The computational complexity of two-state spin systems*

Leslie Ann Goldberg [†]
Dep't of Computer Science
University of Warwick

Mark Jerrum[‡]
Division of Informatics
University of Edinburgh

Mike Paterson[§]
Dep't of Computer Science
University of Warwick

November 29, 2001

Abstract

The subject of this article is spin-systems as studied in statistical physics. We focus on the case of two spins. This case encompasses models of physical interest, such as the classical Ising model (ferromagnetic or antiferromagnetic, with or without an applied magnetic field) and the hard-core gas model. There are three degrees of freedom, corresponding to our parameters β, γ and μ. We wish to study the complexity of (approximately) computing the partition function in terms of these parameters. We pay special attention to the symmetric case $\mu = 1$ for which our results are depicted in Figure 1. Exact computation of the partition function Z is NP-hard except in the trivial case $\beta\gamma = 1$, so we concentrate on the issue of whether Z can be computed within small relative error in polynomial time. We show that there is a fully polynomial randomised approximation scheme (FPRAS) for the partition function in the "ferromagnetic" region $\beta\gamma \geq 1$, but (unless RP = NP) there is no FPRAS in the "antiferromagnetic" region corresponding to the square defined by $0 < \beta < 1$ and $0 < \gamma < 1$. Neither of these "natural" regions — neither the hyperbola nor the square — marks the boundary between tractable and intractable. In one direction, we provide an FPRAS for the partition function within a region which extends well away from the hyperbola. In the other direction, we exhibit two tiny, symmetric, intractable regions extending beyond the antiferromagnetic region. We also extend our results to the asymmetric case $\mu \neq 1$.

*Research Report 386, Department of Computer Science, University of Warwick, Coventry CV4 7AL, UK (Nov 2001). 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).

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

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

[§] `msp@dcs.warwick.ac.uk`, `http://www.dcs.warwick.ac.uk/~msp/`, Department of Computer Science, University of Warwick, Coventry, CV4 7AL, United Kingdom.

1 Introduction

The subject of this article is "spin-systems" as studied in statistical physics. An instance of a spin-system is an n-vertex graph $G = (V, E)$. Let $q \geq 2$ be an integer. A *configuration* of a spin system on G is one of the q^n possible assignments $\sigma : V \to \{0, \dots, q-1\}$ of q *spins* to the vertices of G. (We shall usually refer to spins as *colours*.) Each configuration has an energy $H(\sigma)$ which is the sum of individual contributions from the edges and vertices of G. The contribution of each edge $\{i, j\} \in E$ is a specified function (here assumed symmetric) of the colours $\sigma(i)$ and $\sigma(j)$; likewise, the contribution of vertex $k \in V$ is a function of $\sigma(k)$. Each configuration has weight $w(\sigma) = \exp(-H(\sigma)/T)$, where T is a parameter of the system called *temperature*. The *partition function* of the system is the normalising factor $Z = \sum_\sigma \exp(-H(\sigma)/T)$ that turns the weights into probabilities.[1]

Our goal in this paper is to study the complexity of computing the partition function of spin systems. We shall deal exclusively with two-spin ($q = 2$) systems, since these already seem to present enough of a challenge. Moreover, the case $q = 2$ encompasses models of physical interest, such as the classical Ising model (ferromagnetic or antiferromagnetic, with or without an applied magnetic field), or the hard-core gas model. We refer to the two colours (spins) as "blue" and "green". Since $w(\sigma) = \exp(-H(\sigma)/T)$ and $H(\sigma)$ is a *sum* of contributions from edges and vertices, we can equivalently take a multiplicative view, in which $w(\sigma)$ is defined as a *product* of contributions from the individual edges and vertices. (All this will be set up formally in the next section; however, we hope that this informal account provides an adequate basis for at least a qualitative discussion of the main results of the paper.)

At first sight it seems as though there are three parameters governing edge contributions (corresponding to blue-blue, blue-green and green-green edges), and two governing vertex contributions (corresponding to blue and green vertices). But we may normalise the (multiplicative) blue-green edge contribution to 1, and the blue vertex contribution to 1 also.[2] Thus there are essentially three degrees of freedom. We denote the (multiplicative) blue-blue edge contribution by β, the green-green by γ, and the green vertex contribution by μ. In fact — partly because it is easier to depict a two-dimensional parameter space, and partly because our understanding of the general situation is still incomplete — we shall pay particular attention to the special (symmetric) case $\mu = 1$.

Figure 1 shows the regions in (β, γ)-space as classified by our results when $\mu = 1$. Exact computation of the partition function Z is NP-hard except in the trivial case $\beta\gamma = 1$ so we concentrate on the issue of whether Z can be computed within small relative error in polynomial time. (The precise notion of efficient approximation algorithm used is the "fully polynomial randomised approximation scheme" or FPRAS, which will be defined in §2.) The main features are as follows:

1. To the North-East of the hyperbola is a "ferromagnetic" region $\beta\gamma \geq 1$ within which the partition function may be approximated in the FPRAS sense. This is done by reduction to a ferromagnetic Ising system with external field, whose partition function may be approximated by a Markov chain Monte Carlo algorithm of Jerrum and Sinclair [7]. See §3.

2. The square defined by $0 < \beta < 1$ and $0 < \gamma < 1$ is an "antiferromagnetic" region within which the partition function is hard to approximate (unless RP = NP). Essentially this is because "ground states" (i.e., most likely or most weighty configurations) correspond to maximum cuts in G. This at least is the intuition; the formalisation of it requires some work. See §4.

[1] Readers who do not find the physical setting congenial, may think instead of a weighted version of the graph homomorphism problem. See Section 1.1 of [4].

[2] This is equivalent to normalising the energy contribution of blue-green edges and blue vertices to zero.

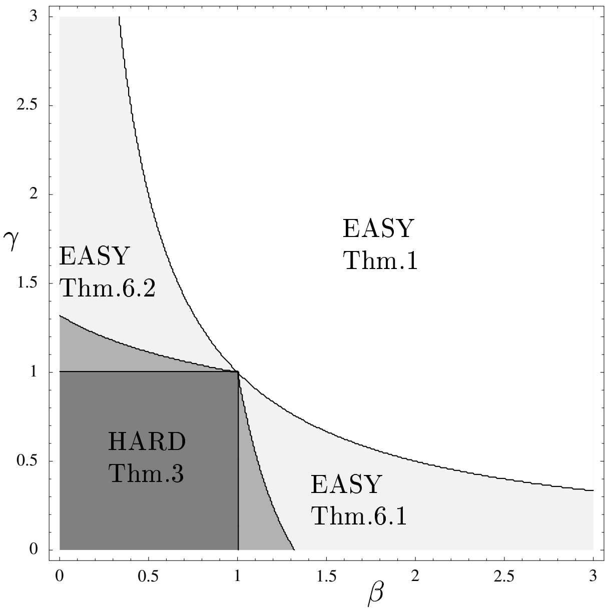

Figure 1: Results for regions of the (β,γ)-plane at $\mu=1$.

3. Neither of these "natural" regions — neither the hyperbola nor the square — marks the boundary between tractable and intractable. In one direction, we provide an FPRAS for the partition function within the light grey region, which extends well away the hyperbola. This FPRAS is based on the Markov chain Monte Carlo method, and its analysis uses the "path-coupling" technique of Bubley and Dyer [2]. See §5.

4. In the other direction, we exhibit two tiny, symmetric, intractable regions extending beyond the "antiferromagnetic" region, close to the points $(0,1)$ and $(1,0)$. This is done by coding up an inapproximable combinatorial optimisation problem following Luby and Vigoda [9]. See §6.

It will be seen that our knowledge even of the $\mu=1$ case is incomplete: specifically, we don't know what happens in the remaining (medium intensity grey) regions. For example, we don't know whether tractability is monotone in β (or γ). In the remaining sections, we prove the results depicted in Figure 1 and also extend these results beyond the symmetric case $\mu=1$.

2 Definitions

To formalise the claims made in the introduction we need to define precisely the terms (two-state) "spin system" and "FPRAS" (our notion of efficient approximate computation).

In order to define the partition function of a two-state spin spin system specified by weights β, γ and μ, it is convenient to identify blue and green with the unit vectors $(1,0)'$ and $(0,1)'$, respectively. (Primes will be used to denote transposition, so spins are column vectors.) Then the partition function for a graph $G=(V,E)$ may be expressed as

$$Z(G)=\sum_{\sigma}\prod_{\{i,j\}\in E}\sigma(i)'A\,\sigma(j)\prod_{k\in V}b'\,\sigma(k),$$

where

$$A = \begin{pmatrix} \beta & 1 \\ 1 & \gamma \end{pmatrix} \quad \text{and} \quad b = \begin{pmatrix} 1 \\ \mu \end{pmatrix},$$

and σ ranges over $\{(1,0)', (0,1)'\}^V$. To see this, note that each of the four possible assignments of unit vectors to $\sigma(i)$ and $\sigma(j)$ picks out a distinct element of the matrix A, and similarly with b.

The problem whose complexity we study is (β, γ, μ)-PARTITION, defined as follows. Let β, γ and μ be non-negative real numbers.

Name. (β, γ, μ)-PARTITION.

Instance. A graph G.

Output. The quantity $Z(G)$, where Z is the partition function with parameters β, γ and μ.

Note that the graph G alone forms the problem instance, which means we have a separate problem for every triple (β, γ, μ). (Our notation is intended to emphasise this.) Our goal is to map out the tractable region of the parameter space. To avoid the issues of specifying and computing with arbitrary real numbers, we assume that β, γ and μ are rational.

In this area, approximation algorithms are usually viewed as computing functions $f : \Sigma^* \to \mathbb{N}$, where Σ is a finite alphabet for encoding problem instances. In the current application, however, the output may be an arbitrary rational number. Rather than redefining a well-established notion of efficient approximate computation, we shall stick with the usual definition, and then explain how to view (β, γ, μ)-PARTITION in this framework.

A *randomised approximation scheme* for a counting problem $f : \Sigma^* \to \mathbb{N}$ (e.g., the number of matchings in a graph) is a randomised 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 $N \in \mathbb{N}$ (a random variable of the "coin tosses" made by the algorithm) such that, for every instance x,

$$\Pr\left[e^{-\varepsilon} f(x) \le N \le e^{\varepsilon} f(x)\right] \ge \frac{3}{4}. \tag{1}$$

We speak of a *fully polynomial randomised approximation scheme*, or *FPRAS*, if the algorithm runs in time bounded by a polynomial in $|x|$ and ε^{-1}. It is a standard result that the number $\frac{3}{4}$ appearing in (1) could be replaced by any number in the open interval $(\frac{1}{2}, 1)$.

To bring the problem (β, γ, μ)-PARTITION within the FPRAS framework, we suggest the following: Assume β, γ and μ are rational, and let L be the least common multiple of their denominators. Then the desired output $Z(G)$ can be expressed as a rational number z with denominator L^{n+m}, where n is the number of vertices in G and m the number of edges. Then our goal is to design an FPRAS computing $L^{n+m} z$.

3 The "ferromagnetic region" is tractable

We argue that the region $\beta\gamma \ge 1$ corresponds fairly directly to the ferromagnetic Ising model with external field. It follows that there is an FPRAS for the partition function Z in this region. When $\beta\gamma = 1$ the partition function is trivially computable in polynomial time.

To make this correspondence explicit, observe that

$$A = \sqrt{\frac{\beta}{\gamma}} \begin{pmatrix} 1 & 0 \\ 0 & \gamma/\alpha \end{pmatrix} \begin{pmatrix} \alpha & 1 \\ 1 & \alpha \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & \gamma/\alpha \end{pmatrix},$$

where $\alpha = \sqrt{\beta\gamma}$, and hence

$$\begin{aligned} \sigma(i)' A\, \sigma(j) &= \sqrt{\frac{\beta}{\gamma}}\, \sigma(i)' \begin{pmatrix} 1 & 0 \\ 0 & \gamma/\alpha \end{pmatrix} \begin{pmatrix} \alpha & 1 \\ 1 & \alpha \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & \gamma/\alpha \end{pmatrix} \sigma(j) \\ &= \sqrt{\beta/\gamma}\, \big[(1, \gamma/\alpha)\, \sigma(i)\big] \big[\sigma(i)' \widehat{A}\, \sigma(j)\big] \big[(1, \gamma/\alpha)\, \sigma(j)\big], \end{aligned}$$

where

$$\widehat{A} = \begin{pmatrix} \alpha & 1 \\ 1 & \alpha \end{pmatrix},$$

and the final equality uses the fact that spins are unit vectors. Thus we obtain the following alternative expression for the partition function:

$$Z(G) = \left(\frac{\beta}{\gamma}\right)^{m/2} \sum_{\sigma} \prod_{\{i,j\} \in E} \sigma(i)' \widehat{A}\, \sigma(j) \prod_{k \in V} \big(1, (\gamma/\alpha)^{d(k)} \mu\big)\, \sigma(k), \tag{2}$$

where $d(k)$ is the degree of vertex k. To verify (2), note that each of the $d(k)$ edges incident at k contributes a factor $(1, \gamma/\alpha)\, \sigma(k)$ to $Z(G)$, in addition to the $(1, \mu)\, \sigma(k)$ already present.

Suppose for the moment that $\mu = 1$. When $\beta\gamma \geq 1$, i.e., when $\alpha \geq 1$, equation (2) is, up to an easily computable factor, the partition function for a ferromagnetic Ising system with external field. Jerrum and Sinclair [7] have exhibited an FPRAS for computing the partition function of such systems, from which it follows that the region $\beta\gamma \geq 1$ is tractable. More precisely:

Theorem 1 *For any fixed β, γ satisfying $\beta\gamma \geq 1$ there is an FPRAS for $(\beta, \gamma, 1)$-*PARTITION. *More generally, there is an FPRAS for (β, γ, μ)-*PARTITION *provided, in addition, $\beta \geq \gamma$ and $\mu \leq \sqrt{\beta/\gamma}$ (or $\beta \leq \gamma$ and $\mu \geq \sqrt{\beta/\gamma}$).*

Proof. Once we have provided a translation between the terminology of the current paper and that of [7], it will be seen that the existence of an FPRAS is immediate from [7, Thm. 5]. (The latter theorem simply asserts the existence of an FPRAS for estimating the partition function of a ferromagnetic Ising system.)

First, a brief description of the Ising model. The *Ising model* is a two-spin model in which interactions are symmetric under interchange of the two colours (spins): in our terminology $\alpha = \beta = \gamma$. In the *ferromagnetic* Ising model, like spins are favoured over unlike, i.e., $\alpha \geq 1$. There may be an *external* (or *applied*) *field*, that causes one colour to be favoured over the other: in our terminology $\mu \neq 1$. The interactions are allowed to vary from edge to edge, provided they are all ferromagnetic. Thus we may have a separate matrix A_{ij} associated with each edge $\{i, j\} \in E$, provided each matrix individually satisfies the conditions stated above (diagonal entries equal and not less than 1.) The interactions with the external field may also vary, i.e., the vector b may vary from vertex to vertex. However, one colour must be uniformly favoured over the other; in other words the parameter μ must be uniformly at least 1, or uniformly at most 1.

Inspecting equation (2), we see that the aforementioned conditions are met, provided only that $(\gamma/\alpha)^{d(k)}\mu = (\gamma/\beta)^{d(k)/2}\mu$ is uniformly at least 1 or at most 1. This will certainly be the case if $\mu = 1$. But it will also hold in the other situations identified in the statement of the theorem.[3]

In order to give the details of the reduction from (β, γ, μ)-PARTITION to Theorem 5 of [7], we need to show how to encode the input, that is, G, α, and the quantities $(\gamma/\beta)^{d(k)/2}\mu$ as binary

[3]Since it is trivial to deal with any isolated vertices of G, we may assume that $d(k) \geq 1$ for all k.

strings of appropriate length. The details of this are routine, and are omitted. (Clearly, only an approximation to α is used, since α itself may be irrational.)

One final technical point concerning [7, Thm. 5]. In the proof of that theorem it is assumed that the interaction of the external field with spins is uniform over all sites, whereas we require here a non-uniform (though consistently oriented) interaction. The proof was organised in this way for simplicity of presentation. The clean fix is to routinely amend the proof by introducing explicit individual interaction strengths at the various sites. However, an alternative fix that does not involve delving into the original proof is to reduce the case of varying interaction strengths to that of fixed. In particular, suppose ε is our desired accuracy parameter and consider an instance G with, for each vertex v, an interaction strength $\mu_v \geq 1$. Let

$$\delta = \frac{\varepsilon}{n\alpha\lceil \max_v \ln \mu_v \rceil},$$

$\mu = 1 + \delta$, and

$$z = \frac{(\alpha - 1)\delta}{1 + \alpha + \delta} = \frac{\mu\alpha + 1}{\mu + \alpha} - 1.$$

The Ising partition function for this instance is closely approximated by the partition function of a new instance in which the graph, G', is formed from G by attaching

$$r_v \doteq \left\lceil \frac{\ln(\mu_v/\mu)}{z} \right\rceil$$

pendant edges to each vertex v and giving each vertex interaction strength μ. To see that the approximation is sufficiently close, note that the relative weight of colouring v green rather than blue in G' is

$$\psi \doteq \mu\left(\frac{\mu\alpha + 1}{\mu + \alpha}\right)^{r_v}.$$

Thus the definitions guarantee

$$e^{-\varepsilon/n}\mu_v \leq \psi \leq e^{\varepsilon/n}\mu_v.$$

as required. □

Remark 1 *When $\beta\gamma = 1$, expression (2) factorises and the (exact) computation of Z is trivial.*

Remark 2 *Another situation in which $(\gamma/\alpha)^{d(k)}\mu$ is assured to be uniformly at least* 1 *or at most* 1 *is when $d(k)$ is constant, i.e., G is regular.*

The parameter values *not* covered by Theorem 1, i.e., $\beta > \gamma$ and $\mu > \sqrt{\beta/\gamma}$ (or $\beta < \gamma$ and $\mu < \sqrt{\beta/\gamma}$) present a conundrum. These correspond to a situation, which may be physically unrealistic, in which some vertices incline to one colour and others to the other. On the one hand, there is no obvious barrier to FPRASability when this occurs. On the other hand, the proof of [7, Thm. 5] certainly breaks down. The issue is that the quantity $\tanh \beta B$ in [7, eq. (2)] will be of inconsistent sign, leading to negative weights $w(X)$ in [7, eq. (3)]. In this situation, the so-called "subgraphs world" process is no longer well defined, as various "probabilities" become negative.

4 The "antiferromagnetic region" is intractable

Let $\leq_{\mathrm{AP}}$ be the approximation-preserving reduction from [5]. Let #SAT and #LARGECUT be defined as follows.

Name. #SAT.

Instance. A Boolean formula φ in conjunctive normal form (CNF).

Output. The number of satisfying assignments to φ.

Name. #LARGECUT.

Instance. A positive integer k and a connected graph G in which every cut[4] has size at most k.

Output. The number of size-k cuts of G.

An AP-reduction from #SAT to #LARGECUT appears in [7].[5] For certain β, γ and μ (see Lemma 2) we will give an AP-reduction from #LARGECUT to (β,γ,μ)-PARTITION. The combination of these reductions implies #SAT $\leq_{\mathrm{AP}}$ (β,γ,μ)-PARTITION which in turn implies that there cannot be an FPRAS for (β,γ,μ)-PARTITION unless NP = RP (see Section 3 of [5]).

Lemma 2 *Let β, γ and μ be fixed parameters satisfying $0<\beta<1$, $0<\gamma<1$ and $\mu>0$. Then* #LARGECUT $\leq_{\mathrm{AP}}$ (β,γ,μ)-PARTITION.

Proof. Let k and $G=(V,E)$ be an instance of #LARGECUT and let n denote $|V|$ and m denote $|E|$. We wish to construct an instance $G'=(V',E')$ of (β,γ,μ)-PARTITION. In order to make the reduction explicit, we will need to define a quantity s which depends upon β, γ, μ and n. The reader should think of s as simply being a sufficiently large polynomial in n. For completeness, let c be a positive integer such that the quantity

$$\rho \doteq \frac{(\min(\beta,\gamma))^2}{(\max(\beta,\gamma))^{c-1}}$$

exceeds 1. It will then suffice to let s be the smallest integer satisfying $s \leq \rho^{s/(2c)}$ which is at least

$$\max\left(\frac{n+6}{\lg\left(\frac{1}{\beta\gamma}\right)}, \frac{c\ln\left(\frac{2^{2n}(\max(1,\mu))^{2n}n2^5}{(\beta\gamma)^m\mu^n}\right)}{\ln\left(\frac{1}{\max(\beta,\gamma)}\right)}, \frac{2c\ln\left(\max(\mu,1/\mu)n2^{n+5}\right)}{\ln\rho}, cn\right).$$

We now give the construction of G'. For every vertex u of G let A_u and B_u be disjoint sets of size s. Let

$$V' = \bigcup_{u\in V} A_u \cup B_u$$

and

$$E' = \left(\bigcup_{u\in V} A_u \times B_u\right) \cup \left(\bigcup_{(u,v)\in E}\ \bigcup_{i\in\{1,\ldots,s\}} \{(A_u[i],A_v[i]),(B_u[i],B_v[i])\}\right).$$

Let $\Omega(G')$ denote the set of all two-spin configurations on G'. For any subset $W\subseteq\Omega(G')$, let $Z_W(G')$ denote the contribution to $Z(G')$ corresponding to configurations in W. A configuration σ

[4]Recall that a "cut" of a graph is an unordered partition of its vertex set into two subsets and that the size of the cut is the number of edges between the two subsets.

[5]The definition of #LARGECUT may seem unnatural because it is not easy in general to verify the promise that no cuts exceeding size-k exist in the input graph. However, the reduction in [7] can be viewed as producing an input graph together with a "witness" which allows the promise to be checked.

is *full* if, for every vertex u of G, all of A_u is coloured with one of the two possible spins and all of B_u is coloured with the other spin. Every cut of G corresponds to exactly two full configurations: If u and v are in the same part of the cut then A_u and A_v are coloured with the same colour. If σ is a full configuration corresponding to a size-j cut then

$$Z_{\{\sigma\}}(G') = \prod_{\{i,j\}\in E'} \sigma(i)' A\, \sigma(j) \prod_{k\in V'} b'\, \sigma(k) = (\beta\gamma)^{s(m-j)}\mu^{sn}.$$

Let N be the number of size-k cuts of G and let C be the set of full configurations which correspond to size-k cuts. Let

$$\Psi = 2(\beta\gamma)^{s(m-k)}\mu^{sn},$$

so

$$Z_C(G') = N\Psi.$$

We will now show

$$Z_{\Omega(G')-C}(G') \le 2^{-4}(\Psi + 2^{-n}Z(G')). \tag{3}$$

Equation (3) implies

$$N \le \frac{Z(G')}{\Psi} \le N + \frac{1}{4}. \tag{4}$$

To see this, consider first the case $N=0$. In this case $Z(G') = Z_{\Omega(G')-C}(G')$, so Equation (3) gives

$$Z(G')(1 - 2^{-(n+4)}) \le 2^{-4}\Psi$$

and therefore $0 \le Z(G')/\Psi \le 1/4$ so Equation (4) holds. If $N>0$ then Equation (3) gives

$$Z(G') = Z_C(G') + Z_{\Omega(G')-C}(G') \le Z_C(G')\left(1 + \frac{1}{16N}\right) + \frac{1}{16N}Z(G').$$

Thus

$$Z(G') \le \frac{1+\frac{1}{16N}}{1-\frac{1}{16N}} Z_C(G') \le (1+\frac{1}{4N})Z_C(G'),$$

which implies Equation (4). From Equation (4), we find that

$$N = \left\lfloor \frac{Z(G')}{\Psi} \right\rfloor. \tag{5}$$

Also, the floor function in Equation (5) does not distort the accuracy overly much: An approximation to $Z(G')$ gives an approximation to N. The details about the accuracy of the approximation are the same as those in the proof of Theorem 3 of [5].

So, to conclude the proof we prove Equation (3). We do this by splitting $\Omega(G') - C$ into several (potentially overlapping) sets and then summing the partition function over these sets. Let F be the set of full configurations corresponding to cuts of size less than k. Then since there are at most 2^n cuts,

$$Z_F(G') \le 2^n(\beta\gamma)^s\Psi \le 2^{-6}\Psi. \tag{6}$$

The second inequality in Equation (6) follows from the fact that s is at least the first term in its definition.

For $u \in V$, let a_u be the set of configurations in which A_u has at least s/c green vertices and at least s/c blue vertices. Let $a = \cup_u a_u$.

$$Z_{a_u}(G') \le 2^{2ns}\max(1,\mu)^{2ns}\max(\beta,\gamma)^{s\cdot s/c} \le 2^{-6}\Psi/n. \tag{7}$$

To see why the first inequality in Equation (7) holds, observe that the number of configurations is at most 2^{2ns}. Each of the $2ns$ vertices has weight at most $\max(1,\mu)$. All edge-weights are at most one, but each of the s vertices in B_u has weight at most $\max(\beta,\gamma)^{s/c}$. The second inequality in Equation (7) follows from the fact that s is at least the second term in its definition.

For $w \in \cup_{u\in V} B_u$, let a'_w be the set of $\sigma \in \Omega(G') - a$ in which at least half of the edges from w to $\cup_u A_u$ are monochromatic. Let $a' = \cup_w a'_w$. We will show

$$Z_{a'_w}(G') \le 2^{-(n+5)} Z(G')/(ns). \tag{8}$$

We will use the following notation to establish (8). For a configuration σ, and a vertex w of G', let $\sigma\backslash w$ be the restriction of σ to $V' - \{w\}$. Let R be the restrictions of configurations in $\Omega(G') - a$ to $V' - \{w\}$. In particular,

$$R = \{\sigma\backslash w \mid \sigma \in \Omega(G') - a\}.$$

For every $\pi \in R$, let

$$Z^R_{\{\pi\}} = \prod_{\substack{\{i,j\}\in E' \\ i\ne w,\ j\ne w}} \pi(i)' A\, \pi(j) \prod_{\substack{k\in V' \\ k\ne w}} b'\, \pi(k).$$

Then

$$Z_{a'_w}(G') \le \sum_{\pi\in R} Z^R_{\{\pi\}} \max(\beta,\gamma)^{(1-1/c)s} \max(1,\mu).$$

Also,

$$Z(G') \ge Z_{\Omega(G')-a-a'_w}(G') \ge \sum_{\pi\in R} Z^R_{\{\pi\}} \min(\beta,\gamma)^{(s/c)+n} \min(1,\mu).$$

So

$$\begin{aligned}
Z_{a'_w}(G') &\le \frac{\max(\beta,\gamma)^{(1-1/c)s}\max(1,\mu)}{\min(\beta,\gamma)^{(s/c)+n}\min(1,\mu)} \sum_{\pi\in R} Z^R_{\{\pi\}} \min(\beta,\gamma)^{(s/c)+n}\min(1,\mu) \\
&\le \frac{\max(\beta,\gamma)^{(1-1/c)s}}{\min(\beta,\gamma)^{(s/c)+n}} \max(\mu,1/\mu) Z(G') \\
&\le \left(\frac{\max(\beta,\gamma)^{c-1}}{\min(\beta,\gamma)^2}\right)^{s/c} \max(\mu,1/\mu) Z(G') \\
&\le 2^{-(n+5)} Z(G')/(ns).
\end{aligned}$$

The second-to-last inequality uses $n \le s/c$ and the final inequality follows from the fact that $s \le \rho^{s/(2c)}$ and the fact that s is at least the third term in its definition. Thus, Equation (8) is established.

For $u \in V$, let b_u be the set of configurations in which B_u has at least s/c green vertices and at least s/c blue vertices. Let $b = \cup_u b_u$. By analogy to Equation (7), we get

$$Z_{b_u}(G') \le 2^{-6}\Psi/n. \tag{9}$$

For $w \in \cup_{u\in V} A_u$, let b'_w be the set of $\sigma \in \Omega(G') - b$ in which at least half of the edges from w to $\cup_u B_u$ are monochromatic. Let $b' = \cup_w b'_w$. By analogy to Equation (8), we get

$$Z_{b'_w}(G') \le 2^{-(n+5)} Z(G')/(ns). \tag{10}$$

Equation (3) follows from Equations (6), (7), (8), (9) and (10) since

$$Z_{\Omega(G')-C}(G') \leq Z_F(G') + Z_a(G') + Z_{a'}(G') + Z_b(G') + Z_{b'}(G'). \tag{11}$$

To see that (11) holds, consider any configuration σ which is not in $a \cup a' \cup b \cup b'$. Consider any vertex u of G. Since $\sigma \notin a_u$, more than $1 - 1/c$ of the nodes in A_u have a certain colour. So, since $\sigma \notin a'_w$ for any $w \in B_u$, all of B_u is coloured with the other colour. Finally, since $\sigma \notin b'_w$ for any $w \in A_u$, all of A_u is coloured with the same colour. We conclude that σ is full, so it is either in F or in C.

□

Lemma 2 has the following consequence

Theorem 3 *Let β, γ and μ be fixed parameters satisfying $0 < \beta < 1$, $0 < \gamma < 1$ and $\mu > 0$. Then there is no FPRAS for (β, γ, μ)-*Partition *unless* NP = RP.

5 An additional tractable region

Theorem 3 showed that there is unlikely to be an FPRAS for (β, γ, μ)-Partition when $0 < \beta < 1$ and $0 < \gamma < 1$. Theorem 1 showed that in the region $\beta\gamma \geq 1$ there is an FPRAS. In this section we will assume that $\beta\gamma < 1$ and either $\beta > 1$ or $\gamma > 1$. Our aim is to identify an additional region where there is still an FPRAS. The FPRAS is based on the simulation of the single-site heat-bath Markov chain, which is studied in Section 5.1.

5.1 Rapid mixing within the region

The single-site heat-bath chain for the two-state partition function works as follows. Given a (connected) n-vertex input graph $G = (V, E)$, $\Omega(G)$ is the state space (the set of configurations, i.e., the set of all 2-colourings of G, including improper colourings). From a configuration $\sigma \in \Omega(G)$, the chain first chooses a vertex $x \in V$ u.a.r. Let $\sigma(x \to g)$ denote the configuration obtained from σ by colouring x green, and $\sigma(x \to b)$ the configuration corresponding to colouring x blue. Let

$$p(x, g)(\sigma) = \frac{Z_{\{\sigma(x\to g)\}}(G)}{Z_{\{\sigma(x\to g)\}}(G) + Z_{\{\sigma(x\to b)\}}(G)}$$

and $p(x, b)(\sigma) = 1 - p(x, g)(\sigma)$. The new state is taken to be $\sigma(x \to g)$ with probability $p(x, g)(\sigma)$ and $\sigma(x \to b)$ otherwise.

We will use path coupling [2] to prove that single-site heat bath is rapidly mixing. We adopt the notation from [3]. Let $S \subseteq \Omega(G)^2$ be the set of pairs of configurations with Hamming-distance 1. If σ and σ' are configurations which disagree only at vertex v then $\Psi(\sigma, \sigma')$ (the *proximity* of σ and σ') is defined to be the degree of v in G, which we denote $\Delta[v]$. The distance function is given in the usual way: For each pair $(\sigma, \sigma') \in \Omega(G)^2$, $\mathcal{P}(\sigma, \sigma')$ is the set of all sequences $\sigma = \sigma_1, \sigma_2, \ldots, \sigma_{r-1}, \sigma_r = \sigma'$ with $(\sigma_i, \sigma_{i+1}) \in S$ for $i \in \{1, \ldots, r-1\}$. The distance function is defined by

$$\delta(\sigma, \sigma') = \min_{\mathcal{P}(\sigma,\sigma')} \sum_{i=1}^{r-1} \Psi(\sigma_i, \sigma_{i+1}), \tag{12}$$

which can be written as

$$\delta(\sigma, \sigma') = \sum_{v\in V} I_v(\sigma, \sigma')\Delta[v], \tag{13}$$

where $I_v(\sigma,\sigma')$ is the indicator for the event that σ and σ' differ at vertex v, i.e., the event $\sigma(v) \neq \sigma'(v)$. Note that if $Z_{\{\sigma\}}(G)$ and $Z_{\{\sigma'\}}(G)$ are both positive, then there is a chain $\sigma = \sigma_1, \sigma_2, \ldots, \sigma_{r-1}, \sigma_r = \sigma'$ which minimises the right-hand-side of (12), and for which each σ_i has $Z_{\{\sigma_i\}}(G) > 0$. For example, if $\beta > 0$ then the chain is constructed by first colouring some green vertices blue and then colouring some blue vertices green.

We will now define a coupling which, for every $(X_0, Y_0) \in S$ and every $(X_1, Y_1) \in \Omega(G)^2$, gives the probability of a joint transition from (X_0, Y_0) to (X_1, Y_1). Suppose that X_0 and Y_0 differ on v. The coupling will be the optimal one, subject to the assumption that the same vertex x is selected in X_0 and in Y_0. First, a vertex x is chosen u.a.r. If x is not a neighbour of v then the same colour is chosen for x in X_1 and in Y_1. If x is a neighbour of v, then with probability $\min(p(x,g)(X_0), p(x,g)(Y_0))$, $X_1 = X_0(x \to g)$ and $Y_1 = Y_0(x \to g)$, and with probability $\min(p(x,b)(X_0), p(x,b)(Y_0))$, $X_1 = X_0(x \to b)$ and $Y_1 = Y_0(x \to b)$. The rest of the coupling is forced by the requirement that the marginals be correct.

The path coupling lemma in [2, 3] guarantees that the chain is rapidly mixing as long as there is an $\varepsilon_n > 1/\mathrm{poly}(n)$ such that for every pair $(X_0, Y_0) \in S$, $E(\delta(X_1, Y_1)) \leq (1-\varepsilon_n)\delta(X_0, Y_0)$. In particular, the total variation distance between the t-step distribution of the chain and the stationary distribution is at most ε after only $\ln(n\varepsilon^{-1})/\varepsilon_n$ steps.

So, suppose that X_0 and Y_0 differ at vertex v. For concreteness, suppose that $X_0(v)$ is blue. For every neighbour w of v, let b_w denote the number of neighbours of w, other than v, which are coloured blue in X_0 (or equivalently, in Y_0). Let g_w denote the number of neighbours of w, other than v which are coloured green in X_0. Thus $\Delta[w] = b_w + g_w + 1 \geq 1$. Let

$$f_w(i) = \frac{\beta^{\Delta[w]-i}}{\beta^{\Delta[w]-i} + \mu\gamma^i} = \frac{1}{1 + \mu(\beta\gamma)^i \beta^{-\Delta[w]}}.$$

Note that $p(w,b)(X_0) = f_w(g_w)$ and $p(w,b)(Y_0) = f_w(g_w + 1)$.

We require that, for every pair $(X_0, Y_0) \in S$ which disagree on vertex v,

$$\left(1 - \frac{1}{n}\right)\Delta[v] + \frac{1}{n}\sum_{w\sim v} |f_w(g_w+1) - f_w(g_w)|\Delta[w] \leq (1-\varepsilon_n)\Delta[v]. \tag{14}$$

Equation (14) follows from (13). The probability that X_1 and Y_1 differ on v is equal to $1 - 1/n$, which is the probability that v is not chosen. If w is chosen then the probability that X_1 and Y_1 differ on w is $|p(w,b)(Y_0) - p(w,b)(X_0)|$. The $\Delta[v]$ on the right-hand-side of (14) represents $\delta(X_0, Y_0)$. In order to establish (14), it suffices to show that for every neighbour w of v,

$$|f_w(g_w+1) - f_w(g_w)|\Delta[w] \leq \theta, \tag{15}$$

for some $\theta < 1$, depending only on β,γ and μ. Then we can take $\varepsilon_n = (1-\theta)/n$. We will identify regions where (15) holds. We start by focusing on the case where $\beta > 1 > \gamma$. (The case $\gamma > 1 > \beta$ is symmetric to this case, and will be handled below.)

Since $\beta\gamma < 1$, $f_w(i)$ is an increasing function of i, and so $|f_w(i+1) - f_w(i)| = f_w(i+1) - f_w(i)$ for all i. To satisfy (15) for all n, G, v and w, it is sufficient to show that, for all integers $i \geq 0$ and all real $\Delta \geq 1$,

$$\left(\frac{1}{1 + y^{i+1}x_\Delta} - \frac{1}{1 + y^i x_\Delta}\right)\Delta \leq \theta. \tag{16}$$

where $y = \beta\gamma$ and $x_\Delta = \mu\beta^{-\Delta}$. Of course, we only really require this inequality for integer values of Δ, but the bounds we obtain are sufficient for our purposes.

For any fixed $i \geq 0$, let $c_i = y^i x_\Delta$. So c_i is a decreasing function of Δ, with derivative $-c_i \ln \beta$. The derivative with respect to Δ of the logarithm of

$$\left(\frac{1}{1+yc_i} - \frac{1}{1+c_i}\right)\Delta = \frac{(1-y)\Delta}{(1+yc_i)(1+c_i^{-1})}$$

is

$$\frac{1}{\Delta} - \ln\beta\left(\frac{1}{1+yc_i} - \frac{1}{1+c_i^{-1}}\right),$$

a decreasing function of Δ. This continuous function tends to $+\infty$ as $\Delta \to 0$, and tends to $-\ln\beta < 0$ as $\Delta \to \infty$. Therefore there is a unique (finite) positive value Δ_i for which this derivative vanishes, and at which the maximum value of the left-hand side of inequality (16) with respect to the real variable Δ is attained.

To map the boundary of the region of (β, γ, μ)-space for which (16) holds, we can therefore solve the following simultaneous pair of relations, expressing the conditions that the maximising value of Δ yields a value at most θ:

$$\frac{1}{\Delta_i} - \ln\beta\left(\frac{1}{1+yc_i} - \frac{1}{1+c_i^{-1}}\right) = 0, \text{ and} \tag{17}$$

$$\left(\frac{1}{1+yc_i} - \frac{1}{1+c_i}\right)\Delta_i \leq \theta. \tag{18}$$

Eliminating the explicit occurrences of Δ_i, we find the following quadratic inequality for c_i:

$$c_i^2 y + \frac{c_i(1-y)}{\theta \ln\beta} - 1 \leq 0. \tag{19}$$

Solving this quadratic for $1/c_i$ implies the inequality

$$c_i \leq \frac{2\theta\ln\beta}{1 - y + \sqrt{(1-y)^2 + y(2\theta\ln\beta)^2}}. \tag{20}$$

Since c_i is decreasing in i, (20) is satisfied for all $i \geq 0$ if and only if it is satisfied for $i = 0$. Since we can choose $\theta < 1$ arbitrarily and the right-hand side of (20) is an increasing function of θ, we can replace θ by 1 in (20) but make the inequality strict, i.e.,

$$c_i \leq c_0 < C(\ln\beta) \tag{21}$$

where

$$C(z) = \frac{2z}{1 - \beta\gamma + \sqrt{(1-\beta\gamma)^2 + 4\beta\gamma z^2}}. \tag{22}$$

Equation (17) yields

$$\ln\beta^{\Delta_0} = D(c_0), \tag{23}$$

where the right-hand side is an increasing function of c_0,

$$D(z) = \left(\frac{1}{1+\beta\gamma z} - \frac{1}{1+z^{-1}}\right)^{-1}. \tag{24}$$

We can use (21) and (23) to derive an upper bound on μ as a function of β and γ. Since $\mu = c_0\beta^{\Delta_0}$, any choice of μ which satisfies Equation (25) (below) also satisfies Equation (21) for the maximising value of Δ_0 given by Equation (17) and for some $\theta < 1$. Hence, it satisfies Equation (18), as required. Our final bound is given by

$$\mu < C(\ln\beta)\, e^{D(C(\ln\beta))}. \tag{25}$$

Note that the left-hand-side of (18) is a decreasing function of γ for the critical case, $i = 0$. Thus we can get a simpler (but worse) bound by considering the extreme case, $\gamma = 0$. Here, $C(z) = z$ and $D(z) = 1 + z$, so (25) gives the bound: $\mu < e\beta\ln\beta$.

The region defined by $\beta\gamma < 1$, $\gamma > 1$ is symmetric to the region that we have just considered. In particular, the blue–green symmetry yields the following relationships:

$$\begin{aligned} p(w,g)(X_0) &= 1 - p(w,b)(X_0) = 1 - f_w(g_w) = 1/(\mu^{-1}\gamma^{-g_w}\beta^{\Delta[w]-g_w} + 1) \\ &= 1/(1 + \mu^{-1}\beta^{b_w+1}\gamma^{b_w+1-\Delta[w]}) = \hat{f}_w(b_w + 1), \end{aligned}$$

and

$$p(w,g)(Y_0) = 1 - f_w(g_w + 1) = 1/(1 + \mu^{-1}\beta^{b_w}\gamma^{b_w-\Delta[w]}) = \hat{f}_w(b_w),$$

where $\hat{f}_w$ is derived from f_w by replacing μ, β, γ by μ^{-1}, γ, β respectively. Our requirement is

$$|\hat{f}_w(b_w + 1) - \hat{f}_w(b_w)|\Delta[w] \le \theta_n,$$

which is symmetric to Equation (15) and this is therefore met within the region defined by $\beta\gamma < 1$, $\gamma > 1$ and $\mu^{-1} < C(\ln\gamma)\, e^{D(C(\ln\gamma))}$ (from Equation (25)).

Thus, we have shown the following result.

Lemma 4 *The single-site heat-bath Markov chain for the two-state partition function is rapidly mixing within the regions defined by: $\beta\gamma < 1$ and either*

1. *$\beta > 1$ and $\mu < C(\ln\beta)e^{D(C(\ln\beta))}$, or*
2. *$\gamma > 1$ and $1/\mu < C(\ln\gamma)e^{D(C(\ln\gamma))}$,*

where functions C and D are defined in Equations (22) and (24).

Corollary 5 *The single-site heat-bath Markov chain for the two-state partition function is rapidly mixing within the regions defined by: $\beta\gamma < 1$ and either*

1. *$\beta > 1$ and $\mu < e\beta\ln\beta$, or*
2. *$\gamma > 1$ and $1/\mu < e\gamma\ln\gamma$.*

Figure 2 shows a portion of the bounding surface of the region described in Lemma 4. The rapid-mixing region lies below this surface, which represents the logarithm of the maximal μ within the region $1 < \beta \le 2$ and $0 \le \gamma < 1$.

5.2 Reducing approximate counting to sampling

Lemma 4 showed that the single-site heat-bath Markov chain is rapidly mixing within the specified region. Thus, this Markov chain provides a fully-polynomial approximate sampler (FPAS) for the two-spin partition function within the region. Such an FPAS is an algorithm which, when it is

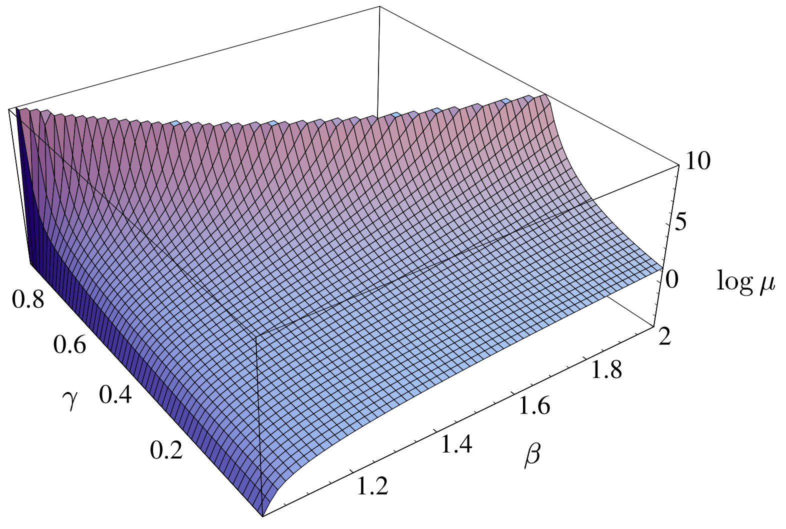

Figure 2: Logarithm of the maximal μ in the region $1 < \beta \le 2$ and $0 \le \gamma < 1$.

given a connected graph $G = (V, E)$ and an accuracy parameter $\varepsilon \in (0, 1]$, outputs a configuration $\sigma \in \Omega(G)$ according to a measure μ_G which satisfies $\mathrm{d}_{\mathrm{TV}}(\mu_G, \pi_G) \le \varepsilon$ where

$$\pi_G(\sigma) = \frac{Z_{\{\sigma\}}(G)}{Z(G)}$$

and d_{TV} denotes total variation distance. Using the method of Jerrum, Valiant and Vazirani [8], it is straightforward to show that the FPAS can be turned into an FPRAS for (β, γ, μ)-Partition within the region. Thus, we obtain the following theorem.

Theorem 6 *There is an FPRAS for* (β, γ, μ)-Partition *when the fixed parameters* β*,* γ *and* μ *are in the regions defined by* $\beta\gamma < 1$ *and either*

1. $\beta > 1$ *and* $\mu < C(\ln \beta) e^{D(C(\ln \beta))}$*, or*

2. $\gamma > 1$ *and* $1/\mu < C(\ln \gamma) e^{D(C(\ln \gamma))}$*,*

where functions C *and* D *are defined in Equations (22) and (24).*

Corollary 7 *There is an FPRAS for* (β, γ, μ)-Partition *when the fixed parameters* β*,* γ *and* μ *are in the regions defined by* $\beta\gamma < 1$ *and either*

1. $\beta > 1$ *and* $\mu < e\beta \ln \beta$*, or*

2. $\gamma > 1$ *and* $1/\mu < e\gamma \ln \gamma$*.*

The special case of $\mu = 1$ *gives a numerical lower bound of about 1.3211 for* β *and* γ *respectively.*

To aid the reader, we provide some details showing how to turn the FPAS into an FPRAS within Region 1. Region 2 can be handled by a similar argument, exploiting the blue–green symmetry as in Section 5.1.

Suppose that the edge set E of G is $E = \{e_1, \ldots, e_m\}$ and let G_i be the graph $(V, \{e_1, \ldots, e_i\})$. Thus, our job is to approximate

$$Z(G) = Z(G_m) = \frac{Z(G_m)}{Z(G_{m-1})} \frac{Z(G_{m-1})}{Z(G_{m-2})} \cdots \frac{Z(G_1)}{Z(G_0)} Z(G_0).$$

The quantity $Z(G_0)$ is easy to compute, so the main task is to estimate the quantity

$$\varrho_i = \frac{Z(G_i)}{Z(G_{i-1})}.$$

Suppose that e_i is the edge (x_i, y_i). For spin s, let Ω_{i-1}^s denote the set of all configurations in $\Omega(G_{i-1})$ in which x_i and y_i are assigned spin s. Let Ω_{i-1}^* denote the set of all configurations in $\Omega(G_{i-1})$ in which x_i and y_i are assigned different spins. Then

$$\varrho_i = \frac{\beta \sum_{\sigma \in \Omega_{i-1}^b} Z_{\{\sigma\}}(G_{i-1}) + \gamma \sum_{\sigma \in \Omega_{i-1}^g} Z_{\{\sigma\}}(G_{i-1}) + \sum_{\sigma \in \Omega_{i-1}^*} Z_{\{\sigma\}}(G_{i-1})}{\sum_{\sigma \in \Omega(G_{i-1})} Z_{\{\sigma\}}(G_{i-1})}. \tag{26}$$

We need a method for estimating ϱ_i. Consider the following experiment (which makes sense, since $\gamma, 1 \le \beta$): Sample a configuration $\sigma \in \Omega(G_{i-1})$ with weight $\pi_{G_{i-1}}(\sigma)$. If $\sigma \in \Omega_{i-1}^b$, output "yes". If $\sigma \in \Omega_{i-1}^g$, output "yes" with probability γ/β and "no" otherwise. If $\sigma \in \Omega_{i-1}^*$, output "yes" with probability $1/\beta$ and "no" otherwise. From (26), we deduce that the probability that the algorithm outputs "yes" is ϱ_i/β. Thus, we can accurately estimate ϱ_i by applying the experiment to several outputs of the FPAS. (We need several outputs because the FPAS has measure $\mu_{G_{i-1}}$ not $\pi_{G_{i-1}}$.) Also, we can conclude that $\varrho_i \le \beta$. It is known [8] that as long as $\varrho_i \ge 1/\mathrm{poly}(n)$, where $n = |V|$, then the required number of samples is only polynomial in n and ε^{-1}, so we get an FPRAS. Details can be found in the proof of Proposition 3.4 of [6].

We conclude this section by showing that, in the region of interest, $\varrho_i \ge 1/(1+\mu)$. To start with, we observe that

$$Z_{\Omega_{i-1}^g}(G_{i-1}) \le \mu Z_{\Omega_{i-1}^*}(G_{i-1}). \tag{27}$$

To see (27), consider the injection which maps every $\sigma \in \Omega_{i-1}^g$ to $\sigma' \in \Omega_{i-1}^*$ by colouring y_i blue. Since $\gamma \le 1$ and $\beta \ge 1$, $\mu Z_{\{\sigma'\}}(G_{i-1}) \ge Z_{\{\sigma\}}(G_{i-1})$. Thus, from (26), since $\beta \ge 1$,

$$\begin{aligned}
\varrho_i &\ge \frac{Z_{\Omega_{i-1}^b}(G_{i-1}) + Z_{\Omega_{i-1}^*}(G_{i-1})}{Z_{\Omega_{i-1}^g}(G_{i-1}) + Z_{\Omega_{i-1}^b}(G_{i-1}) + Z_{\Omega_{i-1}^*}(G_{i-1})} \\
&\ge \frac{Z_{\Omega_{i-1}^b}(G_{i-1}) + Z_{\Omega_{i-1}^*}(G_{i-1})}{Z_{\Omega_{i-1}^b}(G_{i-1}) + (1+\mu) Z_{\Omega_{i-1}^*}(G_{i-1})} \\
&\ge \frac{1}{1+\mu}.
\end{aligned}$$

6 An additional intractable region

In the previous section, we saw that the tractable region extends beyond that defined by the hyperbola $\beta\gamma \ge 1$. The main result of this section is that the intractable region extends beyond the square defined by $\beta < 1$ and $\gamma < 1$. Specifically, we show:

Theorem 8 *Let η be a sufficiently small constant ($\eta = 10^{-7}$ will do), and suppose that $1 \le \beta \le 1 + \eta$, $0 \le \gamma \le \eta$ and $\frac{1}{2} \le \mu \le 2$. Then there is no FPRAS for (β, γ, μ)-*Partition *unless* NP = RP.

By the same symmetry considerations exploited in Section 5, Theorem 8 remains true with the roles of β and γ reversed.

The region covered by Theorem 8 is admittedly small. The estimates in the proof could undoubtedly be tightened with a view to expanding the range of parameter values covered by the theorem. However, since our main aim is to uncover some intractable region of positive volume lying outside the square, we shall instead aim to keep the technical complications to a minimum.

Our starting point is an inapproximability result concerning independent sets in bounded degree graphs. It is well known that that it is NP-hard to determine the size of a maximum independent set in a graph of maximum degree 4. A result of Berman and Karpinski [1, Thm 1(iv)] tells us more:

Proposition 9 *For any $\varepsilon > 0$, it is NP-hard to determine the size of a maximum independent set in a graph G to within ratio $\frac{73}{74} + \varepsilon$, even when G is restricted to have maximum degree 4.*

(By "determining the size… within ratio ρ" we mean computing a number $\hat{k}$ such that $\rho k \leq \hat{k} \leq k$, where k is the size of a maximum independent set in G.) The possibility of establishing results such as Proposition 9 has been opened up by the theory of "polynomially checkable proofs" (PCPs).

Proof of Theorem 8. Our proof strategy is to design a reduction that takes a graph $G = (V, E)$ of maximum degree 4 and forms a graph G' with the following informal property: The partition function $Z(G')$ of the new graph G' determines the size of the largest independent set in G within ratio 0.99. Since such a tight performance guarantee is precluded by Proposition 9, this will be enough to establish the result.

We now describe the construction of G' from G. For every vertex u of G let A_u be a distinct set of size r, where r is a constant to be determined later. Then define

$$V' = \bigcup_{u \in V} A_u$$

and

$$E' = \bigcup_{\{u,v\} \in E} A_u \times A_v.$$

Presently, we shall argue that the partition function of G' is bounded below and above as follows:

$$Z(G') \geq (1+\mu)^{rk} \tag{28}$$

and

$$Z(G') \leq \sum_{i=0}^{k} \binom{n}{i} (1+\mu)^{ri} \sum_{j=0}^{r(n-i)} \binom{r(n-i)}{j} \mu^j (1+\eta)^{r^2 m - j} \eta^j, \tag{29}$$

where $n = |V|$, $m = |E|$ and k is the size of a maximum independent set in G. It transpires that when the parameters β, γ and μ satisfy the conditions of the theorem, these inequalities locate $\ln Z(G')$ rather accurately: see inequality (32). Thus a good estimate for $Z(G')$ provides a good estimate for k.

The lower bound (28) is the easier of the two to justify. Let I be any independent set in G of size k. The lower bound (28) comes from considering just the configurations which assign blue to all vertices in $\bigcup_{u \in V \setminus I} A_u$. Since $\beta \geq 1$ and there are no green-green edges, every such configuration σ contributes at least μ^j to the partition function, where j is the number of green vertices in σ. Since the green vertices are freely selected from a set of size rk, inequality (28) is now immediate.

The upper bound (29) is not much more difficult, if viewed in the right way. A *base* for a configuration σ is an independent set I in G such that:

- for every $u \in I$ the block A_u contains at least one green vertex;

- for every block A_u containing a green vertex, either u is in I or u is adjacent to a vertex in I.

Every configuration has at least one base, since we may take I to be any maximal independent set within the subgraph of G induced by the vertex set

$$\{u \in V : A_u \text{ contains at least one green vertex}\}.$$

It is convenient to think of the term "base" as applying both to the vertex set I in G and the vertex set $\bigcup_{u \in I} A_u$ in G'.

For each base, we shall estimate the total weight of configurations with that base, and then sum over all possible bases. This will lead to overcounting, since each configuration has many bases in general. This is fine, as we are shooting for an upper bound. The key observation is that, in any configuration with base I, each green vertex lying outside the base is adjacent to some green vertex lying inside. Thus the number of green-green edges is at least as large as the number of green non-base vertices.

With these considerations in mind, the formula in (29) may be read left-to-right as follows: (i) i ranges over the possible sizes of a base, k being an upper bound since any base is an independent set in G; (ii) $\binom{n}{i}$ is a bound on the number of bases of size i; (iii) $(1+\mu)^{ri}$ counts colourings of the base-vertices; (iv) j is the number of green vertices among the non-base vertices, ranging from $j = 0$ (no green vertices) to $j = r(n-i)$ (all green); (v) μ^j comes from the j green vertices; and finally (vi) $(1+\eta)^{r^2m-j}\eta^j$ is an upper bound on edge weights, since there must be at least j green-green edges.

Next, we simplify the upper bound (29) by approximating the two sums:

$$\begin{aligned} Z(G') &\le \sum_{i=0}^{k} \binom{n}{i} (1+\mu)^{ri} (1+\eta)^{r^2m} \sum_{j=0}^{rn} \binom{rn}{j} \mu^j \eta^j \\ &= (1+\eta)^{r^2m} \sum_{i=0}^{k} \binom{n}{i} (1+\mu)^{ri} (1+\mu\eta)^{rn} \\ &\le (1+\eta)^{r^2m} (1+\mu\eta)^{rn} 2^n \sum_{i=0}^{k} (1+\mu)^{ri} \\ &\le (1+\eta)^{r^2m} (1+\mu\eta)^{rn} 2^{n+1} (1+\mu)^{rk}, \end{aligned} \tag{30}$$

where the final inequality assumes (as will certainly be the case) that $(1+\mu)^r \ge 2$.

Taking logarithms of (28) and (30) we may sandwich $\ln Z(G')$ as follows:

$$rk \ln(1+\mu) \le \ln Z(G') \le (1 + c_1 + c_2 + c_3)\, rk \ln(1+\mu), \tag{31}$$

where

$$c_1 = \frac{r^2 m \ln(1+\eta)}{rk \ln(1+\mu)}, \quad c_2 = \frac{rn \ln(1+\mu\eta)}{rk \ln(1+\mu)}, \quad \text{and} \quad c_3 = \frac{(n+1)\ln 2}{rk \ln(1+\mu)}.$$

Now $m \le 2n$ since G has maximum degree 4. Furthermore, $k \ge \frac{1}{4}n$ since G is 4-colourable by Brooks' Theorem. (The largest colour class is an independent set). So assuming $r = 1000$,

$\frac{1}{2} \leq \mu \leq 2$ and $0 \leq \eta \leq 10^{-7}$, we have the following bounds on c_1, c_2 and c_3:

$$\begin{aligned}
c_1 &\leq \frac{2rn\eta}{k\ln(1+\mu)} \leq \frac{8r\eta}{\ln(1+\mu)} \leq 0.002 \\
c_2 &\leq \frac{n\mu\eta}{k\ln(1+\mu)} \leq \frac{4\mu\eta}{\ln(1+\mu)} \leq 0.001 \\
c_3 &\leq \frac{(1+o(1))4\ln 2}{r\ln(1+\mu)} \leq 0.007,
\end{aligned}$$

for sufficiently large n.

Thus from (31),

$$rk\ln(1+\mu) \leq \ln Z(G') \leq 1.01\, rk\ln(1+\mu),$$

and hence

$$0.99\, k \leq \frac{0.99\ln Z(G')}{r\ln(1+\mu)} \leq k. \tag{32}$$

Finally, suppose β, γ and μ are as stated in the theorem, and that there is an FPRAS for (β,γ,μ)-Partition. Then we would be able to compute an approximation L to $\ln Z(G')$ within additive error 1 (say), in polynomial time, with high probability. But then $0.99\, L/1000\ln(1+\mu)$ (rounded to the nearest integer) would approximate the size of a maximum independent set in G to within ratio uniformly better than $\frac{73}{74}$. By Proposition 9, this entails RP = NP. □

References

[1] P. Berman and M. Karpinski, On some tighter inapproximability results (extended abstract), *Proceedings of the 26th EATCS International Colloquium on Automata, Languages and Programming* (ICALP), (Springer-Verlag, 1999) 200–209.

[2] R. Bubley and M. Dyer, Path coupling: A technique for proving rapid mixing in Markov chains, *Proceedings of the 38th IEEE Annual Symposium on Foundations of Computer Science*, (IEEE, Los Alamitos, 1997) 223–231.

[3] M. Dyer and C. Greenhill, *Random walks on combinatorial objects*, in J. D. Lamb and D. A. Preece, eds., *Surveys in Combinatorics 1999*, vol. 267 of *London Mathematical Society Lecture Note Series*, (Cambridge University Press, Cambridge, 1999) 101–136.

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

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

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

[7] M. Jerrum and A. Sinclair, Polynomial-time approximation algorithms for the Ising model, *SIAM Journal on Computing* **22** (1993) 1087–1116.

[8] 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.

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