

Algorithms for #BIS-hard problems on expander graphs

Matthew Jenssen* Peter Keevash† Will Perkins‡

July 12, 2018

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.

1 Introduction

There are two natural computational problems associated to a statistical physics spin model on a graph G : 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

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

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

*University of Oxford, jenssen@maths.ox.ac.uk.

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

Research supported in part by ERC Consolidator Grant 647678.

‡University of Birmingham, math@willperkins.org.

Research supported in part by EPSRC grant EP/P009913/1.

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 $(1 - \varepsilon)\hat{Z} \leq Z_G(\lambda) \leq (1 + \varepsilon)\hat{Z}$) and runs in time polynomial in $|V(G)|$ and $1/\varepsilon$. An *efficient 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 $\|\mu_{G,\lambda} - \mu_{\text{alg}}\|_{TV} < \varepsilon$.

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 an efficient sampling algorithm due to Weitz [35]; whereas when $\lambda > \lambda_c(\Delta)$ both computational problems are hard: there is no polynomial-time algorithm unless $\text{NP}=\text{RP}$ [32, 33, 12]. The value $\lambda_c(\Delta)$ is the *uniqueness threshold* of the hard-core model on the infinite Δ -regular tree [22].

On the other hand, if we restrict ourselves to bipartite graphs, then 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 [10], and many interesting approximate counting and sampling problems have been shown to be $\#\text{BIS}$ -hard [17, 8, 13]. In particular, Cai, Galanis, Goldberg, Guo, Jerrum, Štefankovič, and Vigoda [7] 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 [1, 25]). 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 [24] 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 [20] 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 [30] along with the Taylor series truncation method due to Barvinok [2, 3].

Here we adapt the algorithms of [20] 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 $|N(S)| \geq (1 + \alpha)|S|$ for all $S \subseteq \mathcal{O}$ with $|S| \leq |\mathcal{O}|/2$ and all $S \subseteq \mathcal{E}$ with $|S| \leq |\mathcal{E}|/2$.

Theorem 1. *There exists an absolute constant C such that for every $\alpha > 0$, $\Delta \geq 3$, and any $\lambda > C\Delta^{5/\alpha}$, there exists an FPTAS and an efficient sampling algorithm for the hard-core model at fugacity λ on bipartite α -expander graphs of maximum degree Δ .*

As with Weitz's algorithm [35], the running time of the approximation algorithm of 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 .

We can extend the methods of the proof of Theorem 1 to obtain efficient counting and sampling algorithms for the hard-core model on random regular bipartite graphs for all λ larger than an absolute constant, independent of Δ . 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 an absolute constant $\lambda^* > 0$ so that for every $\Delta \geq 3$ and all $\lambda > \lambda^*$, there is an FPTAS and an efficient sampling algorithm for the hard-core model at fugacity λ on almost every Δ -regular bipartite graph.*

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

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).

Galanis, Štefankovič, Vigoda, and Yang [13] showed that approximating the ferromagnetic Potts model partition function for $q \geq 3$ 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 [13] 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 $e(S, S^c) \geq \alpha|S|$ for all subsets $S \subseteq V(G)$ with $|S| \leq |V(G)|/2$.

Theorem 3. *For all $\alpha > 0$, $\Delta \geq 3$, $q \geq 2$ and $\beta > 2 \log(15q\Delta^2)/\alpha$, there is an FPTAS and efficient 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 efficient sampling algorithm for the ferromagnetic Potts model at inverse temperature β on almost every Δ -regular graph.*

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.

1.2 Discussion

We note that 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 [21], turned into a sampling algorithm via self-reducibility by Randall and Wilson [31]; see also the recent proof of Guo and Jerrum [19] 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.

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 [26]. As $\lambda_c(\Delta)$ decays like $1/\Delta$, it is an interesting open problem to find efficient algorithms that work closer to $\lambda_c(\Delta)$. We have made no efforts to optimize the constant λ^* in Theorem 2 (the current proof shows that $(2e)^{250}$ suffices). A natural next step for future research would be to reduce λ^* to a constant below 1. This would furnish us with an FPTAS for counting independent sets and sampling independent sets uniformly in the random regular bipartite graph.

In the low fugacity regime with $\lambda < \lambda_c(\Delta)$, Weitz's algorithm applies to $\mathbf{G}_{n, \Delta}^{\text{bip}}$, and Efthymiou, Hayes, Štefankovic, Vigoda, and Yin [11] have also shown that the Glauber dynamics have mixing time $O(n \log n)$.

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 [13] show 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, Štefankovic, Vigoda, and Yang [5] have recently given an efficient algorithm to obtain an n^{-c} -approximate sample from the Potts model (ferromagnetic and anti-ferromagnetic) on $\mathbf{G}_{n,\Delta}$ when the parameters lie 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 [33] 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 [9] 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.

1.3 Proof ideas

Our main technical contribution is to show that the hard-core and Potts 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 the *ground states*: the all even and all odd occupied configurations in the hard-core model, and the monochromatic configurations in the Potts model.

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

1. First we show that the partition functions of the hard-core model and Potts model on (bipartite) expanders are 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 bound on the Peierls’ constant of such a polymer model, which allows us to verify the Kotecký-Preiss condition for the convergence of the cluster expansion [23] and prove the existence of a zero-free region of the polymer model partition function in the complex plane.
3. This last step allows us to implement the approximate counting algorithm, from [20], based on Barvinok’s method of truncating the Taylor series of the log partition

function. 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, and we state the Kotecký-Preiss condition and the algorithmic results from [20] that we will use.

We prove our results for the Potts model in Section 3 and for the hard-core model in Section 4. We conclude with some discussion and open problems in Section 5.

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 [18, 23].

Let G be a graph and Ω a finite set of spins. A *polymer* $\gamma = (\bar{\gamma}, \omega_{\bar{\gamma}})$ is a connected subgraph $\bar{\gamma}$ of G along with an assignment $\omega_{\bar{\gamma}}$ of spins from Ω to the vertices of $\bar{\gamma}$. A polymer model consists of a set of polymers $\mathcal{C}(G)$ along with a positive, integer-valued *surface energy*, $\|\gamma\|$, for each polymer. We measure the size of a polymer by $|\bar{\gamma}|$, the number of vertices of $\bar{\gamma}$.

We say two polymers γ, γ' are *compatible* if $d(\bar{\gamma}, \bar{\gamma}') > 1$ 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

$$\Xi(G, z) := \sum_{\Gamma \in \mathcal{G}(G)} \prod_{\gamma \in \Gamma} z^{\|\gamma\|}.$$

We think of $\Xi(G, z)$ as a univariate polynomial in the complex variable z .

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 surface energy of every polymer v to be $\|v\| = 1$, then the abstract polymer partition function $\Xi(G, z)$ is exactly the hard-core partition function $Z_G(z)$.

2.2 Convergent cluster expansions

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

For a multiset of polymers Γ , the *incompatibility graph* $H(\Gamma)$ has one vertex for each polymer with an edge between two vertices corresponding to polymers γ, γ' if $d(\bar{\gamma}, \bar{\gamma}') \leq 1$. Let $\mathcal{G}_k^{\text{clust}}(G)$ be the collection of all multisets of k polymers from $\mathcal{C}(G)$ whose incompatibility graph is connected. We call an element of $\mathcal{G}_k^{\text{clust}}(G)$ a *cluster*. The cluster expansion is then the (formal) power series

$$\log \Xi(G, z) = \sum_{k \geq 1} \sum_{\Gamma \in \mathcal{G}_k^{\text{clust}}(G)} \phi(\Gamma) \prod_{\gamma \in \Gamma} z^{\|\gamma\|},$$

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

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

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

Theorem 5 ([23]). *Suppose that for all $|z| < \delta$ and all $\gamma \in \mathcal{C}(G)$*

$$\sum_{\gamma': d(\gamma', \gamma) \leq 1} e^{|\bar{\gamma}'| |z|^{\|\gamma'\|}} \leq |\bar{\gamma}|. \quad (1)$$

Then for every $|z| < \delta$, the cluster expansion converges absolutely and, in particular, $\Xi(G, z) \neq 0$.

Often a crucial ingredient in proving the convergence of the cluster expansion is to show that there exists a constant ρ such that $\|\gamma\| \geq \rho |\bar{\gamma}|$ for all polymers γ . We refer to the supremum over all ρ for which such a bound holds as the *Peierls' constant* of the model.

2.3 Algorithms

Under two mild conditions on the polymer model and the surface energies, 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.

Theorem 6 ([20], Theorem 2.2). *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:*

(i) *There are absolute constants $\rho, C > 0$ so that for every $G \in \mathfrak{G}$ and every $\gamma \in \mathcal{C}(G)$,*

$$\rho |\bar{\gamma}| \leq \|\gamma\| \leq C |\bar{\gamma}|.$$

(ii) *Given $\bar{\gamma}, \omega_{\bar{\gamma}}$, determining whether $\gamma \in \mathcal{C}(G)$ and computing $\|\gamma\|$ can be done in time polynomial in $|\bar{\gamma}|$.*

(iii) *There exists $\delta > 0$ so that for all complex $|z| < \delta$ and all $G \in \mathfrak{G}$,*

$$\Xi(G, z) \neq 0.$$

Then for every $0 < z < \delta$, there is an FPTAS for $\Xi(G, z)$ for $G \in \mathfrak{G}$.

The running time of this algorithm is $(n/\varepsilon)^{O(\log \Delta)}$ (and in fact it applies to approximating $\Xi(G, z)$ for complex z as well).

The algorithm can also be used to sample from a polymer model. We can define a probability measure $\nu_{G,z}$ on $\mathcal{G}(G)$:

$$\nu_{G,z}(\Gamma) = \frac{\prod_{\gamma \in \Gamma} z^{|\gamma|}}{\Xi(G, z)}, \quad \Gamma \in \mathcal{G}(G).$$

Theorem 7 ([20], Theorem 5.1). *Under the conditions of Theorem 6, for all $0 < z < \delta$ there is an efficient sampling algorithm for $\nu_{G,z}$ for all $G \in \mathfrak{G}$.*

The algorithm of Theorem 6 is based on truncating the Taylor series of $\log \Xi(G, z)$ around $z = 0$ (the approach of Barvinok [3] for a broad range of approximation problems). As the cluster expansion and Taylor series for $\log \Xi$ are the same power series, just organized differently, the algorithm computes the low-order coefficients of the Taylor series via the cluster expansion (see also [27] for a different approach to computing coefficients of partition functions of bounded-degree graphs).

To prove Theorems 1 and 3 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, can be approximated well by sums of partition functions of abstract polymer models. We then verify conditions (i) and (ii) of Theorem 6, which is straightforward.

Condition (iii) is non-trivial, and we verify it by showing the Kotecký-Preiss condition (1) holds at sufficiently high fugacity and low temperature.

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 [20]. 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 8. *Let Z be a real number. We call \hat{Z} an ε -relative approximation to Z if*

$$(1 - \varepsilon)\hat{Z} \leq Z \leq (1 + \varepsilon)\hat{Z}.$$

Given a graph $G = (V, E)$ and a set $S \subseteq V$, we let $N(S)$ denote the set of vertices in S^c adjacent to a vertex in S and we let $e(S, S^c)$ denote the number of edges in G with one endpoint in S and the other in S^c . Recall our notion of expansion.

Definition 9. A graph G is an α -expander if $e(S, S^c) \geq \alpha|S|$ 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 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 10. For $\beta > 2 \log(eq)/\alpha$, $Z_G^*(\beta)$ is an e^{-n} -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 $e(S_j, S_j^c) \geq \alpha|S_j|$. 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

$$\frac{1}{2} \sum_{i=1}^q e(S_i, S_i^c) \geq \frac{1}{2} \sum_{i=1}^q \alpha|S_i| = \frac{\alpha n}{2}$$

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 $Z_G^r(\beta)$.

For $\omega \in \Omega_r^n$, let

$$\Gamma(\omega) := \{v \in V(G) : \omega(v) \neq r\}.$$

We call a connected component of the induced subgraph $G[\Gamma(\omega)]$, along with its coloring, a *polymer* of ω . Note that a coloring $\omega \in \Omega_r^n$ is uniquely defined by its collection of polymers. These polymers constitute a polymer model with the following surface energy function:

$$\|\gamma\| := \sum_{\substack{\{i,j\} \in E(G): \\ \{i,j\} \cap \bar{\gamma} \neq \emptyset}} (1 - \delta_{\omega(i), \omega(j)}).$$

Intuitively, the polymers describe the way in which the coloring ω deviates from the ground state coloring in which all vertices are red. The surface energy $\|\gamma\|$ is the number of bichromatic edges incurred by the vertices in γ . Recall that we say that two polymers γ_1, γ_2 are *compatible* if $d(\bar{\gamma}_1, \bar{\gamma}_2) > 1$ where $d(\cdot, \cdot)$ denotes graph distance. We let $\mathcal{C} = \mathcal{C}(G)$ denote the set of polymers and we let $\mathcal{G} = \mathcal{G}(G)$ denote the family of all sets of mutually compatible polymers. Recall that the partition function of the polymer model is defined as

$$\Xi^r(G, z) := \sum_{\Gamma \in \mathcal{G}} \prod_{\gamma \in \Gamma} z^{\|\gamma\|},$$

where we use the superscript r to denote the ground state color.

Our aim is to verify conditions (i)–(iii) of Theorem 6 for this polymer model in order to obtain an FPTAS for $\Xi^r(G, z)$ when z is sufficiently small. Via a connection between $\Xi^r(G, z)$ and $Z_G^r(\beta)$, this will furnish us with an FPTAS for $Z_G^r(\beta)$ and hence also $Z_{G,q}(\beta)$. This connection is as follows. Letting

$$\bar{\mathcal{G}} := \left\{ \Gamma \in \mathcal{G} : \sum_{\gamma \in \Gamma} |\bar{\gamma}| < \frac{n}{2} \right\}$$

we have the following alternative representation of $Z_G^r(\beta)$.

$$Z_G^r(\beta) = e^{\beta e(G)} \sum_{\Gamma \in \bar{\mathcal{G}}} \prod_{\gamma \in \Gamma} e^{-\beta \|\gamma\|}.$$

In this form $Z_G^r(\beta)$ resembles a scaling of $\Xi^r(G, e^{-\beta})$ except for the fact that we have the global constraint that the total size of the polymers in a configuration must not exceed $n/2$. However, for β large this constraint has little effect. Indeed we have the following.

Lemma 11. *For $\beta > 2 \log(eq)/\alpha$, $e^{\beta e(G)} \cdot \Xi^r(G, e^{-\beta})$ is an e^{-n} -relative approximation to $Z_G^r(\beta)$.*

Before proving Lemma 11 we need the following lemma which establishes a lower bound on the Peierls' constant of our polymer model and verifies condition (i) of Theorem 6.

Lemma 12. *Let $\gamma \in \mathcal{C}$, then*

$$\alpha|\bar{\gamma}| \leq \|\gamma\| \leq \Delta|\bar{\gamma}|.$$

Proof. Note that since red is the majority color in ω , all polymers of ω have size at most $n/2$. Since G is an α -expander it follows that

$$e(\bar{\gamma}, \bar{\gamma}^c) \geq \alpha|\bar{\gamma}|.$$

Since γ is a connected component of non-red vertices, it follows that all edges lying between $\bar{\gamma}$ and $\bar{\gamma}^c$ are bichromatic. This gives the lower bound. Since G has maximum degree at most Δ , there are at most $\Delta|\bar{\gamma}|$ edges incident to $\bar{\gamma}$. The upper bound follows. \square

Proof of Lemma 11. Let $\Gamma \in \mathcal{G} \setminus \bar{\mathcal{G}}$. It follows that $\Gamma = \{\gamma_1, \dots, \gamma_k\}$ where $\sum_i |\bar{\gamma}_i| > n/2$. By Lemma 12, $\|\gamma_i\| \geq \alpha|\bar{\gamma}_i|$ for all i . It follows that

$$e^{\beta e(G)} \cdot \Xi^r(G, e^{-\beta}) - Z_\beta^r(G) = e^{\beta e(G)} \sum_{\Gamma \in \mathcal{G} \setminus \bar{\mathcal{G}}} \prod_{\gamma \in \Gamma} e^{-\beta \|\gamma\|} \leq e^{\beta e(G)} q^n e^{-\beta \alpha n/2},$$

where we have used the crude bound $|\mathcal{G} \setminus \bar{\mathcal{G}}| \leq |\mathcal{G}| = q^n$. It follows that

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

where for the second inequality we use the trivial lower bound $\Xi^r(G, e^{-\beta}) \geq 1$. \square

It remains to verify conditions (ii) and (iii) of Theorem 6 for our polymer model in order to obtain an FPTAS for $\Xi^r(G, z)$. Verifying condition (ii) is essentially immediate. Given $\bar{\gamma}, \omega_{\bar{\gamma}}$, determining whether $\gamma \in \mathcal{C}$ amounts to checking whether $G[\bar{\gamma}]$ is connected and consists of non-red vertices. This can be done in $O(|\bar{\gamma}|)$ time by a depth-first search algorithm. There are at most $\Delta \cdot |\bar{\gamma}|$ edges incident to $\bar{\gamma}$ and so computing $\|\gamma\|$ can also be done in $O(|\bar{\gamma}|)$ time.

We now turn our attention to verifying (iii), i.e. that there exists $\delta > 0$ so that for all complex $|z| < \delta$ we have $\Xi^r(G, z) \neq 0$.

3.2 Verifying the Kotecký-Preiss condition

In this section we will show that $\Xi^r(G, z)$ is zero-free in a complex disc containing the origin. By Theorem 5, it suffices to show there exists $\delta > 0$ such that

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

for all $|z| < \delta$ and all $\gamma \in \mathcal{C}(G)$.

Note that we will choose $\delta < 1$ and so by Lemma 12 it suffices to show that

$$\sum_{\gamma': d(\bar{\gamma}', \bar{\gamma}) \leq 1} e^{(1+\alpha \log |z|)|\bar{\gamma}'|} \leq |\bar{\gamma}|.$$

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

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

then by summing this inequality over all $v \in \bar{\gamma} \cup N(\bar{\gamma})$ (noting that $|\bar{\gamma} \cup N(\bar{\gamma})| \leq (\Delta + 1)|\bar{\gamma}|$), we would be done.

In order to establish (2) we borrow the following lemma

Lemma 13 ([14], 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 supports $\bar{\gamma}$ of polymers on t vertices that contain a given vertex v is bounded by $(e\Delta)^t$ and so the total number of polymers γ on t vertices containing v is bounded by $(eq\Delta)^t$. We thus have

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

for $|z| < (15q\Delta^2)^{-1/\alpha}$.

This verifies condition (iii) of Theorem 6 with $\delta = (15q\Delta^2)^{-1/\alpha}$ and so Theorem 6 gives an FPTAS for $\Xi^r(G, z)$ for all complex $|z| < (15q\Delta^2)^{-1/\alpha}$.

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 $O(nq^n)$ 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 6, we can obtain Z_{alg} , an $\varepsilon/3$ -relative approximation to $e^{\beta e(G)} \cdot \Xi^r(G, e^{-\beta})$ in time polynomial in n and $1/\varepsilon$. By Lemma 11, $e^{\beta e(G)} \cdot \Xi^r(G, e^{-\beta})$ is an e^{-n} -relative approximation to $Z_G^r(\beta)$ and so $qe^{\beta e(G)} \cdot \Xi^r(G, e^{-\beta})$ is an e^{-n} -relative approximation to $Z_G^*(\beta)$. By Lemma 10, it follows that $qe^{\beta e(G)} \cdot \Xi^r(G, e^{-\beta})$ is a $\varepsilon/3$ -relative approximation to $Z_{G,q}(\beta)$ and so qZ_{alg} is an ε -relative approximation to $Z_{G,q}(\beta)$ as required.

For the approximate sampling algorithm, we will apply Theorem 7.

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

$$\nu_{G,\beta}^r(\Gamma) = \frac{\prod_{\gamma \in \Gamma} e^{-\beta \|\gamma\|}}{\Xi^r(G, e^{-\beta})}.$$

Then convert Γ into a q -coloring $\omega \in \Omega^n$ by setting $\omega(v) = \omega_{\bar{\gamma}}(v)$ if $v \in \bar{\gamma}$ for some $\gamma \in \Gamma$, and $\omega(v) = r$ otherwise. The resulting distribution of ω is $\hat{\mu}$.

By Lemmas 10 and 11,

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

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 simply appeal to Theorem 7. □

3.4 Proof of Corollary 4

To prove Corollary 4 for the random regular graph, we need the following fact about its expansion properties.

Theorem 14 (Bollobás [6]). *Let $\Delta \geq 3$, and let $\eta \in (0, 1)$ be such that*

$$(1 - \eta) \log_2(1 - \eta) + (1 + \eta) \log_2(1 + \eta) > 4/\Delta$$

then almost every Δ -regular graph is a $(1 - \eta)\Delta/2$ -expander.

Proof of Corollary 4. By Theorem 14, almost every Δ -regular graph is a $c\Delta$ -expander for some absolute constant c . The result follows by Theorem 3. □

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 combination of polymer models, and then we verify the conditions of Theorem 6 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 15. *A bipartite graph $G = (\mathcal{O}, \mathcal{E}, E)$ is a bipartite α -expander if $|N(S)| \geq (1 + \alpha)|S|$ for all $S \subseteq \mathcal{O}$ with $|S| \leq |\mathcal{O}|/2$ and all $S \subseteq \mathcal{E}$ with $|S| \leq |\mathcal{E}|/2$.*

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} .

We first show that the hard-core model partition function $Z_G(\lambda)$ is dominated by independent sets that occupy more than three quarters of one of the partition classes. To make this precise we make the following definitions. Recall that $\mathcal{I} := \mathcal{I}(G)$ denotes the family of all independent sets in G . Let

$$\mathcal{I}_{\mathcal{O}} := \left\{ I \in \mathcal{I} : |I \cap \mathcal{O}| > \frac{3}{4}|\mathcal{O}| \right\}$$

and let

$$Z_{\mathcal{O}}(\lambda) := \sum_{I \in \mathcal{I}_{\mathcal{O}}} \lambda^{|I|}.$$

Define $\mathcal{I}_{\mathcal{E}}$ and $Z_{\mathcal{E}}(\lambda)$ similarly. Note that $\mathcal{I}_{\mathcal{O}}$ and $\mathcal{I}_{\mathcal{E}}$ are disjoint since, supposing that $|\mathcal{O}| \geq |\mathcal{E}|$, if $|I \cap \mathcal{O}| > 3|\mathcal{O}|/4$ then certainly $|N(I \cap \mathcal{O})| > |\mathcal{E}|/4$ and each element of $N(I \cap \mathcal{O}) \subseteq \mathcal{E}$ must be unoccupied (not in the independent set). Let $\mathcal{I}^* := \mathcal{I}_{\mathcal{O}} \cup \mathcal{I}_{\mathcal{E}}$ and set

$$Z_G^*(\lambda) := \sum_{I \in \mathcal{I}^*} \lambda^{|I|} = Z_{\mathcal{O}}(\lambda) + Z_{\mathcal{E}}(\lambda).$$

Lemma 16. *For $\lambda \geq (2e)^{40/\alpha}$, $Z_G^*(\lambda)$ is an e^{-n} -relative approximation to $Z_G(\lambda)$.*

Proof. Let $I \notin \mathcal{I}^*$, then by definition $|I \cap \mathcal{O}| \leq 3|\mathcal{O}|/4$ and $|I \cap \mathcal{E}| \leq 3|\mathcal{E}|/4$. Suppose without loss of generality that $|\mathcal{O}| \geq |\mathcal{E}|$. We consider two cases depending on the size of $I \cap \mathcal{O}$.

Suppose first that $|I \cap \mathcal{O}| \leq |\mathcal{O}|/2$. In this case, since G is a bipartite α -expander, we have

$$|N(I \cap \mathcal{O})| \geq (1 + \alpha)|I \cap \mathcal{O}|$$

and so

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

since each element of $N(I \cap \mathcal{O})$ must be unoccupied. It follows that

$$|I| \leq |\mathcal{E}| - \alpha|I \cap \mathcal{O}| \leq |\mathcal{O}| - \alpha|I \cap \mathcal{O}|.$$

If $|I \cap \mathcal{O}| \geq |\mathcal{O}|/5$, then by the above we would have

$$|I| \leq \left(1 - \frac{\alpha}{5}\right)|\mathcal{O}|.$$

If instead, $|I \cap \mathcal{O}| < |\mathcal{O}|/5$ then, recalling that $|I \cap \mathcal{E}| \leq 3|\mathcal{E}|/4 \leq 3|\mathcal{O}|/4$, we have

$$|I| \leq \frac{19}{20}|\mathcal{O}|.$$

We deduce that

$$|I| \leq |\mathcal{O}| \cdot \max \left\{ 1 - \frac{\alpha}{5}, \frac{19}{20} \right\}. \quad (3)$$

Suppose now that $|I \cap \mathcal{O}| > |\mathcal{O}|/2$. In this case we certainly have $|N(I \cap \mathcal{O})| > |\mathcal{E}|/2$ and so $|I \cap \mathcal{E}| < |\mathcal{E}|/2$. A similar argument to the one above, this time using the expansion of the set $I \cap \mathcal{E}$, again yields the bound (3).

Letting $c_\alpha := \max\{1 - \alpha/5, 19/20\}$ and using the crude bounds $|\mathcal{I} \setminus \mathcal{I}^*| \leq 2^n$ and $Z_G(\lambda) \geq \lambda^{|\mathcal{O}|}$, it follows that

$$Z_G(\lambda) - Z_G^*(\lambda) = \sum_{I \notin \mathcal{I}^*} \lambda^{|I|} \leq 2^n \cdot \lambda^{c_\alpha |\mathcal{O}|}$$

and so

$$\left| 1 - \frac{Z_G^*(\lambda)}{Z_G(\lambda)} \right| < 2^n \cdot \lambda^{(c_\alpha-1)|\mathcal{O}|} \leq 2^n \cdot \lambda^{(c_\alpha-1)n/2} \leq e^{-n}.$$

□

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

Given $I \in \mathcal{I}_{\mathcal{O}}$, let

$$\Gamma(I) := (I^c \cap \mathcal{O}) \cup (I \cap \mathcal{E}).$$

Note that I is uniquely determined by the set $\Gamma(I)$. We think of elements of $\Gamma(I)$ as ‘bad’ vertices which deviate from the maximal independent set \mathcal{O} . The *polymers* of an independent set $I \in \mathcal{I}_{\mathcal{O}}$ are defined to be the connected components $\bar{\gamma}$ of the induced subgraph $G[\Gamma(I)]$ together with an assignment of spin ‘0’ to the vertices of $|\bar{\gamma} \cap \mathcal{O}|$ and spin ‘1’ to the vertices of $|\bar{\gamma} \cap \mathcal{E}|$. Note that, unlike in the previous section, a polymer γ is uniquely defined by its support and so for notational convenience we will simply identify a polymer γ with its support $\bar{\gamma}$.

These polymers constitute a polymer model with the following surface energy function:

$$\|\gamma\| := |\gamma \cap \mathcal{O}| - |\gamma \cap \mathcal{E}|.$$

We let $\mathcal{C}^{\mathcal{O}} = \mathcal{C}^{\mathcal{O}}(G)$ denote the set of possible polymers associated to elements of $\mathcal{I}_{\mathcal{O}}$ and let $\mathcal{G}^{\mathcal{O}} = \mathcal{G}^{\mathcal{O}}(G)$ the family of all sets of mutually compatible polymers from $\mathcal{C}^{\mathcal{O}}$. The polymer model partition function is given by

$$\Xi^{\mathcal{O}}(G, z) := \sum_{\Gamma \in \mathcal{G}^{\mathcal{O}}} \prod_{\gamma \in \Gamma} z^{|\gamma|}.$$

Again our aim is to verify conditions (i)–(iii) of Theorem 6 for this polymer model in order to obtain an FPTAS for $\Xi^{\mathcal{O}}(G, z)$ when $|z|$ is sufficiently small. The following connection between $\Xi^{\mathcal{O}}(G, z)$ and $Z_{\mathcal{O}}(\lambda)$, will furnish us with an FPTAS for $Z_{\mathcal{O}}(\lambda)$ and hence also $Z_G(\lambda)$. Letting

$$\overline{\mathcal{G}^{\mathcal{O}}} := \left\{ \Gamma \in \mathcal{G}^{\mathcal{O}} : \sum_{\gamma \in \Gamma} |\gamma \cap \mathcal{O}| \leq \frac{|\mathcal{O}|}{4} \right\}.$$

we have the following alternative representation of $Z_{\mathcal{O}}(\lambda)$.

$$Z_{\mathcal{O}}(\lambda) = \lambda^{|\mathcal{O}|} \sum_{\Gamma \in \overline{\mathcal{G}^{\mathcal{O}}}} \prod_{\gamma \in \Gamma} \lambda^{-\|\gamma\|}.$$

In this form $Z_{\mathcal{O}}(\lambda)$ resembles a scaling of $\Xi^{\mathcal{O}}(G, 1/\lambda)$ except for the global size constraint on the polymers. However, for λ large the effect of this constraint is small.

Lemma 17. For $\lambda > (2e)^{\frac{8n}{\alpha|\mathcal{O}|}}$, $\lambda^{|\mathcal{O}|}\Xi^{\mathcal{O}}(G, 1/\lambda)$ is an e^{-n} -relative approximation to $Z_{\mathcal{O}}(\lambda)$.

Before proving Lemma 17 we need the following lemma which establishes a lower bound on the Peierls' constant of our polymer model and verifies condition (i) of Theorem 6.

Lemma 18. For each polymer $\gamma \in \mathcal{C}$,

$$\frac{\alpha}{1+\alpha} \cdot |\gamma \cap \mathcal{O}| \leq \|\gamma\| \leq |\gamma|.$$

Proof. First note that when $\gamma \cap \mathcal{E} = \emptyset$ i.e. γ is a single vertex in \mathcal{O} , the result is immediate. We may therefore assume that $\gamma \cap \mathcal{E} \neq \emptyset$.

Suppose that $|\gamma \cap \mathcal{E}| > |\mathcal{E}|/2$. Then since G is a bipartite α -expander $|N(\gamma \cap \mathcal{E})| \geq (1+\alpha)|\mathcal{E}|/2$. However, by our definition of an expander, we must also have $|\mathcal{E}| \geq (1+\alpha)|\mathcal{O}|/2$. It follows that $|N(\gamma \cap \mathcal{E})| > |\mathcal{O}|/4$. Since each element of $N(\gamma \cap \mathcal{E})$ must be unoccupied this contradicts the fact that γ is a polymer of an element of $\mathcal{I}_{\mathcal{O}}$.

We deduce that $|\gamma \cap \mathcal{E}| \leq |\mathcal{E}|/2$. Note that since γ is connected $\gamma \cap \mathcal{O} = N(\gamma \cap \mathcal{E})$. Since G is a bipartite α -expander we then have

$$|\gamma \cap \mathcal{O}| = |N(\gamma \cap \mathcal{E})| = (1+\alpha')|\gamma \cap \mathcal{E}|, \quad (4)$$

where $\alpha' \geq \alpha$. It follows that

$$\|\gamma\| \geq \alpha'|\gamma \cap \mathcal{E}| \geq \frac{\alpha'}{1+\alpha'}|\gamma \cap \mathcal{O}| \geq \frac{\alpha}{1+\alpha}|\gamma \cap \mathcal{O}|.$$

The upper bound is immediate. □

Proof of Lemma 17. Let $\Gamma \in \mathcal{G}^{\mathcal{O}} \setminus \overline{\mathcal{G}^{\mathcal{O}}}$. It follows by definition that $\sum_{\gamma \in \Gamma} |\gamma \cap \mathcal{O}| > |\mathcal{O}|/4$. Using the crude bound $|\mathcal{G}^{\mathcal{O}} \setminus \overline{\mathcal{G}^{\mathcal{O}}}| \leq |\mathcal{G}^{\mathcal{O}}| \leq 2^n$ and Lemma 18, we then have

$$\lambda^{|\mathcal{O}|}\Xi^{\mathcal{O}}(G, 1/\lambda) - Z_{\mathcal{O}}(\lambda) = \lambda^{|\mathcal{O}|} \sum_{\Gamma \in \mathcal{G}^{\mathcal{O}} \setminus \overline{\mathcal{G}^{\mathcal{O}}}} \prod_{\gamma \in \Gamma} \lambda^{-\|\gamma\|} \leq \lambda^{|\mathcal{O}|} \cdot 2^n \cdot \lambda^{-\alpha|\mathcal{O}|/4(1+\alpha)},$$

It follows, using the bound $\Xi^{\mathcal{O}}(G, 1/\lambda) \geq 1$, that

$$\left| 1 - \frac{Z_{\mathcal{O}}(\lambda)}{\lambda^{|\mathcal{O}|}\Xi^{\mathcal{O}}(G, 1/\lambda)} \right| \leq 2^n \cdot \lambda^{-\alpha|\mathcal{O}|/4(1+\alpha)} \leq e^{-n}.$$

In the final inequality we have used the fact that $\alpha \leq 1$. □

It remains to verify conditions (ii) and (iii) of Theorem 6 for the polymer model. Verifying condition (ii) is straightforward. Given a subset $\gamma \subseteq V(G)$ checking whether $G[\gamma]$ is connected and calculating $\|\gamma\| = |\gamma \cap \mathcal{O}| - |\gamma \cap \mathcal{E}|$ can be done in time linear in $|\gamma|$. We now turn our attention to verifying (iii).

4.2 Verifying the Kotecký-Preiss condition

In this section we will show that $\Xi^{\mathcal{O}}(G, z)$ is zero-free in a complex disc containing the origin. Again, using Theorem 5 our task is to show that there exists $\delta > 0$ such that

$$\sum_{\gamma': d(\gamma', \gamma) \leq 1} e^{|\gamma'| |z|^{\|\gamma'\|}} \leq |\gamma|, \quad (\dagger)$$

for all $|z| < \delta$ and all $\gamma \in \mathcal{C}$.

Note that if $d(\gamma', \gamma) \leq 1$ and $|\gamma'| = 1$, then γ' must be a vertex of $\gamma \cap \mathcal{O}$ and so

$$\sum_{\gamma': d(\gamma', \gamma) \leq 1} e^{|\gamma'| |z|^{\|\gamma'\|}} = |\gamma \cap \mathcal{O}| \cdot e \cdot |z| + \sum_{\substack{\gamma': d(\gamma', \gamma) \leq 1 \\ |\gamma'| > 1}} e^{|\gamma'| |z|^{\|\gamma'\|}}.$$

If we choose $|z| < e^{-1}$ then in order to verify (\dagger) it suffices to verify

$$\sum_{\substack{\gamma': d(\gamma', \gamma) \leq 1 \\ |\gamma'| > 1}} e^{|\gamma'| |z|^{\|\gamma'\|}} \leq |\gamma \cap \mathcal{E}|.$$

Now, by (4), for some $\alpha' \geq \alpha$ we have

$$e^{|\gamma'| |z|^{\|\gamma'\|}} = e^{(2+\alpha')|\gamma' \cap \mathcal{E}|} |z|^{\alpha' |\gamma' \cap \mathcal{E}|} \leq e^{(2+\alpha)|\gamma' \cap \mathcal{E}|} |z|^{\alpha |\gamma' \cap \mathcal{E}|},$$

It therefore suffices to show

$$\sum_{\substack{\gamma': d(\gamma', \gamma) \leq 1 \\ |\gamma'| > 1}} e^{(2+\alpha)|\gamma' \cap \mathcal{E}|} |z|^{\alpha |\gamma' \cap \mathcal{E}|} \leq |\gamma \cap \mathcal{E}|.$$

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

$$\sum_{\substack{\gamma': \gamma' \ni v \\ |\gamma'| > 1}} e^{(2+\alpha)|\gamma' \cap \mathcal{E}|} |z|^{\alpha |\gamma' \cap \mathcal{E}|} \leq \frac{1}{(\Delta + 1)^2}, \quad (5)$$

then by summing this inequality over all $v \in \gamma \cup N(\gamma)$ (noting that $|\gamma \cup N(\gamma)| \leq (\Delta + 1)^2 |\gamma \cap \mathcal{E}|$), we would be done.

In order to establish (5), consider the auxiliary graph H on vertex set \mathcal{E} where we join two vertices if they have a common neighbor in G . For a subset $S \subseteq \mathcal{E}$, $H[S]$ is connected if and only if $G[S \cup N(S)]$ is connected. Let us call a polymer with more than one vertex *non-trivial*. A non-trivial polymer γ is determined by its intersection with \mathcal{E} since $\gamma \cap \mathcal{O} = N(\gamma \cap \mathcal{E})$. It follows that there is a one-to-one correspondence between non-trivial polymers and subsets $S \subseteq \mathcal{E}$ for which $H[S]$ is connected. Note also that H has maximum degree at most Δ^2 . It follows from Lemma 13 that the number of polymers γ' containing a vertex $v \in \mathcal{E}$ with $|\gamma' \cap \mathcal{E}| = t$ is at most $(e\Delta^2)^t$. If $v \in \mathcal{O}$, then by considering the polymers containing one of its neighbors we see that the number of polymers γ' containing v with $|\gamma' \cap \mathcal{E}| = t$ is at most $\Delta(e\Delta^2)^t$. We thus have

$$\sum_{\gamma': \gamma' \ni v} e^{(2+\alpha)|\gamma' \cap \mathcal{E}|} |z|^{|\gamma' \cap \mathcal{E}|} \leq \Delta \sum_{t=2}^{\infty} (\Delta^2 e^{3+\alpha} |z|^\alpha)^t \leq \frac{1}{(\Delta+1)^2},$$

when $|z| < e^{-1}(4e^3\Delta^5)^{-1/\alpha}$.

4.3 Proof of Theorem 1

We consider two cases separately. If $\varepsilon < e^{-n/2}$, 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 polynomial in $1/\varepsilon$.

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

$$\lambda > \max \left\{ e(4e^3\Delta^5)^{1/\alpha}, (2e)^{40/\alpha} \right\}.$$

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 $\Xi^{\mathcal{O}}(G, 1/\lambda)$ given by Theorem 6, we may find $Z_{\mathcal{O}}^{\text{alg}}$, an $\varepsilon/7$ -relative approximation to $\lambda^{|\mathcal{O}|} \Xi^{\mathcal{O}}(G, 1/\lambda)$ in time polynomial in n and $1/\varepsilon$. By Lemma 17, $\lambda^{|\mathcal{O}|} \Xi^{\mathcal{O}}(G, 1/\lambda)$ is an $\varepsilon/7$ -relative approximation to $Z_{\mathcal{O}}(\lambda)$ and so $Z_{\mathcal{O}}^{\text{alg}}$ is an $\varepsilon/3$ -relative approximation to $Z_{\mathcal{O}}(\lambda)$.

In identical fashion (noting that $n/|\mathcal{E}| \leq 3$ for the ‘ \mathcal{E} -version’ of Lemma 17) we may find $Z_{\mathcal{E}}^{\text{alg}}$, an $\varepsilon/3$ -relative approximation to $Z_{\mathcal{E}}(\lambda)$ in time polynomial in n and $1/\varepsilon$. It follows that $Z^{\text{alg}} := Z_{\mathcal{O}}^{\text{alg}} + Z_{\mathcal{E}}^{\text{alg}}$ is an $\varepsilon/3$ -relative approximation to $Z_G^*(\lambda)$. By Lemma 16, $Z_G^*(\lambda)$ is an $\varepsilon/3$ -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. Consider the distribution $\hat{\mu}$ on $\mathcal{I}(G)$ defined by choosing \mathcal{O} or \mathcal{E} with probability proportional to $\Xi^{\mathcal{O}}(G, 1/\lambda)$ and $\Xi^{\mathcal{E}}(G, 1/\lambda)$ respectively. Then, supposing we chose \mathcal{O} , sample Γ from the measure

$$\nu_{G, 1/\lambda}^{\mathcal{O}}(\Gamma) = \frac{\prod_{\gamma \in \Gamma} \lambda^{-\|\gamma\|}}{\Xi^{\mathcal{O}}(G, 1/\lambda)}$$

and obtain the independent set I corresponding to the set of contours Γ . The distribution of I is $\hat{\mu}$.

By Lemmas 16 and 17 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 $\nu_{G, 1/\lambda}^{\mathcal{O}}$ and $\nu_{G, 1/\lambda}^{\mathcal{E}}$ in time polynomial in n and $1/\varepsilon$. This is provided by Theorem 7. □

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 19. *A bipartite graph $G = (\mathcal{O}, \mathcal{E}, E)$ is a bipartite (σ, ρ) -expander if $|N(S)| \geq \rho|S|$ 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 20 (Bassalygo [4]). *Almost every Δ -regular bipartite graph is an (σ, ρ) -expander provided*

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

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

Proof of Theorem 2. Note that Theorem 20 implies that almost every Δ -regular bipartite graph is a bipartite 0.16-expander and so the conclusion of Theorem 2 holds for all $\lambda > C\Delta^{32}$ for some absolute constant C (we may therefore assume that $\Delta \geq 4$ in the following). In order to remove the dependence on Δ we exploit the fact that small subsets expand a lot more in the random regular bipartite graph. Indeed, Theorem 20 shows that almost every Δ -regular bipartite graph is also a $(1/\Delta, \Delta/3)$ -expander.

Suppose that in the proof of Theorem 1 we also knew that $G = (\mathcal{O}, \mathcal{E}, E)$ was a $(1/\Delta, \Delta/3)$ -expander and was balanced (i.e. $|\mathcal{O}| = |\mathcal{E}| = m = n/2$ as in a regular bipartite graph).

The key observation is as follows. If γ is a polymer of an independent set $I \in \mathcal{I}_{\mathcal{O}}$ then we must have $|\gamma \cap \mathcal{E}| \leq m/\Delta$ else, by the $(1/\Delta, \Delta/3)$ -expansion property, we would have $|N(\gamma \cap \mathcal{E})| \geq m/3$ contradicting the fact that $|I \cap \mathcal{O}| > 3m/4$ (since each element of $N(\gamma \cap \mathcal{E})$ must be unoccupied). By $(1/\Delta, \Delta/3)$ -expansion we then have

$$|\gamma \cap \mathcal{O}| = |N(\gamma \cap \mathcal{E})| \geq \frac{\Delta}{3} |\gamma \cap \mathcal{E}|,$$

and so

$$\|\gamma\| \geq \left(\frac{\Delta}{3} - 1\right) |\gamma \cap \mathcal{E}|.$$

It follows that for such G , we may replace every instance of α in Subsection 4.2 by $\Delta/3 - 1$. Completing the proof as before, we see that Theorem 2 holds for

$$\lambda > \max \left\{ e(4e^3 \Delta^5)^{\frac{1}{\Delta/3-1}}, (2e)^{40/0.16} \right\} = (2e)^{250}.$$

□

5 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.

Using contour-based techniques Galvin and Tetali [16] showed slow mixing of the Glauber dynamics on Δ -regular bipartite expander graphs. Somewhat counterintuitively, the proof of slow mixing via contour methods is related to verifying the Kotecký-Preiss condition. The expansion condition is weaker and the value of λ smaller in [16] than in Theorem 1 and so one could hope to devise efficient counting and sampling algorithms matching the class of graphs and range of parameters from [16]. We leave for future work the task of adapting some of their ideas and techniques to algorithms.

Moreover, one could hope to find efficient algorithms for other statistical physics models that have been proved to exhibit phase coexistence via contour arguments (e.g. the q -coloring model on a bipartite graph [15, 28, 29]). We state a conjecture in this direction on proper colorings in random regular bipartite graphs, but one could also make conjectures along the same lines for all bipartite expander graphs or for the anti-ferromagnetic Potts model.

Conjecture 21. *For all $q \geq 3$ there is Δ large enough so that there exist an FPTAS and efficient sampling algorithm for proper q -colorings on almost every Δ -regular bipartite graph.*

One could also hope to improve the running time of the algorithms presented here. The 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 [34]. 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 [20] 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.

References

- [1] 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.
- [2] A. Barvinok. Computing the permanent of (some) complex matrices. *Foundations of Computational Mathematics*, 16(2):329–342, 2016.

- [3] A. Barvinok. *Combinatorics and Complexity of Partition Functions*, volume 30 of *Algorithms and Combinatorics*. Springer, 2017.
- [4] L. Bassalygo. Asymptotically optimal switching circuits. *Problems of Information Transmission*, 17(3):206–211, 1981.
- [5] 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. *arXiv preprint arXiv:1804.08111*, 2018.
- [6] B. Bollobás. The isoperimetric number of random regular graphs. *European Journal of Combinatorics*, 9(3):241–244, 1988.
- [7] 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.
- [8] P. Chebolu, L. A. Goldberg, and R. Martin. The complexity of approximately counting stable matchings. *Theoretical Computer Science*, 437:35–68, 2012.
- [9] 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.
- [10] M. Dyer, L. A. Goldberg, C. Greenhill, and M. Jerrum. The relative complexity of approximate counting problems. *Algorithmica*, 38(3):471–500, 2004.
- [11] C. Efthymiou, T. P. Hayes, D. Štefankovic, 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.
- [12] A. Galanis, D. Štefankovič, 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.
- [13] A. Galanis, D. Stefankovic, E. Vigoda, and L. Yang. Ferromagnetic Potts model: Refined #-BIS-hardness and related results. *SIAM Journal on Computing*, 45(6):2004–2065, 2016.
- [14] 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.
- [15] D. Galvin, J. Kahn, D. Randall, and G. B. Sorkin. Phase coexistence and torpid mixing in the 3-coloring model on \mathbf{Z}^d . *SIAM Journal on Discrete Mathematics*, 29(3):1223–1244, 2015.

- [16] D. Galvin and P. Tetali. Slow mixing of Glauber dynamics for the hard-core model on regular bipartite graphs. *Random Structures & Algorithms*, 28(4):427–443, 2006.
- [17] L. A. Goldberg and M. Jerrum. Approximating the partition function of the ferromagnetic Potts model. *Journal of the ACM (JACM)*, 59(5):25, 2012.
- [18] C. Gruber and H. Kunz. General properties of polymer systems. *Communications in Mathematical Physics*, 22(2):133–161, 1971.
- [19] H. Guo and M. Jerrum. Random cluster dynamics for the Ising model is rapidly mixing. *Ann. Appl. Probab.*, 28(2):1292–1313, 04 2018.
- [20] T. Helmuth, W. Perkins, and G. Regts. Algorithmic Pirogov-Sinai theory. *arXiv preprint arXiv:1806.11548*, 2018.
- [21] M. Jerrum and A. Sinclair. Polynomial-time approximation algorithms for the Ising model. *SIAM Journal on computing*, 22(5):1087–1116, 1993.
- [22] F. P. Kelly. Stochastic models of computer communication systems. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 379–395, 1985.
- [23] R. Kotecký and D. Preiss. Cluster expansion for abstract polymer models. *Communications in Mathematical Physics*, 103(3):491–498, 1986.
- [24] 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.
- [25] 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.
- [26] 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.
- [27] 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.
- [28] R. Peled. High-dimensional Lipschitz functions are typically flat. *The Annals of Probability*, 45(3):1351–1447, 2017.
- [29] R. Peled and Y. Spinka. A condition for long-range order in discrete spin systems with application to the antiferromagnetic Potts model. *arXiv preprint arXiv:1712.03699*, 2017.
- [30] S. A. Pirogov and Y. G. Sinai. Phase diagrams of classical lattice systems. *Theoretical and Mathematical Physics*, 25(3):1185–1192, 1975.

- [31] 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.
- [32] 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.
- [33] A. Sly and N. Sun. Counting in two-spin models on d -regular graphs. *The Annals of Probability*, 42(6):2383–2416, 2014.
- [34] R. H. Swendsen and J.-S. Wang. Nonuniversal critical dynamics in Monte Carlo simulations. *Physical review letters*, 58(2):86, 1987.
- [35] 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.