

APPROXIMATELY COUNTING INDEPENDENT SETS IN BIPARTITE GRAPHS VIA GRAPH CONTAINERS

MATTHEW JENSSEN, WILL PERKINS, AND ADITYA POTUKUCHI

ABSTRACT. By implementing algorithmic versions of Sapozhenko’s graph container methods, we give new algorithms for approximating the number of independent sets in bipartite graphs. Our first algorithm applies to d -regular, bipartite graphs satisfying a weak expansion condition: when d is constant, and the graph is a bipartite $\Omega(\log^2 d/d)$ -expander, we obtain an FPTAS for the number of independent sets. Previously such a result for $d > 5$ was known only for graphs satisfying the much stronger expansion conditions of random bipartite graphs. The algorithm also applies to weighted independent sets: for a d -regular, bipartite α -expander, with $\alpha > 0$ fixed, we give an FPTAS for the hard-core model partition function at fugacity $\lambda = \Omega(\log d/d^{1/4})$. Finally we present an algorithm that applies to all d -regular, bipartite graphs, runs in time $\exp\left(O\left(n \cdot \frac{\log^3 d}{d}\right)\right)$, and outputs a $(1 + o(1))$ -approximation to the number of independent sets.

1. INTRODUCTION

Let $\mathcal{I}(G)$ denote the set of independent sets of a graph G and $i(G) = |\mathcal{I}(G)|$ denote the number of independent sets of G . Computing $i(G)$ is a #P-hard problem, even when restricted to bounded degree, bipartite graphs [47]. Even approximating $i(G)$ (to a constant or even subexponential factor) remains NP-hard, even when restricted to d -regular graphs with $d \geq 6$ [53].

Intuitively, one might expect the problem of approximating $i(G)$ to be easier on the class of bipartite graphs; for one, there is a polynomial-time algorithm to find a maximum-size independent set in a bipartite graph while the corresponding problem is NP-hard for general graphs. Dyer, Goldberg, Greenhill, and Jerrum [11] defined the counting problem #BIS (bipartite independent set) and showed that several natural combinatorial counting problems are as hard to approximate as #BIS. These problems include approximately counting stable matchings, approximating the ferromagnetic Potts model partition function ($q \geq 3$) [24, 15], counting the number of q -colorings in a bipartite graph ($q \geq 3$), and approximating the ferromagnetic Ising model partition with non-uniform external fields [23].

The search for approximation algorithms for $i(G)$ that exploit bipartite structure generally falls into two categories. The first approach finds classes of graphs on which polynomial-time approximation algorithms exist. Liu and Lu [41] gave the first such algorithm, providing an FPTAS for the number of independent sets in bipartite graphs in which degrees on one side are bounded by 5, while degrees on the other side are unrestricted. Another line of work in this direction includes [31] in which approximation algorithms are given for the hard-core partition function $Z_G(\lambda)$ (counting weighted independent sets) in bounded-degree, bipartite expander graphs, based on two tools from statistical physics: polymer models and the cluster expansion (following [28]). This work was followed by several improvements, extensions, and

generalizations including [40, 5, 7, 13, 12, 14, 8]. All of these algorithms are ‘low-temperature’ algorithms: they exploit the fact that on a bipartite graph with sufficient expansion, most (weighted) independent sets have few vertices from one side of the bipartition; that is, they are close to one of the two ground states consisting of all subsets of one side. In contrast, the algorithms of Weitz [55] (an FPTAS for $i(G)$ in graphs of maximum degree at most 5) and Liu and Lu [41] are ‘high-temperature’ algorithms, exploiting correlation decay properties of the uniform distribution on $\mathcal{I}(G)$ (or more generally the hard-core model of weighted independent sets) for small vertex degrees.

The second category of bipartite approximation algorithms are those that apply to all bipartite graphs, with running times better than what is known for general graphs. The main result in this direction is the algorithm of Goldberg, Lapinskas, and Richerby [25] which provides an ϵ -relative approximation to $i(G)$ for bipartite graphs, running in time $O(2^{.2372n}(1/\epsilon)^{O(1)})$, beating the best known running time for general graphs of $O(2^{.268n}(1/\epsilon)^{O(1)})$ from the same paper (in comparison, the best known running time for exact counting algorithms for general graphs is $O(2^{.3022n})$ [21]).

1.1. Main results. In this paper we use tools from combinatorics, namely the graph container method of Sapozhenko to give new approximate counting algorithms for independent sets in bipartite graphs. Our first result is an FPTAS for $i(G)$ for weakly expanding, d -regular bipartite graphs, for constant d .

For $z > 0$, we say that $\hat{z}$ is an ϵ -relative approximation to z if $(1 - \epsilon) \leq z/\hat{z} \leq (1 + \epsilon)$. A fully polynomial-time approximation scheme (FPTAS) is an algorithm that for every $\epsilon > 0$ outputs an ϵ -relative approximation to $i(G)$ and runs in time polynomial in $|V(G)|$ and $1/\epsilon$. We let μ_G denote the uniform distribution on $\mathcal{I}(G)$. A polynomial-time sampling scheme for μ_G runs in time polynomial in $|V(G)|$ and $1/\epsilon$ and outputs an independent set with distribution $\hat{\mu}$ within ϵ total variation distance of μ_G .

For $\alpha > 0$, we say that a d -regular bipartite graph G with bipartition X, Y is a *bipartite α -expander* if for every $A \subseteq X$ and $A \subseteq Y$ of size at most $|X|/2$, we have

$$(1) \quad |N(A)| \geq (1 + \alpha)|A|,$$

where $N(A) := \{v \in V(G) : u \sim v \text{ for some } u \in A\}$, the neighborhood of A .

Theorem 1. *There exists a constant $C_1 > 0$ so that for d fixed and sufficiently large, and $\alpha = \frac{C_1 \log^2 d}{d}$, there is an FPTAS for $i(G)$ and a polynomial-time sampling scheme for μ_G for the class of d -regular, bipartite α -expander graphs.*

Previous results on expander graphs include an FPTAS for $i(G)$ in the case that G is a typical d -regular, random bipartite graph [31, 40, 8]. These algorithms exploit the very strong expansion conditions satisfied by a random graph: sets of size $\tilde{O}(n/d)$ on each side of the bipartition expand by a factor $\tilde{\Omega}(d)$.

A natural and powerful generalization of the notion of counting independent sets is to consider weighted independent sets in the form of the hard-core model partition function (also known as the independence polynomial)

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

The corresponding probability measure on independent sets, known as the hard-core model, is given by

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

Taking $\lambda = 1$ gives $i(G)$ and μ_G . In previous works, an FPTAS for $Z_G(\lambda)$ for bounded-degree, bipartite α -expanders was obtained for λ much larger than 1, specifically $\lambda \geq K\Delta^{c/\alpha}$ for constants $c, K > 1$ [31, 7, 12]. In particular, under the expansion conditions of Theorem 1, these algorithms require $\lambda \geq \Delta^{\tilde{\Omega}(\Delta)}$.

Our next result adapts the algorithm of Theorem 1 to work for bipartite α -expanders for λ much smaller than 1.

Theorem 2. *For every $\alpha > 0$ there exists a constant $C_2 > 0$ so that for $d \geq 3$ and $\lambda > \frac{C_2 \log d}{d^{1/4}}$ there is an FPTAS for $Z_G(\lambda)$ and a polynomial-time sampling scheme for $\mu_{G,\lambda}$ for the class of d -regular, bipartite α -expanders.*

In fact in proving Theorems 1 and 2 we can interpolate between the two cases, letting the lower bound on λ shrink as the expansion condition gets stronger. The more general condition obtained is essentially the same as the condition for slow mixing of Glauber dynamics given by Galvin and Tetali [20].

Our next result is an approximation algorithm for $i(G)$ for all (not necessarily expanding) d -regular bipartite graphs G , where d may either be constant or growing with the size of the graph, n . When $d \rightarrow \infty$ as $n \rightarrow \infty$, the algorithm runs in subexponential time. This algorithm estimates $i(G)$ by separating the contribution from non-expanding sets and expanding sets on each side of the bipartition and uses the expander algorithm of Theorem 1 as a subroutine.

Theorem 3. *For every $c > 0$ and $d \geq 2$, there is a randomized algorithm that given a d -regular, n -vertex bipartite graph G outputs an n^{-c} -relative approximation to $i(G)$ with probability at least $2/3$ and runs in time*

$$\exp\left(O\left(\frac{n \log^3 d}{d}\right)\right).$$

Note that while the algorithm of Theorem 3 applies more generally than that of Theorem 1, the algorithmic guarantees are weaker in several senses (in addition to the slower running time): the algorithm uses randomness and the accuracy is limited to being polynomially small in n . Moreover we do not provide a corresponding sampling algorithm for Theorem 3; the problem is not self-reducible for regular graphs, and so there is not a direct reduction of approximate sampling to counting. We can overcome this in Theorem 1, however, using the self-reducibility of polymer models.

1.2. Background. The study of independent sets plays a central role in combinatorics since a broad range of problems can be phrased in terms of independent sets in graphs (and more generally hypergraphs). The container method is one of the most powerful combinatorial tools for studying independent sets. At a high level, the container method exploits a clustering phenomenon exhibited by independent sets which can often be used to deduce useful structural information for typical independent sets in a given graph or hypergraph. For graphs, the method was developed in the early 1980's by Kleitman and Winston [35, 36] and was independently discovered by Sapozhenko who used the method to enumerate independent sets in regular graphs [49, 50]. See the survey of Samotij [48] for background and

examples. The full potential of the container method only recently became apparent with the powerful generalization of the method to the context of hypergraphs developed by Saxton and Thomason [52] and Balogh, Morris and Samotij [4]. These developments have made the container method one of the most influential tools in modern combinatorics.

In this paper, we will only need ideas from the theory of graph containers, and our treatment is most closely related to that of Sapozhenko [49]. A canonical application of Sapozhenko’s version of the container method is his proof that the number of independent sets in the d -dimensional hypercube, Q_d , is asymptotically equal to $2\sqrt{e} \cdot 2^{2^{d-1}}$ (a result originally proved by Korshunov and Sapozhenko [38]). See also Galvin’s exposition of Sapozhenko’s proof [18].

Recently, Jenssen and Perkins [32] used Sapozhenko’s graph containers for Q_d (and Galvin’s extension to weighted independent sets [17]) along with the theory of polymer models and the cluster expansion to deduce refined counting estimates and detailed probabilistic information for independent sets in Q_d . The polymer models they consider closely resemble those used by Jenssen, Keevash, and Perkins [31] to design approximate counting algorithms in bipartite expander graphs. However, the hypercube Q_d is a far weaker expander than the graphs considered in [31], ruling out a direct application of the cluster expansion method. To overcome this obstacle, they showed that the container method arises naturally as a tool for proving cluster expansion convergence. This synthesis of the cluster expansion and the container method has now seen a number of applications to enumeration problems in combinatorics [30, 33, 9, 3].

Graph containers have in fact been used previously in a number of results in theoretical computer science, including [20, 16, 19]. These results use Sapozhenko’s techniques to prove a type of negative algorithmic result: that the Glauber dynamics (a local Markov chain) for sampling independent sets (or q -colorings) in bipartite graphs with sufficient expansion exhibit torpid mixing; that is, the mixing time is exponentially large in the size of the graph.

In this paper, we return to the algorithmic context to prove positive results, showing that the container method can be algorithmically implemented to design efficient approximate counting and sampling algorithms for a broad class of bipartite graphs. There is in fact a connection between the torpid mixing results above and the current algorithmic results: a key step in using polymer models and the cluster expansion as algorithms to approximately count independent sets in bipartite graphs is showing that most (weighted) independent sets can be accounted for by one of two polymer models that capture small deviations from independent sets that are fully contained in one side of the bipartition. This step (Lemmas 16 below) amounts to proving a kind of torpid mixing result, closely related to the result of Galvin and Tetali who proved torpid mixing for d -regular bipartite α -expanders, for $\alpha = \Omega(\log^3 d/d)$ [20, Corollary 1.3], similar to the class of graphs to which Theorem 1 applies.

While Theorem 1 gives efficient approximate counting and sampling algorithms for a larger class of bipartite expander graphs than previous approaches, it does not say much about the tractability of $\#BIS$, beyond ruling out a class of graphs as hard examples. In another prominent approximation problem of undetermined complexity, the Unique Games problem, efficient approximation algorithms for expander graphs [2, 37, 44] were later leveraged via graph decompositions into expanding pieces to find subexponential-time algorithms for all instances [1]. This suggests a very natural goal of finding much faster algorithms for approximating $\#BIS$.

Question 1. *Is there a subexponential-time approximation algorithm for $\#BIS$?*

Theorem 3 makes a small step in this direction, giving subexponential-time algorithms for regular graphs of growing degree, using an expander algorithm as a subroutine to account for the contribution to $i(G)$ from expanding sets. Non-expanding sets are accounted for separately, also using ideas from graph containers. It is tempting to think that algorithmic graph decomposition results (e.g. [45, 6]) could be used in conjunction with expander algorithms for $\#BIS$, but it is not clear to us how to use results that bound the number of edges between expanding pieces to obtain improved approximation algorithms for $\#BIS$, and so this remains an interesting direction for future research.

1.3. Outline. In Section 2 we give a brief warm-up to show how ideas from graph containers can be used for approximately counting independent sets in graphs. In Section 3 we present the graph container results we will need for the main algorithmic results. In Section 4 we define a polymer model which we will use to approximate the number of independent sets in expander graphs. In Section 5 we prove Theorem 1. In Section 6 we prove Theorem 3. In Section 7 we extend the graph container results of Section 3 and the polymer model of Section 4 to the case of weighted independent sets to prove Theorem 2.

2. WARM-UP: ALGORITHMS FROM CONTAINERS

In this section, we demonstrate how ideas from graph containers can be used to design faster algorithms to approximately count independent sets in (not necessarily bipartite) d -regular graphs. This section is intended as a warm-up which introduces the interplay between the container method and counting algorithms, and is not needed for the proofs of Theorems 1, 2 and 3.

Let us assume that we have an algorithm A that runs in time $a(n)$ on a general (not necessarily d -regular) graph G on n vertices and outputs an ϵ -relative approximation to $i(G)$. We will show that if G is d -regular for $d = \omega(1)$, one can obtain an algorithm running in time $a(n/2) \cdot 2^{o(n)}$.

Let $T := n \ln d / d$. We note that throughout the paper, we let $\ln$ denote the natural logarithm and let $\log$ denote $\log_2$. We may then write

$$i(G) = i_{<T}(G) + i_{\geq T}(G)$$

where $i_{<T}(G)$ is the number of independent sets of G of size less than T , and $i_{\geq T}(G)$ is the number of independent sets of size at least T . One can compute $i_{<T}(G)$ by brute-force in time $\binom{n}{T} \cdot \text{poly}(n)$. To compute $i_{\geq T}(G)$, we use a subtle idea of Sapozhenko [51] (also see [29], [34]) with a tighter analysis. First, we fix an ordering $\prec$ of the vertices of the graph G . Given a subset $S \subseteq V(G)$ and vertex $v \in V(G)$, we let $d_S(v)$ denote the number of neighbor of v in S . The following algorithm takes an independent set I of size at least T as input and returns a “certificate” ξ :

- $t, i \leftarrow 0, \xi \leftarrow (0)^n, V_0 \leftarrow V(G)$
- while $t \leq T$ do
 - $v \leftarrow \arg\max_{v \in V_i} d_{V_i}(v)$ with ties broken using $\prec$
 - if $v \in I$
 - * $V_{i+1} \leftarrow V_i \setminus (\{v\} \cup N(v))$
 - * $t \leftarrow t + 1, i \leftarrow i + 1, \xi_i \leftarrow 1$.
 - if $v \notin I$

- * $V_{i+1} \leftarrow V_i \setminus \{v\}$
- * $i \leftarrow i + 1.$
- return ξ

It is worth noting that ξ is not the indicator vector of any subset of $V(G)$, but rather an indicator vector describing the steps at which a vertex $v \in I$ is removed from V_i . Assume that the algorithm runs for $k = k(\xi)$ steps (i.e., k is that final value that i takes in the execution of the algorithm) when the output is ξ and let $V_\xi = V_k$. A key property of ξ is that it determines both V_ξ , and $I \cap (V(G) \setminus V_\xi)$. We may therefore group independent sets according to their certificate ξ and write

$$(2) \quad i_{\geq T}(G) = \sum_{\substack{\xi \in \{0,1\}^n, \\ |\xi| = T}} i(G[V_\xi]).$$

The advantage of this expression is that the number of possible certificates $\binom{n}{T}$ is relatively small and, as we show next, the size of each V_ξ is close to $n/2$. We have therefore exploited the clustering phenomenon of independent sets characteristic of the container method to effectively halve the size of the input graph to our algorithm (albeit at the expense of losing regularity of the graph).

We now give the argument to bound the size of each V_ξ . For any subset $S \subseteq V(G)$, let us denote e_S to be the number of edges in the induced subgraph $G[S]$. A straightforward extension of the Hoffman bound (see e.g. [10, 43, 22, 27]), and using the fact that the smallest eigenvalue of a d -regular graph is at least $-d$, gives us

$$2e_S \geq \frac{2d}{n} |S|^2 - d|S|.$$

So for $i \leq k$, we have

$$(3) \quad \max_{v \in V_i} d_{V_i}(v) \geq \frac{2e_{V_i}}{|V_i|} \geq \frac{d}{n} (2|V_i| - n),$$

and so if $v \in I$, then $|V_{i+1}| \leq |V_i| (1 - \frac{2d}{n}) + d$, and so at the end of the algorithm, we have $|V_k| \leq \frac{n}{2} + O(\frac{n \ln d}{d})$.

Thus, if we have an algorithm A running in time $a(n) = a(n, \epsilon)$ on general graphs on n vertices which outputs an ϵ -relative approximation to $i(G)$ for a general graph G on n vertices, then an ϵ -relative approximation to (2) may be computed in time $\binom{n}{T} \cdot a(\frac{n}{2} + O(\frac{n \ln d}{d}))$. Combining this with the brute-force computation of $i < T$ gives an algorithm running in time

$$\text{poly}(n) \cdot \left(\frac{n}{\frac{n \ln d}{d}} \right)^2 \cdot a\left(\frac{n}{2} + O\left(\frac{n \ln d}{d}\right)\right)$$

that outputs an ϵ -relative approximation to $i(G)$ for a d -regular graph G on n vertices. So in particular, if $d = \omega(1)$, then using the algorithm of Goldberg, Lapinskas, and Richerby [25] as a blackbox, we obtain a $2^{(0.134+o(1))n}$ time algorithm for approximating $i(G)$ in d -regular graphs on n vertices. We will see below in Section 6 that with more sophisticated ideas this running time can be made subexponential provided that the error $\epsilon = n^{-c}$ for some $c > 0$ (Theorem 3).

3. GRAPH CONTAINER LEMMAS

In this section we introduce results from the theory of graph containers that will be key for the algorithms of Theorems 1, 2 and 3. Many of the ideas here have their roots in the aforementioned container method of Sapozhenko [49].

We assume throughout that G is a d -regular bipartite graph on $2n$ vertices with bipartition X, Y , so $|X| = |Y| = n$. For a subset $A \subseteq X$, we use W to denote $N(A)$. Let us define $[A] := \{u \in X \mid N(u) \subseteq W\}$ to be the *closure* of A . Let us call A *2-linked* if the subgraph of G^2 (the square of the graph G) induced by A is connected. We say that A is *expanding* if

$$|W| - |[A]| \geq (C_1/2) \frac{\log^2 d}{d} |W|,$$

where the constant $C_1 > 0$ will be sufficiently large and chosen later. Otherwise, we say that A is *non-expanding*.

Let $\mathcal{G}(v, a, w)$ denote the set of 2-linked expanding sets A such that $A \ni v$, $|[A]| = a$, and $|W| = w$. Let $\mathcal{G}'(v, a)$ be the set of 2-linked non-expanding sets $A \ni v$ such that $A = [A]$, and $|A| = a$.

Remark: Observe that if G is a bipartite α -expander (as defined at (1)) then

$$(4) \quad |N(A)| \geq (1 + \alpha)|[A]|$$

for each $A \subseteq X$ or $A \subseteq Y$ such that $|[A]| \leq n/2$. Indeed, if $|[A]| \leq n/2$, then $|N(A)| = |N([A])| \geq (1 + \alpha)|[A]|$. Since $\alpha \leq 1$, inequality (4) implies that $|[A]| \leq |N(A)|(1 - \alpha/2)$, and so

$$(5) \quad |N(A)| - |[A]| \geq (\alpha/2)|N(A)|$$

for each A such that $|[A]| \leq n/2$. In the following, condition (5) is slightly more convenient to work with and motivates our definition of an expanding set.

We now state our main technical lemmas. The first bounds the number of expanding sets and the second bounds the number of non-expanding sets (and gives an algorithm to enumerate them).

Lemma 4. *There is an absolute constant $c_1 > 0$ such that for every v, a, w we have*

$$|\mathcal{G}(v, a, w)| \leq 2^{w - c_1(w - a)}.$$

Lemma 5. *There is an absolute constant $c_2 > 0$ such that for every v, a , we have*

$$|\mathcal{G}'(v, a)| \leq 2^{c_2 \frac{a \log^2 d}{d}}.$$

Moreover, there is an algorithm running in time $2^{O\left(\frac{a \log^2 d}{d}\right)} \cdot \text{poly}(n)$ that outputs the set $\mathcal{G}'(v, a)$.

Recall that for a subset $A \subseteq X$, we use W to denote $N(A)$, with $|[A]| = a$, $|W| = w$. For every $s > 0$, let $W_s = \{u \in W \mid d_{[A]}(u) \geq s\}$. We next define a notion of an approximation of the set W (which in turn determines $[A]$).

Definition 6. *A set $F \subseteq W$ is an essential subset for A if*

- (1) $F \supseteq W_{d/2}$
- (2) $N(F) \supseteq [A]$.

The next lemma gives a family $\mathcal{C}(v, a, w) \subseteq 2^Y$ that contains an essential subset for each member of $\mathcal{G}(v, a, w)$. Crucially the set of approximating sets $\mathcal{C}(v, a, w)$ is far smaller than $\mathcal{G}(v, a, w)$.

Lemma 7. *There is a family $\mathcal{C}(v, a, w) \subseteq 2^Y$ of size at most*

$$2^{\frac{16w \log^2 d}{d}}$$

such that $\mathcal{C}(v, a, w)$ contains an essential subset for every 2-linked set $A \ni v$ such that $|[A]| = a$ and $|W| = w$. Moreover, there is an algorithm running in time $2^{\frac{16w \log^2 d}{d}} \cdot \text{poly}(n)$ that outputs the set $\mathcal{C}(v, a, w)$.

The following lemma of Park [46] strengthens a result of Sapozhenko [49] (the lemma is proved implicitly in [34] also).

Lemma 8. *There is an absolute constant $c_3 > 0$ such that the following holds: for every $F \subseteq Y$, let $\mathcal{G}(F, a, w)$ be the set of expanding 2-linked sets $A \subseteq X$ such that $|[A]| = a$, $|W| = w$, and F is an essential subset for A . Then*

$$|\mathcal{G}(F, a, w)| \leq 2^{w - c_3(w-a)}.$$

With these lemmas in hand, we now prove Lemma 4 and Lemma 5.

Proof of Lemma 4. First note that

$$(6) \quad \mathcal{G}(v, a, w) \subseteq \bigcup_{F \in \mathcal{C}(v, a, w)} \mathcal{G}(F, a, w)$$

and so by Lemmas 7 and 8, we have that

$$\begin{aligned} |\mathcal{G}(v, a, w)| &\leq \sum_{F \in \mathcal{C}(v, a, w)} |\mathcal{G}(F, a, w)| \\ &\leq |\mathcal{C}(v, a, w)| \cdot \max_{F \in \mathcal{C}(v, a, w)} |\mathcal{G}(F, a, w)| \\ &\leq 2^{\frac{16w \log^2 d}{d}} \cdot 2^{w - c_3(w-a)} \\ &\leq 2^{w - (\frac{c_3}{2})(w-a)} \end{aligned}$$

where for the last inequality we used that $w - a \geq (C_1/2) \frac{\log^2 d}{d} w$ by the definition of an expanding set and assumed that $C_1 > 64/c_3$. $\square$

Proof of Lemma 5. Let F be an essential subset for A where A is non-expanding. Note that each vertex in $W \setminus F$ has at least $d/2$ neighbors in $[A]^c$, and there are at most $d(w - a)$ edges between W and $[A]^c$ since G is d -regular. It follows that

$$(d/2) \cdot |W \setminus F| \leq d(w - a)$$

and so

$$|W \setminus F| \leq 2(w - a) \leq C_1 \frac{\log^2 d}{d} w.$$

Moreover, $W \setminus F \subseteq N^2(F)$, and $|N^2(F)| \leq wd^2$, and so there are at most

$$\left(\frac{wd^2}{\leq C_1 w (\log^2 d)/d} \right) \leq 2^{O\left(\frac{a \log^3 d}{d}\right)}$$

choices for W , each of which determines a $[A]$. Let $\mathcal{G}'(F, a)$ denote the collection of $A \subseteq X$ such that $A = [A]$, $v \in A$, A is 2-linked, non-expanding and F is an essential subset for A . Then by the above,

$$|\mathcal{G}'(F, a)| = 2^{O\left(\frac{a \log^3 d}{d}\right)}.$$

Moreover, we can generate the set $\mathcal{G}'(F, a)$ in time $2^{O\left(\frac{a \log^3 d}{d}\right)} \text{poly}(n)$ by listing each set W that is a union of F with a subset of $N^2(F)$ of size at most $C_1 w(\log^2 d)/d$, generating the corresponding closed set $[A]$ such that $N(A) = W$, and checking it satisfies the required conditions.

Now, by Lemma 7,

$$(7) \quad \mathcal{G}'(v, a) \subseteq \bigcup_{w \in \left[a, a\left(1 + C_1 \frac{\log^2 d}{d}\right)\right]} \bigcup_{F \in \mathcal{C}(v, a, w)} \mathcal{G}'(F, a)$$

and so

$$|\mathcal{G}'(v, a)| \leq \left(C_1 \frac{a \log^2 d}{d}\right) \cdot 2^{O\left(\frac{a \log^3 d}{d}\right)} = 2^{O\left(\frac{a \log^3 d}{d}\right)}.$$

Similarly, by (7), Lemma 7, and the above algorithm for generating $\mathcal{G}'(F, a)$, we may generate $\mathcal{G}'(v, a)$ in time $2^{O\left(\frac{a \log^3 d}{d}\right)} \cdot \text{poly}(n)$. $\square$

We will need a covering result originally due to Lovász [42] and Stein [54].

Theorem 9. *Let H be a bipartite graph on vertex sets P and Q where the degree of each vertex in P is at least a and the degree of each vertex in Q is at most b . Then there is subset $Q' \subseteq Q$ of size at most $\frac{|Q|}{a}(1 + \ln b)$ such that $P \subseteq N(Q')$.*

We record the following corollary of Theorem 9 which we will use in the proof of Lemma 7.

Corollary 10. *Let $A \subseteq X$ be 2-linked and let $A \ni v$. Then the following hold:*

- (1) *There exists a 2-linked subset $A' \subseteq A$ of size at most $2\frac{a}{d} \ln d + 2\frac{w}{d}$ such that $A' \ni v$ and $N(A')$ is an essential subset for A .*
- (2) *There exists a 2-linked subset $A'' \subseteq A$ of size at most $2\frac{a}{d} \ln d + 2\frac{w}{d} + 2(w - a)$ such that $N(A'') = W$.*

Proof. We begin by proving (1). Let $A_0 \subseteq [A]$ be a maximal subset of vertices containing v with pairwise disjoint neighborhoods. Clearly $|A_0| \leq \frac{w}{d}$ and $N^2(A_0) \supseteq A$. Theorem 9 guarantees a subset $A_1 \subseteq A$ of size at most $2\frac{a}{d} \ln d$ such that $W_{d/2} \subseteq N(A_1)$. Suppose $A_0 \cup A_1$ is not 2-linked. We observe that $A_0 \cup A_1$ has at most $\frac{w}{d}$ 2-linked components. Indeed, this is true since $N(A_0 \cup A_1) \subseteq W$ and each two linked component covers at least d vertices of W . Since $[A]$ is 2-linked, it follows that one can choose a subset $A_2 \subseteq [A]$ of size at most $\frac{w}{d}$ such that $A' := A_0 \cup A_1 \cup A_2$ is 2-linked. We note that $|A'| \leq 2\frac{a}{d} \ln d + 2\frac{w}{d}$. To show that $N(A')$ is an essential subset for A observe that

- $N^2(A') \supseteq N^2(A_0) \supseteq [A]$, and
- $W_{d/2} \subseteq N(A_1) \subseteq N(A')$.

We now turn to (2). Note that each vertex in $W \setminus W_{d/2}$ has at least $d/2$ neighbors in $[A]^c$, and there are at most $d(w - a)$ edges between W and $[A]^c$. It follows that

$$(d/2) \cdot |W \setminus W_{d/2}| \leq d(w - a)$$

and so

$$|W \setminus W_{d/2}| \leq 2(w - a).$$

Let $A_3 \subseteq A$ be a minimal cover of $W \setminus W_{d/2}$. We have that $|A_3| \leq |W \setminus W_{d/2}| \leq 2(w - a)$, and every vertex of A_3 is at a distance 2 from some vertex in $A_0 \subseteq A'$ by the maximality of A_0 . Thus $A'' = A' \cup A_3$ is 2-linked, $|A''| \leq 2(w - a) + |A'|$ and $N(A'') \supseteq W$, completing the proof. $\square$

We now prove Lemma 7 using techniques similar to those found in [49].

Proof of Lemma 7. Let $A \ni v$ be a 2-linked subset as in the statement of the lemma. By Corollary 10, there exists a 2-linked subset $A' \subseteq A$ of size at most $4\frac{w}{d} \ln d$ such that $v \in A'$ and $N(A')$ is an essential subset for A . In view of this, we let $\mathcal{B}(v, a, w)$ be the set of all 2-linked subsets of A , containing v , of size at most $4\frac{w}{d} \ln d$ and let

$$\mathcal{C}(v, a, w) = \{N(A) : A \in \mathcal{B}(v, a, w)\}.$$

It remains to upper bound the size of $\mathcal{C}(v, a, w)$ and describe an algorithm that outputs the set. Note that $|\mathcal{C}(v, a, w)| \leq |\mathcal{B}(v, a, w)|$ and $|\mathcal{B}(v, a, w)|$ is at most the number of trees in G^2 containing v as the root with at most $4\frac{w \ln d}{d}$ vertices. Note that the maximum degree in G^2 is at most $d(d - 1)$. So $\mathcal{B}(v, a, w)$ can be enumerated using the following procedure:

- (1) Assume an ordering on the neighbors of every vertex in the graph G^2 . Let us use v_i to denote the i 'th neighbor of a vertex v . Let us also denote $v_0 = v$.
- (2) Generate a list $S \in \{0, \dots, d(d - 1)\}^{8\frac{w \ln d}{d}}$.
- (3) Consider the set $T_S = \{v^{(0)}, \dots, v^{(s)}\}$ where $v^{(0)} = v$ and $v^{(i)} = v_{S_i}^{(i-1)}$.
- (4) If $|T_S| \leq 4\frac{w \ln d}{d}$, then output T_S .

Consider any tree in G^2 with root v and $s \leq 4\frac{w \ln d}{d}$ nodes. There is at least one choice of the list S that causes the above procedure to output the vertices of this tree, namely, if $(S_1, \dots, S_{2s})$ is the DFS traverse order and $S_i = 0$ for $i \geq 2s$.

For each list S , the procedure takes $\text{poly}(n)$ time, and the number of possible lists is $(d(d - 1) + 1)^{8\frac{w}{d} \ln d} \leq 2^{\frac{16w \log^2 d}{d}}$. Therefore there is an algorithm running in time $2^{\frac{16w \log^2 d}{d}} \cdot \text{poly}(n)$ that outputs the set $\mathcal{C}(v, a, w)$. $\square$

4. POLYMER MODELS

In this section we introduce a variant of the polymer models used by Jenssen, Keevash, and Perkins [31] to obtain approximation algorithms for bipartite expander graphs. Polymer models originated in statistical physics (e.g. [26, 39]) as a means to study spin models on lattices. Recently they were used to design algorithms for spin models at low temperatures [28].

Fix a d -regular bipartite graph G with bipartition (X, Y) of size n each. A *polymer* of G is a 2-linked, expanding subset of X . Recall, from the previous section, that a set $A \subseteq X$ is

expanding if

$$(8) \quad |N(A)| - |[A]| \geq (C_1/2) \frac{|N(A)| \log^2 d}{d}$$

where C_1 is the constant from Theorem 1. Let $\mathcal{P}(G)$ denote the set of all polymers of G . The *weight* of a polymer γ is given by $2^{-|N(\gamma)|} =: w_\gamma$. We call two polymers γ and γ' *compatible* if $\gamma \cup \gamma'$ is not 2-linked, and *incompatible* otherwise, in which case we write $\gamma \approx \gamma'$. Let $\Omega(G)$ denote the collection all subsets of mutually compatible polymers (including the empty set). The polymer model partition function is

$$(9) \quad \Xi_G^X := \sum_{\Lambda \in \Omega(G)} \prod_{\gamma \in \Lambda} w_\gamma,$$

and the associated Gibbs distribution ν_G^X on $\Omega(G)$ defined by

$$\nu_G^X(\Lambda) = \frac{\prod_{\gamma \in \Lambda} w_\gamma}{\Xi_G^X} \text{ for } \Lambda \in \Omega(G).$$

When the weights of the polymer model are small enough one can hope to understand the partition function and the Gibbs distribution via perturbative techniques, and in particular, the cluster expansion.

For a tuple Γ of polymers, the *incompatibility graph*, $H(\Gamma)$, is the graph with vertex set Γ and an edge between any two incompatible polymers. A *cluster* Γ is an ordered tuple of polymers so that $H(\Gamma)$ is connected. Let us use $\mathcal{C}(G)$ to denote the set of all clusters of G . The *cluster expansion* of Ξ_G^X is the formal series expansion

$$(10) \quad \ln \Xi_G^X = \sum_{\Gamma \in \mathcal{C}(G)} \phi(H(\Gamma)) \prod_{\gamma \in \Gamma} w_\gamma,$$

where ϕ is the Ursell function defined by

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

We note that $\mathcal{C}(G)$ is an infinite set since it contains, for example, all tuples of the form $(\gamma, \dots, \gamma)$ where γ is a fixed polymer. The series in (10) is therefore an infinite sum. In our application, we will work with a truncated sum after establishing a fast enough rate of convergence. For these convergence bounds, (as in [28, 31]) we use a special case of the Kotecký–Preiss condition [39].

Theorem 11 ([39]). *Fix functions $f : \mathcal{P}(G) \rightarrow [0, \infty)$ and $g : \mathcal{P}(G) \rightarrow [0, \infty)$. Suppose that for every $\gamma \in \mathcal{P}(G)$, we have*

$$(11) \quad \sum_{\gamma' \not\approx \gamma} w_{\gamma'} e^{f(\gamma') + g(\gamma')} \leq f(\gamma).$$

Then the cluster expansion converges absolutely. Moreover, if we let $g(\Gamma) = \sum_{\gamma \in \Gamma} g(\gamma)$, then for every vertex v ,

$$(12) \quad \sum_{\substack{\Gamma \in \mathcal{C}(G) \\ \Gamma \ni v}} \left| \phi(\Gamma) \prod_{\gamma \in \Gamma} w_\gamma \right| e^{g(\Gamma)} \leq 1.$$

We will use this to obtain some relevant structural information about the polymer model. For polymers γ , define modified polymer weights $\tilde{w}_\gamma = 2^{-|N(\gamma)|} \cdot 2^{\frac{|\gamma| \log^2 d}{d}}$ and let

$$(13) \quad \tilde{\Xi}_G^X = \sum_{\Lambda \in \Omega(G)} \prod_{\gamma \in \Lambda} \tilde{w}_\gamma$$

be the modified polymer model partition function. We first show that this polymer model satisfies the Kotecký–Preiss condition for a suitable choice of functions f and g .

Lemma 12. *Consider the modified polymer model where the polymers are all the 2-linked subsets $A \subseteq X$ that are expanding, and the weight of a polymer γ is given by*

$$\tilde{w}_\gamma = 2^{-|N(\gamma)|} \cdot 2^{\frac{|\gamma| \log^2 d}{d}}.$$

Let $f(\gamma) = \ln 2 \cdot \frac{|\gamma| \log^2 d}{d}$, and $g(\gamma) = 2 \ln 2 \cdot \frac{|N(\gamma)| \log^2 d}{d}$. Then for d sufficiently large, every polymer γ satisfies (11). The conclusion also holds for the original polymer model with weights w_γ and the same choice of functions $f(\cdot), g(\cdot)$.

Proof of Lemma 12. It suffices to prove the claim for the modified polymer model since $\tilde{w}_\gamma \geq w_\gamma$ for all polymers γ .

We begin by bounding

$$(14) \quad \sum_{\gamma' \not\supset \gamma} \tilde{w}_{\gamma'} e^{f(\gamma') + g(\gamma')} \leq \sum_{v \in N^2(\gamma)} \sum_{\gamma' \ni v} \tilde{w}_{\gamma'} e^{f(\gamma') + g(\gamma')} \leq d^2 |\gamma| \cdot \max_{v \in N^2(\gamma)} \sum_{\gamma' \ni v} 2^{-|N(\gamma')|} e^{2f(\gamma') + g(\gamma')}.$$

Here we use $N^2(\gamma)$ to denote $N(N(\gamma))$, the second neighborhood of γ , and note that $|N^2(\gamma)| \leq d^2 |\gamma|$. Recall that we let $\mathcal{G}(v, a, w)$ denote the set of all polymers γ such that $\gamma \ni v$, $|\gamma| = a$, and $|N(\gamma)| = w$. Since each polymer γ' in the above sum is expanding by definition, it follows that $\gamma' \in \mathcal{G}(v, w - t, w)$ for some $t \geq (C_1/2) \frac{w \log^2 d}{d}$ and $w \geq d$. The sum right hand side of (14) is therefore at most

$$(15) \quad \sum_{w \geq d} \left(\sum_{t \geq (C/2) \frac{w \log^2 d}{d}} |\mathcal{G}(v, w - t, w)| 2^{-w} \cdot e^{4 \ln 2 \cdot \frac{w \log^2 d}{d}} \right).$$

By Lemma 4, $|\mathcal{G}(v, w - t, w)| \leq 2^{w - c_1 t}$. Moreover, by the proof of Lemma 4 we have that $c_1 > 32/C_1$. It follows that (15) is bounded above by

$$(16) \quad \sum_{w \geq d} e^{4 \ln 2 \cdot \frac{w \log^2 d}{d}} \left(\sum_{t \geq (C/2) \frac{w \log^2 d}{d}} 2^{-32t/C} \right) \leq d \sum_{w \geq d} e^{-\frac{8w \log^2 d}{d}} \leq d^2 e^{-8 \log^2 d}.$$

Combining (14)-(16) we conclude that

$$\sum_{\gamma' \not\supset \gamma} \tilde{w}_{\gamma'} e^{f(\gamma') + g(\gamma')} \leq d^4 |\gamma| e^{-8 \log^2 d} \leq f(\gamma),$$

for d sufficiently large. □

Therefore, Theorem 11 gives us that

$$(17) \quad \sum_{\substack{\Gamma \in \mathcal{C}(G) \\ \Gamma \ni v}} \left| \phi(\Gamma) \prod_{\gamma \in \Gamma} w_\gamma \right| e^{g(\gamma)} \leq 1,$$

and the same with $\tilde{w}_\gamma$ replacing w_γ .

Now define the exponential of the truncated cluster expansion

$$(18) \quad \Xi_G^X(\ell) := \exp \left(\sum_{\substack{\Gamma \in \mathcal{C}(G) \\ \|\Gamma\| \leq \ell}} \phi(\Gamma) \prod_{\gamma \in \Gamma} w_\gamma \right),$$

where $\|\Gamma\| := \sum_{\gamma \in \Gamma} |\gamma|$. The following lemma bounds the error in approximating Ξ_G^X by $\Xi_G^X(\ell)$.

Lemma 13. *We have for every $\ell \geq 1$,*

$$(19) \quad |\ln \Xi_G^X - \ln \Xi_G^X(\ell)| \leq n \cdot 2^{-2 \frac{\ell \log^2 d}{d}}.$$

In particular, if $\ell \geq \frac{d}{2 \log^2 d} \log(n/\epsilon)$, then

$$|\ln \Xi_G^X - \ln \Xi_G^X(\ell)| \leq \epsilon.$$

Proof. First recall that for a cluster Γ ,

$$g(\Gamma) = \sum_{\gamma \in \Gamma} g(\gamma) = 2 \ln 2 \frac{\log^2 d}{d} \sum_{\gamma \in \Gamma} |N(\gamma)| \geq 2 \ln 2 \frac{\log^2 d}{d} \|\Gamma\|.$$

It follows from (17) that

$$\begin{aligned} |\ln \Xi_G^X - \ln \Xi_G^X(\ell)| &= \left| \sum_{\substack{\Gamma \in \mathcal{C}(G) \\ \|\Gamma\| > \ell}} \phi(\Gamma) \prod_{\gamma \in \Gamma} w_\gamma \right| \\ &\leq \sum_v \sum_{\substack{\Gamma \in \mathcal{C}(G) \\ \Gamma \ni v \\ \|\Gamma\| > \ell}} \left| \phi(\Gamma) \prod_{\gamma \in \Gamma} w_\gamma \right| \\ &\leq n e^{-2 \ln 2 \frac{\ell \log^2 d}{d}} \\ &= n \cdot 2^{-2 \frac{\ell \log^2 d}{d}}. \end{aligned} \quad \square$$

Let $\|\mathbf{\Lambda}\| = \sum_{\gamma \in \mathbf{\Lambda}} |\gamma|$ for $\mathbf{\Lambda} \sim \nu_G^X$. We have the following large deviation result for $\|\mathbf{\Lambda}\|$, following [32, Lemma 16].

Lemma 14. *For any $\delta \in (0, 1)$, there is a $d_0 = d_0(\delta)$ such for $d \geq d_0$, we have*

$$\mathbb{P}(\|\mathbf{\Lambda}\| \geq \delta n) \leq 2^{-\left(\frac{\delta n \log^2 d}{2d}\right)}.$$

Proof. With $\tilde{\Xi}_G^X$ as defined at (13), we have that $\ln \tilde{\Xi}_G^X - \ln \Xi_G^X = \ln \mathbf{E} e^{\zeta \cdot \|\mathbf{A}\|}$ for $\zeta = \frac{\ln 2 \log^2 d}{d}$. Summing (17) over all v , and using the fact that $|N(\gamma)| \geq d$ for every γ , we get

$$(20) \quad \ln \tilde{\Xi}_G^X \leq \sum_v \sum_{\substack{\Gamma \in \mathcal{C}(G) \\ \Gamma \ni v}} \left| \phi(\Gamma) \prod_{\gamma \in \Gamma} \tilde{w}_\gamma \right| \leq n \cdot 2^{-2 \log^2 d}.$$

Therefore, we have

$$\ln \mathbf{E} e^{\zeta \cdot \|\mathbf{A}\|} \leq \ln \tilde{\Xi}_G^X \leq n \cdot 2^{-2 \log^2 d},$$

and so by Markov's inequality we have

$$\begin{aligned} \mathbb{P}(\|\mathbf{A}\| \geq \delta n) &\leq \exp\left(-\zeta \delta n + n \cdot 2^{-2 \log^2 d}\right) \\ &\leq 2^{-\frac{\delta n \log^2 d}{2d}}, \end{aligned}$$

where the last inequality follows because $(1/2) \ln 2 \geq (\delta \log^2 d)^{-1} \cdot d \cdot 2^{-2 \log^2 d}$ for large enough d . $\square$

Approximate counting and sampling for polymer models. Here, we will use an algorithm from [31] to approximate Ξ_G^X and approximately sample from ν_G^X .

Lemma 15. *Then there is an FPTAS for Ξ_G^X that runs in time $(n/\epsilon)^{O(d)}$. Moreover, for every $\epsilon > 0$ there is a randomized algorithm that runs in time $(n/\epsilon)^{O(d)}$ that outputs a configuration $\Lambda \in \Omega(G)$ with distribution ν_{alg}^X so that $\|\nu_{alg}^X - \nu_G^X\|_{TV} < \epsilon$.*

To give a sense of the above algorithms, we briefly describe the FPTAS for Ξ_G^X . Using Lemma 13, it is enough to compute $\Xi_G^X(\ell)$ (as defined in (18)) where $\ell = \left\lceil \frac{d}{2 \log^2 d} \log(n/\epsilon) \right\rceil$.

This may be done by

- (1) listing all clusters Γ of size at most ℓ ,
- (2) computing $\phi(\Gamma)$,
- (3) computing $\prod_{\gamma \in \Gamma} w_\gamma$, and
- (4) evaluating $\Xi_G^X(\ell)$ by exponentiating the truncated cluster expansion.

To approximately sample from ν_G^X we appeal to the self-reducibility of the abstract polymer model, see [28, Theorem 10].

Remark 1. *In the proof of Theorem 3 in Section 6, we will in fact need to approximate the partition function of various subgraphs of G . These subgraphs, $H \subseteq G$, all have the following form: H is the subgraph induced on vertex set $X' \cup Y'$ for some $X' \subseteq X$, $Y' \subseteq Y$ such that $N(X') \subseteq Y'$. For such a subgraph, we define a polymer of H to be an expanding 2-linked subset of X' and define the partition function $\Xi_H^{X'}$ in the obvious way. We note that the polymers of H are a subset of the polymers of G . In particular, since the polymer model on G satisfies the hypothesis of Theorem 11, the polymer model on H also satisfies the hypothesis of Theorem 11 with the same functions f and g . In particular, Lemmas 13, 14 and 15 also apply to the polymer model on H .*

5. AN ALGORITHM FOR EXPANDER GRAPHS: PROOF OF THEOREM 1

In this section, we prove Theorem 1. We assume throughout that G is a d -regular bipartite α -expander with $\alpha = C_1 \frac{\log d}{d}$. We let X, Y denote the vertex classes of G and let $n = |X| = |Y|$. We note that by (5), we have that $|N(A)| - |A| \geq (C_1/2) \frac{|N(A)| \log^2 d}{d}$ for every $A \subseteq X$ or $A \subseteq Y$ such that $|A| \leq n/2$, i.e. each such set A is expanding.

First we show that $i(G)$ can be approximated well by a linear combination of the polymer model partition functions Ξ_G^X and Ξ_G^Y (as defined in (9)). We may then use the algorithm of Section 4 to approximate Ξ_G^X and Ξ_G^Y . Recall that we let $\mathcal{I} = \mathcal{I}(G)$ denote the set of independent sets of G . For the sampling algorithms, we show that μ_G can be approximated by a mixture of probability distributions on $\mathcal{I}$ derived from the polymer measures ν_G^X, ν_G^Y . We let $\hat{\nu}_G^X$ denote the probability distribution on $\mathcal{I}$ defined as follows:

- (1) Sample a collection of compatible polymers Λ from the measure ν_G^X .
- (2) Set $I = J \cup \bigcup_{\gamma \in \Lambda} \gamma$ where J is a uniformly random subset of $Y \setminus \bigcup_{\gamma \in \Lambda} N(\gamma)$.

We define $\hat{\nu}_G^Y$ analogously and define the mixture

$$\hat{\mu}_G = \frac{\Xi_G^X}{\Xi_G^X + \Xi_G^Y} \hat{\nu}_G^X + \frac{\Xi_G^Y}{\Xi_G^X + \Xi_G^Y} \hat{\nu}_G^Y.$$

Lemma 16. *For all d sufficiently large and $n = n(d)$ sufficiently large,*

$$2^n \cdot (\Xi_G^X + \Xi_G^Y)$$

is an ϵ -relative approximation to $i(G)$ where $\epsilon = 2^{-\frac{n \log^2 d}{60d}}$. Moreover,

$$\|\mu_G - \hat{\mu}_G\|_{TV} \leq 2\epsilon.$$

Proof. Let us define

$$\mathcal{I}_X := \{I \in \mathcal{I} \mid \text{every 2-linked component of } I \cap X \text{ is expanding}\}$$

i.e. the set of all I such that $\nu_G^X(I) > 0$. We note that since the independence number of G is n , we have that for each $I \in \mathcal{I}$,

$$\min(|I \cap X|, |I \cap Y|) \leq n/2.$$

In particular, every component of either $I \cap X$ or $I \cap Y$ is expanding. Therefore $\mathcal{I} = \mathcal{I}_X \cup \mathcal{I}_Y$ and so

$$(21) \quad i(G) = |\mathcal{I}_X| + |\mathcal{I}_Y| - |\mathcal{I}_X \cap \mathcal{I}_Y|.$$

Note that

$$(22) \quad \mathcal{I}_X = 2^n \cdot \Xi_G^X$$

and moreover, the set of 2-linked components of $I \cap X$ for a uniformly chosen $I \in \mathcal{I}_X$ is distributed exactly accordingly to ν_G^X . It will suffice to bound the size of $\mathcal{I}_X \cap \mathcal{I}_Y$.

Letting $\mathcal{I}_X^\delta = \{I \in \mathcal{I}_X : |I \cap X| \leq \delta n\}$, for $\delta \in (0, 1)$, it follows by Lemma 14 that

$$(23) \quad |\mathcal{I}_X \setminus \mathcal{I}_X^\delta| \leq |\mathcal{I}_X| \cdot 2^{-\frac{\delta n \log^2 d}{2d}}.$$

Defining $\mathcal{I}_Y$, $\mathcal{I}_Y^\delta$, analogously and taking $\delta = 1/30$ we have

$$(24) \quad |\mathcal{I}_X^\delta \cap \mathcal{I}_Y^\delta| \leq \sum_{i=0}^{2\delta n} \binom{2n}{i} \leq 2^n \cdot 2^{-\frac{n \log^2 d}{60d} - 1} \leq |\mathcal{I}_X| \cdot 2^{-\frac{n \log^2 d}{60d} - 1}.$$

By (22), (23) and (24) we conclude that

$$(25) \quad |\mathcal{I}_X \cap \mathcal{I}_Y| \leq |\mathcal{I}_X \setminus \mathcal{I}_X^\delta| + |\mathcal{I}_Y \setminus \mathcal{I}_Y^\delta| + |\mathcal{I}_Y^\delta \cap \mathcal{I}_X^\delta| \leq 2^n (\Xi_G^X + \Xi_G^Y) 2^{-\frac{n \log^2 d}{60d}},$$

and therefore, by (21),

$$(26) \quad 2^n (\Xi_G^X + \Xi_G^Y) \cdot \left(1 - 2^{-\frac{n \log^2 d}{60d}}\right) \leq i(G) \leq 2^n (\Xi_G^X + \Xi_G^Y).$$

□

Now we prove Theorem 1.

Proof of Theorem 1. Set $\epsilon_0 = 2^{-\frac{n \log^2 d}{60d}}$. First suppose $\epsilon \leq 2\epsilon_0$, then $i(G)$ may be computed exactly and a uniformly random independent set can be sampled by brute-force in time $2^{n+o(n)} = (1/\epsilon)^{O(d/\log^2 d)}$.

Now suppose $\epsilon > 2\epsilon_0$. By Lemma 16, it is enough to compute an $(\epsilon/4)$ -relative approximation to both Ξ_G^X and Ξ_G^Y . By using the algorithm given by Lemma 15, this takes time $(n/\epsilon)^{O(d)}$.

For the approximate sampling algorithm, we note that by Lemma 16, it is enough to obtain an $\epsilon/2$ -approximate sample from $\hat{\mu}_G$. We do this as follows: we first compute $\epsilon/8$ -relative approximations to Ξ_G^X and Ξ_G^Y by computing $\Xi_G^X(\ell)$ and $\Xi_G^Y(\ell)$ respectively, with ℓ chosen as in Lemma 15. We then pick X or Y with respective probabilities $\Xi_G^X(\ell)/(\Xi_G^X(\ell) + \Xi_G^Y(\ell))$ and $\Xi_G^Y(\ell)/(\Xi_G^X(\ell) + \Xi_G^Y(\ell))$, and then use the polymer sampling algorithm of Lemma 15 to approximately sample a configuration of compatible polymers Λ from X (resp. Y), accurate to within total variation distance $\epsilon/8$. Given the polymer configuration Λ we then independently select each vertex of $Y \setminus N(\Lambda)$ (resp. $X \setminus N(\Lambda)$) with probability $1/2$ and add these to the independent set. The distribution of the output is then within total variation distance $\epsilon/2$ of $\hat{\mu}_G$. See the sampling algorithm of [31] for details on the calculation of this bound. □

6. AN ALGORITHM FOR GENERAL REGULAR BIPARTITE GRAPHS: PROOF OF THEOREM 3

In this section we prove Theorem 3, giving an algorithm that, for any constant $c > 0$, returns an n^{-c} -relative approximation for the number of independent sets in a general d -regular bipartite graph. As in the previous sections we let G denote a d -regular graph on $2n$ vertices with vertex classes X and Y . The algorithm proceeds by separating the contribution of expanding and non-expanding 2-linked sets of X to the independent set count. To estimate the contribution from non-expanding components, we use a simple argument inspired by the container method (see Lemma 17 below). To estimate the contribution from expanding components, we appeal to the algorithm of Lemma 15.

We begin with the the following lemma which will allow us to group non-expanding components according to their closure. We say that a set $A \subseteq X$ is *closed* if $A = [A]$.

Lemma 17. *Let $A \subseteq X$ be a 2-linked, closed, non-expanding set. Then there is a randomized $\text{poly}(n) \cdot \epsilon^{-2} \ln(1/\delta) \cdot 2^{O\left(\frac{|A| \log^2 d}{d}\right)}$ -time algorithm that outputs an ϵ -relative approximation to the number of 2-linked $B \subseteq A$ such that $N(B) = N(A)$ with probability at least $1 - \delta$.*

Proof. Let $W = N(A)$, $w = |W|$ and $a = |[A]|$. Let

$$\mathcal{D} := \{B \subseteq A \mid N(B) = W \text{ and } B \text{ is 2-linked}\}$$

be the set whose size we would like to estimate. By Corollary 10 there exists a set $A' \in \mathcal{D}$ of size at most

$$2 \frac{a}{d} \ln d + 2 \frac{w}{d} + 2(w - a) = O\left(\frac{a \log^2 d}{d}\right).$$

It follows that

$$|\mathcal{D}| \geq 2^{a - O\left(\frac{a \log^2 d}{d}\right)}.$$

Indeed every superset $B \supseteq A'$ satisfies $N(B) = W$ and B is 2-linked. The first property is clear, since $N(B) \supseteq N(A') = W$. The second property holds because every vertex in B is either in A' or at a distance 2 from some vertex in A' , which is itself 2-linked. It follows that $|\mathcal{D}|$ can be estimated to relative error ϵ by sampling

$$\frac{1}{\epsilon^2} \ln(1/\delta) \cdot 2^{O\left(\frac{a \log^2 d}{d}\right)}$$

subsets of A . □

The algorithm. For a closed, 2-linked subset $A \subseteq X$, let us denote

$$\mathcal{D}(A) := \#\{B \subseteq A \mid B \text{ is 2-linked, } N(B) = N(A)\}.$$

We now define an algorithm with inputs a graph G on n vertices and an accuracy parameter $\epsilon > 0$. Let $L := \lceil \frac{d}{2 \log^2 d} \log(2n/\epsilon) \rceil$. If $d \leq \sqrt{n}$, the algorithm is as follows:

- (1) List all vectors $(a_1, \dots, a_\ell)$ of positive integers such that $\ell \leq n/d$ and $\sum a_i \leq n$.
- (2) For each vector $(a_1, \dots, a_\ell)$ from Step 1, list all sets $\{A_1, \dots, A_\ell\}$ such that the $N(A_i)$'s are pairwise disjoint and $A_i \in \mathcal{G}'(v_i, a_i)$ for some $v_i \in X$ for each i .
- (3) For each set $\{A_1, \dots, A_\ell\}$ from Step 2, compute $\tilde{\mathcal{D}}(A_i)$, which is a $(\epsilon/2n)$ -relative approximation to $\mathcal{D}(A_i)$ using Lemma 17, setting $\delta = (1/3) \cdot 2^{-n^2}$ for all i .
- (4) For each set $A = \{A_1, \dots, A_\ell\}$ from Step 2, let $Y_A = Y \setminus N(\cup_{i=1}^\ell A_i)$, $X_A = X \setminus N^2(\cup_{i=1}^\ell A_i)$, $G_A = G[X_A \cup Y_A]$ and compute $\Xi_{G_A}^{X_A}(L)$.
- (5) Output

$$\sum_{\ell=0}^{n/d} \sum_{\substack{A_1, \dots, A_\ell \text{ compatible} \\ \forall i, A_i \text{ non-expanding} \\ \text{2-linked, closed}}} \left(\prod_{i=1}^{\ell} \tilde{\mathcal{D}}(A_i) \right) \cdot \left(2^{|Y| - \sum_{i=1}^{\ell} N(A_i)} \cdot \Xi_{G_A}^{X_A}(L) \right).$$

If $d > \sqrt{n}$, the algorithm is to run Steps 1-3 above and output

$$(27) \quad \sum_{\ell=0}^{n/d} \sum_{\substack{A_1, \dots, A_\ell \text{ compatible} \\ \forall i, A_i \text{ non-expanding} \\ \text{2-linked, closed}}} \left(\prod_{i=1}^{\ell} \tilde{\mathcal{D}}(A_i) \right) \cdot \left(2^{|Y| - \sum_{i=1}^{\ell} N(A_i)} \right).$$

Proof of Theorem 3. We first prove the correctness of the algorithm: for any $c > 0$ and $\epsilon = n^{-c}$, the output is an ϵ -relative approximation to $i(G)$. As before we let X, Y denote the vertex classes of G . Suppose first that $d \leq \sqrt{n}$. We then have

$$\begin{aligned}
i(G) &= 2^{|Y|} \cdot \sum_{t=0}^{n/d} \sum_{\substack{A_1, \dots, A_t \text{ compatible} \\ A_i \text{ 2-linked } \forall i}} \left(\prod_{i=1}^t 2^{-N(A_i)} \right) \\
&= 2^{|Y|} \cdot \sum_{t=0}^{n/d} \sum_{\ell=0}^t \sum_{\substack{A_1, \dots, A_\ell \text{ compatible} \\ \forall i, A_i \text{ non-expanding,} \\ \text{2-linked}}} \sum_{\substack{A_{\ell+1}, \dots, A_t \\ \subseteq X \setminus N^2(\cup_{j=1}^\ell A_j) \\ \text{compatible} \\ \forall i A_i \text{ expanding,} \\ \text{2-linked}}} \left(\prod_{i=1}^\ell 2^{-N(A_i)} \cdot \prod_{i=\ell+1}^t 2^{-N(A_i)} \right) \\
&= 2^{|Y|} \cdot \sum_{t=0}^{n/d} \sum_{\ell=0}^t \sum_{\substack{A_1, \dots, A_\ell \text{ compatible} \\ \forall i, A_i \text{ non-expanding} \\ \text{2-linked, closed}}} \sum_{\substack{B_1, \dots, B_\ell \\ \forall i, B_i \subseteq A_i, \\ N(B_i) = N(A_i), \\ B_i \text{ 2-linked}}} \left(\prod_{i=1}^\ell 2^{-N(A_i)} \cdot \sum_{\substack{A_{\ell+1}, \dots, A_t \\ \subseteq X \setminus N^2(\cup_{j=1}^\ell A_j) \\ \text{compatible} \\ \forall i A_i \text{ expanding,} \\ \text{2-linked}}} \prod_{i=\ell+1}^t 2^{-N(A_i)} \right) \\
&= \sum_{\ell=1}^{n/d} \sum_{\substack{A_1, \dots, A_\ell \text{ compatible} \\ \forall i, A_i \text{ non-expanding} \\ \text{2-linked, closed}}} \sum_{\substack{B_1, \dots, B_\ell \\ \forall i, B_i \subseteq A_i, \\ N(B_i) = N(A_i), \\ B_i \text{ 2-linked}}} \left(2^{|Y| - \sum_{i=1}^\ell N(A_i)} \cdot \sum_{t=1}^{n/d - \ell} \sum_{\substack{A'_1, \dots, A'_t \\ \subseteq X \setminus N^2(\cup_{j=1}^\ell A_j) \\ \text{compatible} \\ \forall i A'_i \text{ expanding,} \\ \text{2-linked}}} \prod_{i=1}^t 2^{-N(A'_i)} \right) \\
&= \sum_{\ell=1}^{n/d} \sum_{\substack{A_1, \dots, A_\ell \text{ compatible} \\ \forall i, A_i \text{ non-expanding} \\ \text{2-linked, closed}}} \left(\prod_{i=1}^\ell \mathcal{D}(A_i) \right) \cdot \left(2^{|Y| - \sum_{i=1}^\ell N(A_i)} \cdot \Xi_{G_A}^{X_A} \right) \\
&= (1 \pm \epsilon) \sum_{\ell=1}^{n/d} \sum_{\substack{A_1, \dots, A_\ell \text{ compatible} \\ \forall i, A_i \text{ non-expanding} \\ \text{2-linked, closed}}} \left(\prod_{i=1}^\ell \tilde{\mathcal{D}}(A_i) \right) \cdot \left(2^{|Y| - \sum_{i=1}^\ell N(A_i)} \cdot \Xi_{G_A}^{X_A}(L) \right).
\end{aligned}$$

For the last equality we used that by Remark 1, we may apply Lemma 13 to G_A , and so $\Xi_{G_A}^{X_A}(L)$ is an ϵ -relative approximation to $\Xi_{G_A}^{X_A}$. Observe that there are at most 2^{n^2} many choices of nonexpanding $A_1, \dots, A_\ell$. So by union bounding over all these choices, the last summation is exactly what the algorithm outputs with probability at least $1 - \delta \cdot 2^{n^2} = 2/3$, we have the required approximation guarantee. If $d \geq \sqrt{n}$, then we note that since the polymers

of G_A are a subset of the polymers of G , we have by (20) that $\ln \Xi_{G_A}^{X_A} \leq \ln \Xi_G^X \leq 2^{-\Omega(\log^2 n)}$. It follows that 1 is trivially an ϵ -relative approximation to $\Xi_{G_A}^{X_A}$ (recall that $\epsilon = n^{-c}$) and so (27) is an ϵ -relative approximation to $i(G)$.

We now show that if $\epsilon = n^{-c}$, with $c > 0$ fixed, the above algorithm runs in time $2^{O\left(\frac{\log^3 d}{d}n\right)}$. We consider the algorithm step by step.

Step 1. For $\ell \leq n/d$, and $k \leq n$ the number of vectors $(a_1, \dots, a_\ell)$ of positive integers such that $\sum a_i = k$ (i.e. the number of ordered partitions of k with ℓ parts) is

$$\binom{k-1}{\ell-1} \leq \binom{n}{n/d} = 2^{O\left(\frac{\log d}{d}n\right)}.$$

Moreover, it is clear that the set of all such partitions can be listed in time $2^{O\left(\frac{\log d}{d}n\right)}$ and so Step 1 takes time $2^{O\left(\frac{\log d}{d}n\right)}$.

Step 2. Let $(a_1, \dots, a_\ell)$ be a vector of positive integers such that $\ell \leq n/d$ and $\sum a_i \leq n$. We first list all tuples vertices $\{v_1, \dots, v_\ell\} \subseteq X$ which takes time $\binom{n}{\ell} = 2^{O\left(\frac{\log d}{d}n\right)}$. For each $\{v_1, \dots, v_\ell\}$, we then appeal to Lemma 5 to output the tuple $(\mathcal{G}'(v_1, a_1), \dots, \mathcal{G}'(v_\ell, a_\ell))$ in time $2^{O\left(\frac{\log^3 d}{d}n\right)}$. We note that by Lemma 5

$$|\mathcal{G}'(v_1, a_1) \times \dots \times \mathcal{G}'(v_\ell, a_\ell)| \leq \prod_{i=1}^{\ell} 2^{O\left(\frac{\log^3 d}{d}a_i\right)} = 2^{O\left(\frac{\log^3 d}{d}n\right)}.$$

We may therefore check each element of $\mathcal{G}'(v_1, a_1) \times \dots \times \mathcal{G}'(v_\ell, a_\ell)$ to see if it satisfies the required conditions and output the desired list in time $2^{O\left(\frac{\log^3 d}{d}n\right)}$.

Step 3. Given a set $\{A_1, \dots, A_\ell\}$ from Step 2, we use Lemma 17 to compute an ϵ/n -relative approximation $\tilde{\mathcal{D}}(A_i)$ to $\mathcal{D}(A_i)$ for all i . This takes time

$$(\epsilon/n)^{-2} \ln(1/\delta) 2^{O\left(\frac{\log^2 d}{d}n\right)} = 2^{O\left(\frac{\log^2 d}{d}n\right)}$$

where we recall that $\epsilon = n^{-C}$ and $\delta = (1/3) \cdot 2^{-n^2}$.

Step 4. If $d < \sqrt{n}$, then given any set $A = \{A_1, \dots, A_\ell\}$ from Step 2, we may compute $\Xi_{G_A}^{X_A}(L)$ in time $(n/\epsilon)^{O(d)} = 2^{O\left(\frac{\log^2 d}{d}n\right)}$ by the algorithm in Section 4 restricted to polymers of G_A . If $d \geq \sqrt{n}$, we skip Step 4.

We conclude that the algorithm takes time $2^{O\left(\frac{\log^3 d}{d}n\right)}$ in total.

7. WEIGHTED INDEPENDENT SETS: PROOF OF THEOREM 2

The proof of Theorem 2 will follow the same lines as that of Theorem 1, with the main difference being that polymer weights will now be $w_\gamma = \frac{\lambda^{|\gamma|}}{(1+\lambda)^{|\mathcal{N}(\gamma)|}}$, generalizing the $\lambda = 1$ case of Theorem 1.

We assume throughout this section that G is a d -regular, bipartite α -expander with bipartition X, Y of size n each.

Define a polymer model with polymers consisting of the small 2-linked, subsets of X (resp. Y) with two polymers compatible if their union is not 2-linked (recall that $\gamma \subseteq X$ is small if $|\gamma| \leq n/2$). The weight of a polymer γ is $w_\gamma = \frac{\lambda^{|\gamma|}}{(1+\lambda)^{|N(\gamma)|}}$. Let $\Xi_G^X(\lambda)$ be the polymer model partition function and $\nu_{G,\lambda}^X$ be the corresponding Gibbs measure on collections of compatible polymers.

Theorem 2 follows from the following two lemmas. Lemma 18 below is the analogue of Lemma 15 and Lemma 19 is the analogue of Lemma 16.

Lemma 18. *For every $\alpha > 0$, there exists $C_2 > 0$ so that if $\lambda \geq \frac{C_2 \log d}{d^{1/4}}$ then there is an FPTAS to compute $\Xi_G^X(\lambda)$ and $\Xi_G^Y(\lambda)$ and a polynomial-time sampling scheme for $\nu_{G,\lambda}^X$ and $\nu_{G,\lambda}^Y$.*

As in Section 5, we define a probability measure on independent sets as a mixture of measures derived from the two polymer models. Define the distribution $\hat{\nu}_{G,\lambda}^X$ on $\mathcal{I}(G)$ as follows.

- (1) Sample a collection of compatible polymers Λ from the measure $\nu_{G,\lambda}^X$.
- (2) Set $I = J \cup \bigcup_{\gamma \in \Lambda} \gamma$ where J is a random subset of $Y \setminus \bigcup_{\gamma \in \Lambda} N(\gamma)$ formed by including each vertex independently with probability $\lambda/(1+\lambda)$.

Define $\hat{\nu}_{G,\lambda}^Y$ analogously and define the mixture

$$\hat{\mu}_{G,\lambda} = \frac{\Xi_G^X(\lambda)}{\Xi_G^X(\lambda) + \Xi_G^Y(\lambda)} \hat{\nu}_{G,\lambda}^X + \frac{\Xi_G^Y(\lambda)}{\Xi_G^X(\lambda) + \Xi_G^Y(\lambda)} \hat{\nu}_{G,\lambda}^Y.$$

Lemma 19. *For every $\alpha > 0$, there exists $C_2 > 0$ so that if $\lambda \geq \frac{C_2 \log d}{d^{1/4}}$ then for n sufficiently large,*

$$(1+\lambda)^n (\Xi_G^X(\lambda) + \Xi_G^Y(\lambda))$$

is an ϵ -relative approximation to $Z_G(\lambda)$ and

$$\|\mu_{G,\lambda} - \hat{\mu}_{G,\lambda}\|_{TV} < \epsilon$$

where $\epsilon = \exp(-\Omega(n))$, with the implicit constant a function of d .

To prove Lemmas 18 and 19 we extend the estimates of Sections 3 and 4 to the more general case. The main graph container lemma comes from the paper of Galvin and Tetali [20]. Let us define

$$\mathcal{W}_\lambda(v, a, w) := \sum_{\substack{A \subseteq X \\ |A| \leq n/2, \\ A \text{ is 2-linked} \\ |N(A)| = w}} \lambda^{|A|} (1+\lambda)^{-w},$$

and

$$\beta(\lambda) := \frac{\log^2(1+\lambda)}{\log(1+\lambda) + \log(2d^5/\alpha)}.$$

We use the following lemma from [20].

Lemma 20. *There are constants c_4 and c_5 such that the following holds: Let G be a bipartite α -expander and suppose $\lambda > 0$. If $\beta(\lambda)$ satisfies*

$$\beta(\lambda) \geq c_4 \max \left\{ \frac{\log(d^5/\alpha)}{\sqrt{d}}, \frac{2 \log^2 d}{\alpha d} \right\}.$$

Then

$$\mathcal{W}_\lambda(v, a, w) \leq 2^{-c_5(w-a)\beta(\lambda)}.$$

The hypothesis of the above lemma says that $\alpha \cdot \beta(\lambda) \geq 2c_4 \cdot \frac{\log^2 d}{d}$. However, in our application, we will also assume that $\beta(\lambda)$ is also large enough to ensure

$$(28) \quad \alpha \cdot \beta(\lambda) \geq \frac{4000}{c_5} \cdot \frac{\log^2 d}{d}.$$

This holds whenever $\lambda \geq \frac{C_2 \log d}{d^{1/4}}$ for a large enough constant C_2 .

We now prove Lemma 18.

Proof of Lemma 18. As in the proof of Lemma 15, the FPTAS and polynomial-time sampling scheme will follow from verifying the Kotecký–Preiss condition for the polymer model.

We will first show that the polymer model satisfies (11) with $f(\gamma) = c_5 \alpha \ln 2 \frac{\beta(\lambda)|\gamma|}{8}$ and $g(\gamma) = c_5 \alpha \ln 2 \frac{\beta(\lambda)|N(\gamma)|}{8}$. We begin by bounding

$$(29) \quad \sum_{\gamma' \not\sim \gamma} w_{\gamma'} e^{f(\gamma') + g(\gamma')} \leq \sum_{v \in N^2(\gamma)} \sum_{\gamma' \ni v} w_{\gamma'} e^{f(\gamma') + g(\gamma')}.$$

Again, we use $N^2(\gamma)$ to denote $N(N(\gamma))$, the second neighborhood of γ , and note that $|N^2(\gamma)| \leq d^2 |\gamma|$. Since each polymer γ' in the above sum is expanding by definition, the right hand side of (29) is therefore at most

$$(30) \quad \sum_{w \geq d} \left(\sum_{t \geq (\alpha/2)w} \mathcal{W}_\lambda(v, w-t, w) \cdot e^{c_5 \alpha \ln 2 \frac{\beta(\lambda)w}{4}} \right).$$

By Lemma 20, it follows that (30) is bounded above by

$$(31) \quad \sum_{w \geq d} 2^{c_5 \alpha \frac{\beta(\lambda)w}{4}} \left(\sum_{t \geq (\alpha/2)w} 2^{-c_5 \beta(\lambda)t} \right) \leq 2d \sum_{w \geq d} 2^{-c_5 \alpha \frac{\beta(\lambda)w}{4}} \leq 4d^2 |\gamma| 2^{-c_5 \alpha \frac{\beta(\lambda)d}{4}}.$$

Combining (28)-(31) we conclude that

$$\begin{aligned} \sum_{\gamma' \not\sim \gamma} w_{\gamma'} e^{f(\gamma') + g(\gamma')} &\leq 4d^4 |\gamma| 2^{-c_5 \alpha \frac{\beta(\lambda)d}{4}} \\ &\leq 2d^4 |\gamma| 2^{-500 \log^2 d} \cdot 2^{-c_5 \alpha \frac{\beta(\lambda)d}{8}} \\ &\leq |\gamma| \cdot c_5 \alpha \frac{\beta(\lambda)}{16} \\ &\leq f(\gamma). \end{aligned}$$

The second last inequality follows because for $d \geq 3$, we have $2^{500 \log^2 d} > 2d^4$ and for any $x, y > 0$, we have $2^{-x} \leq \frac{x}{2y}$ if $x \geq 500 \log^2 y$.

Define $\Xi_G^X(\ell, \lambda)$ to be the exponential of the truncated cluster expansion as in (18). By the calculation similar to that in the proof of Lemma 13, we have for every $\ell \geq 1$,

$$(32) \quad |\ln \Xi_G^X(\lambda) - \ln \Xi_G^X(\ell, \lambda)| \leq n \cdot 2^{-500 \frac{\log^2 d \cdot \ell}{d}}.$$

In particular, if $\ell \geq \frac{d}{500 \log^2 d} \log(n/\epsilon)$, then

$$|\ln \Xi_G^X - \ln \Xi_G^X(\ell, \lambda)| \leq \epsilon. \quad \square$$

We now prove Lemma 19.

Proof of Lemma 19. Consider a modified polymer model with weights $\tilde{w}_\gamma(\lambda) = \lambda^{|\gamma|}(1 + \lambda)^{-|N(\gamma)|} 2^{c_5 \alpha \frac{\beta(\lambda)}{16}}$. The calculation in the proof of Lemma 18 shows that the modified polymer model with weights $\tilde{w}_\gamma(\lambda)$ satisfies (11) with $f(\gamma) = c_5 \alpha \ln 2 \frac{\beta(\lambda)|\gamma|}{16}$ and $g(\gamma) = c_5 \alpha \ln 2 \frac{\beta(\lambda)|N(\gamma)|}{8}$.

Let $\tilde{\Xi}_G^X(\lambda)$ denote the modified polymer model partition function as in (13). We then have

$$\ln \tilde{\Xi}_G^X(\lambda) - \ln \Xi_G^X(\lambda) = \ln \mathbf{E} e^{\zeta \cdot \|\mathbf{A}\|}$$

for $\zeta = c_5 \alpha \ln 2 \frac{\beta(\lambda)}{16}$. Summing (17) over all v , and using the fact that $|N(\gamma)| \geq d$ for every γ , we get (as in (20))

$$(33) \quad \ln \tilde{\Xi}_G^X \leq \sum_v \sum_{\substack{\Gamma \in \mathcal{C}(G) \\ \Gamma \ni v}} \left| \phi(\Gamma) \prod_{\gamma \in \Gamma} \tilde{w}_\gamma \right| \leq n \cdot 2^{-c_5 \alpha \beta(\lambda) \frac{d}{8}}.$$

Therefore, we have

$$\ln \mathbf{E} e^{\zeta \cdot \|\mathbf{A}\|} \leq \ln \tilde{\Xi}_G^X \leq n \cdot 2^{-c_5 \alpha \beta(\lambda) \frac{d}{8}}.$$

Fix any $\delta \geq \frac{10}{\sqrt{d}}$. By Markov's inequality, we have

$$(34) \quad \begin{aligned} \mathbb{P}(\|\mathbf{A}\| \geq \delta n) &\leq \exp\left(-\zeta \delta n + n \cdot 2^{-c_5 \alpha \beta(\lambda) \frac{d}{8}}\right) \\ &\leq 2^{-c_5 \alpha \frac{\beta(\lambda)}{32} \cdot \delta n} \\ &\leq 2^{-\frac{100 \log^2 d}{d} \cdot \delta n}, \end{aligned}$$

where the penultimate inequality follows because for any $x, y > 0$, we have $2^{-x \cdot y} \leq \frac{x}{2\sqrt{y}}$ if $x \geq 500 \log^2 y$, and therefore

$$2^{-c_5 \alpha \beta(\lambda) \frac{d}{8}} \leq \frac{c_5 \alpha \beta(\lambda)}{16\sqrt{d}} \leq \frac{\zeta \delta}{2}.$$

The final inequality (34) follows from (28).

As in the proof of Lemma 16, let $\mathcal{I}_X = \{I \in \mathcal{I} : \nu_{G,\lambda}^X(I) > 0\}$ i.e. the set of all I such that each 2-linked component of $I \cap X$ is small. For $\delta > 0$, let $\mathcal{I}_X^\delta = \{I \in \mathcal{I}_X : |I \cap X| \leq \delta n\}$ and define $\mathcal{I}_Y, \mathcal{I}_Y^\delta$ similarly. As in Lemma 16, $\mathcal{I} = \mathcal{I}_X \cup \mathcal{I}_Y$ and so

$$(35) \quad Z_G(\lambda) = \sum_{I \in \mathcal{I}_X} \lambda^{|I|} + \sum_{I \in \mathcal{I}_Y} \lambda^{|I|} - \sum_{I \in \mathcal{I}_X \cap \mathcal{I}_Y} \lambda^{|I|} = (1 + \lambda)^n (\Xi_G^X(\lambda) + \Xi_G^Y(\lambda)) - \sum_{I \in \mathcal{I}_X \cap \mathcal{I}_Y} \lambda^{|I|}.$$

Now, let I be a random independent set chosen from the distribution $\nu_{G,\lambda}^X$. It follows from (34) that for $\delta \geq \frac{10}{\sqrt{d}}$,

$$(36) \quad \mathbb{P}(|I \cap X| > \delta n) = \sum_{I \in \mathcal{I}_X \setminus \mathcal{I}_X^\delta} \frac{\lambda^{|I|}}{(1+\lambda)^n \Xi_G^X(\lambda)} \leq 2^{-\frac{100 \log^2 d}{d} \cdot \delta n}.$$

Furthermore, we have

$$\begin{aligned} \sum_{I \in \mathcal{I}_X^\delta \cap \mathcal{I}_Y^\delta} \frac{\lambda^{|I|}}{(1+\lambda)^n \Xi_G^X(\lambda)} &\leq \frac{\sum_{i,j \leq \delta n} \binom{n}{i} \binom{n}{j} \lambda^{i+j}}{(1+\lambda)^n} \\ &= \frac{\left(\sum_{i \leq \delta n} \binom{n}{i} \lambda^i \right)^2}{(1+\lambda)^n} \\ &= (1+\lambda)^n \mathbb{P} \left(\text{Bin} \left(n, \frac{\lambda}{1+\lambda} \right) \leq \delta n \right)^2 \end{aligned}$$

This quantity can be made to be at most $e^{-\frac{\delta n}{16}}$ for $\delta = \frac{\lambda}{100(1+\lambda)}$. We note that $\delta \geq \frac{C \log^2 d}{d^{1/4}} \geq \frac{10}{\sqrt{d}}$ for a large enough C , and so

$$\begin{aligned} \sum_{I \in \mathcal{I}_X \cap \mathcal{I}_Y} \lambda^{|I|} &\leq \sum_{I \in \mathcal{I}_X \setminus \mathcal{I}_X^\delta} \lambda^{|I|} + \sum_{I \in \mathcal{I}_Y \setminus \mathcal{I}_Y^\delta} \lambda^{|I|} + \sum_{I \in \mathcal{I}_X^\delta \cap \mathcal{I}_Y^\delta} \lambda^{|I|} \\ &\leq (1+\lambda)^n (\Xi_G^X(\lambda) + \Xi_G^Y(\lambda)) e^{-\Omega(n)}. \end{aligned}$$

By (35), we conclude that

$$\left| \frac{Z_G(\lambda)}{(1+\lambda)^n (\Xi_G^X(\lambda) + \Xi_G^Y(\lambda))} - 1 \right| = e^{-\Omega(n)}.$$

The bound on total variation distance follows in the same manner as in the proof of Lemma 16. $\square$

ACKNOWLEDGEMENTS

WP is supported in part by NSF grant DMS-1847451. AP is supported in part by NSF grant CCF-1934915.

REFERENCES

- [1] Sanjeev Arora, Boaz Barak, and David Steurer, *Subexponential algorithms for unique games and related problems*, Journal of the ACM (JACM) **62** (2015), 1–25.
- [2] Sanjeev Arora, Subhash A Khot, Alexandra Kolla, David Steurer, Madhur Tulsiani, and Nisheeth K Vishnoi, *Unique games on expanding constraint graphs are easy*, Proceedings of the fortieth annual ACM Symposium on Theory of Computing (STOC), 2008, pp. 21–28.
- [3] József Balogh, Ramon I García, and Lina Li, *Independent sets in the middle two layers of Boolean lattice*, Journal of Combinatorial Theory, Series A **178** (2021), 105341.
- [4] József Balogh, Robert Morris, and Wojciech Samotij, *Independent sets in hypergraphs*, Journal of the American Mathematical Society **28** (2015), 669–709.

- [5] Sarah Cannon and Will Perkins, *Counting independent sets in unbalanced bipartite graphs*, Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, 2020, pp. 1456–1466.
- [6] Charles Carlson, Ewan Davies, and Alexandra Kolla, *Efficient algorithms for the Potts model on small-set expanders*, arXiv preprint arXiv:2003.01154 (2020).
- [7] Zongchen Chen, Andreas Galanis, Leslie A Goldberg, Will Perkins, James Stewart, and Eric Vigoda, *Fast algorithms at low temperatures via Markov chains*, Random Structures & Algorithms **58** (2021), 294–321.
- [8] Zongchen Chen, Andreas Galanis, Daniel Štefankovič, and Eric Vigoda, *Sampling colorings and independent sets of random regular bipartite graphs in the non-uniqueness region*, arXiv preprint arXiv:2105.01784 (2021).
- [9] Ewan Davies, Matthew Jenssen, and Will Perkins, *A proof of the Upper Matching Conjecture for large graphs*, Journal of Combinatorial Theory, Series B **151** (2021), 393–416.
- [10] Philippe Delsarte, *An algebraic approach to the association schemes of coding theory*, Philips Res. Rep. Suppl. **10** (1973).
- [11] Martin Dyer, Leslie Ann Goldberg, Catherine Greenhill, and Mark Jerrum, *The relative complexity of approximate counting problems*, Algorithmica **38** (2004), 471–500.
- [12] Tobias Friedrich, Andreas Göbel, Martin S Krejca, and Marcus Pappik, *Polymer dynamics via cliques: New conditions for approximations*, arXiv preprint arXiv:2007.08293 (2020).
- [13] Andreas Galanis, Leslie Ann Goldberg, and James Stewart, *Fast algorithms for general spin systems on bipartite expanders*, 45th International Symposium on Mathematical Foundations of Computer Science (MFCS 2020), Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020.
- [14] Andreas Galanis, Leslie Ann Goldberg, and James Stewart, *Fast mixing via polymers for random graphs with unbounded degree*, arXiv preprint arXiv:2105.00524 (2021).
- [15] Andreas Galanis, Daniel Štefanković, Eric Vigoda, and Linji Yang, *Ferromagnetic Potts model: Refined #BIS-hardness and related results*, SIAM Journal on Computing **45** (2016), 2004–2065.
- [16] David Galvin, *Sampling 3-colourings of regular bipartite graphs*, Electronic Journal of Probability **12** (2007), 481–497.
- [17] David Galvin, *A threshold phenomenon for random independent sets in the discrete hypercube*, Combinatorics, Probability and Computing **20** (2011), 27–51.
- [18] David Galvin, *Independent sets in the discrete hypercube*, arXiv preprint arXiv:1901.01991 (2019).
- [19] David Galvin and Dana Randall, *Torpid mixing of local Markov chains on 3-colorings of the discrete torus*, Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms, 2007, pp. 376–384.
- [20] David J. Galvin and Prasad Tetali, *Slow mixing of Glauber dynamics for the hard-core model on regular bipartite graphs*, Random Struct. Algorithms **28** (2006), 427–443.
- [21] Serge Gaspers and Edward J Lee, *Faster graph coloring in polynomial space*, International Computing and Combinatorics Conference, Springer, 2017, pp. 371–383.
- [22] Chris D Godsil and Mike W Newman, *Eigenvalue bounds for independent sets*, Journal of Combinatorial Theory, Series B **98** (2008), 721–734.
- [23] Leslie Ann Goldberg and Mark Jerrum, *The complexity of ferromagnetic Ising with local fields*, Combinatorics, Probability and Computing **16** (2007), 43–61.
- [24] Leslie Ann Goldberg and Mark Jerrum, *Approximating the partition function of the ferromagnetic Potts model*, Journal of the ACM (JACM) **59** (2012), 1–31.
- [25] Leslie Ann Goldberg, John Lapinskas, and David Richerby, *Faster exponential-time algorithms for approximately counting independent sets*, arXiv preprint arXiv:2005.05070 (2020).
- [26] Christian Gruber and Hervé Kunz, *General properties of polymer systems*, Communications in Mathematical Physics **22** (1971), 133–161.
- [27] Willem H Haemers, *Hoffman’s ratio bound*, Linear Algebra and its Applications **617** (2021), 215–219.
- [28] Tyler Helmuth, Will Perkins, and Guus Regts, *Algorithmic Pirogov–Sinai theory*, Probability Theory and Related Fields **176** (2020), 851–895.
- [29] L. Ilinca and Jeff Kahn, *Counting maximal antichains and independent sets*, Order **30** (2013), 427–435.
- [30] Matthew Jenssen and Peter Keevash, *Homomorphisms from the torus*, arXiv preprint arXiv:2009.08315 (2020).
- [31] Matthew Jenssen, Peter Keevash, and Will Perkins, *Algorithms for #BIS-hard problems on expander graphs*, SIAM Journal on Computing **49** (2020), 681–710.
- [32] Matthew Jenssen and Will Perkins, *Independent sets in the hypercube revisited*, Journal of the London Mathematical Society **102** (2020), 645–669.

- [33] Matthew Jenssen, Will Perkins, and Aditya Potukuchi, *Independent sets of a given size and structure in the hypercube*, arXiv preprint arXiv:2106.09709 (2021).
- [34] Jeff Kahn and Jinyoung Park, *The number of maximal independent sets in the Hamming cube*, arXiv preprint arXiv:1909.04283 (2019).
- [35] Daniel J Kleitman and Kenneth J Winston, *On the number of graphs without 4-cycles*, Discrete Mathematics **41** (1982), 167–172.
- [36] DJ Kleitman and KJ Winston, *The asymptotic number of lattices*, Combinatorial Mathematics, Optimal Designs and their Applications (J. Srivastava, ed.), Ann. Discrete Math **6** (1980), 243–249.
- [37] Alexandra Kolla, *Spectral algorithms for unique games*, Computational Complexity **20** (2011), 177–206.
- [38] AD Korshunov and AA Sapozhenko, *The number of binary codes with distance 2*, Problemy Kibernet **40** (1983), 111–130.
- [39] Roman Kotecký and David Preiss, *Cluster expansion for abstract polymer models*, Communications in Mathematical Physics **103** (1986), 491–498.
- [40] Chao Liao, Jiabao Lin, Pinyan Lu, and Zhenyu Mao, *Counting independent sets and colorings on random regular bipartite graphs*, Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2019), Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2019.
- [41] Jingcheng Liu and Pinyan Lu, *FPTAS for #BIS with degree bounds on one side*, Proceedings of the forty-seventh annual ACM symposium on Theory of Computing, 2015, pp. 549–556.
- [42] L. Lovász, *On the ratio of optimal integral and fractional covers*, Discrete Mathematics **13** (1975), 383 – 390.
- [43] László Lovász, *On the Shannon capacity of a graph*, IEEE Transactions on Information theory **25** (1979), 1–7.
- [44] Konstantin Makarychev and Yury Makarychev, *How to play unique games on expanders*, International Workshop on Approximation and Online Algorithms, Springer, 2010, pp. 190–200.
- [45] Shayan Oveis Gharan and Luca Trevisan, *Partitioning into expanders*, Proceedings of the twenty-fifth annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, 2014, pp. 1256–1266.
- [46] Jinyoung Park, *Note on the number of balanced independent sets in the Hamming cube*, arXiv preprint arXiv:2103.11198 (2021).
- [47] J Scott Provan and Michael O Ball, *The complexity of counting cuts and of computing the probability that a graph is connected*, SIAM Journal on Computing **12** (1983), 777–788.
- [48] Wojciech Samotij, *Counting independent sets in graphs*, European Journal of Combinatorics **48** (2015), 5–18.
- [49] AA Sapozhenko, *On the number of connected subsets with given cardinality of the boundary in bipartite graphs*, Metody Diskret Analiz **45** (1987), 42–70.
- [50] Aleksandr Antonovich Sapozhenko, *On the number of independent sets in extenders*, Diskretnaya Matematika **13** (2001), 56–62.
- [51] Aleksandr Antonovich Sapozhenko, *The number of independent sets in graphs*, Moscow University Mathematics Bulletin **62** (2007), 116–118.
- [52] David Saxton and Andrew Thomason, *Hypergraph containers*, Inventiones Mathematicae **201** (2015), 925–992.
- [53] Allan Sly, *Computational transition at the uniqueness threshold*, 2010 IEEE 51st Annual Symposium on Foundations of Computer Science, IEEE, 2010, pp. 287–296.
- [54] S.K Stein, *Two combinatorial covering theorems*, Journal of Combinatorial Theory, Series A **16** (1974), 391 – 397.
- [55] Dror Weitz, *Counting independent sets up to the tree threshold*, Proceedings of the thirty-eighth annual ACM Symposium on Theory of Computing, 2006, pp. 140–149.

SCHOOL OF MATHEMATICS, UNIVERSITY OF BIRMINGHAM

DEPARTMENT OF MATHEMATICS, STATISTICS, AND COMPUTER SCIENCE, UNIVERSITY OF ILLINOIS AT CHICAGO

Email address: m.jenssen@bham.ac.uk, math@willperkins.org, adityap@uic.edu