

Algorithms for #BIS-hard problems on expander graphs

Jenssen, Matthew; Perkins, Will; Keevash, Peter

Document Version
Peer reviewed version

Citation for published version (Harvard):
Jenssen, M, Perkins, W & Keevash, P 2020, 'Algorithms for #BIS-hard problems on expander graphs', *SIAM Journal on Computing*.

[Link to publication on Research at Birmingham portal](#)

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Algorithms for #BIS-hard problems on expander graphs

Matthew Jenssen* Peter Keevash† Will Perkins‡

March 23, 2020

Abstract

We give an FPTAS and an efficient sampling algorithm for the high-fugacity hard-core model on bounded-degree bipartite expander graphs and the low-temperature ferromagnetic Potts model on bounded-degree expander graphs. The results apply, for example, to random (bipartite) Δ -regular graphs, for which no efficient algorithms were known for these problems (with the exception of the Ising model) in the non-uniqueness regime of the infinite Δ -regular tree. We also find efficient counting and sampling algorithms for proper q -colorings of random Δ -regular bipartite graphs when q is sufficiently small as a function of Δ .

1 Introduction

There are two natural computational problems associated to a statistical physics spin model on a graph: the *approximate counting* problem of approximating the partition function of the model and the *sampling* problem of obtaining a random spin configuration approximately distributed according to the model.

A prominent example is the hard-core model of weighted independent sets. For a graph G and fugacity parameter $\lambda > 0$, the hard-core model is the probability distribution $\mu_{G,\lambda}$ on the collection $\mathcal{I}(G)$ of independent sets of G given by

$$\mu_{G,\lambda}(I) = \frac{\lambda^{|I|}}{Z_G(\lambda)}$$

where

$$Z_G(\lambda) = \sum_{I \in \mathcal{I}(G)} \lambda^{|I|}$$

*University of Birmingham, m.jenssen@bham.ac.uk.

†University of Oxford, keevash@maths.ox.ac.uk.

Research supported in part by ERC Consolidator Grant 647678.

‡University of Illinois at Chicago, math@willperkins.org.

Research supported in part by EPSRC grant EP/P009913/1 and NSF Career award DMS-1847451.

is the hard-core partition function (also known as the independence polynomial in graph theory).

A *fully polynomial-time approximation scheme* (FPTAS) is an algorithm that for every $\varepsilon > 0$ outputs an ε -relative approximation to $Z_G(\lambda)$ (that is, a number \hat{Z} so that $e^{-\varepsilon} \hat{Z} \leq Z_G(\lambda) \leq e^\varepsilon \hat{Z}$) and runs in time polynomial in $|V(G)|$ and $1/\varepsilon$. A *polynomial-time sampling algorithm* is a randomized algorithm that for every $\varepsilon > 0$ runs in time polynomial in $|V(G)|$ and $1/\varepsilon$ and outputs an independent set I with distribution μ_{alg} so that $\|\mu_{G,\lambda} - \mu_{\text{alg}}\|_{TV} < \varepsilon$ (this terminology is not completely standard, but it defines the natural sampling equivalent of the running-time guarantees of an FPTAS).

The computational complexity of the approximate counting and sampling problems for the hard-core model is well understood for bounded-degree graphs. For graphs of maximum degree at most Δ , when $\lambda < \lambda_c(\Delta) = \frac{(\Delta-1)^{\Delta-1}}{(\Delta-2)^\Delta}$, there is an FPTAS and a polynomial-time sampling algorithm due to Weitz [47] (see also the recent [2]); whereas when $\lambda > \lambda_c(\Delta)$ both computational problems are hard: there is no polynomial-time algorithm unless $\text{NP}=\text{RP}$ [43, 44, 21]. The value $\lambda_c(\Delta)$ is the *uniqueness threshold* of the hard-core model on the infinite Δ -regular tree [31].

On the other hand, if we restrict ourselves to bipartite graphs, then classifying the computational complexity of these tasks are open problems. The class $\#\text{BIS}$ is the class of problems polynomial-time equivalent to approximating the number of independent sets of a bipartite graph [18], and many interesting approximate counting and sampling problems have been shown to be $\#\text{BIS}$ -hard [24, 14, 22] (that is, at least as hard as approximating the number of independent sets in a bipartite graph). In particular, Cai, Galanis, Goldberg, Guo, Jerrum, Štefankovič, and Vigoda [11] showed that for all $\Delta \geq 3$ and all $\lambda > \lambda_c(\Delta)$, it is $\#\text{BIS}$ -hard to approximate the hard-core partition function at fugacity λ on a bipartite graph of maximum degree Δ . Resolving the complexity of $\#\text{BIS}$ is a major open problem in the field of approximate counting.

One direction for partial progress on any intermediate complexity class is to find subclasses of instances for which the problem is tractable (e.g. results showing that the Unique Games problem is tractable on expander graphs [3, 35]). For $\#\text{BIS}$, we would like to find subclasses of bipartite graphs on which we can efficiently approximate the number of independent sets or the hard-core partition function. One example is the algorithm of Liu and Lu [34] which works when $\lambda < \lambda_c(\Delta)$ and one side of the bipartition has maximum degree Δ but the other side of the bipartition is allowed unbounded degree.

Recently, Helmuth, Perkins, and Regts [27] gave efficient algorithms for the hard-core model at high fugacity on the torus $(\mathbb{Z}/n\mathbb{Z})^d$ and subsets of the lattice \mathbb{Z}^d with certain boundary conditions. The algorithms are based on contour models from Pirogov-Sinai theory [40] along with the cluster expansion and Barvinok's polynomial interpolation method [4, 5].

Here we extend this approach and use abstract polymer models and the cluster expansion to give efficient approximate counting and sampling algorithms for all bounded-degree, bipartite expander graphs at sufficiently high fugacity. We say that a bipartite graph $G = (\mathcal{O}, \mathcal{E}, E)$ is a *bipartite α -expander* if $|\partial S| \geq (1 + \alpha)|S|$ for all $S \subseteq \mathcal{O}$ with $|S| \leq |\mathcal{O}|/2$ and all $S \subseteq \mathcal{E}$ with $|S| \leq |\mathcal{E}|/2$. Here ∂S denotes the *vertex boundary* of S ,

the set of all vertices in $V(G)\setminus S$ with a neighbor in S .

Theorem 1. *There exists an absolute constant C such that for every $\alpha > 0$, $\Delta \geq 3$, and any $\lambda > C\Delta^{4/\alpha}$, there exists an FPTAS and a polynomial-time sampling algorithm for the hard-core model at fugacity λ on bipartite α -expander graphs of maximum degree Δ .*

In general the running time of our approximate counting algorithm in Theorem 1 (and the theorems that follow) is $(n/\varepsilon)^{O(\log \Delta)}$ for an ε -relative approximation to $Z_G(\lambda)$ for an n -vertex graph G . Since Δ is fixed this running time is polynomial in n and $1/\varepsilon$, but one might hope to improve the dependence on Δ . In fact, the larger λ is and the larger α is, the faster the running time, and for large enough values of the parameters the algorithm runs in almost linear time. We discuss this more below. See also [15] in which the authors find faster sampling algorithms in similar settings using abstract polymer models and Markov chains.

We can extend our methods to obtain efficient counting and sampling algorithms for the hard-core model on random regular bipartite graphs for much smaller values of λ , all the way down to $\lambda = \Omega\left(\frac{\log^2 \Delta}{\Delta}\right)$.

Let $\mathcal{G}^{\text{bip}}(n, \Delta)$ be the set of all Δ -regular bipartite graphs on n vertices (where n is even), and let $\mathbf{G}_{n, \Delta}^{\text{bip}}$ be a uniformly chosen graph from $\mathcal{G}^{\text{bip}}(n, \Delta)$. We say that a property holds for *almost every* Δ -regular bipartite graph if the property holds with probability $\rightarrow 1$ as $n \rightarrow \infty$ for $\mathbf{G}_{n, \Delta}^{\text{bip}}$.

Theorem 2. *There exists a constant Δ_0 so that for every $\Delta \geq \Delta_0$ and all $\lambda > \frac{50 \log^2 \Delta}{\Delta}$, there is an FPTAS and an efficient sampling algorithm for the hard-core model at fugacity λ on almost every Δ -regular bipartite graph. The running time of the approximate counting algorithm is $(n/\varepsilon)^{1+O(\log^2 \Delta/\Delta)}$.*

In particular, by setting $\lambda = 1$ this gives an FPTAS for counting the total number of independent sets in random Δ -regular bipartite graphs for large enough Δ .

Remark 1. *The guarantees of the algorithm of Theorem 2 (and of those below in Corollary 4 and Theorem 5) are of the following form: the counting algorithm will output an estimate of the partition function for every Δ -regular bipartite graph. With probability $1 - o(1)$ over the choice of graph $\mathbf{G}_{n, \Delta}^{\text{bip}}$ the estimate will satisfy the guarantees of an FPTAS. In fact what is needed is that the graph satisfy some expansion conditions which hold with probability $1 - o(1)$ for $\mathbf{G}_{n, \Delta}^{\text{bip}}$. It would be desirable to find algorithms with a stronger type of guarantee: an algorithm could give no answer on a vanishing fraction of graphs but if it gives an answer, the answer must satisfy the stated guarantees (e.g. the sampling algorithm in [9]). We find such an algorithm for the Potts model on random regular graphs below (Corollary 4) by certifying expansion using spectral methods. We believe such an algorithm can be found for the hard-core model as well (also using spectral results, e.g. [46, 10]), but certifying the specific expansion properties used in Lemma 24 below seems more difficult, and so we do not pursue it here.*

To the best of our knowledge, no efficient counting or sampling algorithms for random regular bipartite graphs were previously known for any $\lambda > \lambda_c(\Delta)$.

Moreover, for any $\delta > 0$, taking $\Delta = \Omega(\log^2(1/\delta)/\delta)$ gives an approximate counting algorithm running in time $O((n/\varepsilon)^{1+\delta})$.

1.1 The Potts Model

Given a graph G and $q \in \mathbb{N}$, let $\Omega = [q] = \{1, \dots, q\}$ and let $\Omega^{V(G)}$ be the set of all colorings $\omega : V(G) \rightarrow \Omega$. Given $\omega \in \Omega^{V(G)}$ let $m(G, \omega)$ denote the number of monochromatic edges induced by the coloring ω , that is

$$m(G, \omega) := \sum_{\{i,j\} \in E(G)} \delta_{\omega(i), \omega(j)}$$

where δ is the Kronecker delta function. The q -color Potts model on G at inverse temperature β is the probability distribution on $\Omega^{V(G)}$ defined by

$$\mu_{G,q,\beta}(\omega) = \frac{e^{\beta \cdot m(G,\omega)}}{Z_{G,q}(\beta)}, \quad \omega \in \Omega^{V(G)}$$

where

$$Z_{G,q}(\beta) := \sum_{\omega \in \Omega^{V(G)}} e^{\beta \cdot m(G,\omega)}$$

is the Potts model partition function. When $\beta > 0$ the model is *ferromagnetic* (monochromatic edges preferred) and when $\beta < 0$ the model is *antiferromagnetic* (bichromatic edges preferred).

Goldberg and Jerrum [24] showed that approximating the partition function of the ferromagnetic Potts model for $q \geq 3$ is #BIS-hard. Later, Galanis, Štefankovič, Vigoda, and Yang [22] refined this result by showing that this task is #BIS-hard on graphs of maximum degree Δ when $\beta > \beta_o(q, \Delta)$, the order/disorder threshold of the infinite Δ -regular tree (see [22] for a precise definition of $\beta_o(q, \Delta)$; in particular, $\beta_o(q, \Delta) > \beta_c(q, \Delta)$, the uniqueness threshold on the infinite Δ -regular tree).

Our next theorem gives efficient counting and sampling algorithms for the ferromagnetic Potts model at low enough temperatures on expander graphs. We say that a graph G is an α -*expander* if $|\partial_e(S)| \geq \alpha|S|$ for all subsets $S \subseteq V(G)$ with $|S| \leq |V(G)|/2$. Here $\partial_e(S)$ denotes the *edge boundary* of S , the set of edges of G with one endpoint in S and the other in $V(G) \setminus S$.

Theorem 3. *For all $\alpha > 0$, $\Delta \geq 3$, $q \geq 2$ and $\beta \geq \frac{4+2\log(q\Delta)}{\alpha}$, there is an FPTAS and polynomial-time sampling algorithm for the q -color ferromagnetic Potts model at inverse temperature β on all α -expander graphs of maximum degree Δ .*

Our algorithms apply to the Potts model on the random Δ -regular graph as well. Let $\mathcal{G}(n, \Delta)$ be the set of all Δ -regular graphs on n vertices, and let $\mathbf{G}_{n,\Delta}$ be a uniformly chosen graph from $\mathcal{G}(n, \Delta)$ (as long as this set is non-empty). We say that a property holds for *almost every* Δ -regular graph if the property holds with probability $\rightarrow 1$ as $n \rightarrow \infty$ for $\mathbf{G}_{n,\Delta}$.

Corollary 4. *There is an absolute constant $C > 0$ so that for every $\Delta \geq 3$, $q \geq 2$, and all $\beta > \frac{C \log(q\Delta)}{\Delta}$, there is an FPTAS and polynomial-time sampling algorithm for the ferromagnetic Potts model at inverse temperature β on almost every Δ -regular graph. Moreover, there is a polynomial-time algorithm to certify conditions on a Δ -regular graph G that suffice for the guarantees of the FPTAS and sampling algorithm, and these conditions hold with probability $1 - o(1)$ for $\mathbf{G}_{n,\Delta}$.*

Again to the best of our knowledge no efficient counting or sampling algorithms were known previously for the $q \geq 3$ Potts model on random regular graphs for β above the uniqueness threshold $\beta_c(q, \Delta)$ of the infinite Δ -regular tree. The lower bound on β needed in Corollary 4 is necessarily above the ordering threshold $\beta_o(q, \Delta)$ on the infinite tree since our proof involves a phase coexistence result (Lemma 12). We note however that our lower bound on β is within a constant factor of β_o for Δ fixed and q large (since $\beta_o(q, \Delta) = (1 + o_q(1)) \frac{2 \log q}{\Delta}$ as $q \rightarrow \infty$ [22]).

1.2 Counting and sampling proper colorings on bipartite graphs

For $q \geq 3$, let $\mathcal{X}_q(G)$ be the set of all proper q -colorings of G . Let $Z_G(q) = |\mathcal{X}_q(G)|$ and let $\mu_{G,q}$ be the uniform distribution on $\mathcal{X}_q(G)$. In particular, if G is bipartite then $\mathcal{X}_q(G)$ is guaranteed to be non-empty. We would like to approximate $Z_G(q)$ and sample from $\mu_{G,q}$. Galanis, Stefankovic, Vigoda, and Yang [22] show that these problems are #BIS-hard on bipartite graphs of maximum degree Δ when $q \leq \Delta/(2 \log \Delta)$. Here we show that under the stricter condition that $q \leq c\sqrt{\Delta}/\log^2 \Delta$ (or equivalently $\Delta \geq Cq^2 \log^2 q$), these problems are tractable on random regular bipartite graphs. This solves Conjecture 1 from the earlier extended abstract of this paper [29].

Theorem 5. *There is an absolute constant $C > 0$ so that for all $q \geq 3$, and for $\Delta \geq Cq^2 \log^2 q$, there is an FPTAS and polynomial-time sampling algorithm for proper colorings on almost every Δ -regular bipartite graph.*

Adapting techniques from [27, 29], Liao, Lin, Lu, and Mao [33] have recently and independently proved similar results to Theorems 2 and 5 with slightly stronger conditions needed on λ and Δ respectively to obtain efficient algorithms.

1.3 Discussion

We take $\Delta \geq 3$ in all of our theorems since computing the relevant partition functions exactly on paths and cycles takes linear time.

We also note that for the $q = 2$ case of the Potts model (the Ising model), efficient algorithms are known for *all* graphs and all temperatures: the approximate counting algorithm of Jerrum and Sinclair [30], turned into a sampling algorithm via self-reducibility by Randall and Wilson [41]; see also the recent proof of Guo and Jerrum [26] showing polynomial-time mixing of the random-cluster dynamics.

The approximate counting and sampling problems for the hard-core and Potts models on random graphs have received considerable attention, with positive algorithmic results in the low-fugacity, high-temperature uniqueness regimes of the infinite Δ -regular

tree, and some negative algorithmic results, in the form of torpid mixing of certain Markov chains, in the high-fugacity and low-temperature regimes. Our results are novel in providing positive algorithmic results in the high-fugacity and low-temperature regimes.

In the low fugacity regime with $\lambda < \lambda_c(\Delta)$, Weitz's algorithm applies to $\mathbf{G}_{n,\Delta}^{\text{bip}}$. Efthymiou, Hayes, Štefanković, Vigoda, and Yin [19] have also shown that the Glauber dynamics have mixing time $O(n \log n)$ on $\mathbf{G}_{n,\Delta}^{\text{bip}}$ for $\lambda < (1 - \varepsilon(\Delta))\lambda_c(\Delta)$ for some $\varepsilon(\Delta) \rightarrow 0$ as $\Delta \rightarrow \infty$. For $\lambda > \lambda_c(\Delta)$, the Glauber dynamics for the hard-core model on $\mathbf{G}_{n,\Delta}^{\text{bip}}$ are known to mix slowly [36], and perhaps Theorem 2 can be improved to work for all $\lambda > \lambda_c(\Delta)$, though this would likely require new ideas.

For the Potts model, a natural conjecture for the optimal bound on β for the particular polymer-based algorithm we use here is the order/disorder transition point $\beta_o(q, \Delta) = \log \frac{q-2}{(q-1)^{1-2/\Delta}-1}$. Galanis, Štefankovič, Vigoda, and Yang [22] have shown that the Swendsen-Wang dynamics mix slowly at $\beta_o(q, \Delta)$ on the random Δ -regular graph (for q large enough). In fact their analysis shows that the approximation lemmas we use below fail for $\mathbf{G}_{n,\Delta}$ and $\beta \leq \beta_o(q, \Delta)$. The bound we obtain in Corollary 4 is at worst a factor of order $\log \Delta$ away from this natural barrier and matches up to a constant factor when q and Δ are polynomially related.

In the high-temperature regime, Blanca, Galanis, Goldberg, Štefanković, Vigoda, and Yang [9] have recently given an efficient algorithm to obtain an n^{-c} -approximate sample (that is, a sample within total variation distance n^{-c} for n -vertex graphs for some constant $c > 0$) from the Potts model (ferromagnetic and anti-ferromagnetic) on $\mathbf{G}_{n,\Delta}$ when the parameters lie in the uniqueness regime for the infinite Δ -regular tree.

While efficient counting and sampling algorithms for these problems on random regular graphs were previously only known for the uniqueness regime, the probabilistic properties of these models are well understood at all fugacities and temperatures. Sly and Sun [44] showed that the limiting free energy (the normalized log partition function) of *any* sequence of locally tree-like bipartite graphs converges to the replica symmetric solution predicted by the cavity method from statistical physics. This result applies in particular to the hard-core model on random bipartite Δ -regular graphs. Dembo, Montanari, Sly, and Sun [16] then showed that the limiting free energy of the ferromagnetic Potts model on a sequence of graphs converging locally to the infinite Δ -regular tree is given by the replica symmetric solution from the cavity method.

For other recent work on algorithms for the Potts model and low temperature and the hard-core model at high fugacity, see [6, 15, 12]. For further algorithmic applications of abstract polymer models and the cluster expansion see [13].

1.4 Proof ideas

Our main technical contribution is to show that the hard-core, Potts and coloring models are well approximated by mixtures of *polymer models* with convergent cluster expansions in the relevant range of parameters. These polymer models each represent deviations from one of a collection of *ground states*. For example, in the case of the Potts model, the

ground states are the monochromatic configurations, while in the case of the hard-core model the ground states are the two collections of independent sets with no occupied even or odd vertices respectively.

The main steps in the proofs of Theorems 1 and 3 are as follows.

1. First we show that the partition function of the relevant model on (bipartite) expanders is dominated by configurations that are ‘close’ to one of the ground states.
2. For each ground state we define a polymer model representing deviations from the given state. We show that α -expansion implies a strong upper bound on the polymer weights, which allows us to verify the Kotecký-Preiss condition for the convergence of the cluster expansion [32].
3. This last step allows us to implement a version of the approximate counting algorithm from [27], based on truncating the cluster expansion. This algorithm is inspired by Barvinok’s method of truncating the Taylor series of the log partition function, but here we can work directly with the cluster expansion and avoid any use of complex analysis.
4. The sampling algorithm is based on a form of self-reducibility for abstract polymer models.

In Section 2 we define polymer models and the cluster expansion, we state the Kotecký-Preiss condition, and describe the counting algorithm. We prove our results for the Potts model in Section 3, for the hard-core model in Section 4, and for proper colorings in Section 5. We conclude with some discussion and open problems in Section 6.

2 Polymer models

2.1 Abstract polymer models

We define polymer models in sufficient generality for the purposes of this paper. A more general treatment can be found, for example, in [25, 32].

A *polymer* γ is a connected subgraph of G . A polymer model consists of a set of allowed polymers $\mathcal{C}(G)$ along with a complex-valued weight function w_γ for each polymer. We measure the size of a polymer by $|\gamma|$, the number of vertices of γ .

We say two polymers γ, γ' are *compatible* if $d(\gamma, \gamma') > 1$ and *incompatible* otherwise, where $d(\cdot, \cdot)$ is the graph distance. Let $\mathcal{G}(G)$ be the collection of all finite subsets (including the empty set) of $\mathcal{C}(G)$ consisting of mutually compatible polymers.

We can then define the polymer model partition function

$$\Xi(G) := \sum_{\Gamma \in \mathcal{G}(G)} \prod_{\gamma \in \Gamma} w_\gamma.$$

The prototypical example of a polymer model is the low-fugacity hard-core model on a graph G : the set of polymers $\mathcal{C}(G)$ is simply the set of vertices $V(G)$. The collection of sets of mutually compatible polymers $\mathcal{G}(G)$ is exactly $\mathcal{I}(G)$, the collection of independent sets of G . If we set the weight function of every polymer to be $w_\gamma = \lambda$, then the abstract polymer partition function $\Xi(G)$ is exactly the hard-core partition function $Z_G(\lambda)$.

2.2 Convergent cluster expansions

A detailed probabilistic understanding of a polymer model can be obtained by showing that the *cluster expansion* of its log partition function converges.

For a multiset of polymers Γ , the *incompatibility graph* $H(\Gamma)$ has one vertex for each polymer in the multiset (with multiplicity) with an edge between two vertices corresponding to polymers γ, γ' that are incompatible. A *cluster* is an ordered list of polymers from $\mathcal{C}(G)$ (with repetitions allowed) whose incompatibility graph is connected. The size of a cluster is $|\Gamma| = \sum_{\gamma \in \Gamma} |\gamma|$. Let $\mathcal{G}^{\text{clust}}(G)$ be the collection of all clusters. The cluster expansion is then the (formal) power series in the variables w_γ , $\gamma \in \mathcal{C}(G)$,

$$\log \Xi(G) = \sum_{\Gamma \in \mathcal{G}^{\text{clust}}(G)} \phi(\Gamma) \prod_{\gamma \in \Gamma} w_\gamma,$$

where $\phi(\Gamma)$ is the *Ursell function* of the incompatibility graph $H = H(\Gamma)$ defined by

$$\phi(H) = \frac{1}{|V(H)|!} \sum_{\substack{A \subseteq E(H) \\ \text{spanning, connected}}} (-1)^{|A|}.$$

(The Ursell function $\phi(H)$ is an evaluation of the Tutte polynomial of H). In fact the cluster expansion is simply the multivariate Taylor series for $\log \Xi(G)$ in the variables w_γ , as observed by Dobrushin [17]. See also Scott and Sokal [42] for a derivation of the cluster expansion and much more.

A sufficient condition for the convergence of the cluster expansion is given by the following specialization of a result of Kotecký and Preiss.

Theorem 6 ([32]). *Fix a function $g : \mathcal{C}(G) \rightarrow [0, \infty)$, and extend g to clusters by defining*

$$g(\Gamma) = \sum_{\gamma \in \Gamma} g(\gamma).$$

Suppose that for all $\gamma \in \mathcal{C}(G)$

$$\sum_{\gamma' : d(\gamma', \gamma) \leq 1} |w_{\gamma'}| e^{|\gamma'| + g(\gamma')} \leq |\gamma|. \quad (1)$$

Then the cluster expansion converges absolutely and, moreover, for every $v \in V(G)$,

$$\sum_{\substack{\Gamma \in \mathcal{G}^{\text{clust}}(G) \\ \Gamma \ni v}} \left| \phi(\Gamma) \prod_{\gamma \in \Gamma} w_\gamma \right| e^{g(\Gamma)} \leq 1, \quad (2)$$

where we write $v \in \Gamma$ if there exists $\gamma \in \Gamma$ so that $v \in \gamma$.

We remark that the theorem appearing in [32] is more general. For instance it allows a general choice of function $f : \mathcal{C}(G) \rightarrow [0, \infty)$ in lieu of the function $|\cdot|$ appearing in (1). A more careful choice of the function f might lead to better dependencies between the parameters in our results, although as far as we can tell only constant factor improvements would be gained.

2.3 Algorithms

We can use Theorem 6 to approximate the partition function. For a given function $g : \mathcal{C}(G) \rightarrow [0, \infty)$ for which (1) holds, define the truncated cluster expansion

$$T_m(G) = \sum_{\substack{\Gamma \in \mathcal{G}^{\text{clust}}(G) \\ g(\Gamma) < m}} \phi(\Gamma) \prod_{\gamma \in \Gamma} w_\gamma.$$

We refer to g as a *decay function* for the polymer model. Under the Kotecký-Preiss condition, truncating the cluster expansion with large enough m gives a good approximation to the log partition function.

Fact 7. *If condition (1) holds and $m \geq \log(|V(G)|/\varepsilon)$, then*

$$|\log \Xi(G) - T_m(G)| \leq \varepsilon.$$

This fact follows from summing (2) over all $v \in V(G)$.

Under mild conditions on the polymer model, and under the non-trivial condition of zero-freeness of the partition function in a disc in the complex plane, Helmuth, Perkins, and Regts gave an efficient approximation algorithm for the partition function [27, Theorem 6]. The following theorem is an adaptation of that theorem.

Theorem 8. *Fix Δ and let \mathfrak{G} be some class of graphs of maximum degree at most Δ . Suppose the following hold for a given polymer model with decay function $g(\cdot)$:*

- (i) *There exists constants $c_1, c_2 > 0$ such that given a connected subgraph γ , determining whether $\gamma \in \mathcal{C}(G)$ and then computing w_γ and $g(\gamma)$ can be done in time $O(|\gamma|^{c_1} e^{c_2|\gamma|})$.*
- (ii) *There exists $\rho(\Delta) = \rho > 0$ so that for every $G \in \mathfrak{G}$ and every $\gamma \in \mathcal{C}(G)$, $g(\gamma) \geq \rho|\gamma|$.*
- (iii) *The Kotecký-Preiss condition (1) holds with the given function $g(\cdot)$.*

Then there is an FPTAS for $\Xi(G)$ for $G \in \mathfrak{G}$. The running time of this algorithm is $n \cdot (n/\varepsilon)^{O((\log \Delta + c_2)/\rho)}$.

The algorithm is as follows. Given $\varepsilon > 0$, and a graph G on n vertices, let $m = \log(n/\varepsilon)$ and:

1. Enumerate all clusters $\Gamma \in \mathcal{G}^{\text{clust}}(G)$ with $|\Gamma| < m/\rho$; call the list of such clusters \mathcal{L} .
2. For each cluster $\Gamma \in \mathcal{L}$ compute $\phi(\Gamma)$, $\prod_{\gamma \in \Gamma} w_\gamma$, and $g(\Gamma)$.
3. Compute $T_m(G)$ by summing:

$$T_m(G) = \sum_{\substack{\Gamma \in \mathcal{L} \\ g(\Gamma) < m}} \phi(\Gamma) \prod_{\gamma \in \Gamma} w_\gamma.$$

4. Output $\exp(T_m(G))$.

Proof of Theorem 8. Fact 7 along with condition (iii) of the theorem tells us that $\exp(T_m(G))$ is an ε -relative approximation to $\Xi(G)$ and so we just need to show that steps 1–3 can be done in time $n \cdot (n/\varepsilon)^{O((\log \Delta + c_2)/\rho)}$.

By condition (ii) of the theorem, we need to list clusters Γ with $|\Gamma| < m/\rho$, compute the Ursell function for each cluster, and compute the product of polymer weights in each cluster. Using the algorithm of [27, Theorem 6], we can list all clusters of size at most m/ρ in time $n\Delta^{O(m/\rho)}$ and compute their associated Ursell functions (this result relies on the algorithms of [38] for enumerating connected subgraphs in a bounded-degree graph and [8] for evaluating the Tutte polynomial, and hence the Ursell function, of a graph in time exponential in its number of vertices). Computing $\prod w_\gamma$ for a single cluster can be done in time $(m/\rho)^{c_1} e^{O(c_2 m/\rho)}$ using condition (i). Putting this together yields an algorithm with running time $n \cdot (n/\varepsilon)^{O((\log \Delta + c_2)/\rho)}$. \square

We can also sample efficiently from a polymer model satisfying condition (1). We define a probability measure ν_G on $\mathcal{G}(G)$:

$$\nu_G(\Gamma) = \frac{\prod_{\gamma \in \Gamma} w_\gamma}{\Xi(G)}, \quad \Gamma \in \mathcal{G}(G).$$

Following [27, Theorem 10], the approximate counting algorithm also provides a polynomial-time sampling algorithm for ν_G .

Theorem 9 ([27], Theorem 10). *Under the conditions of Theorem 8, there is a polynomial-time sampling algorithm for ν_G for all $G \in \mathfrak{G}$.*

Theorem 9 is an application of self-reducibility to the abstract polymer model. We briefly summarize the algorithm here for completeness and since we will use some of its properties later.

Let $n = |V(G)|$, and given $\varepsilon > 0$, let $m = \log(2n/\varepsilon)$. Let $\mathcal{C}^{(m)}(G)$ be the set of all polymers γ so that $g(\gamma) < m$, and let $\mathcal{G}^{(m)}(G)$ be the collection of pairwise compatible sets of polymers from $\mathcal{C}^{(m)}(G)$. By Fact 7, with probability at least $1 - \varepsilon/2$, Γ drawn from ν_G belongs to $\mathcal{G}^{(m)}(G)$, and so it suffices to sample from $\mathcal{G}^{(m)}(G)$. The algorithm builds such a configuration one polymer at a time.

Fix an arbitrary ordering of the vertices of G , v_1, \dots, v_n . Let $\Gamma_0 = \emptyset$, and let $\mathcal{C}_1 = \mathcal{C}^{(m)}(G)$. Then for $j = 1, \dots, n$ do the following:

1. List all polymers $\gamma \in \mathcal{C}_j$ so that $v_j \in \gamma$; call this list \mathcal{L}_j .
2. For each $\gamma \in \mathcal{L}_j$, approximate the probability $\Pr[\gamma \in \Gamma | \Gamma_{j-1} \subseteq \Gamma]$.
3. Choose at most one polymer from \mathcal{L}_j according to these approximate probabilities and add it to Γ_{j-1} to form Γ_j ; if no polymer is chosen let $\Gamma_j = \Gamma_{j-1}$.
4. Let $\mathcal{C}_{j+1} = \mathcal{C}_j \setminus \mathcal{L}_j$.

Output $\Gamma = \Gamma_n$.

The algorithm is simply self-reducibility applied to the polymer model. Its correctness relies on two facts: 1) at most one polymer from \mathcal{L}_j can appear in Γ since any two polymers in \mathcal{L}_j are incompatible with each other 2) for each polymer $\gamma \in \mathcal{C}^{(m)}(G)$ there is exactly one step at which it can potentially be added to Γ : the step j corresponding to the first vertex v_j in γ according to the arbitrary vertex ordering. The efficiency of the sampling algorithm relies on the fact that \mathcal{L}_j can be formed efficiently and that the conditional probabilities can be computed efficiently. The latter is done by approximating partition functions of polymer models with various subsets of polymers. Crucially, if condition (1) holds for $\mathcal{C}(G)$ then it holds for any subset $\mathcal{C}' \subseteq \mathcal{C}(G)$.

While the algorithm of Theorem 8 does not use complex analysis or refer the zeros of a partition function, it is closely related to and inspired by the approach of Barvinok [5] for a broad range of approximation problems that involves truncating the Taylor series of the log partition function and using the absence of zeros of the partition function in the complex plane to deduce convergence. The cluster expansion itself is the multivariate Taylor series of $\log \Xi$ and the fact that it is supported on clusters which are connected objects is related to the method for efficient computation of the coefficients of graph polynomials in [37]. The algorithm of [27] for approximating a polymer model partition function uses the cluster expansion in an indirect way: the cluster expansion is used to prove that the partition function does not vanish in a disc around 0 in the complex plane; then Barvinok's algorithm of truncating the univariate Taylor series is applied. Here we use the cluster expansion directly, truncating it to approximate the log partition function, and using the guarantees of (2) to bound the approximation error. The technical reason zero-freeness and the Taylor series were needed in [27] is that the weight functions of the more complicated contour models used there are ratios of partition functions, instead of the explicit polymer weights used here.

To prove Theorems 1, 2, 3, and 5 we will take \mathfrak{G} to be the class of (bipartite) α -expander graphs of maximum degree Δ . We will show that the hard-core and Potts partition functions, at sufficiently high fugacity and low temperature respectively, and

the number of proper q -colorings for q sufficiently small, can be approximated well by sums of partition functions of abstract polymer models. This involves showing that the expansion condition implies that almost all the weight of the given partition function comes from configurations close to one of a small number of ground states.

We then verify in each case conditions (i) of Theorem 8, which is straightforward.

Conditions (ii) and (iii), choosing the function g and showing the Kotecký-Preiss condition holds, are non-trivial, and it is here that the dependence of the temperature and fugacity on the degree and expansion is determined.

We note that although we are working with low-temperature models, we are able to use the polymer model formulation instead of the more complex *contour model* formulation of Pirogov-Sinai theory used for the algorithms on \mathbb{Z}^d in [27]. The reason polymer models suffice is that the strong expansion condition allows us to express our partition functions in terms of deviations from the ground states directly and not in the recursive fashion of a contour model.

3 The Potts model

3.1 Approximation by a polymer model

In this section we show that the Potts model partition function of an expander graph can be well-approximated by the partition function of a certain polymer model. Recall that our measure of approximation is the following.

Definition 10. *Let Z be a real number. We call \hat{Z} an ε -relative approximation to Z if*

$$e^{-\varepsilon} \hat{Z} \leq Z \leq e^{\varepsilon} \hat{Z}.$$

Given a graph G and a set $S \subseteq V(G)$, recall that we let ∂S denote the set of vertices in S^c (the complement of S in $V(G)$) adjacent to a vertex in S and we let $\partial_e(S)$ denote the set of edges in G with one endpoint in S and the other in S^c . We let $G[S]$ denote the subgraph of G induced by the vertex set S . Let $S^+ = S \cup \partial S$ and let $\nabla(S) = E(G[S]) \cup \partial_e(S)$, the set of edges of G that are incident to a vertex in S . Recall our notion of expansion.

Definition 11. *A graph G is an α -expander if $|\partial_e(S)| \geq \alpha|S|$ for all subsets $S \subseteq V(G)$ with $|S| \leq |V(G)|/2$.*

Let $\mathfrak{G}(\alpha, \Delta)$ denote the class of all α -expander graphs with maximum degree at most Δ . For the remainder of this section we fix a graph $G \in \mathfrak{G}(\alpha, \Delta)$ on n vertices.

First we first show that the main contribution to the Potts model partition function of G comes from colorings where one color dominates. To make this precise we make a few definitions. First let $\Omega = [q]$ and let Ω^n be the set of all (not necessarily proper) colorings $\omega : V(G) \rightarrow [q]$, and recall that $m(G, \omega)$ denotes the number of monochromatic edges of G induced by ω . The Potts model partition function is then

$$Z_{G,q}(\beta) := \sum_{\omega \in \Omega^n} e^{\beta \cdot m(G, \omega)}.$$

For $j \in [q]$, let

$$\Omega_j^n = \{\omega \in \Omega^n : |\omega^{-1}(\{j\})| > n/2\},$$

let

$$Z_G^j(\beta) := \sum_{\omega \in \Omega_j^n} e^{\beta \cdot m(G, \omega)},$$

and let

$$Z_G^*(\beta) = \sum_{j=1}^q Z_G^j(\beta).$$

Lemma 12. For $\beta > 2 \log(eq)/\alpha$, $Z_G^*(\beta)$ is an e^{-n} -relative approximation to $Z_{G,q}(\beta)$.

Proof. Let $\Omega_*^n = \bigcup_{i=1}^q \Omega_i^n$ and note that this is a disjoint union. Let $\omega \notin \Omega_*^n$, then for each $j \in [q]$ we have

$$|\omega^{-1}(\{j\})| \leq \frac{n}{2}.$$

Letting $S_j = \omega^{-1}(\{j\})$ it follows that $|\partial_e(S_j)| \geq \alpha|S_j|$. The set S_j consists of all vertices of G with the color j and so every edge lying between S_j and S_j^c is bichromatic. Summing over colors j we thus have at least

$$\frac{1}{2} \sum_{j=1}^q |\partial_e(S_j)| \geq \frac{1}{2} \sum_{j=1}^q \alpha|S_j| = \frac{\alpha n}{2}$$

bichromatic edges and so

$$m(G, \omega) \leq e(G) - \frac{\alpha n}{2}.$$

Using the crude bound $|\Omega^n \setminus \Omega_*^n| \leq |\Omega^n| = q^n$ we then have

$$Z_{G,q}(\beta) - Z_G^*(\beta) = \sum_{\omega \notin \Omega_*^n} e^{\beta \cdot m(G, \omega)} \leq q^n e^{\beta(e(G) - \alpha n/2)},$$

and so

$$\left| 1 - \frac{Z_G^*(\beta)}{Z_{G,q}(\beta)} \right| \leq \frac{q^n e^{\beta(e(G) - \alpha n/2)}}{Z_{G,q}(\beta)} \leq q^{n-1} e^{-\beta \alpha n/2} \leq e^{-n},$$

where for the second inequality we use the trivial lower bound $Z_{G,q}(\beta) > qe^{\beta e(G)}$. \square

This allows us to focus on approximating $Z_G^*(\beta)$. Henceforth, let us fix $r \in [q]$. We will refer to r as the color ‘red’. By symmetry

$$Z_G^*(\beta) = q \cdot Z_G^r(\beta),$$

and so we may in fact focus on approximating $Z_G^r(\beta)$.

We now define a polymer model, whose partition function will serve as an approximation to $Z_G^r(\beta)$.

Define a polymer to be a set $\gamma \subseteq V(G)$ such that $G[\gamma]$ is connected and $|\gamma| \leq n/2$. Following the set-up in Section 2, we say that two polymers γ_1, γ_2 are compatible if

$d(\gamma_1, \gamma_2) > 1$. We let $\mathcal{C} = \mathcal{C}(G)$ denote the set of all polymers of G and let $\mathcal{G} = \mathcal{G}(G)$ denote the family of all sets of mutually compatible polymers from \mathcal{C} . To each polymer $\gamma \in \mathcal{C}$, we assign the weight

$$w_\gamma := e^{-\beta|\nabla(\gamma)|} Z_{G[\gamma], q-1}(\beta).$$

This defines a polymer model with partition function

$$\Xi(G) = \sum_{\Gamma \in \mathcal{G}} \prod_{\gamma \in \Gamma} w_\gamma.$$

We remark that since \mathcal{G} includes the empty set we have $\Xi(G) \geq 1$.

Lemma 13. For $\beta > 2 \log(eq)/\alpha$

$$\tilde{Z}_G(\beta) := e^{\beta \cdot e(G)} \cdot \Xi(G)$$

is an e^{-n} -relative approximation to $Z_G^r(\beta)$.

Proof. For $S \subseteq V(G)$, let $\Omega^n(S)$ denote the set of colorings $\omega \in \Omega^n$ such that $\omega(S) \subseteq [q] \setminus \{r\}$ and $\omega(S^c) = \{r\}$. Note that for $\omega \in \Omega^n(S)$, we have

$$m(G, \omega) = e(G) - |\nabla(S)| + m(G[S], \omega). \quad (3)$$

Note also that if $\{\gamma_1, \dots, \gamma_k\}$ are the connected components of $G[S]$ then

$$e^{-\beta|\nabla(S)|} Z_{G[S], q-1}(\beta) = \prod_{i \in [k]} e^{-\beta|\nabla(\gamma_i)|} Z_{G[\gamma_i], q-1}(\beta). \quad (4)$$

We call a subset $S \subseteq V(G)$ *small* if $|S| \leq n/2$ and *large* otherwise. We call S *sparse* if each of the connected components of $G[S]$ is small. Note that there is a one to one correspondence between sparse subsets of $V(G)$ and elements of \mathcal{G} . We thus have by (3) and (4) that

$$\begin{aligned} \tilde{Z}_G(\beta) &= e^{\beta \cdot e(G)} \sum_{\Gamma \in \mathcal{G}} \prod_{\gamma \in \Gamma} e^{-\beta|\nabla(\gamma)|} Z_{G[\gamma], q-1}(\beta), \\ &= e^{\beta \cdot e(G)} \sum_{S \text{ sparse}} e^{-\beta|\nabla(S)|} Z_{G[S], q-1}(\beta), \\ &= \sum_{S \text{ sparse}} \sum_{\omega \in \Omega^n(S)} e^{\beta \cdot m(G, \omega)}. \end{aligned}$$

On the other hand

$$Z_G^r(\beta) = \sum_{S \text{ small}} \sum_{\omega \in \Omega^n(S)} e^{\beta \cdot m(G, \omega)},$$

and so

$$\tilde{Z}_G(\beta) - Z_G^r(\beta) = \sum_{\substack{S \text{ sparse,} \\ \text{large}}} \sum_{\omega \in \Omega^n(S)} e^{\beta \cdot m(G, \omega)}.$$

Observe that if S is large and sparse then $|\partial_e(S)| \geq \alpha|S| \geq \alpha n/2$, and so by (3) $m(G, \omega) \leq e(G) - \alpha n/2$ for all $\omega \in \Omega^n(S)$. It follows that

$$\left| 1 - \frac{\tilde{Z}_G(\beta)}{Z_G^r(\beta)} \right| \leq e^{-\beta \cdot e(G)} \left(\tilde{Z}_G(\beta) - Z_G^r(\beta) \right) \leq q^n e^{-\beta \alpha n/2} \leq e^{-n},$$

where we have the lower bound $\tilde{Z}_G(\beta) \geq e^{\beta \cdot e(G)}$ (since $\Xi(G) \geq 1$) and the crude upper bound of q^n on the number of ω such that $\omega \in \Omega^n(S)$ for some large and sparse set S . \square

Our aim is to apply Theorem 8 to this polymer model in order to obtain an FPTAS for $\Xi(G)$ when β is sufficiently large which, by Lemmas 12 and 13, will furnish us with an FPTAS for $Z_{G,q}(\beta)$.

To this end we aim to verify conditions (i)–(iii) of Theorem 8 for our polymer model. Verifying condition (i) is essentially immediate. Given γ , determining whether $\gamma \in \mathcal{C}$ amounts to checking whether $G[\gamma]$ is connected and of size at most $n/2$. This can be done in $O(|\gamma|)$ time by a depth-first search algorithm. Computing the weight function w_γ can be done in time $O(\Delta|\gamma|(q-1)^{|\gamma|})$ by calculating $m(G[\gamma], \omega)$ for all assignments ω of $(q-1)$ colors to γ .

We now turn our attention to verifying conditions (ii) and (iii) for an appropriate choice of function g .

3.2 Verifying the Kotecký-Preiss condition

We choose

$$g(\gamma) = |\gamma|.$$

Condition (ii) of Theorem 8 holds trivially with $\rho = 1$. It remains to show the Kotecký-Preiss condition holds. That is,

$$\sum_{\gamma': d(\gamma', \gamma) \leq 1} w_{\gamma'} e^{2|\gamma'|} \leq |\gamma|$$

for all $\gamma \in \mathcal{C}(G)$. We begin by bounding w_γ :

$$\begin{aligned} w_\gamma &= e^{-\beta|\nabla(\gamma)|} Z_{G[\gamma], q-1}(\beta) \\ &\leq e^{-\beta|\partial_e \gamma|} (q-1)^{|\gamma|} \\ &\leq e^{-\beta\alpha|\gamma|} (q-1)^{|\gamma|}. \end{aligned}$$

For the first inequality we used that $|\nabla(\gamma)| = |\partial_e \gamma| + e(G[\gamma])$ and that $Z_{G[\gamma], q-1}(\beta) \leq (q-1)^{|\gamma|} e^{\beta \cdot e(G[\gamma])}$. For the second inequality we used that G is an α -expander. It thus suffices to show that

$$\sum_{\gamma': d(\gamma', \gamma) \leq 1} e^{(2-\alpha\beta+\log(q-1))|\gamma'|} \leq |\gamma|.$$

If we could show that for each $v \in V(G)$

$$\sum_{\gamma': \gamma' \ni v} e^{(2-\alpha\beta+\log(q-1))|\gamma'|} \leq \frac{1}{\Delta+1}, \quad (5)$$

then by summing this inequality over all $v \in \gamma^+ (= \gamma \cup \partial\gamma)$ and noting that $|\gamma^+| \leq (\Delta+1)|\gamma|$, we would be done.

In order to establish (5) we borrow the following lemma:

Lemma 14 ([23], Lemma 2.1). *In a graph of maximum degree at most Δ , the number of connected, induced subgraphs of order t containing a fixed vertex v is at most $(e\Delta)^t$.*

It follows that the number of polymers γ of size t that contain a given vertex v is bounded by $(e\Delta)^t$, and so

$$\sum_{\gamma': \gamma' \ni v} e^{(2-\alpha\beta+\log(q-1))|\gamma'|} \leq \sum_{t=1}^{\infty} \left((q-1)\Delta \cdot e^{(3-\alpha\beta)} \right)^t \leq \frac{1}{\Delta+1},$$

for $\beta \geq \frac{4+2\log(q\Delta)}{\alpha}$.

This verifies condition (iii) of Theorem 8 and so Theorem 8 gives an FPTAS for $\Xi(G)$ for all $\beta \geq \frac{4+2\log(q\Delta)}{\alpha}$.

3.3 Proof of Theorem 3

We consider two cases separately. If $\varepsilon \leq e^{-n/2}$, then we proceed by brute force, calculating $m(G, \omega)$ for each of the q^n possible colorings of G . In this way we can calculate the partition function $Z_{G,q}(\beta)$ exactly in time $O(n\Delta q^n)$ and therefore in time polynomial in $1/\varepsilon$. Similarly we can obtain an exact sample from $\mu_{G,q,\beta}$ by brute force in time polynomial in $1/\varepsilon$.

Now we assume $\varepsilon > e^{-n/2}$. Using Theorem 8, we can obtain Z_{alg} , an $\varepsilon/2$ -relative approximation to $e^{\beta e(G)} \cdot \Xi(G)$ in time polynomial in n and $1/\varepsilon$ (here we use the fact that for a polymer γ , w_γ can be computed in time $O(|\gamma| \Delta q^{|\gamma|})$). By Lemma 13, $e^{\beta e(G)} \cdot \Xi(G)$ is an e^{-n} -relative approximation to $Z_G^r(\beta)$ and so $q e^{\beta e(G)} \cdot \Xi(G)$ is an e^{-n} -relative approximation to $Z_G^*(\beta)$. By Lemma 12, it follows that $q e^{\beta e(G)} \cdot \Xi(G)$ is an $\varepsilon/2$ -relative approximation to $Z_{G,q}(\beta)$ and so $q Z_{\text{alg}}$ is an ε -relative approximation to $Z_{G,q}(\beta)$ as required. Lemmas 12 and 13 apply since $\frac{4+2\log(q\Delta)}{\alpha} > 2\log(eq)/\alpha$.

For the approximate sampling algorithm, we will apply Theorem 9.

Consider the following distribution $\hat{\mu}$ on Ω^n . Choose $r \in [q]$ uniformly at random and then sample Γ from the measure $\nu_{G,\beta}$ on $\mathcal{G}(G)$ defined by

$$\nu_{G,\beta}(\Gamma) = \frac{\prod_{\gamma \in \Gamma} w_\gamma}{\Xi(G)}.$$

Then convert Γ into a q -coloring $\omega \in \Omega^n$ as follows: for each $\gamma \in \Gamma$, choose a $(q-1)$ coloring (excluding color r) ω_γ according to $\mu_{G[\gamma],q-1,\beta}$. Then set $\omega(v) = \omega_\gamma(v)$ if $v \in \gamma$ for some $\gamma \in \Gamma$, and $\omega(v) = r$ otherwise. The resulting distribution of ω is $\hat{\mu}$.

By Lemmas 12 and 13 and the condition on ε ,

$$\|\hat{\mu} - \mu_{G,q,\beta}\|_{TV} \leq \varepsilon/2.$$

Thus to complete the proof of Theorem 3 it suffices to obtain an $\varepsilon/2$ -approximate sample from $\nu_{G,\beta}^r$ in time polynomial in n and $1/\varepsilon$, and to do this we appeal to Theorem 9. Crucially, because of the truncation involved in the polymer sampling algorithm of Theorem 9, we have $|\gamma| = O(\log(n/\varepsilon))$ for each $\gamma \in \Gamma$ and so for each γ a sample from $\mu_{G[\gamma],q-1,\beta}$ can be obtained by brute force in time polynomial in n and $1/\varepsilon$. □

3.4 Proof of Corollary 4

In order to prove Corollary 4 we require a very brief review of spectral graph theory. For a graph G on n vertices we let

$$h(G) := \min_{\{S:|S|\leq n/2\}} \frac{|\partial_e S|}{|S|}.$$

In other words, $h(G)$ is the largest value of α for which G is an α -expander. The exact determination of $h(G)$, given G , is known to be **coNP**-complete. However, as shown by Alon [1], one can efficiently approximate the expansion properties of a graph using its *spectrum*. For a graph G on n vertices, we let $\lambda_1 \geq \dots \geq \lambda_n$ denote the eigenvalues of the adjacency matrix of G (that is, the $n \times n$ matrix A , with rows and columns indexed by $V(G)$, where $A_{ij} = 1$ if $\{i, j\} \in E(G)$ and $A_{ij} = 0$ otherwise).

The following lemma, one of Cheeger's inequalities (see e.g. [28]), bounds $h(G)$ in terms of the second largest eigenvalue of G when G is regular.

Lemma 15. *If G is a Δ -regular graph, then*

$$h(G) \geq \frac{\Delta - \lambda_2}{2}.$$

We will also use the following celebrated result of Friedman on the spectral gap of the random regular graph. For a graph G we let $\lambda(G) = \max(|\lambda_2|, |\lambda_n|)$.

Theorem 16 (Friedman [20]). *For every fixed $\varepsilon > 0$, almost every Δ -regular graph G satisfies*

$$\lambda(G) \leq 2\sqrt{\Delta - 1} + \varepsilon. \tag{6}$$

Proof of Corollary 4. Fix $\varepsilon = 1/100$. Since the eigenvalues of a graph G can be calculated in polynomial time, there is a polynomial time algorithm to determine whether a graph G satisfies (6). If G satisfies (6) then by Lemma 15, G is a $\Delta/40$ -expander (since for $\Delta \geq 3$, $\Delta - 2\sqrt{\Delta - 1} - 1/100 \geq \Delta/20$). By Theorem 3 it follows that if $\beta > 200 \frac{\log(q\Delta)}{\Delta}$, there is an FPTAS and polynomial-time sampling algorithm for the ferromagnetic Potts model at inverse temperature β on G . We conclude by noting that almost every Δ -regular graph satisfies (6) by Theorem 16. \square

4 The hard-core model

4.1 Approximation by a polymer model

In this section we prove Theorem 1 following the same strategy as in the previous section. First we approximate the hard-core partition function by a sum of polymer model partition functions, and then we verify the conditions of Theorem 8 for these models.

We let $G = (\mathcal{O}, \mathcal{E}, E)$ denote a bipartite graph with partition classes \mathcal{O} , \mathcal{E} and edge set E . We will refer to vertices of \mathcal{O} and \mathcal{E} as ‘odd’ and ‘even’ vertices respectively.

Recall our notion of expansion for a bipartite graph.

Definition 17. For $\alpha > 0$, a bipartite graph $G = (\mathcal{O}, \mathcal{E}, E)$ is a bipartite α -expander if $|\partial S| \geq (1 + \alpha)|S|$ for all $S \subseteq \mathcal{O}$ with $|S| \leq |\mathcal{O}|/2$ and all $S \subseteq \mathcal{E}$ with $|S| \leq |\mathcal{E}|/2$.

We remark that for regular bipartite graphs the expansion property of Definition 17 also follows from a spectral gap condition. Indeed if G is a Δ -regular connected bipartite graph whose second largest eigenvalue is λ_2 , then a classical result of Tanner [46] implies that G is a bipartite α -expander where $\alpha = (\Delta^2 - \lambda_2^2)/(\Delta^2 + \lambda_2^2)$.

Let $\mathfrak{G}^{\text{bip}}(\alpha, \Delta)$ denote the class of all bipartite α -expander graphs with maximum degree at most Δ . From this point on, let us fix a graph $G \in \mathfrak{G}^{\text{bip}}(\alpha, \Delta)$ on n vertices with partition classes \mathcal{O}, \mathcal{E} . Recall that $\mathcal{I}(G)$ denotes the family of all independent sets in G .

Let us call a set $S \subseteq \mathcal{O}$ *small* if $|S| \leq |\mathcal{O}|/2$. We define small subsets of \mathcal{E} similarly.

Lemma 18. Let $I \in \mathcal{I}(G)$, then at least one of the sets $I \cap \mathcal{O}, I \cap \mathcal{E}$ is small.

Proof. Without loss of generality let $|\mathcal{O}| \geq |\mathcal{E}|$ and suppose that $|I \cap \mathcal{O}| > |\mathcal{O}|/2$. Since G is a bipartite α -expander, by considering a subset of $I \cap \mathcal{O}$ of size $|\mathcal{O}|/2$, we have $|\partial(I \cap \mathcal{O})| \geq (1 + \alpha)|\mathcal{O}|/2$ and so

$$|I \cap \mathcal{E}| < |\mathcal{E}| - (1 + \alpha)|\mathcal{O}|/2 < |\mathcal{E}|/2.$$

\square

We now define two distinct polymer models whose partition functions we will use to approximate $Z_G(\lambda)$.

First let us introduce some notation and terminology. We let G^k denote the k th power of the graph G , that is the graph on vertex set $V(G)$ where two vertices are adjacent if and only if they are at distance at most k in G .

Definition 19. We say a subset $S \subseteq V(G)$ is G^k -connected if the induced subgraph $G^k[S]$ is connected. We call the connected components of $G^k[S]$, the G^k -connected components of S .

Throughout this section we will be concerned only with G^2 -connected sets. Note that a set $S \subseteq \mathcal{O}, \mathcal{E}$ is G^2 -connected if and only if $G[S^+]$ is connected.

A *polymer* is any *small* G^2 -connected set γ which lies entirely in \mathcal{E} or entirely in \mathcal{O} ; in the first case we say it is an even polymer and in the second an odd polymer. The size of a polymer γ , $|\gamma|$, is again the number of vertices in γ . We define a compatibility relation on the set of even (odd) polymers, with γ_1, γ_2 compatible if $d_G(\gamma_1, \gamma_2) > 2$ (that is if $\gamma_1 \cup \gamma_2$ is not G^2 -connected); otherwise γ_1 and γ_2 are incompatible. Note that this is consistent with the definition of a polymer model in Section 2: a polymer is a connected subgraph in G^2 and two polymers γ, γ' are compatible if their graph distance in G^2 is greater than 1.

For each polymer γ we assign a weight

$$w_\gamma := \frac{\lambda^{|\gamma|}}{(1 + \lambda)^{|\partial\gamma|}}.$$

We let $\mathcal{C}^\mathcal{E} = \mathcal{C}^\mathcal{E}(G)$ denote the set of all even polymers of G and let $\mathcal{G}^\mathcal{E} = \mathcal{G}^\mathcal{E}(G)$ denote the family of all sets of mutually compatible polymers from $\mathcal{C}^\mathcal{E}$. We define $\mathcal{C}^\mathcal{O}, \mathcal{G}^\mathcal{O}$ similarly.

The set of polymers $\mathcal{C}^\mathcal{E} (\mathcal{C}^\mathcal{O})$ constitutes a polymer model with partition function $\Xi^\mathcal{E}(G) (\Xi^\mathcal{O}(G))$. That is,

$$\Xi^\mathcal{E}(G) = \sum_{\Gamma \in \mathcal{G}^\mathcal{E}} \prod_{\gamma \in \Gamma} w_\gamma \quad \text{and} \quad \Xi^\mathcal{O}(G) = \sum_{\Gamma \in \mathcal{G}^\mathcal{O}} \prod_{\gamma \in \Gamma} w_\gamma.$$

Note that $(1 + \lambda)^{|\mathcal{O}|} \Xi^\mathcal{E}(G)$ represents the contribution to $Z_G(\lambda)$ from independent sets that are dominated by odd occupied vertices, and vice-versa.

We now show that a certain linear combination of the partition functions $\Xi^\mathcal{E}(G), \Xi^\mathcal{O}(G)$ serves as a good approximation to $Z_G(\lambda)$.

Lemma 20. For $\lambda > e^{11/\alpha}$, the polynomial

$$\tilde{Z}_G(\lambda) = (1 + \lambda)^{|\mathcal{O}|} \Xi^\mathcal{E}(G) + (1 + \lambda)^{|\mathcal{E}|} \Xi^\mathcal{O}(G)$$

is an e^{-n} -relative approximation to $Z_G(\lambda)$.

Proof. Let us call a set $A \subseteq \mathcal{E}, \mathcal{O}$ *sparse* if its G^2 -connected components are all small. Note that $\mathcal{G}^\mathcal{E} (\mathcal{G}^\mathcal{O})$ is in one to one correspondence with sparse subsets of $\mathcal{E} (\mathcal{O})$ and so

$$\begin{aligned} \tilde{Z}_G(\lambda) &= (1 + \lambda)^{|\mathcal{O}|} \sum_{\Gamma \in \mathcal{G}^\mathcal{E}} \prod_{\gamma \in \Gamma} \frac{\lambda^{|\gamma|}}{(1 + \lambda)^{|\partial\gamma|}} + (1 + \lambda)^{|\mathcal{E}|} \sum_{\Gamma \in \mathcal{G}^\mathcal{O}} \prod_{\gamma \in \Gamma} \frac{\lambda^{|\gamma|}}{(1 + \lambda)^{|\partial\gamma|}} \\ &= (1 + \lambda)^{|\mathcal{O}|} \sum_{\text{sparse } A \subseteq \mathcal{E}} \frac{\lambda^{|A|}}{(1 + \lambda)^{|\partial A|}} + (1 + \lambda)^{|\mathcal{E}|} \sum_{\text{sparse } A \subseteq \mathcal{O}} \frac{\lambda^{|A|}}{(1 + \lambda)^{|\partial A|}}. \end{aligned} \quad (7)$$

Consider an independent $I \in \mathcal{I}(G)$ for which $I \cap \mathcal{E} = A$ where A is some fixed subset of \mathcal{E} . The possible intersections $I \cap \mathcal{O}$ are then precisely the subsets of $\mathcal{O} \setminus \partial A$. The contribution to $Z_G(\lambda)$ from independent sets I such that $I \cap \mathcal{E} = A$ is therefore $\lambda^{|A|}(1 + \lambda)^{|\mathcal{O}| - |\partial A|}$. Similarly the contribution to $Z_G(\lambda)$ from independent sets I such that $I \cap \mathcal{O} = B$ is $\lambda^{|B|}(1 + \lambda)^{|\mathcal{E}| - |\partial B|}$. Let us call an independent set I sparse if both $I \cap \mathcal{E}, I \cap \mathcal{O}$ are sparse. Since by Lemma 18, for any independent set I , at least one of $I \cap \mathcal{E}, I \cap \mathcal{O}$ is small (and therefore sparse), the sums in (7) contain the contribution to $Z_G(\lambda)$ from all $I \in \mathcal{I}(G)$ and double count the contribution from precisely the sparse independent sets i.e.

$$\tilde{Z}_G(\lambda) = Z_G(\lambda) + \sum_{I \text{ sparse}} \lambda^{|I|}.$$

Let I be a sparse independent set. Since $I \cap \mathcal{E}, I \cap \mathcal{O}$ are composed of small G^2 -connected components it follows that $|\partial(I \cap \mathcal{E})| > (1 + \alpha)|I \cap \mathcal{E}|$ and $|\partial(I \cap \mathcal{O})| > (1 + \alpha)|I \cap \mathcal{O}|$. Since each vertex in $\partial(I \cap \mathcal{E})$ and $\partial(I \cap \mathcal{O})$ must be unoccupied, it follows that

$$|I \cap \mathcal{E}| < |\mathcal{E}| - (1 + \alpha)|I \cap \mathcal{O}|$$

and

$$|I \cap \mathcal{O}| < |\mathcal{O}| - (1 + \alpha)|I \cap \mathcal{E}|.$$

By summing these two inequalities we conclude that

$$|I| < \frac{n}{2 + \alpha}.$$

It follows that

$$\sum_{I \text{ sparse}} \lambda^{|I|} \leq 2^n \lambda^{\frac{n}{2 + \alpha}}.$$

Using the crude bound $Z_G(\lambda) \geq \lambda^{n/2}$ we have

$$\left| 1 - \frac{\tilde{Z}_G(\lambda)}{Z_G(\lambda)} \right| < 2^n \frac{\lambda^{\frac{n}{2 + \alpha}}}{\lambda^{n/2}} = 2^n \lambda^{-\frac{\alpha}{4 + 2\alpha}n} \leq e^{-n}.$$

□

In order to approximate $Z_G(\lambda)$ we approximate $\Xi^{\mathcal{E}}(G)$ and $\Xi^{\mathcal{O}}(G)$ separately. We focus on approximating $\Xi^{\mathcal{E}}(G)$, noting that the approximation algorithm for $\Xi^{\mathcal{O}}(G)$ will be identical up to a change of notation.

We will verify conditions (i)–(iii) of Theorem 8 for the polymer model defined on $\mathcal{C}^{\mathcal{E}}$ in order to obtain an FPTAS for $\Xi^{\mathcal{E}}(G)$. Verifying condition (i) is essentially immediate. Given $\gamma \subseteq V(G)$, determining whether $\gamma \in \mathcal{C}^{\mathcal{E}}$ amounts to checking whether $\gamma \subseteq \mathcal{E}$ and whether $G[\gamma^+]$ is connected. This can be done in $O(\Delta|\gamma|)$ time by a depth-first search algorithm. Computing $|\gamma|, |\partial\gamma|$ and thus w_γ can also clearly be done in $O(\Delta|\gamma|)$ time.

We now turn our attention verifying conditions (ii) and (iii) for an appropriate choice of function g .

4.2 Verifying the Kotecký-Preiss condition

We choose

$$g(\gamma) = |\gamma|.$$

Condition (ii) of Theorem 8 holds trivially with $\rho = 1$. It remains to show the Kotecký-Preiss condition holds. That is,

$$\sum_{\gamma': d(\gamma', \gamma) \leq 1} w_{\gamma'} e^{g(\gamma') + |\gamma'|} \leq |\gamma|$$

for all $\gamma \in \mathcal{C}^{\mathcal{E}}$. Note that since G is a bipartite α -expander, for $\gamma \in \mathcal{C}^{\mathcal{E}}$ we have

$$w_{\gamma} = \frac{\lambda^{|\gamma|}}{(1+\lambda)^{|\partial\gamma|}} \leq \frac{\lambda^{|\gamma|}}{(1+\lambda)^{(1+\alpha)|\gamma|}} \leq (1+\lambda)^{-\alpha|\gamma|}.$$

It thus suffices to show that

$$\sum_{\gamma': d(\gamma', \gamma) \leq 1} (1+\lambda)^{-\alpha|\gamma'|} \cdot e^{2|\gamma'|} \leq |\gamma|.$$

If we could show that for each $v \in V(G)$

$$\sum_{\gamma': \gamma' \ni v} (1+\lambda)^{-\alpha|\gamma'|} \cdot e^{2|\gamma'|} \leq \frac{1}{\Delta^2}, \quad (8)$$

then by summing this inequality over all $v \in \mathcal{E}$ at distance at most 2 from γ (noting that there are at most $(\Delta(\Delta-1)+1)|\gamma| \leq \Delta^2|\gamma|$ such vertices), we would be done.

In order to establish (8), first observe that the graph G^2 has maximum degree at most Δ^2 and so by Lemma 14 the number of G^2 -connected sets of size t containing vertex v is at most $(e\Delta^2)^t$. We thus have

$$\sum_{\gamma': \gamma' \ni v} (1+\lambda)^{-\alpha|\gamma'|} \cdot e^{2|\gamma'|} \leq \sum_{t=1}^{\infty} (e^3\Delta^2(1+\lambda)^{-\alpha})^t \leq \frac{1}{\Delta^2}$$

provided $\lambda > (2e^3\Delta^4)^{1/\alpha}$.

4.3 Proof of Theorem 1

We consider two cases separately. If $\varepsilon < 2^{-n}$, then we proceed by brute force, checking all subsets of $V(G)$ to see if they are independent. In this way we can calculate the partition function $Z_G(\lambda)$ exactly in time $O(n2^n)$ and therefore count and sample in time n/ε .

Now we assume $\varepsilon > 2^{-n}$ and take

$$\lambda > \max \left\{ (2e^3\Delta^4)^{1/\alpha}, e^{11/\alpha} \right\}.$$

Assume without loss of generality that $|\mathcal{O}| \geq |\mathcal{E}|$. Note that by the definition of a bipartite expander we also have $|\mathcal{E}| \geq |\mathcal{O}|/2$.

Using the FPTAS for $\Xi^\mathcal{E}(G)$ given by Theorem 8, we may find $Z_{\text{alg}}^\mathcal{E}$, an $\varepsilon/2$ -relative approximation to $(1 + \lambda)^{|\mathcal{O}|} \Xi^\mathcal{E}(G)$ in $(n/\varepsilon)^{O(\log \Delta)}$ -time.

In identical fashion we may find $Z_{\text{alg}}^\mathcal{O}$, an $\varepsilon/2$ -relative approximation to $(1 + \lambda)^{|\mathcal{E}|} \Xi^\mathcal{O}(G)$ in $(n/\varepsilon)^{O(\log \Delta)}$ -time. It follows that $Z^{\text{alg}} := Z_{\text{alg}}^\mathcal{O} + Z_{\text{alg}}^\mathcal{E}$ is an $\varepsilon/2$ -relative approximation to $\tilde{Z}_G(\lambda)$ (as defined in Lemma 20). By Lemma 20, $\tilde{Z}_G(\lambda)$ is an $\varepsilon/2$ -relative approximation to $Z_G(\lambda)$ and so Z^{alg} is an ε -relative approximation to $Z_G(\lambda)$ as required.

The proof for the sampling algorithm is much like that for the Potts model except that we lack the exact symmetry between ground states. Consider the distribution $\hat{\mu}$ on $\mathcal{I}(G)$ defined as follows. First choose \mathcal{E} or \mathcal{O} with probability proportional to $(1 + \lambda)^{|\mathcal{O}|} \Xi^\mathcal{E}(G)$ and $(1 + \lambda)^{|\mathcal{E}|} \Xi^\mathcal{O}(G)$ respectively. Then, supposing we chose \mathcal{O} , sample Γ from the measure

$$\nu_G^\mathcal{O}(\Gamma) = \frac{\prod_{\gamma \in \Gamma} w_\gamma}{\Xi^\mathcal{O}(G)}.$$

We then set

$$I = J \cup \bigcup_{\gamma \in \Gamma} \gamma$$

where we sample J from the set $\mathcal{E} \setminus \bigcup_{\gamma \in \Gamma} \partial\gamma$ by including each vertex independently with probability $\frac{\lambda}{1+\lambda}$. The distribution of I is $\hat{\mu}$.

By Lemma 20 we have

$$\|\hat{\mu} - \mu_{G,\lambda}\|_{TV} = O(e^{-n})$$

and so to obtain an ε -approximate sample from $\mu_{G,\lambda}$ efficiently, it suffices to obtain an $\varepsilon/2$ -approximate sample from $\hat{\mu}$. We do this as follows:

1. Compute $Z_{\text{alg}}^\mathcal{E}$, an $\varepsilon/8$ -relative approximation to $(1 + \lambda)^{|\mathcal{O}|} \Xi^\mathcal{E}(G)$, and $Z_{\text{alg}}^\mathcal{O}$, an $\varepsilon/8$ -relative approximation to $(1 + \lambda)^{|\mathcal{E}|} \Xi^\mathcal{O}(G)$.
2. Choose \mathcal{O} with probability $\frac{Z_{\text{alg}}^\mathcal{O}}{Z_{\text{alg}}^\mathcal{E} + Z_{\text{alg}}^\mathcal{O}}$ and \mathcal{E} otherwise.
3. Then, supposing we chose \mathcal{O} , take Γ , an $\varepsilon/4$ approximate sample from $\nu_G^\mathcal{O}$, and let

$$I = J \cup \bigcup_{\gamma \in \Gamma} \gamma$$

where we sample J from the set $\mathcal{E} \setminus \bigcup_{\gamma \in \Gamma} \partial\gamma$ by including each vertex independently with probability $\frac{\lambda}{1+\lambda}$.

The resulting distribution on independent sets is within $\varepsilon/2$ total variation distance of $\hat{\mu}$, and we can obtain the sample in time polynomial in n and $1/\varepsilon$: the computation of $Z_{\text{alg}}^\mathcal{E}$ and $Z_{\text{alg}}^\mathcal{O}$ is done as above, and the approximate sample from $\nu_G^\mathcal{O}$ or $\nu_G^\mathcal{E}$ is obtained efficiently by applying Theorem 9. □

4.4 Proof of Theorem 2

To prove Theorem 2 we need a result on the expansion of random regular bipartite graphs. In order to state the result we first generalize our notion of expansion slightly.

Definition 21. For $\rho > 0$ and $\sigma \in (0, 1)$, a bipartite graph $G = (\mathcal{O}, \mathcal{E}, E)$ is a bipartite (σ, ρ) -expander if $|\partial S| \geq \rho|S|$ for all $S \subseteq \mathcal{O}$ with $|S| \leq \sigma|\mathcal{O}|$ and all $S \subseteq \mathcal{E}$ with $|S| \leq \sigma|\mathcal{E}|$.

Note that our previous definition of a bipartite α -expander is the same notion as a $(1/2, 1 + \alpha)$ -expander.

Theorem 22 (Bassalygo [7]). *Almost every Δ -regular bipartite graph is an (σ, ρ) -expander provided*

$$\Delta > \frac{H(\sigma) + H(\sigma\rho)}{H(\sigma) - \sigma\rho H(1/\rho)},$$

where $H(p) = -p \log_2(p) - (1-p) \log_2(1-p)$ is the binary entropy function.

We will take advantage of the fact that small sets in the random regular bipartite graph expand by a lot.

Lemma 23. *There exists Δ_0 such that for all $\Delta \geq \Delta_0$, almost every Δ -regular bipartite graph is a $\left(\frac{4 \log \Delta}{\Delta}, \frac{\Delta}{4 \log \Delta} - \frac{1}{2}\right)$ -expander.*

Proof. By Theorem 22 it suffices to verify that

$$\Delta > \frac{H(\sigma) + H(\sigma\rho)}{H(\sigma) - \sigma\rho H(1/\rho)}, \tag{9}$$

where $\sigma = 4 \log(\Delta)/\Delta$ and $\rho = 1/\sigma - 1/2$. Using the fact

$$H(x) = \frac{x \log(1/x) + x - x^2/2}{\log 2} + O(x^3), \tag{10}$$

(noting that this expansion is also valid for $H(1-x)$ by symmetry of the entropy function) we have

$$H(\sigma) + H(\sigma\rho) = \frac{6}{\log 2} \cdot \frac{\log^2 \Delta}{\Delta} + O\left(\frac{\log \Delta}{\Delta}\right)$$

and

$$H(\sigma) - \sigma\rho H(1/\rho) = \frac{8}{\log 2} \cdot \frac{\log^2 \Delta}{\Delta^2} + O\left(\frac{\log^3 \Delta}{\Delta^3}\right).$$

It follows that there exists Δ_0 such that (9) holds for all $\Delta \geq \Delta_0$. \square

Henceforth we will assume that $G = (\mathcal{O}, \mathcal{E}, E)$ is a Δ -regular $\left(\frac{4 \log \Delta}{\Delta}, \frac{\Delta}{4 \log \Delta} - \frac{1}{2}\right)$ -expander on $n = 2m$ vertices. This strong expansion condition allows us to prove a strengthened version of Lemma 18. Let us update our notion of a small set and say that a set $S \subseteq V(G)$ is *tiny* if $|S| \leq \frac{4 \log \Delta}{\Delta} m$.

Lemma 24. *Let $I \in \mathcal{I}(G)$, then at least one of the sets $I \cap \mathcal{O}, I \cap \mathcal{E}$ is tiny.*

Proof. Suppose that $|I \cap \mathcal{E}| \geq \frac{4 \log \Delta}{\Delta} m$, then by the expansion property

$$|\partial(I \cap \mathcal{E})| \geq \left(1 - \frac{2 \log \Delta}{\Delta}\right) m$$

and so $|I \cap \mathcal{O}| \leq \frac{2 \log \Delta}{\Delta} m$. □

We define two polymer models in identical fashion to Section 4.1: An even (odd) *polymer* is any *tiny* G^2 -connected set γ which lies entirely in \mathcal{E} (\mathcal{O}). We say two even (odd) polymers γ_1, γ_2 are *compatible* if $d_G(\gamma_1, \gamma_2) > 2$ and for each polymer γ we assign a weight

$$w_\gamma := \frac{\lambda^{|\gamma|}}{(1 + \lambda)^{|\partial\gamma|}}.$$

We let $\mathcal{C}^\mathcal{E}$ denote the set of all even polymers of G and let $\mathcal{G}^\mathcal{E}$ denote the family of all sets of mutually compatible polymers from $\mathcal{C}^\mathcal{E}$. We define $\mathcal{C}^\mathcal{O}, \mathcal{G}^\mathcal{O}$ similarly. Let us again denote the partition functions associated to these polymer models by $\Xi^\mathcal{E}(G), \Xi^\mathcal{O}(G)$ respectively. As before a linear combination of these two partition functions serves a good approximation to $Z_G(\lambda)$.

Lemma 25. *There exists Δ_0 such that for $\Delta \geq \Delta_0, \lambda > 20 \frac{\log^2 \Delta}{\Delta}$, the polynomial*

$$\tilde{Z}_G(\lambda) = (1 + \lambda)^m (\Xi^\mathcal{E}(G) + \Xi^\mathcal{O}(G))$$

is a $(1 + \lambda)^{-n/4}$ -relative approximation to $Z_G(\lambda)$.

Proof. As in the proof of Lemma 20, we have

$$\tilde{Z}_G(\lambda) = Z_G(\lambda) + \sum_{I \text{ sparse}} \lambda^{|I|}, \tag{11}$$

where we call an independent set I *sparse* if the G^2 -connected components of $I \cap \mathcal{E}$ and $I \cap \mathcal{O}$ are all tiny.

Let I be a sparse independent set. Since $I \cap \mathcal{E}, I \cap \mathcal{O}$ are composed of tiny G^2 -connected components it follows that $|\partial(I \cap \mathcal{E})| > \left(\frac{\Delta}{4 \log \Delta} - \frac{1}{2}\right) |I \cap \mathcal{E}|$ and $|\partial(I \cap \mathcal{O})| > \left(\frac{\Delta}{4 \log \Delta} - \frac{1}{2}\right) |I \cap \mathcal{O}|$. Since each element of $\partial(I \cap \mathcal{E})$ and $\partial(I \cap \mathcal{O})$ must be unoccupied, it follows that

$$|I \cap \mathcal{E}| < m - \left(\frac{\Delta}{4 \log \Delta} - \frac{1}{2}\right) |I \cap \mathcal{O}|$$

and

$$|I \cap \mathcal{O}| < m - \left(\frac{\Delta}{4 \log \Delta} - \frac{1}{2}\right) |I \cap \mathcal{E}|.$$

By summing these two inequalities we conclude that

$$|I| < \frac{n}{\frac{\Delta}{4 \log \Delta} + \frac{1}{2}}. \quad (12)$$

We use the following well-known estimate

$$\binom{n}{\leq cn} := \sum_{i=0}^{\lfloor cn \rfloor} \binom{n}{i} \leq 2^{H(c)n} \quad \text{for } c \leq \frac{1}{2}.$$

It follows from (12) and (10) that for Δ sufficiently large

$$\begin{aligned} \sum_{I \text{ sparse}} \lambda^{|I|} &\leq \binom{n}{\leq \frac{4 \log \Delta}{\Delta} n} (1 + \lambda)^{\frac{4 \log \Delta}{\Delta} n} \\ &\leq \exp \left\{ \log 2 \cdot H \left(\frac{4 \log \Delta}{\Delta} \right) n + \frac{4 \log \Delta}{\Delta} \log(1 + \lambda) n \right\} \\ &\leq \exp \left\{ 5 \log 2 \cdot \frac{\log^2 \Delta}{\Delta} n + \frac{4 \log \Delta}{\Delta} \log(1 + \lambda) n \right\} \\ &\leq \exp \{ \log(1 + \lambda) n / 4 \} \\ &= (1 + \lambda)^{n/4}. \end{aligned}$$

It follows from (11) and the crude bound $Z_{\mathcal{O}}^*(\lambda) \geq (1 + \lambda)^{n/2}$ that

$$\left| 1 - \frac{\tilde{Z}_G(\lambda)}{Z_G(\lambda)} \right| < (1 + \lambda)^{-n/4}.$$

□

4.5 Verifying the Kotecký-Preiss condition

We will verify condition (1) with the function $g(\gamma) = |\gamma| \frac{\Delta}{10 \log \Delta} \log(1 + \lambda)$. In particular, we will show

$$\sum_{\gamma \ni v} w_\gamma e^{|\gamma| + g(\gamma)} \leq \frac{1}{\Delta^2}.$$

We have

$$w_\gamma = \frac{\lambda^{|\gamma|}}{(1 + \lambda)^{|\partial \gamma|}} \leq \lambda^{|\gamma|} (1 + \lambda)^{-\frac{\Delta}{5 \log \Delta} |\gamma|},$$

and so proceeding as in Section 4.2 we have

$$\sum_{\gamma \ni v} w_\gamma e^{|\gamma| + g(\gamma)} \leq \sum_{k \geq 1} \exp \left[k \left(2 \log \Delta + 2 + \frac{\Delta}{10 \log \Delta} \log(1 + \lambda) + \log \lambda - \frac{\Delta}{5 \log \Delta} \log(1 + \lambda) \right) \right]$$

$$= \sum_{k \geq 1} \exp \left[k \left(2 \log \Delta + 2 - \frac{\Delta}{10 \log \Delta} \log(1 + \lambda) + \log \lambda \right) \right]$$

which, for Δ large enough and $\lambda \geq \frac{50 \log^2 \Delta}{\Delta}$, is at most $1/\Delta^2$. We may now finish as in Section 4.3, noting that this time we can take $\rho = \Delta/(10 \log \Delta)$, $c_1 = 1$, and $c_2 = 0$ in the application of Theorem 8 so that the run-time of the approximate counting algorithm is $(n/\varepsilon)^{1+O(\log^2 \Delta/\Delta)}$. This completes the proof of Theorem 2.

5 Proper colorings

In this section we prove Theorem 5. Let $G \in \mathcal{G}^{\text{bip}}(n, \Delta)$. As in the proof of Theorem 2, the only property we require of G is that it is a $\left(\frac{4 \log \Delta}{\Delta}, \frac{\Delta}{4 \log \Delta} - \frac{1}{2}\right)$ -expander which holds with high probability. As before we denote the two partition classes of G by \mathcal{O}, \mathcal{E} and let $V = V(G)$. Let $m = n/2$ (so that $|\mathcal{O}| = |\mathcal{E}| = m$) and fix an integer $q \geq 3$.

Throughout this section all colorings will be proper vertex q -colorings. Let $\mathcal{X} = \mathcal{X}_{G,q}$ be the set of all proper colorings $f : V(G) \rightarrow [q]$. For a set $S \subseteq V(G)$ and $f \in \mathcal{X}$, we let $f(S) := \{f(v) : v \in S\}$ and we let $f|_S$ denote the restriction of f to S , that is, the map $f|_S : S \rightarrow [q]$ where $f|_S(v) = f(v)$ for all $v \in S$.

Our aim is to obtain an FPTAS for $Z_G(q) = |\mathcal{X}|$ and an efficient sampling algorithm for $\mu_{G,q}$, the uniform distribution over \mathcal{X} .

We have the following important class of colorings of G , which play the role of ground states.

Definition 26. *Let A, B be disjoint subsets of $[q]$ such that $A \cup B = [q]$. We call a coloring $f \in \mathcal{X}$ an (A, B) -coloring if $f(\mathcal{O}) \subseteq A$, $f(\mathcal{E}) \subseteq B$. We call the pair (A, B) a pattern.*

This notion was inspired by the work of Peled and Spinka [39] where such patterns play a similar role.

Let \mathcal{P} denote the set of all patterns. Given a subset $S \subseteq V$, a coloring $f \in \mathcal{X}$ and a pattern (A, B) , we say that f agrees with (A, B) at $v \in V$ if $v \in \mathcal{O}$ and $f(v) \in A$ or if $v \in \mathcal{E}$ and $f(v) \in B$. We say that f disagrees with (A, B) at v otherwise. Let $\chi_{A,B}(S)$ be the set of colorings $f \in \mathcal{X}$ such that f disagrees with (A, B) at each $v \in S$ and agrees at each $v \in V \setminus S$.

Suppose $(A, B) \in \mathcal{P}$, $S \subseteq V$. We record the following simple bound for future use. Recall that for $S \subseteq V$, ∂S denotes the set of vertices in $V \setminus S$ that are adjacent to a vertex of S .

Lemma 27. *For $(A, B) \in \mathcal{P}$ and $S \subseteq V$ we have*

$$|\chi_{A,B}(S)| \leq |A|^m |B|^m \left(1 - \frac{1}{|A|}\right)^{|\partial S \cap \mathcal{O}|} \left(1 - \frac{1}{|B|}\right)^{|\partial S \cap \mathcal{E}|} \left(\frac{|A|}{|B|}\right)^{|S \cap \mathcal{E}| - |S \cap \mathcal{O}|}.$$

Proof. For $f \in \chi_{A,B}(S)$, there are at most $|A|^{|S \cap \mathcal{E}|} |B|^{|S \cap \mathcal{O}|}$ choices for $f|_S$ and given any such choice there are then at most $(|A| - 1)^{|\partial S \cap \mathcal{O}|} (|B| - 1)^{|\partial S \cap \mathcal{E}|}$ choices for $f|_{\partial S}$.

Finally, given any choice of $f|_{S^+}$ there are at most $|A|^{m-|S^+\cap\mathcal{O}}|B|^{m-|S^+\cap\mathcal{E}|}$ choices for $f|_{V\setminus S^+}$. The result follows. \square

Henceforth, let us call a set $S \subseteq V$ *little* if $|S| \leq 4q \frac{\log \Delta}{\Delta} m$.

Lemma 28. *For every $f \in \mathcal{X}$, there is a pattern (A, B) and a little set $S \subseteq V$ such that $f \in \chi_{A, B}(S)$.*

Proof. Given $f \in \mathcal{X}$ let

$$A' = \{i \in [q] : |f^{-1}(\{i\}) \cap \mathcal{O}| > 4 \frac{\log \Delta}{\Delta} m\},$$

$$B' = \{j \in [q] : |f^{-1}(\{j\}) \cap \mathcal{E}| > 4 \frac{\log \Delta}{\Delta} m\}.$$

Note that since G is a $\left(\frac{4 \log \Delta}{\Delta}, \frac{\Delta}{4 \log \Delta} - \frac{1}{2}\right)$ -expander, we have that $E(X, Y) \neq \emptyset$ for any two subsets $X \subseteq \mathcal{O}, Y \subseteq \mathcal{E}$ such that $|X|, |Y| > 4 \frac{\log \Delta}{\Delta} m$. It follows that A' and B' are disjoint. Letting (A, B) be any pattern such that $A' \subseteq A$ and $B' \subseteq B$, we have $f \in \chi_{A, B}(S)$ for some little set S by construction. \square

We now define a collection of polymer models and partition functions, a linear combination of which will serve as a good approximation to $Z_G(q) = |\mathcal{X}|$. We define a polymer to be a *little*, G^3 -connected subset of G (see Definition 19). We note that a subset $S \subseteq V$ is G^3 -connected if and only if $G[S^+]$ is connected. We say that two polymers γ_1, γ_2 are *compatible* if $d_G(\gamma_1, \gamma_2) > 3$ (i.e. $\gamma_1 \cup \gamma_2$ is not G^3 -connected). We let $\mathcal{C} = \mathcal{C}(G)$ denote the set of all polymers of G and let $\mathcal{G} = \mathcal{G}(G)$ denote the family of all sets of mutually compatible polymers from \mathcal{C} . Let us now fix a pattern $(A, B) \in \mathcal{P}$. To each polymer $\gamma \in \mathcal{C}$, we assign a weight

$$w_{A, B}(\gamma) := \frac{|\chi_{A, B}(\gamma)|}{|A|^m |B|^m}.$$

This defines a polymer model with partition function

$$\Xi_{A, B}(G) = \sum_{\Gamma \in \mathcal{G}} \prod_{\gamma \in \Gamma} w_{A, B}(\gamma).$$

Lemma 29. *There is an absolute constant C such that if $\Delta \geq Cq^2 \log^2 q$, then the sum*

$$\tilde{Z}_G(q) := \sum_{(A, B) \in \mathcal{P}} |A|^m |B|^m \cdot \Xi_{A, B}(G)$$

is a $e^{-m/4q}$ -relative approximation to $Z_G(q)$.

Proof. Let us fix $(A, B) \in \mathcal{P}$ and let $S \subseteq V$. Let $\hat{\chi}_{A, B}(S)$ be the set of colorings $f|_{S^+}$ where $f \in \chi_{A, B}(S)$. Given any $g \in \hat{\chi}_{A, B}(S)$, we may extend g to an element of $\chi_{A, B}(S)$

by arbitrarily assigning vertices of $\mathcal{O} \setminus S^+$ with colors from A and arbitrarily assigning vertices of $\mathcal{E} \setminus S^+$ with colors from B . It follows that

$$\frac{|\chi_{A,B}(S)|}{|A|^m |B|^m} = \frac{|\hat{\chi}_{A,B}(S)|}{|A|^{|S^+ \cap \mathcal{O}|} |B|^{|S^+ \cap \mathcal{E}|}}. \quad (13)$$

If the G^3 -connected components of S are $\gamma_1, \dots, \gamma_k$ we therefore have

$$\frac{|\chi_{A,B}(S)|}{|A|^m |B|^m} = \prod_{i \in [k]} w_{A,B}(\gamma_i).$$

We call a set $S \subseteq V$ *sparse* if all of its G^3 -connected components are little. We note that there is a one-one correspondence between sparse subsets of V and collections of mutually compatible polymers (i.e. elements of \mathcal{G}). It follows that

$$\sum_{S \text{ sparse}} |\chi_{A,B}(S)| = |A|^m |B|^m \sum_{\Gamma \in \mathcal{G}} \prod_{\gamma \in \Gamma} w_{A,B}(\gamma) = |A|^m |B|^m \cdot \Xi_{A,B}(G)$$

and so

$$\tilde{Z}_G(q) = \sum_{(A,B) \in \mathcal{P}} \sum_{S \text{ sparse}} |\chi_{A,B}(S)|.$$

We first show that $\tilde{Z}_G(q)$ is a good approximation to the sum

$$\hat{Z}_G(q) := \sum_{(A,B) \in \mathcal{P}} \sum_{S \text{ little}} |\chi_{A,B}(S)|.$$

To this end we need the following claim.

Claim 30. *If S is sparse then $|S| \leq 12q \frac{\log \Delta}{\Delta} m$.*

Proof of Claim 30. Let S be sparse and let $\Gamma = \{\gamma_1, \dots, \gamma_k\}$ be the G^3 -connected components of S . Suppose that $|S| > 12q \frac{\log \Delta}{\Delta} m$. Since each γ_i is little by assumption, we may partition $\Gamma = \Gamma_1 \cup \Gamma_2$ in such a way that $\sum_{\gamma \in \Gamma_j} |\gamma| \geq \frac{4q \log \Delta}{\Delta} m$ for $j = 1, 2$. Let $S_j = \bigcup_{\gamma \in \Gamma_j} \gamma$ for $j = 1, 2$. Suppose without loss of generality that $|S_1 \cap \mathcal{O}| \geq \frac{2q \log \Delta}{\Delta} m$. Since G is a $\left(\frac{4 \log \Delta}{\Delta}, \frac{\Delta}{4 \log \Delta} - \frac{1}{2}\right)$ -expander, by considering a subset of $S_1 \cap \mathcal{O}$ of size $\frac{4 \log \Delta}{\Delta} m$, we have

$$|\partial(S_1 \cap \mathcal{O})| > \left(1 - \frac{2 \log \Delta}{\Delta}\right) m \quad (14)$$

and so $|S_1^+| > m$. Similarly $|S_2^+| > m$ so that $S_1^+ \cap S_2^+ \neq \emptyset$. It follows that $\gamma_i^+ \cap \gamma_j^+ \neq \emptyset$ for some $\gamma_i \in \Gamma_1, \gamma_j \in \Gamma_2$, contradicting the fact that γ_i, γ_j are distinct G^3 -connected components of S . \square

Suppose now that $S \subseteq V$ is sparse and not little, and suppose without loss of generality that $|S \cap \mathcal{O}| \geq 2q \frac{\log \Delta}{\Delta} m$. As in (14) we have $|\partial(S \cap \mathcal{O})| \geq \left(1 - \frac{2 \log \Delta}{\Delta}\right) m$ and so by Claim 30

$$|\partial S \cap \mathcal{E}| > \left(1 - \frac{14q \log \Delta}{\Delta}\right) m > \frac{m}{2}.$$

By Lemma 27 we then have

$$\begin{aligned} |\chi_{A,B}(S)| &\leq |A|^m |B|^m \left(1 - \frac{1}{|A|}\right)^{|\partial S \cap \mathcal{O}|} \left(1 - \frac{1}{|B|}\right)^{|\partial S \cap \mathcal{E}|} \left(\frac{|A|}{|B|}\right)^{|S \cap \mathcal{E}| - |S \cap \mathcal{O}|} \\ &\leq |A|^m |B|^m \left(1 - \frac{1}{|B|}\right)^{m/2} q^{12q \frac{\log \Delta}{\Delta} m}, \end{aligned}$$

where for the final inequality we have used the crude bounds $|A|/|B| \in [1/q, q]$, $\|S \cap \mathcal{E}\| - \|S \cap \mathcal{O}\| \leq |S|$ and Claim 30. We thus have

$$\begin{aligned} \tilde{Z}_G(q) - \hat{Z}_G(q) &= \sum_{(A,B) \in \mathcal{P}} \sum_{\substack{S \text{ sparse,} \\ \text{not little}}} |\chi_{A,B}(S)|, \\ &\leq \sum_{(A,B) \in \mathcal{P}} \binom{2m}{12q \frac{\log \Delta}{\Delta} m} |A|^m |B|^m \left(1 - \frac{1}{|B|}\right)^{m/2} q^{12q \frac{\log \Delta}{\Delta} m}, \\ &\leq \sum_{(A,B) \in \mathcal{P}} |A|^m |B|^m \exp \left\{ 2 \log 2 \cdot H \left(6q \frac{\log \Delta}{\Delta}\right) m + 12q \log q \frac{\log \Delta}{\Delta} m - \frac{m}{2q} \right\}, \\ &\leq e^{-m/(3q)} \sum_{(A,B) \in \mathcal{P}} |A|^m |B|^m, \end{aligned}$$

where for the first inequality we used the fact that there are at most $\binom{2m}{12q m \log \Delta / \Delta}$ sparse subsets of V and for the final inequality we used the assumed lower bound on Δ . Using the crude bound $\tilde{Z}_G(q) \geq \sum_{(A,B) \in \mathcal{P}} |A|^m |B|^m$ we see that $\hat{Z}_G(q)$ is a $e^{-m/(3q)}$ -relative approximation to $\tilde{Z}_G(q)$.

By Lemma 28 we have

$$\mathcal{X} = \bigcup_{(A,B) \in \mathcal{P}} \bigcup_{S \text{ little}} \chi_{A,B}(S).$$

Suppose now that $(A, B), (C, D) \in \mathcal{P}$, $S, T \subseteq V$ and $\chi_{A,B}(S) \cap \chi_{C,D}(T) \neq \emptyset$. Unless $(A, B, S) = (C, D, T)$ we must have $(A, B) \neq (C, D)$. WLOG suppose that $|A| = \max\{|A|, |B|, |C|, |D|\}$. Then $|A \cap C| \leq |A| - 1$ and $|B \cap D| \leq |B|$. If S, T are both little we then have

$$|\chi_{A,B}(S) \cap \chi_{C,D}(T)| \leq (|A| - 1)^m |B|^m q^{8q \frac{\log \Delta}{\Delta} m}$$

$$\leq \left(\frac{q-1}{2}\right)^{2m} \exp\left\{8q \log q \frac{\log \Delta}{\Delta} m\right\}.$$

The first inequality comes from the fact that there are at most $(|A| - 1)^m$ ways to color the vertices of \mathcal{O} that agree with both (A, B) and (C, D) (and so are colored with elements of $A \cap C$), there are at most $|B|^m$ ways to color the vertices of \mathcal{E} that agree with both (A, B) and (C, D) , and there are at most $q^{8q \frac{\log \Delta}{\Delta} m}$ ways to color the vertices of S and T . We may therefore bound $|\mathcal{X}|$ by inclusion-exclusion as follows:

$$\hat{Z}_G(q) - 2^{2q} \left(4q \frac{\log \Delta}{\Delta} m\right)^2 \left(\frac{q-1}{2}\right)^{2m} \exp\left\{8q \log q \frac{\log \Delta}{\Delta} m\right\} \leq |\mathcal{X}| \leq \hat{Z}_G(q),$$

where we have used the bound $|\mathcal{P}| \leq 2^q$ and that there are at most $\binom{2m}{4qm \log \Delta / \Delta}$ little subsets of V . Using the bound $\hat{Z}_G(q) \geq \left[\frac{q}{2}\right]^m \left[\frac{q}{2}\right]^m$, it follows that $\hat{Z}_G(q)$ is a $e^{-m/q}$ -relative approximation to $|\mathcal{X}|$. The result follows. \square

We may therefore focus on approximating each partition function $\Xi_{A,B}(G)$ individually. Henceforth let us fix a pattern $(A, B) \in \mathcal{P}$.

We will verify conditions (i)–(iii) of Theorem 8 for the polymer model defined on \mathcal{C} with weight function $w_{A,B}$ in order to obtain an FPTAS for $\Xi_{A,B}(G)$.

Verifying condition (i) is essentially immediate. Given $\gamma \subseteq V(G)$, determining whether $\gamma \in \mathcal{C}$ amounts to checking whether γ is G^3 -connected or equivalently whether $G[\gamma^+]$ is connected. This can be done in $O(\Delta |\gamma|)$ time by a depth-first search algorithm. By (13), $w_{A,B}(\gamma)$ can be calculated in $e^{O(\Delta \log q |\gamma|)}$ time by checking all possible colorings of γ^+ .

We now turn our attention verifying conditions (ii) and (iii) for an appropriate choice of function g .

5.1 Verifying the Kotecký-Preiss condition

For brevity we denote $w_{A,B}$ simply by w . We choose

$$g(\gamma) = \frac{\Delta}{10q^2 \log \Delta} |\gamma|.$$

It remains to show the Kotecký-Preiss condition holds. That is,

$$\sum_{\gamma': d(\gamma', \gamma) \leq 3} w(\gamma') e^{|\gamma'| + g(\gamma')} \leq |\gamma|$$

for all $\gamma \in \mathcal{C}$. If we could show that for each $v \in V(G)$

$$\sum_{\gamma': \gamma' \ni v} w(\gamma') \cdot e^{|\gamma'| + g(\gamma')} \leq \frac{1}{\Delta^3},$$

then by summing this inequality over all v at distance at most 3 from γ in G (noting that there are at most $\Delta(\Delta - 1)^2 + 1 \leq \Delta^3|\gamma|$ such vertices) we would be done.

The graph G^3 has maximum degree at most Δ^3 and so by Lemma 14 the number of G^3 -connected sets of size t containing vertex v is at most $(e\Delta^3)^t$. Suppose now that $S \subseteq V$ is little. Since G is a $\left(\frac{4\log \Delta}{\Delta}, \frac{\Delta}{4\log \Delta} - \frac{1}{2}\right)$ -expander, by considering a subset $T \subseteq S \cap \mathcal{O}$ such that $|T| = \frac{|S \cap \mathcal{O}|}{q} \leq \frac{4\log \Delta}{\Delta}m$ we have

$$|\partial(S \cap \mathcal{O})| \geq \left(\frac{\Delta}{4\log \Delta} - \frac{1}{2}\right) \frac{|S \cap \mathcal{O}|}{q}$$

and similarly

$$|\partial(S \cap \mathcal{E})| \geq \left(\frac{\Delta}{4\log \Delta} - \frac{1}{2}\right) \frac{|S \cap \mathcal{E}|}{q}.$$

By summing the above two inequalities we obtain

$$|\partial S| \geq |\partial(S \cap \mathcal{O})| + |\partial(S \cap \mathcal{E})| - |S| \geq \left(\frac{\Delta}{4\log \Delta} - q - \frac{1}{2}\right) \frac{|S|}{q}.$$

If $\gamma \in \mathcal{C}$, so that in particular γ is little, by Lemma 27 we then have

$$w(\gamma) \leq \left(1 - \frac{1}{q}\right)^{|\partial\gamma|} q^{|\gamma|} \leq \exp\left\{-\frac{\Delta}{5q^2 \log \Delta} |\gamma|\right\},$$

where we've assumed that $\Delta > Cq^2 \log^2 q$ for a large absolute constant C as in Lemma 29. Putting everything together we have

$$\sum_{\gamma': \gamma' \ni v} w(\gamma') \cdot e^{|\gamma'| + g(\gamma')} \leq \sum_{t=1}^{\infty} (e^2 \Delta^3)^t \exp\left\{-\frac{\Delta}{10q^2 \log \Delta} t\right\} \leq \frac{1}{\Delta^3}. \quad (15)$$

5.2 Proof of Theorem 5

We consider two cases separately. If $\varepsilon < e^{-n/(8q)}$, then we proceed by brute force, checking all possible elements of $[q]^{V(G)}$ to see if they are a proper q -coloring. In this way we can calculate $Z_G(q)$ exactly in time $O(n\Delta q^n)$ and therefore count and sample in time polynomial in $1/\varepsilon$.

Now we assume $\varepsilon > e^{-n/(8q)}$ and take

$$\Delta > Cq^2 \log^2 q,$$

where C is chosen large enough so that Lemmas 23, 29 and inequality (15) all hold.

For each $(A, B) \in \mathcal{P}$, using the FPTAS for $\Xi_{A,B}(G)$ given by Theorem 8, we may find $Z_{A,B}^{\text{alg}}$, an $\varepsilon/3$ -relative approximation to $|A|^m |B|^m \Xi_{A,B}(G)$ in time polynomial in n and $1/\varepsilon$. It follows that $Z^{\text{alg}} := \sum_{(A,B) \in \mathcal{P}} Z_{A,B}^{\text{alg}}$ is an $\varepsilon/3$ -relative approximation to

$\tilde{Z}_G(q)$ (as defined in Lemma 29). By Lemma 29, $\tilde{Z}_G(q)$ is an $\varepsilon/3$ -relative approximation to $Z_G(q)$ and so Z^{alg} is an ε -relative approximation to $Z_G(q)$ as required.

We now turn our attention to the sampling algorithm. Let μ denote the uniform distribution on \mathcal{X} . Consider the distribution $\hat{\mu}$ on \mathcal{X} defined as follows. First choose $(A, B) \in \mathcal{P}$ with probability proportional to $|A|^m |B|^m \Xi_{A,B}(G)$. Then, sample $\Gamma \in \mathcal{G}$ from the measure

$$\nu_{A,B}(\Gamma) = \frac{\prod_{\gamma \in \Gamma} w_{A,B}(\gamma)}{\Xi_{A,B}(G)}.$$

Let $S = \bigcup_{\gamma \in \Gamma} \gamma$. Uniformly choose a coloring of S^+ which disagrees with (A, B) on S and agrees on ∂S . Then extend this coloring to a coloring of G by coloring vertices of $\mathcal{O} \setminus S^+$ with elements of A and vertices of $\mathcal{E} \setminus S^+$ with elements of B uniformly at random. The distribution of the resulting coloring is $\hat{\mu}$. By Lemma 29 we have

$$\|\hat{\mu} - \mu\|_{TV} \leq \varepsilon/2.$$

Then to obtain an $\varepsilon/2$ -approximate sample from $\hat{\mu}$ efficiently, we proceed as in Section 4.3- first we approximate $\Xi_{A,B}(G)$ for $(A, B) \in \mathcal{P}$ in order to approximately sample $(A, B) \in \mathcal{P}$ and then we apply Theorem 9 to obtain an approximate sample from $\nu_{A,B}$. As in Section 3, the polymer sampling algorithm of Theorem 9 guarantees that $|S^+| = O(\log(n/\varepsilon))$ and so a uniformly random coloring of S^+ which disagrees with (A, B) on S and agrees on ∂S can be obtained by brute force in time polynomial in n and $1/\varepsilon$.

□

6 Concluding Remarks

The algorithms presented here are the first provably efficient counting and sampling algorithms for #BIS-hard problems for the class of expander graphs. However, they are presumably not optimal in terms of either their running time or the range of parameters for which they are provably efficient.

One natural choice for more efficient algorithms would be those based on Markov chains. A candidate algorithm for the Potts model is the Swendsen-Wang dynamics [45]. It is natural to conjecture that the Swendsen-Wang dynamics are rapidly mixing on expander graphs at sufficiently low temperatures, and this would give a more efficient sampling algorithm than the one presented here. Similarly, for the hard-core model on a bipartite expander graph with symmetry between the sides of the bipartition, one could follow the suggestion of [27] and start the Glauber dynamics in either the all even or all odd occupied state with equal probability. Proving that such sampling algorithms are indeed efficient is left as an open problem. Recently Chen, Galanis, Goldberg, Perkins, Stewart, and Vigoda [15] have made some progress in this direction by showing that a restricted version of Glauber dynamics for the Potts and hard-core models mixes in polynomial time under similar conditions to those of Theorems 1 and 3.

For random graphs, we can ask if efficient algorithms for the problems addressed in this paper exist for *all* ranges of parameters.

Question 31. *Is there a polynomial-time sampling algorithm for the Potts model on random regular graphs at all inverse temperatures β ? Are there efficient sampling algorithms for the hard-core model and proper colorings on random regular bipartite graphs at all fugacities λ and for all q ?*

References

- [1] N. Alon. Eigenvalues and expanders. *Combinatorica*, 6(2):83–96, 1986.
- [2] N. Anari, K. Liu, and S. O. Gharan. Spectral independence in high-dimensional expanders and applications to the hardcore model. *arXiv preprint arXiv:2001.00303*, 2020.
- [3] S. Arora, S. A. Khot, A. Kolla, D. Steurer, M. Tulsiani, and N. K. Vishnoi. Unique games on expanding constraint graphs are easy. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, STOC 2008, pages 21–28. ACM, 2008.
- [4] A. Barvinok. Computing the permanent of (some) complex matrices. *Foundations of Computational Mathematics*, 16(2):329–342, 2016.
- [5] A. Barvinok. *Combinatorics and Complexity of Partition Functions*, volume 30 of *Algorithms and Combinatorics*. Springer, 2017.
- [6] A. Barvinok and G. Regts. Weighted counting of solutions to sparse systems of equations. *Combinatorics, Probability and Computing*, 28(5):696–719, 2019.
- [7] L. Bassalygo. Asymptotically optimal switching circuits. *Problems of Information Transmission*, 17(3):206–211, 1981.
- [8] A. Björklund, T. Husfeldt, P. Kaski, and M. Koivisto. Computing the Tutte polynomial in vertex-exponential time. In *Proceedings of the Forty-ninth Annual Symposium on Foundations of Computer Science*, FOCS 2008, pages 677–686. IEEE, 2008.
- [9] A. Blanca, A. Galanis, L. A. Goldberg, D. Stefankovic, E. Vigoda, and K. Yang. Sampling in uniqueness from the Potts and random-cluster models on random regular graphs. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
- [10] G. Brito, I. Dumitriu, and K. D. Harris. Spectral gap in random bipartite biregular graphs and applications. *arXiv preprint arXiv:1804.07808*, 2018.
- [11] J.-Y. Cai, A. Galanis, L. A. Goldberg, H. Guo, M. Jerrum, D. Štefankovič, and E. Vigoda. # BIS-hardness for 2-spin systems on bipartite bounded degree graphs in the tree non-uniqueness region. *Journal of Computer and System Sciences*, 82(5):690–711, 2016.

- [12] S. Cannon and W. Perkins. Counting independent sets in unbalanced bipartite graphs. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1456–1466. SIAM, 2020.
- [13] K. Casel, P. Fischbeck, T. Friedrich, A. Göbel, and J. Lagodzinski. Zeros and approximations of Holant polynomials on the complex plane. *arXiv preprint arXiv:1905.03194*, 2019.
- [14] P. Chebolu, L. A. Goldberg, and R. Martin. The complexity of approximately counting stable matchings. *Theoretical Computer Science*, 437:35–68, 2012.
- [15] Z. Chen, A. Galanis, L. A. Goldberg, W. Perkins, J. Stewart, and E. Vigoda. Fast algorithms at low temperatures via Markov chains. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2019.
- [16] A. Dembo, A. Montanari, A. Sly, and N. Sun. The replica symmetric solution for Potts models on d-regular graphs. *Communications in Mathematical Physics*, 327(2):551–575, 2014.
- [17] R. Dobrushin. Estimates of semi-invariants for the Ising model at low temperatures. *Translations of the American Mathematical Society-Series 2*, 177:59–82, 1996.
- [18] M. Dyer, L. A. Goldberg, C. Greenhill, and M. Jerrum. The relative complexity of approximate counting problems. *Algorithmica*, 38(3):471–500, 2004.
- [19] C. Efthymiou, T. P. Hayes, D. Štefanković, E. Vigoda, and Y. Yin. Convergence of MCMC and loopy BP in the tree uniqueness region for the hard-core model. In *Proceedings of the Fifty-seventh Annual Symposium on Foundations of Computer Science*, FOCS 2016, pages 704–713. IEEE, 2016.
- [20] J. Friedman. A proof of Alon’s second eigenvalue conjecture. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 720–724, 2003.
- [21] A. Galanis, D. Štefanković, and E. Vigoda. Inapproximability of the partition function for the antiferromagnetic Ising and hard-core models. *Combinatorics, Probability and Computing*, 25(4):500–559, 2016.
- [22] A. Galanis, D. Štefanković, E. Vigoda, and L. Yang. Ferromagnetic Potts model: Refined $\#$ -BIS-hardness and related results. *SIAM Journal on Computing*, 45(6):2004–2065, 2016.
- [23] D. Galvin and J. Kahn. On phase transition in the hard-core model on \mathbb{Z}^d . *Combinatorics, Probability and Computing*, 13(2):137–164, 2004.
- [24] L. A. Goldberg and M. Jerrum. Approximating the partition function of the ferromagnetic Potts model. *Journal of the ACM (JACM)*, 59(5):25, 2012.

- [25] C. Gruber and H. Kunz. General properties of polymer systems. *Communications in Mathematical Physics*, 22(2):133–161, 1971.
- [26] H. Guo and M. Jerrum. Random cluster dynamics for the Ising model is rapidly mixing. *Ann. Appl. Probab.*, 28(2):1292–1313, 04 2018.
- [27] T. Helmuth, W. Perkins, and G. Regts. Algorithmic Pirogov-Sinai theory. *Probability Theory and Related Fields*, to appear. An extended abstract appeared at STOC 2019.
- [28] S. Hoory, N. Linial, and A. Wigderson. Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43(4):439–561, 2006.
- [29] M. Jenssen, P. Keevash, and W. Perkins. Algorithms for #BIS-hard problems on expander graphs. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2235–2247. SIAM, 2019.
- [30] M. Jerrum and A. Sinclair. Polynomial-time approximation algorithms for the Ising model. *SIAM Journal on computing*, 22(5):1087–1116, 1993.
- [31] F. P. Kelly. Stochastic models of computer communication systems. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 379–395, 1985.
- [32] R. Kotecký and D. Preiss. Cluster expansion for abstract polymer models. *Communications in Mathematical Physics*, 103(3):491–498, 1986.
- [33] C. Liao, J. Lin, P. Lu, and Z. Mao. Counting independent sets and colorings on random regular bipartite graphs. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2019.
- [34] J. Liu and P. Lu. FPTAS for #BIS with degree bounds on one side. In *Proceedings of the Forty-seventh Annual ACM Symposium on Theory of Computing*, STOC 2015, pages 549–556. ACM, 2015.
- [35] K. Makarychev and Y. Makarychev. How to play unique games on expanders. In *International Workshop on Approximation and Online Algorithms*, pages 190–200. Springer, 2010.
- [36] E. Mossel, D. Weitz, and N. Wormald. On the hardness of sampling independent sets beyond the tree threshold. *Probability Theory and Related Fields*, 143(3-4):401–439, 2009.
- [37] V. Patel and G. Regts. Deterministic polynomial-time approximation algorithms for partition functions and graph polynomials. *SIAM Journal on Computing*, 46(6):1893–1919, 2017.
- [38] V. Patel and G. Regts. Computing the number of induced copies of a fixed graph in a bounded degree graph. *Algorithmica*, 81(5):1844–1858, 2019.

- [39] R. Peled and Y. Spinka. Rigidity of proper colorings of \mathbb{Z}^d . *arXiv preprint arXiv:1808.03597*, 2018.
- [40] S. A. Pirogov and Y. G. Sinai. Phase diagrams of classical lattice systems. *Theoretical and Mathematical Physics*, 25(3):1185–1192, 1975.
- [41] D. Randall and D. Wilson. Sampling spin configurations of an Ising system. In *Proceedings of the Tenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 1999*, pages 959–960. Society for Industrial and Applied Mathematics, 1999.
- [42] A. D. Scott and A. D. Sokal. The repulsive lattice gas, the independent-set polynomial, and the Lovász local lemma. *Journal of Statistical Physics*, 118(5-6):1151–1261, 2005.
- [43] A. Sly. Computational transition at the uniqueness threshold. In *Proceedings of the Fifty-first Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010*, pages 287–296. IEEE, 2010.
- [44] A. Sly and N. Sun. Counting in two-spin models on d-regular graphs. *The Annals of Probability*, 42(6):2383–2416, 2014.
- [45] R. H. Swendsen and J.-S. Wang. Nonuniversal critical dynamics in Monte Carlo simulations. *Physical review letters*, 58(2):86, 1987.
- [46] R. M. Tanner. Explicit concentrators from generalized n-gons. *SIAM Journal on Algebraic Discrete Methods*, 5(3):287–293, 1984.
- [47] D. Weitz. Counting independent sets up to the tree threshold. In *Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing, STOC 2006*, pages 140–149. ACM, 2006.