

Fast Algorithms for General Spin Systems on Bipartite Expanders

Andreas Galanis

Department of Computer Science, University of Oxford, UK
andreas.galanis@cs.ox.ac.uk

Leslie Ann Goldberg

Department of Computer Science, University of Oxford, UK
leslie.goldberg@cs.ox.ac.uk

James Stewart

Department of Computer Science, University of Oxford, UK
james.stewart@cs.ox.ac.uk

Abstract

A spin system is a framework in which the vertices of a graph are assigned spins from a finite set. The interactions between neighbouring spins give rise to weights, so a spin assignment can also be viewed as a weighted graph homomorphism. The problem of approximating the partition function (the aggregate weight of spin assignments) or of sampling from the resulting probability distribution is typically intractable for general graphs.

In this work, we consider arbitrary spin systems on bipartite expander Δ -regular graphs, including the canonical class of bipartite random Δ -regular graphs. We develop fast approximate sampling and counting algorithms for general spin systems whenever the degree and the spectral gap of the graph are sufficiently large. Our approach generalises the techniques of Jenssen et al. and Chen et al. by showing that typical configurations on bipartite expanders correspond to “bicliques” of the spin system; then, using suitable polymer models, we show how to sample such configurations and approximate the partition function in $\tilde{O}(n^2)$ time, where n is the size of the graph.

2012 ACM Subject Classification Theory of computation \rightarrow Design and analysis of algorithms; Theory of computation \rightarrow Randomness, geometry and discrete structures; Mathematics of computing \rightarrow Discrete mathematics

Keywords and phrases bipartite expanders, approximate counting, spin systems

Digital Object Identifier 10.4230/LIPIcs.MFCS.2020.37

Related Version A full version with all proofs can be found at <https://arxiv.org/abs/2004.13442>. The theorem numbers here match those in the full version. Proof sketches here refer to lemmas in the full version without restating them in detail.

1 Introduction

Spin systems are general frameworks that encompass sampling and counting problems in computer science, graph homomorphism problems in combinatorics, and phase transition phenomena in statistical physics. In this paper, we provide algorithms for general spin systems on bounded-degree bipartite expander graphs.

A q -spin system is specified by a set of spins $[q] = \{1, 2, \dots, q\}$ and a symmetric interaction matrix $H \in \mathbb{R}_{\geq 0}^{q \times q}$. Given a graph $G = (V_G, E_G)$, a spin configuration is an assignment $\sigma : V_G \rightarrow [q]$ and the weight of σ is given by $w_{G,H}(\sigma) = \prod_{\{u,v\} \in E_G} H_{\sigma(u),\sigma(v)}$. The Gibbs distribution of the system, denoted by $\mu_{G,H}$, is a probability distribution on the set $\Sigma_{G,H}$ which denotes the set of all possible spin configurations, given by $\mu_{G,H}(\sigma) = \frac{w_{G,H}(\sigma)}{Z_{G,H}}$, where $Z_{G,H} := \sum_{\tau \in \Sigma_{G,H}} w_{G,H}(\tau)$ is the so-called partition function. Well-known examples of spin systems are the Ising/Potts models, where the matrix H has all diagonal entries equal to



© Andreas Galanis, Leslie Ann Goldberg, and James Stewart;
licensed under Creative Commons License CC-BY

45th International Symposium on Mathematical Foundations of Computer Science (MFCS 2020).

Editors: Javier Esparza and Daniel Král'; Article No. 37; pp. 37:1–37:14

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

some parameter $\beta > 0$ and off-diagonal entries equal to 1; the case $q = 2$ is the Ising model, and $q > 2$ is the Potts model. Apart from statistical physics systems, graph homomorphisms also fit naturally into this framework, whenever H has 0-1 entries.

Henceforth, we will normalise H so that its largest entry is equal to 1. More formally, we will consider δ -matrices, defined as follows.

► **Definition 1.** Let $q \geq 2$ and let $\delta \in (0, 1)$. A symmetric matrix $H \in \mathbb{R}_{\geq 0}^{q \times q}$ is called a δ -matrix if $\max_{i,j \in [q]} H_{i,j} = 1$ and $\max_{i,j \in [q]: H_{i,j} \neq 1} H_{i,j} \leq \delta$. ◀

Note that, apart from trivial cases¹, we can always normalise the interaction matrix of a spin system to satisfy Definition 1 for some $\delta \in (0, 1)$.

Approximately sampling from the Gibbs distribution of a spin system and approximating its partition function are well-studied problems in computer science, since they appear in various applications. However, even for the most canonical models, such as the Potts model or graph homomorphisms, these computational problems are hard in general, even on bounded-degree graphs [10, 4, 18, 17, 13, 16, 15].

In light of these hardness results, it is natural to consider whether efficient algorithms can be developed for more restricted classes of graphs. Recently, Jenssen, Keevash, and Perkins [23] (see also [21]) showed a new framework that such algorithms are possible for three canonical models (ferromagnetic Potts model, colourings, and independent sets) on bounded-degree expander graphs, in a range of parameters where the problems are otherwise hard for general bounded-degree graphs. In this paper, we show that this framework can further be used to obtain approximation algorithms of *any* spin system on *bipartite* expander graphs whenever the degree is sufficiently large.

More precisely, we will consider regular bipartite graphs whose second eigenvalue is bounded by a small constant [22]. Let G be an n -vertex Δ -regular bipartite graph. Let $\lambda_1(G) \geq \lambda_2(G) \geq \dots \geq \lambda_n(G)$ denote the eigenvalues of the adjacency matrix of G . It is well-known that $\lambda_1(G) = \Delta$ and $\lambda_n(G) = -\lambda_1(G)$. We define $\lambda(G) = \lambda_2(G)$.

► **Definition 2.** Let $\Delta \geq 3$ be an integer and λ be a positive real strictly less than Δ . We let $\mathcal{G}_{\Delta, \lambda}^{\text{bip}}$ denote the set of all connected Δ -regular bipartite graphs G , for which $\lambda(G) \leq \lambda$. ◀

One of the primary examples of bipartite expander graphs, and one of the main motivations behind this work, are random bipartite regular graphs. It is known [3] that, for any fixed $\varepsilon > 0$, with high probability² over the choice of a random bipartite Δ -regular graph G , it holds that $\lambda(G) \leq 2\sqrt{\Delta - 1} + \varepsilon$. From a counting/sampling perspective these graphs are particularly interesting since they have been key ingredients in obtaining inapproximability results [26, 5, 14]. Somewhat surprisingly, while we know constant factor estimates of the partition function via (non-algorithmic) probabilistic methods that hold with probability $1 - o(1)$ over the choice of the graph [14], it is not known how to approximately sample from the Gibbs distribution efficiently. In fact, even obtaining more refined estimates on the partition function is an open problem. As a corollary of our main result, we address both of these questions, provided that the degree Δ is sufficiently large relative to H .

¹ If H has all of its entries equal to some $c > 0$, then $Z_{G,H} = q^{|V_G|} c^{|E_G|}$ for any graph $G = (V_G, E_G)$. Similarly, if H is the all zeros matrix, then $Z_{G,H} = 0$ for any non-empty graph G . So suppose that H has at least two entries with distinct values and let $h_{\max} = \max_{i,j \in [q]} H_{i,j}$. Then, the matrix $H' = \frac{1}{h_{\max}} H$ is a δ -matrix, for any $\delta \in (0, 1)$ which is bigger than the second largest entry in H' . Moreover, for any graph $G = (V_G, E_G)$ we have that $Z_{G,H} = h_{\max}^{|E_G|} \cdot Z_{G,H'}$.

² Here and throughout the paper, “with high probability” means with probability tending to 1 as the size of the graph tends to infinity.

To formally state our results, we will need some definitions. First, the following computational problem that we will study.

Parameters: H , a symmetric matrix in $\mathbb{R}_{\geq 0}^{q \times q}$, integer $\Delta \geq 3$, and a real $\lambda \in (0, \Delta)$.

Name: $\text{SPIN}_{H, \Delta, \lambda}$.

Input: A graph $G \in \mathcal{G}_{\Delta, \lambda}^{\text{bip}}$.

Output: The value of $Z_{G, H}$.

In particular, we consider the problem of approximating $Z_{G, H}$ and sampling from $\mu_{G, H}$. Given an accuracy parameter $\varepsilon > 0$, we say that \hat{Z} is an ε -approximation to Z if $(1 - \varepsilon)Z \leq \hat{Z} \leq (1 + \varepsilon)Z$. For a distribution μ , we say that a random variable X is an ε -sample from μ if the total variation distance between the distribution of X and μ is at most ε . A *fully polynomial randomised approximation scheme* (FPRAS) for $\text{SPIN}_{H, \Delta, \lambda}$ is a randomised algorithm that, given $\varepsilon > 0$ and $G = (V, E) \in \mathcal{G}_{\Delta, \lambda}^{\text{bip}}$ as input, outputs a random variable that is an ε -approximation to $Z_{G, H}$ with probability at least $3/4$, in time $\text{poly}(|V|, 1/\varepsilon)$.³

We prove the following result. For a bipartite graph G , we use (V_G^0, V_G^1) to denote the bipartition of the vertex set of G , and all logarithms throughout the paper are with base e .

► **Theorem 3.** *Let $q \geq 2$ be an integer, δ be a real in $(0, 1)$, $H \in \mathbb{R}_{\geq 0}^{q \times q}$ be a symmetric δ -matrix. Suppose that Δ, λ satisfy $\frac{\Delta}{\lambda} \geq \frac{100}{1-\delta} q^2 \log(q\Delta)$ and $\Delta \geq \left(\frac{10}{1-\delta} q \log(q\Delta)\right)^4$. Then, there is an FPRAS for $\text{SPIN}_{H, \Delta, \lambda}$.*

In fact, there is a randomised algorithm that, given a graph $G \in \mathcal{G}_{\Delta, \lambda}^{\text{bip}}$ with $n = |V_G^0| = |V_G^1|$ vertices and an accuracy parameter $\varepsilon^ \geq e^{-n/(5q)}$, outputs an ε^* -approximation to $Z_{G, H}$ and an ε^* -sample from the Gibbs distribution $\mu_{G, H}$ in time $O((n/\varepsilon^*)^2 \log^3(n/\varepsilon^*))$.*

As a corollary of Theorem 3, we have the following for random bipartite Δ -regular graphs.

► **Corollary 4.** *Let $q \geq 2$ be an integer, δ be a real in $(0, 1)$, and $H \in \mathbb{R}_{\geq 0}^{q \times q}$ be a symmetric δ -matrix. Then, for all integers $\Delta \geq \left(\frac{10}{1-\delta} q \log(q\Delta)\right)^4$, there is a randomised algorithm such that the following holds with high probability over the choice of a random Δ -regular bipartite graph G with $n = |V_G^0| = |V_G^1|$.*

The algorithm, on input G and an accuracy parameter $\varepsilon^ \geq e^{-n/(5q)}$, outputs in time $O((n/\varepsilon^*)^2 \log^3(n/\varepsilon^*))$ an ε^* -approximation to the partition function $Z_{G, H}$ and an ε^* -sample from the Gibbs distribution $\mu_{G, H}$.*

Proof. Using the result in [3, Theorem 4], we have that, with high probability over the choice of G , it holds that $\lambda(G) \leq 2\sqrt{\Delta}$. It follows that $\frac{\Delta}{\lambda(G)} \geq \frac{1}{2}\sqrt{\Delta}$ and hence the result follows by applying Theorem 3. ◀

Our algorithms apply to a larger class of graphs when δ is small so that the interactions between spins are weak – this corresponds to the so-called “low-temperature” regime – by contrast, approaches such as MCMC and correlation decay typically apply in high-temperature regimes. Since our results concern regular graphs, they easily extend to models with external fields – the fields can be incorporated in the entries of the interaction matrix H .

³ Note, the error probability can be reduced to any $\eta > 0$ by calling the original FPRAS $O(\log(1/\eta))$ times.

1.1 Proof Outline

In order to prove our main Theorem 3, we appeal to what are known as polymer models, as introduced by Gruber and Kunz [19] in the statistical physics literature. Recently, polymer models have been used as a tool in the development of efficient counting and sampling algorithms [21, 23, 25, 8, 6, 7] for problems that are not amenable to traditional approaches such as local-update Markov chains. Our approach is inspired by, and generalises, the approaches in [23, 8], where counting and sampling algorithms are given for the hard-core and ferromagnetic Potts models on expander graphs at low temperatures.

The main idea behind the use of polymer models is that, for graphs with good expansion properties, the partition function and the Gibbs distribution are dominated by configurations which are highly ordered, i.e., whose weight is large. As we shall see in detail in Section 3, these large-weight configurations correspond to “biclques” of H , as defined below.⁴

► **Definition 5.** Let $q \geq 2$ be an integer, $\delta \in [0, 1)$ be a real and $H \in \mathbb{R}_{>0}^{q \times q}$ be a symmetric δ -matrix. A *biclque* of H is a pair (B_0, B_1) with $B_0, B_1 \subseteq [q]$, such that $H_{i,j} = 1$ for all $i \in B_0$ and all $j \in B_1$. We use \mathcal{K}_H to denote the set of all biclques of H and we use \mathcal{K}_H^{\max} to denote the set of all inclusion maximal biclques of H .

Note that, for a bipartite graph G , configurations σ which assign vertices in V_G^0 a spin from B_0 and vertices in V_G^1 a spin from B_1 have weight 1, the largest possible weight (since H is a δ -matrix). Polymer models allow us to capture deviations from such configurations and approximate their contribution to the partition function, see Section 4.

Using the results of Sections 3 and 4, we give the proof of Theorem 3 in Section 5.

2 Preliminaries

Let G be a bipartite graph. We will write G as (V_G^0, V_G^1, E_G) , where (V_G^0, V_G^1) denotes the bipartition of the vertex set of G and E_G its edge set; we will use $V_G = V_G^0 \cup V_G^1$ to denote the vertex set of G .

For a vertex subset $S \subseteq V_G$, let $\partial_G S$ denote the set of vertices of $V_G \setminus S$ that have a neighbour in S , and let $S_G^+ = S \cup \partial_G S$. When $S = \{u\}$, we simply write $\partial_G u$ instead of $\partial_G S$. For vertex subsets $S, T \subseteq V_G$, let $E_G(S, T)$ denote the set of edges of G that have one endpoint in S and the other in T , and let $e_G(S, T) = |E_G(S, T)|$; when $S = T$, we simply write $E_G(S), e_G(S)$ instead of $E_G(S, S), e_G(S, S)$, respectively. We will omit G from all of the above notation where it is obvious from the context.

2.1 Bipartite Expander Graphs

It is well-known that graphs in $\mathcal{G}_{\Delta, \lambda}^{\text{bip}}$ have good expansion properties, and in this section we will review certain edge and vertex expansion properties that will be relevant for us.

The following result relates the spectrum of a regular bipartite graph to its edge-expansion properties. This result was first proven in [20, Theorem 5.1], though the version we state below is taken from [9].

⁴ The “biclque” terminology comes from the homomorphism problem (where H corresponds to a graph), but our interpretation here is analogous which justifies its use.

► **Lemma 6** ([9, Lemma 8]). *Let $G = (V_G^0, V_G^1, E_G) \in \mathcal{G}_{\Delta, \lambda}^{\text{bip}}$ with $n = |V_G^0| = |V_G^1|$. Then, for sets $S_0 \subseteq V_G^0, S_1 \subseteq V_G^1$, we have that*

$$\left| e_G(S_0, S_1) - \frac{\Delta |S_0| |S_1|}{n} \right| \leq \lambda \sqrt{|S_0| |S_1| \left(1 - \frac{|S_0|}{n}\right) \left(1 - \frac{|S_1|}{n}\right)}.$$

The following simple consequence of the above result gives a lower bound on the edge expansion of $G \in \mathcal{G}_{\Delta, \lambda}^{\text{bip}}$ when λ is sufficiently small.

► **Corollary 7.** *Let $G = (V_G^0, V_G^1, E_G) \in \mathcal{G}_{\Delta, \lambda}^{\text{bip}}$ with $n = |V_G^0| = |V_G^1|$. Then, for sets $S_0 \subseteq V_G^0, S_1 \subseteq V_G^1$ such that $\lambda \leq \frac{\Delta}{2n} \sqrt{|S_0| |S_1|}$, it holds that $e_G(S_0, S_1) \geq \frac{\Delta}{2n} |S_0| |S_1|$.*

Proof. Lemma 6 implies that

$$\left| e_G(S_0, S_1) - \frac{\Delta |S_0| |S_1|}{n} \right| \leq \frac{\Delta |S_0| |S_1|}{2n}, \text{ therefore } e_G(S_0, S_1) \geq \frac{\Delta |S_0| |S_1|}{2n}. \quad \blacktriangleleft$$

A second combinatorial notion of expansion is vertex expansion. A well-known result from Tanner [27] relates the spectrum of a graph to its vertex expansion properties (see also [24] for a more refined estimate). Here we state a version from [22, Theorem 4.15]; there, the result is stated and proved for non-bipartite graphs, but a minor adaptation of the proof in [22], which we give for completeness in the full version, also applies to bipartite graphs.

► **Lemma 8.** *Let $G = (V_G^0, V_G^1, E_G) \in \mathcal{G}_{\Delta, \lambda}^{\text{bip}}$, $\rho > 0$ be a real number and $i \in \{0, 1\}$. Then, for all $S \subseteq V_G^i$ with $|S| \leq \rho |V_G^i|$, it holds that $|\partial S| \geq |S| / (\rho + \frac{\lambda^2}{\Delta^2} (1 - \rho))$.*

3 Ground states for spin configurations

In this section, we show that the partition function of a spin system is dominated by configurations which are “close to maximal bicliques”, cf. Definition 5. Let $q \geq 2$, $\Delta \geq 3$ be integers and λ, δ be reals with $\lambda \in (0, \Delta)$ and $\delta \in (0, 1)$. Let $G \in \mathcal{G}_{\Delta, \lambda}^{\text{bip}}$ and let $H \in \mathbb{R}_{\geq 0}^{q \times q}$ be a symmetric δ -matrix.

We next describe more precisely the configurations which are “close” to some maximal biclique of H . Given $\sigma : V_G \rightarrow [q]$ and a spin $i \in [q]$, we write $\sigma^{-1}(i)$ for the set of vertices of G whose image under σ is i . More generally, for a subset of spins $Q \subseteq [q]$, we let $\sigma^{-1}(Q) = \{v \in V_G \mid \sigma(v) \in Q\}$.

► **Definition 9.** *Let $\varepsilon \in (0, 1)$. For $(B_0, B_1) \in \mathcal{K}_H^{\text{max}}$, define $\Sigma_{G, H, \varepsilon}^{B_0, B_1}$ to be the set of spin configurations $\sigma \in \Sigma_{G, H}$ for which $|\sigma^{-1}(B_0) \cap V_G^0| + |\sigma^{-1}(B_1) \cap V_G^1| \geq (1 - \varepsilon) |V_G|$. We define $\Sigma_{G, H, \varepsilon}$ to be the union of the sets $\Sigma_{G, H, \varepsilon}^{B_0, B_1}$ over all bicliques $(B_0, B_1) \in \mathcal{K}_H^{\text{max}}$ and define $Z_{G, H, \varepsilon} := \sum_{\sigma \in \Sigma_{G, H, \varepsilon}} w_{G, H}(\sigma)$.* ◀

The following result shows that $Z_{G, H, \varepsilon}$ gives a close approximation to $Z_{G, H}$ whenever ε is sufficiently large relative to λ, Δ, q .⁵

► **Lemma 10.** *Let $\varepsilon \in (0, 1)$ be such that $\varepsilon \geq 2q\lambda/\Delta$ and $\varepsilon^2 \geq \frac{8q^2 \log q}{\Delta \log(1/\delta)}$. Then, for $G \in \mathcal{G}_{\Delta, \lambda}^{\text{bip}}$ with $n = |V_G^0| = |V_G^1|$, we have that $Z_{G, H, \varepsilon}$ is an e^{-n} -approximation to $Z_{G, H}$.*

⁵ Note, in Lemma 10, as in other lemmas as well, our assumed inequalities for ε impose some restrictions on Δ, λ, q to ensure that such an ε exists. These restrictions will be carefully accounted for when we apply these lemmas; namely, in the proof of Theorem 3.

Proof. We associate each spin configuration $\sigma \in \Sigma_{G,H}$ with a pair of spin subsets $(B_0(\sigma), B_1(\sigma))$ by setting for $i \in \{0, 1\}$, $B_i(\sigma) = \left\{ j \in [q] : |\sigma^{-1}(j) \cap V_G^i| \geq \frac{\varepsilon n}{q} \right\}$. Note that for $\sigma \in \Sigma_{G,H}$, there are fewer than εn vertices of V_G^i which are *not* assigned spins from $B_0(\sigma) \cup B_1(\sigma)$. Also note that, since $\varepsilon \in (0, 1)$, we have that $B_i(\sigma) \neq \emptyset$ for $i \in \{0, 1\}$.

Fix arbitrary $\sigma \in \Sigma_{G,H}$. We first show that

$$\text{either } (B_0(\sigma), B_1(\sigma)) \in \mathcal{K}_H \text{ or } w_{G,H}(\sigma) \leq \delta^{\Delta \varepsilon^2 n / (2q^2)}, \quad (4)$$

i.e., either $(B_0(\sigma), B_1(\sigma))$ is a biclique of H or σ has small weight. For $i \in \{0, 1\}$, consider arbitrary $j_i \in B_i(\sigma)$ and let $S_i = \sigma^{-1}(j_i) \cap V_G^i$. Since $|S_i| \geq \varepsilon n / q$, we have that $\frac{\Delta}{2n} \sqrt{|S_0||S_1|} \geq \Delta \varepsilon / (2q) \geq \lambda$, thus it follows from Corollary 7 that $e(S_0, S_1) \geq \frac{\Delta |S_0||S_1|}{2n} \geq \Delta \varepsilon^2 n / (2q^2)$. It follows that, if $H_{j_0, j_1} \leq \delta$, then $w_{G,H}(\sigma) \leq \delta^{\Delta \varepsilon^2 n / (2q^2)}$; otherwise, $H_{j_0, j_1} = 1$. Since j_0, j_1 were arbitrary spins in $B_0(\sigma), B_1(\sigma)$, respectively, we conclude (4).

Let σ be such that $(B_0(\sigma), B_1(\sigma)) \in \mathcal{K}_H$. Then there exists $(B_0, B_1) \in \mathcal{K}_H^{\max}$ such that $B_i(\sigma) \subseteq B_i$ for $i \in \{0, 1\}$. Moreover, for $i \in \{0, 1\}$ and $j \in [q] \setminus B_i$, we have that $|\sigma^{-1}(j) \cap V_G^i| < \varepsilon n / q$ and therefore that $|\sigma^{-1}([q] \setminus B_i) \cap V_G^i| < \varepsilon n$. Hence, we conclude that $|\sigma^{-1}(B_0) \cap V_G^0| + |\sigma^{-1}(B_1) \cap V_G^1| \geq (1 - \varepsilon)|V_G|$. Combining this with (4), we obtain that for all $\sigma \in \Sigma_{G,H} \setminus \Sigma_{G,H,\varepsilon}$ it holds that $w_{G,H}(\sigma) \leq \delta^{\Delta \varepsilon^2 n / (2q^2)}$, and hence

$$Z_{G,H} - Z_{G,H,\varepsilon} \leq \sum_{\sigma \in \Sigma_{G,H} \setminus \Sigma_{G,H,\varepsilon}} \delta^{\frac{\Delta \varepsilon^2 n}{2q^2}} \leq q^{2n} \delta^{\frac{\Delta \varepsilon^2 n}{2q^2}} \leq q^{-2n},$$

where in the last inequality we used that $\Delta \geq \frac{8q^2 \log(q)}{\varepsilon^2 \log(1/\delta)}$. The result follows since $Z_{G,H} \geq 1$; this bound can be seen by considering the configuration that maps V_G^0 to j_0 and V_G^1 to j_1 , where $j_0, j_1 \in [q]$ are such that $H_{j_0, j_1} = 1$. \blacktriangleleft

For our approximation algorithms, it will be useful to consider the following quantities.

► **Definition 11.** For $\varepsilon \in (0, 1)$, let

$$\widehat{Z}_{G,H,\varepsilon} = \sum_{(B_0, B_1) \in \mathcal{K}_H^{\max}} \sum_{\sigma \in \Sigma_{G,H,\varepsilon}^{B_0, B_1}} w_{G,H}(\sigma) \quad \text{and} \quad Z_{G,H,\varepsilon}^{\text{overlap}} := \sum_{\sigma \in \Sigma_{G,H,\varepsilon}^{\text{overlap}}} w_{G,H}(\sigma),$$

where $\Sigma_{G,H,\varepsilon}^{\text{overlap}} := \bigcup_{(B_0, B_1) \in \mathcal{K}_H^{\max}} \bigcup_{(C_0, C_1) \in \mathcal{K}_H^{\max} \setminus \{(B_0, B_1)\}} (\Sigma_{G,H,\varepsilon}^{B_0, B_1} \cap \Sigma_{G,H,\varepsilon}^{C_0, C_1})$.

The following lemma shows, given a value of ε that is sufficiently small, that $\widehat{Z}_{G,H,3\varepsilon}$ is a close approximation to $Z_{G,H,\varepsilon}$ by showing that $Z_{G,H,3\varepsilon}^{\text{overlap}}$ is small relative to $Z_{G,H,\varepsilon}$.

► **Lemma 12.** Let $\varepsilon \in (0, \frac{1}{240q \log q}]$ be such that $\varepsilon^2 \geq \frac{8q^2 \log(q)}{\Delta \log(1/\delta)}$ and $\varepsilon \geq 2q \frac{\lambda}{\Delta}$. Then, for all $G \in \mathcal{G}_{\Delta, \lambda}^{\text{bip}}$ with $n = |V_G^0| = |V_G^1|$ sufficiently large, we have that $Z_{G,H,3\varepsilon}^{\text{overlap}} \leq e^{-n/(3q)} Z_{G,H,\varepsilon}$ and that $\widehat{Z}_{G,H,3\varepsilon}$ is an $e^{-n/(4q)}$ -approximation to $Z_{G,H,\varepsilon}$.

Proof (Sketch). The main overcounting in $\widehat{Z}_{G,H,3\varepsilon}$ in comparison to $Z_{G,H,\varepsilon}$ comes from double-counting configurations in $\Sigma_{G,H,3\varepsilon}^{\text{overlap}}$ (we also have to account for configurations in $\Sigma_{G,H,3\varepsilon} \setminus \Sigma_{G,H,\varepsilon}$, but this can be done using Lemma 10). Consider such a configuration σ ; then, there exist distinct maximal bicliques $(B_0, B_1), (C_0, C_1) \in \mathcal{K}_H^{\max}$ such that $S := \bigcup_{i \in \{0,1\}} (V_G^i \setminus \sigma^{-1}(B_i))$ and $T := \bigcup_{i \in \{0,1\}} (V_G^i \setminus \sigma^{-1}(C_i))$ satisfy $|S|, |T| \leq 6\varepsilon n$. The key observation is that for $i \in \{0, 1\}$, the vertices in $V_G^i \setminus (S \cup T)$ can be coloured in at most $|B_i \cap C_i|^n$ ways, since they must have a colour in $B_i \cap C_i$. Observe now that at least one of the inequalities $|B_0 \cap C_0| \leq |B_0| - 1, |B_1 \cap C_1| \leq |B_1| - 1$ must hold since otherwise $B_0 \subseteq C_0, B_1 \subseteq C_1$,

contradicting the maximality of (B_0, B_1) . Since $(|B_i| - 1)^n / |B_i|^n \leq (1 - 1/q)^n \leq e^{-n/q}$ for $i \in \{0, 1\}$, we have that the number of such σ (and hence their aggregate weight) is roughly at most $e^{-n/(3q)} \sum_{(B_0, B_1) \in \mathcal{K}_H^{\max}} |B_0|^n |B_1|^n$; note, the factor $e^{-n/(3q)}$ also accounts for the number of spin assignments to the sets S, T , see the full version for details. By considering surjective homomorphisms to bicliques in \mathcal{K}_H^{\max} , we obtain a lower bound on $Z_{G, H, \varepsilon}$ of the order $\sum_{(B_0, B_1) \in \mathcal{K}_H^{\max}} |B_0|^n |B_1|^n$, yielding the result. \blacktriangleleft

4 Using polymer models to estimate the partition function

In this section, we first define abstract polymer models, which will be important in obtaining our approximation algorithms, and then review the algorithmic results of [8] in Section 4.2. We then define a polymer model for spin systems in Section 4.3, and obtain approximation/sampling algorithms for it in Section 4.4.

4.1 Abstract polymer models

Our presentation of abstract polymer models follows mostly [23], but is slightly modified so that a polymer model is a function of both an underlying graph G and a host graph J_G .

Let \mathcal{G} be a class of graphs. Given an underlying graph $G \in \mathcal{G}$ and a set of spins $[q] = \{1, \dots, q\}$, we construct a “host” graph J_G for the polymer model; in our case, we will set for example J_G to be G^3 (see Section 4.3 for more details), but in general other choices are obviously possible. We assign to each vertex $v \in V_{J_G}$ a set of “ground state” spins $g_v \subseteq [q]$. A *polymer* is a pair $\gamma = (V_\gamma, \sigma_\gamma)$ consisting of a J_G -connected set of vertices V_γ and an assignment $\sigma_\gamma : V_\gamma \rightarrow [q]$ such that $\sigma_\gamma(v) \in [q] \setminus g_v$ for all $v \in V_\gamma$. Let \mathcal{P}_G be the set of all polymers.

A *polymer model* for an underlying graph G and host graph J_G is defined by a set of allowed polymers $\mathcal{C}_G \subseteq \mathcal{P}_G$, and a weight function $w_G : \mathcal{C}_G \rightarrow \mathbb{R}_{\geq 0}$. For polymers $\gamma, \gamma' \in \mathcal{P}_G$, we write $\gamma \sim \gamma'$ to denote that γ, γ' are *compatible*, i.e., if $d_{J_G}(V_\gamma, V_{\gamma'}) > 1$ where $d_{J_G}(\cdot, \cdot)$ denotes the graph distance in J_G and for $S, T \subseteq V_G$ we let $d_{J_G}(S, T) = \min_{u \in S, v \in T} d_{J_G}(u, v)$. We define $\Omega_G = \{\Gamma \subseteq \mathcal{C}_G \mid \forall \gamma, \gamma' \in \Gamma, \gamma \sim \gamma'\}$ to be the set of all sets of mutually compatible polymers of \mathcal{C}_G ; elements of Ω_G are called polymer configurations. The polymer model induces the following Gibbs distribution on Ω_G , defined for $\Gamma \in \Omega_G$ by $\mu_G(\Gamma) = \frac{\prod_{\gamma \in \Gamma} w_G(\gamma)}{Z_G}$, where $Z_G := \sum_{\Gamma' \in \Omega_G} \prod_{\gamma \in \Gamma'} w_G(\gamma)$ is the partition function of the model. We use $(\mathcal{C}_G, w_G, J_G)$ to denote the polymer model and $\{(\mathcal{C}_G, w_G, J_G) \mid G \in \mathcal{G}\}$ to denote the family of polymer models corresponding to the class of graphs \mathcal{G} ; we say that the family has degree bound Δ if, for every $G \in \mathcal{G}$, both G and the host graph J_G have maximum degree at most Δ .

4.2 Algorithms for polymer models

Given a family of polymer models, Chen et. al. [8], building upon work of [23], give sufficient conditions under which the partition function of the polymer model can be efficiently approximated using Markov chains. We will briefly describe these conditions, as well as the key results from [8] that will be later important for us.

The first condition is known as *computational feasibility* and is defined as follows.

Definition 13. [8, Definition 3] *Let \mathcal{G} be a class of graphs. A family of polymer models $\{(\mathcal{C}_G, w_G, J_G) \mid G \in \mathcal{G}\}$ is computationally feasible if, for all $G \in \mathcal{G}$ and all $\gamma \in \mathcal{P}_G$, we can determine whether $\gamma \in \mathcal{C}_G$ and, if so, compute $w_G(\gamma)$ in time polynomial in $|V_\gamma|$.* \blacktriangleleft

The second condition is called the *polymer sampling condition* and is defined as follows.

► **Definition 14** ([8, Definition 4]). *Let $q \geq 2, \Delta \geq 3$ be integers, and \mathcal{G} be a class of graphs. A family of polymer models $\{(C_G, w_G, J_G) \mid G \in \mathcal{G}\}$ with q spins and degree bound Δ satisfies the polymer sampling condition with constant $\tau \geq 5 + 3 \log((q-1)\Delta)$ if $w_G(\gamma) \leq e^{-\tau|V_\gamma|}$ for all $\gamma \in \mathcal{C}_G$ and all $G \in \mathcal{G}$.*

The following result from [8] asserts that if a family of polymer models satisfies the above two conditions, then there are efficient approximation and sampling algorithms for the partition function and the Gibbs distribution of the polymer model, respectively.

► **Theorem 15** ([8, Theorems 5 & 6]). *Let $q \geq 2, \Delta \geq 3$ be integers, and \mathcal{G} be a class of graphs. Suppose that $\{(C_G, w_G, J_G) \mid G \in \mathcal{G}\}$ is a family of computationally feasible polymer models with q spins and degree bound Δ that satisfies the polymer sampling condition.*

Then, there is a randomised algorithm that takes as input an n -vertex graph $G \in \mathcal{G}$ and an accuracy parameter $\varepsilon^ \in (0, 1)$ and outputs an ε^* -approximation to Z_G in time $O((n/\varepsilon^*)^2 \log^2(n/\varepsilon^*))$ with probability $\geq 3/4$. Moreover, there is a randomised algorithm that on input G and ε^* as before, outputs an ε^* -approximate sample from μ_G in time $O(n \log(n/\varepsilon^*))$.*

4.3 Polymer model for spin systems

In this section we define a polymer model for spin systems that captures the deviations of spin configurations from maximal bicliques. The polymer model that we propose is a generalisation to arbitrary spin systems of a polymer model that was used in [23, Section 5] in the case of proper colourings.

Let $H \in \mathbb{R}_{\geq 0}^{q \times q}$ be a symmetric matrix and $(B_0, B_1) \in \mathcal{K}_H^{\max}$ be a maximal biclique of H . Let $G \in \mathcal{G}_{\Delta, \lambda}^{\text{bip}}$ be a graph, and let $\varepsilon \in (0, 1)$. The host graph for the polymer model is $J_G = G^3$, where G^3 is the graph defined on V_G with two vertices connected by an edge if the distance between them in G is at most 3. For $v \in V_G^i$ with $i \in \{0, 1\}$, the set of ground state spins g_v is B_i . Let $\mathcal{P}_{G, H}^{B_0, B_1}$ denote the set of all polymers, i.e., all pairs $\gamma = (V_\gamma, \sigma_\gamma)$ consisting of a G^3 -connected set of vertices V_γ and an assignment $\sigma_\gamma : V_\gamma \rightarrow [q]$ such that $\sigma_\gamma(v) \in [q] \setminus g_v$ for all $v \in V_\gamma$. We define the set of allowed polymers as

$$\mathcal{C}_{G, H, \varepsilon}^{B_0, B_1} = \left\{ \gamma \in \mathcal{P}_{G, H}^{B_0, B_1} : |V_\gamma| \leq \varepsilon |V_G| \right\} \quad (9)$$

and let $\Omega_{G, H, \varepsilon}^{B_0, B_1}$ denote the set of all sets of mutually compatible polymers. We define the weight of a polymer $\gamma = (V_\gamma, \sigma_\gamma) \in \mathcal{C}_{G, H, \varepsilon}^{B_0, B_1}$ as

$$w_{G, H}^{B_0, B_1}(\gamma) = \frac{\prod_{\{u, v\} \in E_G(V_\gamma)} H_{\sigma_\gamma(u), \sigma_\gamma(v)} \prod_{u \in \partial V_\gamma} F_u}{\prod_{i \in \{0, 1\}} |B_i|^{|V_G^i \cap V_\gamma^+|}}, \quad (10)$$

where for $u \in V_G^i$ with $i \in \{0, 1\}$, $F_u := \sum_{j \in B_i} \prod_{v \in V_\gamma \cap \partial u} H_{j, \sigma_\gamma(v)}$.

We let $Z_{G, H, \varepsilon}^{B_0, B_1}$ and $\mu_{G, H, \varepsilon}^{B_0, B_1}$ denote the partition function and the Gibbs distribution of the polymer model $(\mathcal{C}_{G, H, \varepsilon}^{B_0, B_1}, w_{G, H}^{B_0, B_1}, J_G)$, as defined in Section 4.1.

The next lemma shows the motivation behind the definition of the weight of a polymer. For a polymer configuration $\Gamma \in \Omega_{G, H, \varepsilon}^{B_0, B_1}$, let $\cup \Gamma = \bigcup_{\gamma \in \Gamma} V_\gamma$, and σ_Γ denote the assignment to vertices in $\cup \Gamma$ obtained by combining all of the assignments σ_γ , for $\gamma \in \Gamma$.

► **Definition 16.** *For $\Gamma \in \Omega_{G, H, \varepsilon}^{B_0, B_1}$, define $\Sigma_{G, H}^{B_0, B_1}(\Gamma)$ to be the set of configurations τ such that $\tau|_{\cup \Gamma} = \sigma_\Gamma$ and which map, for $i \in \{0, 1\}$, $V_G^i \setminus (\cup \Gamma)$ to B_i . ◀*

► **Lemma 17.** *Let $n = |V_G^0| = |V_G^1|$. For all $\varepsilon \in (0, 1)$, and all polymer configurations $\Gamma \in \Omega_{G,H,\varepsilon}^{B_0,B_1}$, we have that*

$$|B_0|^n |B_1|^n \prod_{\gamma \in \Gamma} w_{G,H}^{B_0,B_1}(\gamma) = \sum_{\tau \in \Sigma_{G,H}^{B_0,B_1}(\Gamma)} w_{G,H}(\tau).$$

Proof (Sketch). Let γ, γ' be two distinct polymers in Γ . Since the polymers are compatible, they correspond to distinct G^3 -connected components; in particular, $V_{\gamma'}^+ \cap V_{\gamma}^+ = \emptyset$. It follows that $|(\cup \Gamma)^+| = \sum_{\gamma \in \Gamma} |V_{\gamma}^+|$. Hence, by the definition in (10), we have that

$$\begin{aligned} |B_0|^n |B_1|^n \prod_{\gamma \in \Gamma} w_{G,H}^{B_0,B_1}(\gamma) &= |B_0|^n |B_1|^n \prod_{\gamma \in \Gamma} \frac{\prod_{\{u,v\} \in E_G(V_{\gamma})} H_{\sigma_{\gamma}(u), \sigma_{\gamma}(v)} \prod_{u \in \partial V_{\gamma}} F_u}{\prod_{i \in \{0,1\}} |B_i|^{|V_G^i \cap V_{\gamma}^+|}} \\ &= \prod_{i \in \{0,1\}} |B_i|^{|V_G^i \setminus (\cup \Gamma)^+|} \prod_{\{u,v\} \in E_G(\cup \Gamma)} H_{\sigma_{\Gamma}(u), \sigma_{\Gamma}(v)} \prod_{u \in \partial(\cup \Gamma)} F_u. \end{aligned} \quad (11)$$

On the other hand, for each $\tau \in \Sigma_{G,H}^{B_0,B_1}(\Gamma)$, we have that

$$w_{G,H}(\tau) = \prod_{\{u,v\} \in E_G(\cup \Gamma)} H_{\sigma_{\Gamma}(u), \sigma_{\Gamma}(v)} \prod_{u \in \partial(\cup \Gamma)} \prod_{v \in (\cup \Gamma) \cap \partial u} H_{\sigma_{\Gamma}(v), \tau(u)},$$

i.e., given Γ and that $\tau \in \Sigma_{G,H}^{B_0,B_1}(\Gamma)$, the weight of τ depends only on the assignment of $\partial(\cup \Gamma)$. Using this, we show in the full version that

$$\sum_{\tau \in \Sigma_{G,H}^{B_0,B_1}(\Gamma)} w_{G,H}(\tau) = \prod_{i \in \{0,1\}} |B_i|^{|V_G^i \setminus (\cup \Gamma)^+|} \prod_{\{u,v\} \in E_G(V_{\Gamma})} H_{\sigma_{\Gamma}(u), \sigma_{\Gamma}(v)} \prod_{u \in \partial(\cup \Gamma)} F_u,$$

which combined with (11) gives the desired equality. ◀

The following quantity combines the partition functions of the polymer models corresponding to maximal bicliques of H ; we will use this as our approximation to $Z_{G,H}$.

► **Definition 18.** *For $\varepsilon \in (0, 1)$, let $Z_{G,H,\varepsilon}^{\text{polymer}} = \sum_{(B_0, B_1) \in \mathcal{K}_H^{\max}} |B_0|^n |B_1|^n \cdot Z_{G,H,\varepsilon}^{B_0,B_1}$.*

Using a minor adaptation of [23, Claim 29] in our setting, we show in the full version that the aggregate size of polymer configurations is at most $6\varepsilon n$ (cf. Lemma 19), so combining Lemmas 12 and 17 we obtain the following result (proof in the full version).

► **Lemma 20.** *Let $\varepsilon \in (0, \frac{1}{240q \log q}]$ be such that $\Delta \geq \frac{8q^2 \log(q)}{\varepsilon^2 \log(1/\delta)}$ and $\varepsilon \geq 2q \frac{\lambda}{\Delta}$. Then, for all $G \in \mathcal{G}_{\Delta,\lambda}^{\text{bip}}$ with $n = |V_G^0| = |V_G^1|$ sufficiently large, we have that $Z_{G,H,\varepsilon}^{\text{polymer}}$ is an $e^{-n/(4q)}$ -approximation to $Z_{G,H,\varepsilon}$.*

4.4 Sampling from the polymer model

Let $q \geq 2, \Delta \geq 3$ be integers and $\delta \in (0, 1), \lambda \in (0, \Delta)$ be reals. Let $H \in \mathbb{R}_{\geq 0}^{q \times q}$ be a symmetric δ -matrix and $(B_0, B_1) \in \mathcal{K}_H^{\max}$. For $\varepsilon \in (0, 1)$, we now show that the family of polymer models $\{(\mathcal{C}_{G,H,\varepsilon}^{B_0,B_1}, w_{G,H,\varepsilon}^{B_0,B_1}, J_G) \mid G \in \mathcal{G}_{\Delta,\lambda}^{\text{bip}}\}$ which was defined in the previous subsection is computationally feasible and satisfies the polymer sampling condition.

► **Lemma 21.** *Let $\varepsilon \in (0, 1)$ be such that $\varepsilon \geq \lambda^2/\Delta^2$ and $\varepsilon \leq \frac{1-\delta}{40q \log(q\Delta)}$. The family of polymer models $\{(\mathcal{C}_{G,H,\varepsilon}^{B_0,B_1}, w_{G,H,\varepsilon}^{B_0,B_1}, J_G) \mid G \in \mathcal{G}_{\Delta,\lambda}^{\text{bip}}\}$ is computationally feasible and satisfies the polymer sampling condition with constant $\tau \geq 5 + 3 \log((q-1)\Delta^3)$.*

37:10 Fast Algorithms for General Spin Systems on Bipartite Expanders

Proof (Sketch). Consider an arbitrary polymer $\gamma = (V_\gamma, \sigma_\gamma) \in \mathcal{P}_{G,H}^{B_0, B_1}$. We focus on verifying the polymer sampling condition, since the computational feasibility follows using the definitions. From (10), using that the entries of H are at most 1, we have the bound

$$w_{G,H}^{B_0, B_1}(\gamma) \leq \frac{\prod_{u \in \partial V_\gamma \cap V_G} F_u}{\prod_{i \in \{0,1\}} |B_i|^{|V_G^i \cap V_\gamma^+|}}, \quad (13)$$

where, recall that, for $u \in V_G^i$ with $i \in \{0,1\}$, $F_u := \sum_{j \in B_i} \prod_{v \in V_\gamma \cap \partial u} H_{j, \sigma_\gamma(v)}$. Let v be a neighbour of u in $V_\gamma \cap V_G^{i \oplus 1}$; such v exists since $u \in \partial V_\gamma \cap V_G^i$. Then, there exists $j \in B_i$ such that $H_{j, \sigma_\gamma(v)} \leq \delta$; otherwise $(B_i, B_{i \oplus 1} \cup \{\sigma_\gamma(v)\})$ would also be a biclique of H , contradicting the maximality of (B_0, B_1) (since $\sigma_\gamma(v) \notin B_{i \oplus 1}$). It follows that $F_u \leq |B_i| - 1 + \delta$. Using this in (13), we get that

$$w_{G,H}^{B_0, B_1}(\gamma) \leq \frac{\prod_{i \in \{0,1\}} (|B_i| - 1 + \delta)^{|\partial V_\gamma \cap V_G^i|}}{\prod_{i \in \{0,1\}} |B_i|^{|V_G^i \cap V_\gamma^+|}} \leq \left(1 - \frac{1 - \delta}{q}\right)^{|\partial V_\gamma|} \leq e^{-|\partial V_\gamma| \left(\frac{1 - \delta}{q}\right)}, \quad (14)$$

where in the second to last inequality we used that $|V_\gamma^+ \cap V_G^i| \geq |\partial V_\gamma \cap V_G^i|$ for $i \in \{0,1\}$.

We next lower bound $|\partial V_\gamma|$ in terms of $|V_\gamma|$. For $i \in \{0,1\}$, let $\rho_i = |V_\gamma \cap V_G^i|/n$ and $\rho = |V_\gamma|/n$. Applying Lemma 8 to the set $V_\gamma \cap V_G^i$ and using the inequality $\frac{a}{x} + \frac{b}{y} \geq \frac{(a+b)^2}{ax+by}$ which holds for all $a, b, x, y \geq 0$, we obtain

$$\sum_{i \in \{0,1\}} |\partial(V_\gamma \cap V_G^i)| \geq \sum_{i \in \{0,1\}} \frac{|V_\gamma \cap V_G^i|}{\rho_i + \frac{\lambda^2}{\Delta^2}} \geq \frac{|V_\gamma|}{\rho + \frac{\lambda^2}{\Delta^2}} \geq \frac{|V_\gamma|}{3\varepsilon},$$

using that $\rho \leq 2\varepsilon$ and $\varepsilon \geq \frac{\lambda^2}{\Delta^2}$. Since $\varepsilon \in (0, 1/20)$, we have that $|\partial V_\gamma| \geq \frac{|V_\gamma|}{3\varepsilon} - |V_\gamma| \geq \frac{|V_\gamma|}{4\varepsilon}$.

Plugging this into (14), we obtain $w_{G,H}^{B_0, B_1}(\gamma) \leq e^{-|V_\gamma| \left(\frac{1 - \delta}{4\varepsilon q}\right)} \leq e^{-\tau |V_\gamma|}$, where $\tau := \frac{1 - \delta}{4\varepsilon q} \geq 10 \log(q\Delta) \geq 5 + 3 \log((q-1)\Delta^3)$ using that $\varepsilon \leq \frac{1 - \delta}{40q \log(q\Delta)}$ and $q \geq 2, \Delta \geq 3$. This finishes the proof of Lemma 21, after observing that the degree bound for the family $\{(C_{G,H,\varepsilon}^{B_0, B_1}, w_{G,H}^{B_0, B_1}, J_G) \mid G \in \mathcal{G}_{\Delta, \lambda}^{\text{bip}}\}$ is Δ^3 . \blacktriangleleft

Finally, we can apply Theorem 15, which gives us an efficient algorithm for approximating the partition function of the polymer model.

► **Corollary 22.** *Let $q \geq 2$, $\delta \in (0, 1)$, $H \in \mathbb{R}_{\geq 0}^{q \times q}$ be a symmetric δ -matrix, and $(B_0, B_1) \in \mathcal{K}_H^{max}$. Let $\varepsilon \in (0, 1)$ be such that $\varepsilon \geq \lambda^2/\Delta^2$ and $\varepsilon \leq \frac{1 - \delta}{40q \log(q\Delta)}$.*

Then, there is a randomised algorithm that takes as input an n -vertex graph $G \in \mathcal{G}_{\Delta, \lambda}^{\text{bip}}$ and an accuracy parameter $\varepsilon^ \in (0, 1)$ and outputs an ε^* -approximation to $Z_{G,H,\varepsilon}^{B_0, B_1}$ in time $O((n/\varepsilon^*)^2 \log^2(n/\varepsilon^*))$ with probability $\geq 3/4$. Moreover, there is a randomised algorithm that on input G and ε^* as before, outputs an ε^* -approximate sample from $\mu_{G,H,\varepsilon}^{B_0, B_1}$ in time $O(n \log(n/\varepsilon^*))$.*

5 Proof of Theorem 3

In this section, we combine the results of Sections 3 and 4 to prove Theorem 3.

► **Theorem 3.** *Let $q \geq 2$ be an integer, δ be a real in $(0, 1)$, $H \in \mathbb{R}_{\geq 0}^{q \times q}$ be a symmetric δ -matrix. Suppose that Δ, λ satisfy $\frac{\Delta}{\lambda} \geq \frac{100}{1 - \delta} q^2 \log(q\Delta)$ and $\Delta \geq \left(\frac{10}{1 - \delta} q \log(q\Delta)\right)^4$. Then, there is an FPRAS for $\text{SPIN}_{H, \Delta, \lambda}$.*

In fact, there is a randomised algorithm that, given a graph $G \in \mathcal{G}_{\Delta, \lambda}^{\text{bip}}$ with $n = |V_G^0| = |V_G^1|$ vertices and an accuracy parameter $\varepsilon^* \geq e^{-n/(5q)}$, outputs an ε^* -approximation to $Z_{G,H}$ and an ε^* -sample from the Gibbs distribution $\mu_{G,H}$ in time $O((n/\varepsilon^*)^2 \log^3(n/\varepsilon^*))$.

Proof. As input to the FPRAS, we are given a graph $G \in \mathcal{G}_{\Delta, \lambda}^{\text{bip}}$ and an accuracy parameter $\varepsilon^* \in (0, 1)$. We may assume that $n = |V_G^0| = |V_G^1|$ is sufficiently large, otherwise we can compute $Z_{G,H}$ exactly in constant time, by brute force. Similarly, we may assume that $\varepsilon^* \geq 9e^{-n/(4q)}$, otherwise we can compute $Z_{G,H}$ exactly in $O(nq^{2n}) = \text{poly}(1/\varepsilon^*)$ time, by brute force.

Let $\varepsilon = \frac{1-\delta}{50q \log(q\Delta)}$ and observe that, using the lower bounds on $\frac{\Delta}{\lambda}$ and Δ , we have that

$$\varepsilon \leq \frac{1}{240q \log q}, \quad \varepsilon \geq 2q \frac{\lambda}{\Delta} \geq \frac{\lambda^2}{\Delta^2}, \quad \varepsilon^2 \geq \frac{8q^2 \log q}{\Delta \log(1/\delta)},$$

where the first inequality follows from $\Delta \geq q^4$, the second from rearranging $\frac{\Delta}{\lambda} \geq 2q/\varepsilon$ (using the lower bound on $\frac{\Delta}{\lambda}$), and the last inequality from using that $\log(1/\delta) \geq 1 - \delta$ for all $\delta \in (0, 1)$. In particular, the assumption of Lemmas 10, 20 and Corollary 22 are satisfied.

By Corollary 22, for an arbitrary biclique $(B_0, B_1) \in \mathcal{K}_H^{\text{max}}$, we can obtain an $(\varepsilon^*/8)$ -approximation to $Z_{G,H,\varepsilon}^{B_0, B_1}$ in $O((n/\varepsilon^*)^2 \log^2(n/\varepsilon^*))$ time, with probability at least $3/4$. Taking the median of $O(\log(1/\varepsilon^*))$ runs of this algorithm, we therefore obtain $\widehat{Z}_{G,H,\varepsilon}^{B_0, B_1}$ which is an $(\varepsilon^*/8)$ -approximation to $Z_{G,H,\varepsilon}^{B_0, B_1}$ with probability at least $1 - \varepsilon^*/(16|\mathcal{K}_H^{\text{max}}|)$. By a union bound over the bicliques in $\mathcal{K}_H^{\text{max}}$, it follows that $\widehat{Z}_{G,H,\varepsilon}^{\text{polymer}} := \sum_{(B_0, B_1) \in \mathcal{K}_H^{\text{max}}} |B_0|^n |B_1|^n \cdot \widehat{Z}_{G,H,\varepsilon}^{B_0, B_1}$ is an $(\varepsilon^*/8)$ -approximation to $Z_{G,H,\varepsilon}^{\text{polymer}}$ (cf. Definition 18) with probability at least $1 - \varepsilon^*/16$. By Lemma 20, $Z_{G,H,\varepsilon}^{\text{polymer}}$ is an $(\varepsilon^*/8)$ -approximation to $Z_{G,H,\varepsilon}$ which, by Lemma 10, is an $(\varepsilon^*/8)$ -approximation to $Z_{G,H}$. It therefore follows that $\widehat{Z}_{G,H,\varepsilon}^{\text{polymer}}$ is a $(3\varepsilon^*/8)$ -approximation, and hence an ε^* -approximation, to $Z_{G,H}$ with probability at least $1 - \varepsilon^*/16$. The total run-time of the algorithm is therefore $O((n/\varepsilon^*)^2 \log^3(n/\varepsilon^*))$.

For the sampling algorithm, we assume again that $\varepsilon^* \geq 9e^{-n/(4q)}$. We first sample a biclique $\widehat{\mathbf{B}} = (B_0, B_1) \in \mathcal{K}_H^{\text{max}}$ with probability $\frac{|B_0|^n |B_1|^n \widehat{Z}_{G,H,\varepsilon}^{B_0, B_1}}{\widehat{Z}_{G,H,\varepsilon}^{\text{polymer}}}$, where $\widehat{Z}_{G,H,\varepsilon}^{B_0, B_1}$ and $\widehat{Z}_{G,H,\varepsilon}^{\text{polymer}}$ are as before. Then, using Corollary 22, we sample a polymer configuration $\widehat{\Gamma}$ whose distribution is at distance at most $\varepsilon^*/6$ from $\mu_{G,H,\varepsilon}^{B_0, B_1}$. We output $\widehat{\sigma} = \text{Spin}_{\widehat{\mathbf{B}}}(\widehat{\Gamma})$, where for a biclique $B = (B_0, B_1)$ and a polymer configuration Γ , $\text{Spin}_B(\Gamma)$ is a random configuration τ obtained as follows:

- For every vertex $u \in \cup \Gamma$, we set $\tau(u) = \sigma_\Gamma(u)$.
- For $u \in V_G^i \setminus (\cup \Gamma)^+$ with $i \in \{0, 1\}$, we assign a random spin from B_i uniformly at random.
- For $u \in \partial(\cup \Gamma) \cap V_G^i$ with $i \in \{0, 1\}$, for $j \in B_i$ we set $\tau(u) = j$ with probability $\frac{1}{F_u} \prod_{v \in \partial u \cap (\cup \Gamma)} H_{j, \sigma_\Gamma(v)}$ where $F_u := \sum_{j \in B_i} \prod_{v \in \partial u \cap (\cup \Gamma)} H_{j, \sigma_\Gamma(v)}$.

We claim that $\widehat{\sigma}$ is an ε^* -approximate sample from the Gibbs distribution $\mu_{G,H}$. To prove this, let $\mathbf{B}, \Gamma, \sigma$ be the analogues of $\widehat{\mathbf{B}}, \widehat{\Gamma}, \widehat{\sigma}$, respectively, when there is no error, i.e.,:

1. Set $\mathbf{B} = (B_0, B_1)$ with probability $|B_0|^n |B_1|^n \frac{Z_{G,H,\varepsilon}^{B_0, B_1}}{Z_{G,H,\varepsilon}^{\text{polymer}}}$.
2. Conditioned on $\mathbf{B} = (B_0, B_1)$, set Γ to be a sample from $\mu_{G,H,\varepsilon}^{B_0, B_1}$ and $\sigma = \text{Spin}_{\mathbf{B}}(\Gamma)$.

With probability $1 - \varepsilon^*/16$, we have that, for all bicliques $(B_0, B_1) \in \mathcal{K}_H^{\text{max}}$, $\widehat{Z}_{G,H,\varepsilon}^{B_0, B_1}$ and $\widehat{Z}_{G,H,\varepsilon}^{\text{polymer}}$ are $(\varepsilon^*/8)$ -approximations to $Z_{G,H,\varepsilon}^{B_0, B_1}$ and $Z_{G,H,\varepsilon}^{\text{polymer}}$ respectively, and conditioned on this the total variation distance between the distributions of $\widehat{\mathbf{B}}$ and \mathbf{B} is at most $\frac{1+\varepsilon^*/8}{1-\varepsilon^*/8} - 1 \leq 3\varepsilon^*/8$. It follows that the total variation distance between the distributions of $\widehat{\mathbf{B}}$ and \mathbf{B} is at

37:12 Fast Algorithms for General Spin Systems on Bipartite Expanders

most $3\varepsilon^*/8 + \varepsilon^*/16 \leq \varepsilon^*/2$, and so there is a coupling so that $\Pr[\widehat{\mathbf{B}} \neq \mathbf{B}] \leq \varepsilon^*/2$. Conditioned on $\widehat{\mathbf{B}} = \mathbf{B}$, we have that the total variation distance between $\widehat{\Gamma}$ and Γ is at most $\varepsilon^*/6$, so there is a coupling which further satisfies $\Pr[\widehat{\Gamma} \neq \Gamma \mid \widehat{\mathbf{B}} = \mathbf{B}] \leq \varepsilon^*/6$. Finally, conditioned on $\widehat{\mathbf{B}} = \mathbf{B}$ and $\widehat{\Gamma} = \Gamma$, we can clearly couple $\widehat{\sigma}$ and σ so that they agree. It follows that the total variation distance between $\widehat{\sigma}$ and σ is at most $2\varepsilon^*/3$, so the result will follow by showing that the distribution of σ and $\mu_{G,H}$ are at distance at most $3e^{-n/(4q)} \leq \varepsilon^*/3$.

For this, we will consider the set of configurations $\widehat{\Sigma} := \Sigma_{G,H,\varepsilon} \setminus \Sigma_{G,H,3\varepsilon}^{\text{overlap}}$, where recall from Definition 9 that $\Sigma_{G,H,\varepsilon}$ is the set of configurations τ with $\sum_{i \in \{0,1\}} |\tau^{-1}(B_i) \cap V_G^i| \geq (1-\varepsilon)|V_G|$, and from Definition 11 that $\Sigma_{G,H,3\varepsilon}^{\text{overlap}}$ is the set of configurations τ such that $\tau \in \Sigma_{G,H,3\varepsilon}^{B_0,B_1} \cap \Sigma_{G,H,3\varepsilon}^{C_0,C_1}$ for some distinct maximal bicliques $(B_0, B_1), (C_0, C_1)$.

Consider arbitrary $\tau \in \widehat{\Sigma}$ and let (B_0, B_1) be such that $\tau \in \Sigma_{G,H,\varepsilon}^{B_0,B_1}$. For $i \in \{0,1\}$, let $T_i = V_G^i \cap \tau^{-1}([q] \setminus B_i)$, and S_1, \dots, S_k denote the G^3 -connected components of $T := T_0 \cup T_1$; note, since $\tau \in \widehat{\Sigma}$, we have $|T| \leq \varepsilon|V_G|$. Consider the polymer configuration Γ_τ which is the union of the polymers $(S_1, \tau|_{S_1}), \dots, (S_k, \tau|_{S_k})$. In the full version, we show that

$$\sigma = \tau \text{ iff } \mathbf{B} = (B_0, B_1), \mathbf{\Gamma} = \Gamma_\tau, \sigma|_{V_G \setminus (\cup \Gamma)} = \tau|_{V_G \setminus T}. \quad (15)$$

From (15) and the sampling procedure for σ , we therefore have that $\sigma = \tau$ with probability

$$\frac{|B_0|^n |B_1|^n Z_{G,H,\varepsilon}^{B_0,B_1}}{Z_{G,H,\varepsilon}^{\text{polymer}}} \cdot \frac{w_{G,H}^{B_0,B_1}(\Gamma_\tau)}{Z_{G,H,\varepsilon}^{B_0,B_1}} \cdot \prod_{i \in \{0,1\}} \frac{1}{|B_i|^{|V_G^i \setminus T^+|}} \prod_{u \in \partial T} \left(\frac{1}{F_u} \prod_{v \in T \cap \partial u} H_{\tau(u), \tau(v)} \right). \quad (16)$$

From (11) applied to the polymer configuration Γ_τ , we have that (using $\cup \Gamma_\tau = T$ and $\sigma_{\Gamma_\tau} = \tau$)

$$|B_0|^n |B_1|^n w_{G,H}^{B_0,B_1}(\Gamma_\tau) = \prod_{i \in \{0,1\}} |B_i|^{|V_G^i \setminus T^+|} \prod_{\{u,v\} \in E_G(T)} H_{\tau(u), \tau(v)} \prod_{u \in \partial T} F_u,$$

and hence we obtain that the expression in (16) equals

$$\frac{\prod_{\{u,v\} \in E_G(T)} H_{\tau(u), \tau(v)} \prod_{u \in \partial T} \prod_{v \in T \cap \partial u} H_{\tau(u), \tau(v)}}{Z_{G,H,\varepsilon}^{\text{polymer}}} = \frac{w_{G,H}(\tau)}{Z_{G,H,\varepsilon}^{\text{polymer}}},$$

where the last equality follows by noting that edges that are not in $E_G(T) \cup E_G(T, \partial T)$ contribute a factor of 1 in the weight of T (since their endpoints are assigned spins of the biclique). So, we have shown that $\sigma = \tau$ with probability $w_{G,H}(\tau)/Z_{G,H,\varepsilon}^{\text{polymer}}$.

Let p_σ be the probability that $\sigma \in \widehat{\Sigma}$ and p be the aggregate weight of configurations in the Gibbs distribution $\mu_{G,H}$ in $\widehat{\Sigma}$. Then, using that $Z_{G,H,3\varepsilon}^{\text{overlap}} \leq e^{-n/(3q)} Z_{G,H,\varepsilon}$ from Lemma 12 and $Z_{G,H,\varepsilon}^{\text{polymer}} \leq (1 + e^{-n/(4q)}) Z_{G,H,\varepsilon}$ from Lemma 20, we have that

$$p_\sigma \geq \frac{1}{Z_{G,H,\varepsilon}^{\text{polymer}}} \sum_{\tau \in \widehat{\Sigma}} w_{G,H}(\tau) \geq \frac{Z_{G,H,\varepsilon} - Z_{G,H,3\varepsilon}^{\text{overlap}}}{Z_{G,H,\varepsilon}^{\text{polymer}}} \geq 1 - 2e^{-n/(4q)}, \quad (17)$$

while for p , using that $Z_{G,H} \leq (1 + e^{-n}) Z_{G,H,\varepsilon}$ from Lemma 10, we have the bound

$$p \geq \frac{1}{Z_{G,H}} \sum_{\tau \in \widehat{\Sigma}} w_{G,H}(\tau) \geq \frac{Z_{G,H,\varepsilon} - Z_{G,H,3\varepsilon}^{\text{overlap}}}{Z_{G,H}} \geq 1 - 2e^{-n/(3q)}. \quad (18)$$

It follows that the total variation distance between the distribution of σ and $\mu_{G,H}$ is bounded above by $D := \frac{1}{2}((1 - p_\sigma) + (1 - p) + M)$, where $M := \sum_{\tau \in \widehat{\Sigma}} w_{G,H}(\tau) \left| \frac{1}{Z_{G,H,\varepsilon}^{\text{polymer}}} - \frac{1}{Z_{G,H}} \right|$.

Using Lemma 10 and Lemma 20, we have the bound

$$M \leq \left| \frac{Z_{G,H}}{Z_{G,H,\varepsilon}^{\text{polymer}}} - 1 \right| \leq 2e^{-n/(4q)}. \quad (19)$$

Combining (17), (18) and (19), we obtain that $D \leq 3e^{-n/(4q)}$, i.e., the distance between the distribution of σ and $\mu_{G,H}$ is at most $3e^{-n/(4q)}$, finishing the proof of Theorem 3. ◀

References

- 1 N. Alon. Eigenvalues, geometric expanders, sorting in rounds, and Ramsey theory. *Combinatorica*, 6(3):207–219, 1986.
- 2 N. Alon and F. R. K. Chung. Explicit construction of linear sized tolerant networks. *Discrete Mathematics*, 72(1-3):15–19, 1988.
- 3 G. Brito, I. Dumitriu, and K. D. Harris. Spectral gap in random bipartite biregular graphs and applications. *arXiv:1804.07808*, 2018.
- 4 A. Bulatov and M. Grohe. The complexity of partition functions. *Theoretical Computer Science*, 348(2):148–186, 2005.
- 5 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.
- 6 S. Cannon and W. Perkins. Counting independent sets in unbalanced bipartite graphs. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1456–1466. SIAM, 2020.
- 7 C. Carlson, E. Davies, and A. Kolla. Efficient algorithms for the Potts model on small-set expanders. *CoRR*, abs/2003.01154, 2020.
- 8 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)*, volume 145 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 41:1–41:14, 2019.
- 9 S. De Winter, J. Schillewaert, and J. Verstraete. Large incidence-free sets in geometries. *The electronic journal of Combinatorics*, 19(4):P24, 2012.
- 10 M. Dyer and C. Greenhill. The complexity of counting graph homomorphisms. *Random Structures & Algorithms*, 17(3-4):260–289, 2000.
- 11 M. E. Dyer, L. A. Goldberg, C. S. Greenhill, and M. Jerrum. The relative complexity of approximate counting problems. *Algorithmica*, 38(3):471–500, 2004.
- 12 J. Friedman. A proof of Alon’s second eigenvalue conjecture and related problems. *CoRR*, cs.DM/0405020, 2004.
- 13 A. Galanis, L. A. Goldberg, and M. Jerrum. Approximately counting H -colorings is #BIS-hard. *SIAM Journal on Computing*, 45(3):680–711, 2016.
- 14 A. Galanis, D. Štefankovič, and E. Vigoda. Inapproximability for antiferromagnetic spin systems in the tree nonuniqueness region. *J. ACM*, 62(6), 2015.
- 15 A. Galanis, D. Štefankovič, E. Vigoda, and L. Yang. Ferromagnetic Potts model: Refined #BIS-hardness and related results. *SIAM Journal on Computing*, 45(6):2004–2065, 2016.
- 16 L. A. Goldberg and M. Jerrum. Approximating the partition function of the ferromagnetic Potts model. *J. ACM*, 59(5), 2012.
- 17 L. A. Goldberg, S. Kelk, and M. Paterson. The complexity of choosing an H -coloring (nearly) uniformly at random. *SIAM Journal on Computing*, 33(2):416–432, 2004.
- 18 A. Govorov, J.-Y. Cai, and M. Dyer. A dichotomy for bounded degree graph homomorphisms with nonnegative weights. *arXiv:2002.02021*, 2020.
- 19 C. Gruber and H. Kunz. General properties of polymer systems. *Communications in Mathematical Physics*, 22(2):133–161, 1971.

- 20 W. H. Haemers. Interlacing eigenvalues and graphs. *Linear Algebra and its applications*, 226(228):593–616, 1995.
- 21 Tyler Helmuth, Will Perkins, and Guus Regts. Algorithmic Pirogov–Sinai theory. *Probability Theory and Related Fields*, pages 1–45, 2019.
- 22 S. Hoory, N. Linial, and A. Wigderson. Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43(4):439–561, 2006.
- 23 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 2019)*, pages 2235–2247, 2019.
- 24 N. Kahale. Eigenvalues and expansion of regular graphs. *J. ACM*, 42(5):1091–1106, 1995. doi:10.1145/210118.210136.
- 25 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)*, volume 145 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 34:1–34:12, 2019.
- 26 A. Sly and N. Sun. Counting in two-spin models on d -regular graphs. *Ann. Probab.*, 42(6):2383–2416, November 2014.
- 27 R. M. Tanner. Explicit concentrators from generalized N -gons. *SIAM Journal on Algebraic Discrete Methods*, 5(3):287–293, 1984.