{\alpha \in [1,t], \beta \in [1,\ell-1]} v_{\alpha,\beta}$$

and edge set

$$E_{i,j} = \bigcup_{\alpha \in [1,t]} (i, v_{\alpha,1}) \cup \bigcup_{\alpha \in [1,t], \beta \in [1,\ell-2]} (v_{\alpha,\beta}, v_{\alpha,\beta+1}) \cup \bigcup_{\alpha \in [1,t]} (v_{\alpha,\ell-1}, j).$$

Let all edges in $E_{i,j}$ have weight $\frac{4}{3}$ and all vertices in $V_{i,j}$ have weight 1. Let

$$\lambda_{\text{eff}} = \left(\frac{f_\ell}{a_\ell}\right)^t = \left[1 + \frac{2}{7^\ell - 1}\right]^t.$$

Then the contribution to the partition function $Z(G_{i,j}, \lambda, 1)$ for each possible assignment of spins to i and j is as shown in Table 3.

https://doi.org/10.1017/S096354830600767X Published online by Cambridge University Press

We want to choose $\ell > 1$ and $t \geqslant 0$ so that

$$e^{-\varepsilon/(2m)} \leqslant \frac{\lambda_{\text{eff}}}{\lambda_{i,j}} \leqslant e^{\varepsilon/(2m)}. \tag{3.1}$$

We can do this by first setting ℓ to be the smallest integer, greater than 1, such that

$$1 + \frac{2}{7^\ell - 1} \leqslant e^{\varepsilon/(2m)};$$

then, with ℓ fixed to this value, setting t to be the largest integer such that

$$\left[1 + \frac{2}{7^\ell - 1}\right]^t \leqslant \lambda_{i,j}.$$

(The condition $\ell > 1$ is just there to ensure that we end up with a graph, and not a multigraph. If $t = 0$, then the vertices i and j simply lose their direct connection.) It is clear that we have achieved inequality (3.1), even the stronger one with $e^{\varepsilon/(2m)}$ replaced by 1. Furthermore, $\ell = O(\log(2m/\varepsilon))$ and $t = O((2m/\varepsilon)\log\lambda_{i,j})$, so the total number of edges in $G_{i,j}$ is $O((2m/\varepsilon)\log(\lambda_{i,j})\log(2m/\varepsilon))$. Note that $G_{i,j}$ is not the most efficient gadget – we would be better off using different length stretches on different branches of the thickening – but it is good enough for the purposes of constructing an FPRAS.

Vertex weights may be handled similarly. Consider a vertex i with weight $\mu_i > 1$. To construct the gadget G_i that replaces i in $\widehat{G}$, we first attach t bristles to i, and then perform an ℓ-stretch on the bristles. All vertices have weight 1 except the end-vertices (of degree 1) which have weight 2 (one of our standard weights). From (1.2), the contribution to the partition function of $Z(G_i, \lambda, \hat{\mu}_v)$ from assigning $\sigma(i) = +1$ is $(f_\ell 2^{1/2} + a_\ell 2^{-1/2})^t$. and the contribution from $\sigma(i) = -1$ is $(f_\ell 2^{-1/2} + a_\ell 2^{1/2})^t$. Let

$$\mu_{\text{eff}} = \left(\frac{f_\ell 2^{1/2} + a_\ell 2^{-1/2}}{f_\ell 2^{-1/2} + a_\ell 2^{1/2}}\right)^t = \left[1 + \frac{2}{3 \times 7^\ell - 1}\right]^t.$$

As before, we may achieve $e^{-\varepsilon/(2n)} \leqslant \mu_{\text{eff}}/\mu_i \leqslant e^{\varepsilon/(2n)}$ using $O((2n/\varepsilon)\log(\mu_i)\log(2n/\varepsilon))$ vertices and edges in total. The case $\mu_i < 1$ is handled similarly. Thus, we have established Lemma 3.2. □

4. The Potts model

This section contains the proof of Theorem 1.2, which says that for $q > 2$,

$$\text{Ferromagnetic Potts}(q) \equiv_{\text{AP}} \#\text{SAT}.$$

To establish this theorem, we need to show (i) Ferromagnetic Potts$(q) \leqslant_{\text{AP}} \#$SAT, and (ii) $\#\text{SAT} \leqslant_{\text{AP}}$ Ferromagnetic Potts(q). The first of these is straightforward. The paper [7] shows that every problem in #P is AP-reducible to #SAT. While Ferromagnetic Potts(q) is not itself in #P, it is easy to see how to reduce it to a #P problem. The idea is to multiply all of the rational edge weights $\lambda_{i,j}$ and vertex weights $\mu_{v,c}$ by the same positive integer so that, after multiplication, they are all positive integers. This only changes the partition function by an easily computable factor. If the weights $\lambda_{i,j}$ and $\mu_{v,c}$ are positive

https://doi.org/10.1017/S096354830600767X Published online by Cambridge University Press

integers, then the problem of computing the quantity

$$\sum_{\sigma:V\to\{1,\ldots,q\}} \prod_{(i,j)\in E:\sigma(i)=\sigma(j)} \lambda_{i,j} \prod_{v\in V} \mu_{v,\sigma(v)}$$

is in #P, and, from equation (1.5), the partition function of the original problem can be approximated using this quantity. In the remainder of the section, we show that #SAT $\leqslant_{\mathrm{AP}}$ Ferromagnetic Potts(q).

We start by considering the multiterminal cut problem from [4]. Given a graph $G = (V, E)$ and a set $\{s_1, \ldots, s_q\}$ of vertices in V (which we refer to as terminals), a *multiterminal cut* is a set $E' \subseteq E$ whose removal disconnects the terminals in the sense that the graph $(V, E - E')$ does not contain a path between any two distinct terminals. The size of the multiterminal cut is the number of edges in E'. The problem Multiterminal Cut(q) is defined as follows.

Name. Multiterminal Cut(q).
Instance. A positive integer b, a connected graph $G = (V, E)$, and q distinct vertices $s_1, \ldots, s_q$ from V. The input is only valid if every multiterminal cut for $G, s_1, \ldots, s_q$ has size at least b.
Output. Is there a multiterminal cut for $G, s_1, \ldots, s_q$ of size b?

For fixed $q > 2$, Multiterminal Cut(q) is NP-complete. This result is due to Dahlhaus, Johnson, Papadimitriou, Seymour and Yannakakis [4]. The way that we have stated the problem Multiterminal Cut(q) is slightly unusual, and some comments are in order. The result that we use from [4] is Theorem 3, which says that following problem is NP-hard.

Name. Multiterminal Cut.
Instance. A positive integer b, a graph $G = (V, E)$, and 3 distinct vertices s_1, s_2, s_3 from V.
Output. Is there a multiterminal cut for G, s_1, s_2, s_3 of size at most b?

The NP-hardness proof given in [4] is a reduction from Simple Max Cut.

Name. Simple Max Cut.
Instance. A graph G and a positive integer k
Output. Is there a partition of the vertices of G into two sets V_1 and V_2 such that there are at least k edges between V_1 and V_2?

In particular, Dahlhaus, Johnson, Papadimitriou, Seymour and Yannakakis show how to take a graph $G = (V, E)$ and construct a graph F with terminals s_1, s_2, s_3 such that, for any K, G has a cut of size at least K if and only if F, s_1, s_2, s_3 has a multiterminal cut of size at most $28|E| - K$. Now if we go back to the proof of NP-completeness of Simple Max Cut in [8], we find that we can constrain the input G so that (i) G is connected, and (ii) every cut of G has size at most k. By construction, the graph F is also connected, and every multiterminal cut of F, s_1, s_2, s_3 has size at least $b = 28|E| - k$. Thus,

https://doi.org/10.1017/S096354830600767X Published online by Cambridge University Press

we have established[2] the NP-hardness of MULTITERMINAL CUT(3). As Dahlhaus, Johnson, Papadimitriou, Seymour and Yannakakis mention, it is easy to reduce MULTITERMINAL CUT(3) to MULTITERMINAL CUT(q) for $q > 3$. In our case, since we want the input graph to be connected, we join each of the terminals $s_4, \ldots, s_q$ by a single edge to s_3.

Now consider the counting version of MULTITERMINAL CUT(q), as follows.

Name. #MULTITERMINAL CUT(q).
Instance. A positive integer b, a connected graph $G = (V, E)$ and q distinct vertices $s_1, \ldots, s_q$ from V. Every multiterminal cut for $G, s_1, \ldots, s_q$ has size at least b.
Output. The number of size-b multiterminal cuts for $G, s_1, \ldots, s_q$.

Theorem 1 of [7] states that if a decision problem is NP-complete then the corresponding counting problem is AP-interreducible with #SAT. This implies that #MULTITERMINAL CUT(q) is AP-interreducible with #SAT. We complete the proof of Theorem 1.2 by giving an AP-reduction from #MULTITERMINAL CUT(q) to FERROMAGNETIC POTTS(q).

Fix q and let b, $G = (V, E)$, and $s_1, \ldots, s_q$ be an input to #MULTITERMINAL CUT(q). Let $n = |V|$ and $m = |E|$. Construct an instance of FERROMAGNETIC POTTS(q) as follows. The graph is G. Let $\lambda = 8q^n$ and let $\mu = 8\lambda^b q^n$. For all $(i, j) \in E$, let $\lambda_{i,j} = \lambda$. For all $c \in \{1, \ldots, q\}$, let $\mu_{s_c,c} = \mu$. For every other (v, c), let $\mu_{v,c} = 1$.

Say that a Potts configuration $\sigma : V \to \{1, \ldots, q\}$ is *separating* if, for every $c \in \{1, \ldots, q\}$, $\sigma(s_c) = c$. A separating configuration induces a multiterminal cut E' for $G, s_1, \ldots, s_q$ in which E' is the set of edges that are bichromatic in σ. Every size-b multiterminal cut E' corresponds to exactly one separating configuration σ. The removal of E' from E splits G into exactly q components (if there were more components there would be a smaller multiterminal cut since G is connected). For every spin c, the connected component containing s_c is assigned colour c by σ.

Let N be the number of size-b multiterminal cuts. Each of these corresponds to a separating configuration σ which contributes $\lambda^{m-b}\mu^c$ to $Z(G, \lambda_{i,j}, \mu_{v,c})$.

The contribution to $Z(G, \lambda_{i,j}, \mu_{v,c})$ from separating configurations that induce larger multiterminal cuts is at most $q^n\lambda^{m-b-1}\mu^c$.

Finally, the contribution to $Z(G, \lambda_{i,j}, \mu_{v,c})$ from configurations which are not separating is at most $q^n\lambda^m\mu^{c-1}$.

We conclude that

$$N \leqslant \frac{Z(G, \lambda_{i,j}, \mu_{v,c})}{\lambda^{m-b}\mu^c} \leqslant N + \frac{q^n\lambda^{m-b-1}\mu^c}{\lambda^{m-b}\mu^c} + \frac{q^n\lambda^m\mu^{c-1}}{\lambda^{m-b}\mu^c} = N + \frac{1}{4}.$$

Thus (see the proof of Theorem 3 of [7]), the construction is an AP-reduction from #MULTITERMINAL CUT(q) to FERROMAGNETIC POTTS(q) and we have completed the proof of Theorem 1.2. □

[2] It is possible to equip the input to SIMPLE MAX CUT with a 'witness', which could be used to check that the instance has no cuts of size exceeding k. This could be translated into a witness for the input to MULTITERMINAL CUT(3).

https://doi.org/10.1017/S096354830600767X Published online by Cambridge University Press

The parameters λ and μ that we used in the above reduction are exponential in n. The reduction was written this way for easy presentation, but the large weights can be eliminated using the method of Section 3.2.

In particular, the reduction requires the edges of G to have an effective weight $\lambda_{\text{eff}} \geqslant 8q^n$. This can be achieved by replacing the edge with t parallel 2-edge paths, and giving each of the $2t$ new edges a weight $\lambda' = 2$. The effective weight is given by

$$\lambda_{\text{eff}} = \left(\frac{\lambda'^2 + q - 1}{2\lambda' + q - 2}\right)^t = \left(1 + \frac{1}{q+2}\right)^t,$$

so it suffices to make t equal to the ceiling of the logarithm (base $1 + 1/(q-2)$) of $8q^n$. This gives a λ_{eff} which is between $8q^n$ and $8q^n(1 + 1/(q+2))$. Similarly, the reduction requires each vertex s_c to have an effective weight $\mu_{\text{eff}\,s_c,c} \geqslant 8\lambda_{\text{eff}}^b q^n$. This can be achieved by attaching the vertex s_c to t new vertices $w_1, \ldots, w_t$. Each of the new edges is given weight $\lambda' = 2$ and each of the new vertices w_i is given weight $\mu_{w_i,c} = \mu' = 2$. All of the weights $\mu_{s_c,c'}$ are set to 1 (even for $c' = c$). Then

$$\mu_{\text{eff}\,s_c,c} = \left(\frac{\lambda'\mu' + q - 1}{\lambda' + \mu' + q - 2}\right)^t = \left(1 + \frac{1}{q+2}\right)^t,$$

so t can be chosen appropriately, as before.

References

[1] Alava, M. and Rieger, H. (1998) Chaos in the random field Ising model. *Phys. Review E* **58** 4284–4287.

[2] Baxter, R. J. (1982) *Exactly Solved Models in Statistical Mechanics*, Academic Press.

[3] Biskup, M., Borgs, C. Chayes, J. T. and Kotecky, R. (2000) Gibbs states of graphical representations in the Potts model with external fields. *J. Math. Phys.* **41** 1170–1210.

[4] Dahlhaus, E., Johnson, D. S., Papadimitriou, C. H., Seymour, P. D. and Yannakakis, M. (1994) The complexity of multiterminal cuts. *SIAM J. Comput.* **23** 864–894.

[5] Dietzfelbinger, M. and Kunde, M. (2003) A case against using Stirling's formula. *Bull. European Association for Theoretical Computer Science* (EATCS) **80**.

[6] Dimakos, X. K. (2001) A guide to exact simulation. *International Statistical Review* **69** 27–48.

[7] Dyer, M., Goldberg, L. A., Greenhill, C. and Jerrum, M. (2003) On the relative complexity of approximate counting problems. *Algorithmica* **38** 471–500.

[8] Garey, M. R., Johnson, D. S. and Stockmeyer, L. (1976) Some simplified NP-complete graph problems. *Theoret. Comput. Sci.* **1** 237–267.

[9] Goldberg, L. A. and Jerrum, M. (1999) Randomly sampling molecules. *SIAM J. Comput.* **29** 834–853.

[10] Goldberg, L. A., Jerrum, M. and Paterson, M. (2003) The computational complexity of two-state spin systems. *Random Struct. Alg.* **23** 133–154.

[11] Gore, V. K. and Jerrum, M. R. (1999) The Swendsen–Wang process does not always mix rapidly. *J. Statist. Phys.* **97** 67–86.

[12] Jaeger, F., Vertigan, D. L. and Welsh, D. J. A. (1990) On the computational complexity of the Jones and Tutte polynomials. *Math. Proc. Cambridge Philos. Soc.* **108** 35–53.

https://doi.org/10.1017/S096354830600767X Published online by Cambridge University Press

[13] Jerrum, M. and Sinclair, A. (1993) Polynomial-time approximation algorithms for the Ising model. *SIAM J. Comput.* **22** 1087–1116.
[14] Mizuno, J. (1974) Effects of random local magnetic fields in the Ising model. *J. Physics C: Solid State Phys.* **7** 3755–3757.
[15] Welsh, D. J. A. (1993) *Complexity: Knots, Colourings and Counting*, Vol. 186 of *London Mathematical Society Lecture Note Series*, Cambridge University Press.
[16] Zuckerman, D. (1996) On unapproximable versions of NP-complete problems. *SIAM J. Comput.* **25** 1293–1304.

https://doi.org/10.1017/S096354830600767X Published online by Cambridge University Press