

APPROXIMATE COUNTING AND SAMPLING VIA LOCAL CENTRAL LIMIT THEOREMS

VISHESH JAIN, WILL PERKINS, ASHWIN SAH, AND MEHTAAB SAWHNEY

ABSTRACT. We give an FPTAS for computing the number of matchings of size k in a graph G of maximum degree Δ on n vertices, for all $k \leq (1 - \delta)m^*(G)$, where $\delta > 0$ is fixed and $m^*(G)$ is the matching number of G , and an FPTAS for the number of independent sets of size $k \leq (1 - \delta)\alpha_c(\Delta)n$, where $\alpha_c(\Delta)$ is the NP-hardness threshold for this problem. We also provide quasi-linear time randomized algorithms to approximately sample from the uniform distribution on matchings of size $k \leq (1 - \delta)m^*(G)$ and independent sets of size $k \leq (1 - \delta)\alpha_c(\Delta)n$.

Our results are based on a new framework for exploiting local central limit theorems as an algorithmic tool. We use a combination of Fourier inversion, probabilistic estimates, and the deterministic approximation of partition functions at complex activities to extract approximations of the coefficients of the partition function. For our results for independent sets, we prove a new local central limit theorem for the hard-core model that applies to all fugacities below $\lambda_c(\Delta)$, the uniqueness threshold on the infinite Δ -regular tree.

1. INTRODUCTION

Counting matchings and independent sets in graphs are central problems in the study of exact and approximate counting algorithms. Exact counting of the total number of matchings of a graph, the number of perfect matchings, the number of matchings of a given size, the number of independent sets, and the number of independent sets of a given size are all #P-hard problems [43], even for many restricted classes of input graphs (bipartite graphs, graphs of bounded degree). A singular exception is the classical algorithm of Kasteleyn for counting the number of perfect matchings of a planar graph [28].

Turning to approximate counting, a stark difference emerges between matchings and independent sets. The landmark work of Jerrum and Sinclair gave an FPRAS (fully polynomial-time randomized approximation scheme) for counting (weighted) matchings in general graphs as well as counting matchings of any given size bounded away from the maximum matching [26]. For the special case of bipartite graphs, Jerrum, Sinclair, and Vigoda [27] gave an FPRAS for the number of matchings of any given size, including perfect matchings. Recent work of Alimohammadi, Anari, Shiragur, and Vuong [2] provides an FPRAS for the number of matchings of any given size in planar graphs. On the other hand, for counting (weighted) independent sets and independent sets of a given size, there is a threshold (in terms of degree, weighing factor, or density) above which the approximation problems are NP-hard and below which efficient approximation algorithms exist [3, 14, 44]. We mention that approximating the number of perfect matchings in general graphs is an outstanding open problem (see [42]).

Randomization has played a crucial role in the aforementioned algorithmic results, and especially for matchings, there is a wide gap between what is known to be achievable deterministically and with randomness. For graphs of maximum degree Δ , Bayati, Gamarnik, Katz, Nair, and Tetali [5] gave an FPTAS (fully polynomial-time approximation scheme) for the number of matchings and for weighted matchings with a bounded weighing factor; the running time of their algorithm is

Perkins was supported in part by NSF grants DMS-1847451 and CCF-1934915. Sah and Sawhney were supported by NSF Graduate Research Fellowship Program DGE-1745302.

polynomial in n and $1/\epsilon$, where ϵ is the desired accuracy and where the exponent of the polynomial depends on Δ and the weighing factor. However, it is not known, for instance, how to deterministically approximate the number of near-maximum matchings in bounded degree graphs. Our first main result addresses this by achieving, for bounded degree graphs, what was previously only known to be possible with randomness: we provide an FPTAS for the number of matchings of any given size bounded away from the maximum matching. Let $m_k(G)$ be the number of matchings of G of size k and let $m^*(G)$ be the size of the maximum matching of G .

Theorem 1.1. *Let $\Delta \geq 3$ and $\delta \in (0, 1)$. There exists a deterministic algorithm which, on input a graph $G = (V, E)$ on n vertices of maximum degree at most Δ , an integer $1 \leq k \leq (1 - \delta)m^*(G)$, and an error parameter $\epsilon \in (0, 1)$, outputs an ϵ -relative approximation to $m_k(G)$ in time $(n/\epsilon)^{O_{\delta, \Delta}(1)}$.*

Here, by an ϵ -relative approximation to $m_k(G)$, we mean that the output A satisfies $e^{-\epsilon}m_k(G) \leq A \leq e^{\epsilon}m_k(G)$.

We next prove the corresponding result for independent sets. Recall that the hard-core model on a graph $G = (V, E)$ at fugacity $\lambda \in \mathbb{R}_{\geq 0}$ is the probability distribution on $\mathcal{I}(G)$, the independent sets of G , defined by

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

where $Z_G(\lambda) = \sum_{I \in \mathcal{I}(G)} \lambda^{|I|}$ is the independence polynomial of G . For $\Delta \geq 3$, let

$$\lambda_c(\Delta) = \frac{(\Delta - 1)^{\Delta - 1}}{(\Delta - 2)^{\Delta}};$$

this is the uniqueness threshold for the hard-core model on the infinite Δ -regular tree. For $0 \leq \lambda < \lambda_c(\Delta)$, Weitz gave an FTPAS for $Z_G(\lambda)$ on the class of graphs of maximum degree Δ [44]. Sly [39], Sly and Sun [40], and Galanis, Štefankovič, and Vigoda [19] complemented this by showing that for $\lambda > \lambda_c(\Delta)$, no FPRAS for $Z_G(\lambda)$ exists unless $\text{NP} = \text{RP}$.

Recently, Davies and Perkins [14] showed an analogous threshold for counting independent sets of a given size in bounded degree graphs. Let

$$\alpha_c(\Delta) = \frac{\lambda_c(\Delta)}{1 + (\Delta + 1)\lambda_c(\Delta)} = \frac{(\Delta - 1)^{\Delta - 1}}{(\Delta - 2)^{\Delta} + (\Delta + 1)(\Delta - 1)^{\Delta - 1}};$$

this is the occupancy fraction (i.e. the expected density of an independent set) for the hard-core model on the clique on $\Delta + 1$ vertices at the critical fugacity $\lambda_c(\Delta)$. They showed that for $\alpha < \alpha_c(\Delta)$, there is an FPRAS for $i_k(G)$ (the number of independent sets in G of size k) for any G of maximum degree Δ on n vertices and any $k \leq \alpha n$; conversely, no FPRAS exists for $k \geq \alpha n$ for $\alpha > \alpha_c(\Delta)$ unless $\text{NP} = \text{RP}$. The algorithm of [14] uses randomness in an essential way and the authors conjectured the existence of an FPTAS for $i_k(G)$ for $k \leq \alpha n$ where $\alpha < \alpha_c(\Delta)$ [14, Conjecture 1]. We prove this conjecture.

Theorem 1.2. *Let $\Delta \geq 3$ and $\delta \in (0, 1)$. There exists a deterministic algorithm which, on input a graph $G = (V, E)$ on n vertices of maximum degree at most Δ , an integer $1 \leq k \leq (1 - \delta)n\alpha_c(\Delta)$, and an error parameter $\epsilon \in (0, 1)$, outputs an ϵ -relative approximation to $i_k(G)$ in time $(n/\epsilon)^{O_{\delta, \Delta}(1)}$.*

We remark that a deterministic algorithm for approximating m_k and i_k via the cluster expansion is implicit in [13], but only for k much smaller than the bounds above.

Next, we turn to the problem of uniformly sampling matchings and independent sets of a given size. For fixed $\Delta \geq 3$, $\delta \in (0, 1)$, Davies and Perkins [14] gave an algorithm for ϵ -approximately sampling (i.e. within ϵ in total variation distance) from the uniform distribution on $\mathcal{I}_k(G)$ (the independent sets of size k in G) for a graph $G = (V, E)$ on n vertices with maximum degree at most Δ , for all $1 \leq k \leq (1 - \delta)n\alpha_c(\Delta)$. The running time of their algorithm is $\tilde{O}_{\delta, \Delta}(n^3)$, where

the \tilde{O} conceals polylogarithmic factors in n and $1/\epsilon$. For the more restricted range $1 \leq k \leq (1 - \delta)n/(2(\Delta + 1))$, it was already shown by Buble and Dyer [8] that the down-up walk on $\mathcal{I}_k(G)$ has ϵ -mixing time $O_\delta(n \log(n/\epsilon))$, which is optimal up to constants (see also [1] which gave fast mixing for a larger range of k in graphs satisfying a spectral condition). The down-up walk is the following Markov chain on $\mathcal{I}_k(G)$: at each step, given the current independent set $I \in \mathcal{I}_k(G)$, choose a uniformly random vertex $v \in I$ and a uniformly random vertex $w \in V$. Let $I' = (I \setminus v) \cup w$. If $I' \in \mathcal{I}_k(G)$, then move to I' ; else, stay at I .

For matchings, a polynomial time algorithm for approximately sampling from the uniform distribution on matchings of size k , for all $1 \leq k \leq (1 - \delta)m^*(G)$, is present in the work of Jerrum and Sinclair [26]. The running time of their algorithm scales at least as $n^{7/2}$. In the case of graphs of maximum degree at most Δ , a recent result of Chen, Liu, and Vigoda [11] combined with a rejection sampling procedure (and Lemma 5.1 below) provides an algorithm for this task running in time $\tilde{O}_{\delta,\Delta}(n^{3/2})$.

It was conjectured in [14, Conjecture 2] that the down-up walk on $\mathcal{I}_k(G)$ mixes rapidly for all $1 \leq k \leq (1 - \delta)n\alpha_c(\Delta)$ on all graphs on n vertices of maximum degree at most Δ . A stronger conjecture is that the ϵ -mixing time of the chain for this range of k is $O_{\delta,\Delta}(n \log(n/\epsilon))$. While not resolving this specific conjecture, our next main result provides an approximate sampling algorithm for $1 \leq k \leq (1 - \delta)n\alpha_c(\Delta)$ (and $1 \leq k \leq (1 - \delta)m^*(G)$ in the case of matchings) running in quasi-linear time, which matches (up to a small polylogarithmic factor) the conjectured mixing time of the down-up walk.

Theorem 1.3. *Let $\Delta \geq 3$ and $\delta \in (0, 1)$. There is a randomized algorithm which, on input a graph $G = (V, E)$ on n vertices of maximum degree at most Δ , an integer $1 \leq k \leq (1 - \delta)n\alpha_c(\Delta)$, and an error parameter $\epsilon \in (0, 1)$ outputs a random independent set $I \in \mathcal{I}_k(G)$ such that the total variation distance between the distribution of I and the uniform distribution on $\mathcal{I}_k(G)$ is at most ϵ . The running time of the algorithm is $O_{\delta,\Delta}(n \log(n/\epsilon)(\log n)^3 + n \log(n/\epsilon) \log n \log(1/\epsilon)^{3/2})$.*

There is also a randomized algorithm with the same guarantee and running time for matchings of size k for all $1 \leq k \leq (1 - \delta)m^(G)$.*

A natural extension of the unresolved conjecture of [14] is that the down-up walk for matchings of size k is rapidly mixing in the setting of Theorem 1.3.

Conjecture 1.4. *The down-up walk for matchings of size k mixes in time $O_{\Delta,\delta}(n \log(n/\epsilon))$ for graphs G of maximum degree Δ and $1 \leq k \leq (1 - \delta)m^*(G)$.*

1.1. Local central limit theorems. In previous works on approximately counting matchings [26] and independent sets [14] of a given size, a common approach is followed: sample a matching or independent set from the monomer-dimer model (i.e. hard-core model on the line graph) or hard-core model where the fugacity λ is chosen so that the average size is close to the desired size k . If the fugacity λ is such that sampling from the corresponding monomer-dimer or hard-core model can be done efficiently, and if one can further show that the probability of obtaining a matching or independent set of size exactly k is only polynomially small, then naïve rejection sampling gives an efficient sampling algorithm (which can be converted into an FPRAS for $m_k(G)$ or $i_k(G)$ by standard self-reducibility techniques). In the proofs of Theorems 1.1 and 1.2, we show how to implement a version of this idea deterministically.

Recall that a sequence of random variables X_n with mean μ_n and variance σ_n^2 is said to satisfy a central limit theorem (CLT) if for all $a \leq b \in \mathbb{R}$,

$$\mathbb{P} \left[a \leq \frac{X_n - \mu_n}{\sigma_n} \leq b \right] = \frac{1}{\sqrt{2\pi}} \int_a^b e^{-x^2/2} dx + o_n(1).$$

In particular, central limit theorems provide control on the probability that X_n lies in an interval of length $\Theta(\sigma_n)$. A much more precise notion is that of a local central limit theorem (LCLT). We

say that a sequence of integer-valued random variables X_n with mean μ_n and variance σ_n^2 satisfies an LCLT if for all integers k ,

$$\mathbb{P}[X_n = k] = \frac{1}{\sqrt{2\pi\sigma_n}} e^{-(k-\mu_n)^2/(2\sigma_n^2)} + o_n(\sigma_n^{-1}).$$

Returning to the discussion in the previous paragraph, suppose we could deterministically find a fugacity λ such that:

- (a) The expected size (to the nearest integer) of a matching or independent set drawn from the monomer-dimer or hard-core model is k .
- (b) There is an FPTAS for the the partition function, expectation, and variance of the monomer-dimer or hard-core model at λ .
- (c) The size of a matching or independent set drawn from the monomer-dimer or hard-core model at λ satisfies an LCLT.

Then, from (a) and (c), we have that (as long as $\sigma_n = \omega_n(1)$)

$$\frac{m_k(G_n)\lambda^k}{\sum_{j=0}^n m_j(G_n)\lambda^j} = \mathbb{P}[Y_n = k] = (1 + o_n(1)) \frac{1}{\sqrt{2\pi\sigma_n}} e^{-(k-\mu_n)^2/(2\sigma_n^2)},$$

and similarly for independent sets. Together with (b), this immediately gives a deterministic algorithm for approximating $m_k(G_n)$ or $i_k(G_n)$ to within a factor of $(1 + \epsilon)(1 + o_n(1))$ in time which is polynomial in n and $1/\epsilon$.

For the range of k covered by [Theorems 1.1](#) and [1.2](#), a fugacity λ satisfying (a) and (b) does indeed exist and can be found deterministically using (b) along with a binary search procedure. In particular, for the hard-core model, such a λ satisfies $\lambda < \lambda_c(\Delta)$. Moreover, by the Heilmann–Lieb theorem [\[22\]](#), the roots of the partition function of the monomer-dimer model are restricted to the negative real line, from which it follows that the size of a random matching is distributed as the sum of independent Bernoulli random variables and thus, for λ not too small, satisfies an LCLT [\[20\]](#). However, for the hard-core model, the corresponding LCLT for all $\lambda < \lambda_c(\Delta)$ was not known. In fact, even the much weaker statement that $\mathbb{P}_\lambda[|I| = \lfloor \mu_\lambda \rfloor] = \Omega(\sigma_\lambda^{-1})$ (here, the probability \mathbb{P}_λ , the mean μ_λ , and the variance σ_λ^2 are with respect to the hard-core model with fugacity λ) was unavailable for all $\lambda < \lambda_c(\Delta)$ and is precisely the content of [\[14, Conjecture 3\]](#). A key step in our proof is the resolution of this conjecture in the much stronger form of an LCLT.

Theorem 1.5. *Fix $\Delta \geq 3$ and $\delta \in (0, 1)$. Then for any sequence of graphs G_n on n vertices of maximum degree Δ , and any sequence $\lambda_n \in \mathbb{R}^+$ so that $n\lambda_n \rightarrow \infty$ and $\lambda_n \leq (1 - \delta)\lambda_c(\Delta)$, the size of a random independent set drawn from the hard-core model on G_n at fugacity λ_n satisfies a local central limit theorem.*

In [Section 3](#), we state and prove a quantitative version of this result ([Theorem 3.1](#)). A critical part in the proof of the LCLT is the existence of a suitable zero-free region (in the complex plane) for the independence polynomial, which has previously been exploited using Barvinok’s method (cf. [\[4\]](#)) to devise an FPTAS for the independence polynomial evaluated at λ in a certain complex region containing the interval $[0, \lambda_c(\Delta))$ on graphs of maximum degree Δ [\[34, 35\]](#).

Given the LCLT for λ satisfying (a) and (b), the above discussion immediately leads to a $(1 \pm \epsilon)(1 \pm o_n(1))$ -factor approximation of $m_k(G)$ or $i_k(G)$ in time which is polynomial in n and $1/\epsilon$, where the degree of the polynomial is allowed to depend on Δ and δ . As such, this is only an EPTAS (efficient polynomial-time approximation scheme) since the finest approximation one can achieve with this method is $(1 \pm o_n(1))$ (see [Section 4.1](#) for a more detailed discussion). Nevertheless, we show that the *proof* of the LCLT using Fourier inversion can be converted into an FPTAS, thereby providing a (perhaps surprising) connection between the deterministic approximation of the matching polynomial or independence polynomial at *complex* fugacities and the deterministic

approximation of suitable coefficients of the matching polynomial or independence polynomial. We note that the computational complexity of evaluating partition functions at complex parameters has received much recent attention [7, 9, 21, 31, 35].

In a nutshell, the idea is the following: given a graph G with maximum degree Δ , an error parameter $\epsilon \in (0, 1)$, and k as in [Theorem 1.2](#), consider $\lambda < \lambda_c(\Delta)$ satisfying (a) and (b). Let $Y = |I|$ denote the size of an independent set drawn from the hard-core model on G at fugacity λ , let μ and σ^2 denote the mean and variance of Y , and let $X = (Y - \mu)/\sigma$. By the Fourier inversion formula for lattices, for all $x \in \sigma^{-1} \cdot \mathbb{Z} - \sigma^{-1}\mu$,

$$\mathbb{P}[X = x] = \frac{1}{2\pi\sigma} \int_{-\pi\sigma}^{\pi\sigma} \mathbb{E}[e^{itX}] e^{-itx} dt,$$

where the probability and expectations are with respect to the hard-core model with fugacity λ .

In order to approximate $\mathbb{P}[X = x]$ we consider $\mathbb{E}[e^{itX}]$. First, observe that

$$\mathbb{E}[e^{itX}] = e^{-it\mu/\sigma} \cdot \frac{Z_G(\lambda e^{it/\sigma})}{Z_G(\lambda)}.$$

If t/σ is sufficiently small so that $\lambda e^{-it/\sigma}$ is in the zero-free region of $Z_G(\lambda)$ (viewed as a univariate polynomial of a single complex variable), then there is an FPTAS for $Z_G(\lambda e^{it/\sigma})$ by Barvinok's method [4, 34, 35]. However this method does not handle all t . But over the range of t sufficiently large, our proof of the LCLT shows that the integral is bounded by $(\epsilon/2)\mathbb{P}[X = x]$. It turns out that these two regimes overlap. Therefore using a Riemann sum approximation for the small regime gives an FPTAS for $i_k(G)$.

Towards the proof of [Theorem 1.3](#), given $k \leq (1 - \delta)\alpha_c(\Delta)n$, we provide in [Lemma 5.1](#) a $\tilde{O}_{\delta,\Delta}(n)$ time randomized algorithm for finding λ satisfying $|\mathbb{E}_\lambda Y - k| \leq \sqrt{\text{Var}_\lambda Y}$. As mentioned earlier, this can be combined with the $\tilde{O}_{\Delta,\delta}(n)$ mixing of the Glauber dynamics for the hard-core model at fugacity λ [11] and rejection sampling to approximately sample from the uniform distribution on $\mathcal{I}_k(G)$ in time $\tilde{O}_{\Delta,\delta}(n^{3/2})$, since the acceptance probability is $\tilde{\Omega}_{\Delta,\delta}(1/\sqrt{\text{Var}_\lambda |I|})$. The main idea underlying our algorithm is that the one may instead perform rejection sampling with the base distribution (effectively) being the hard-core distribution conditioned on $Y \equiv k \pmod p$, with $p = \tilde{\Theta}(\sqrt{\text{Var}_\lambda Y})$, while still ensuring that samples from the base distribution can be obtained in time $\tilde{O}(n)$. Given this, the assertion of [Theorem 1.3](#) follows since by the LCLT, the acceptance probability is now $\mathbb{P}_\lambda[Y = k]/\mathbb{P}_\lambda[Y \equiv k \pmod p] = \tilde{\Omega}_{\Delta,\delta}(1)$. Thus, the key step is to show that samples from the hard-core distribution conditioned on $Y \equiv k \pmod p$ may still be obtained in time $\tilde{O}(n)$. For this, we use a multi-stage view of sampling from the hard-core model, motivated directly by the proof of the LCLT.

1.2. Additional results. In [Theorems 1.1](#) and [1.2](#), we gave an FPTAS for $i_k(G)$ and $m_k(G)$ with running times of the form $(n/\epsilon)^{O_{\Delta,\delta}(1)}$. Now, we provide substantially faster randomized algorithms for the same problems.

In [14], an FPRAS (fully polynomial-time randomized approximation scheme) for $i_k(G)$ was given, for all $1 \leq k \leq (1 - \delta)n\alpha_c(\Delta)$, with running time $\tilde{O}_{\Delta,\delta}(n^6\epsilon^{-2})$. Combining this algorithm with our near-optimal sampling algorithm improves the running time to $\tilde{O}_{\Delta,\delta}(n^4\epsilon^{-2})$. In contrast, it is known that there is an FPRAS for the independence polynomial $Z_G(\lambda)$, for all $0 \leq \lambda \leq (1 - \delta)\lambda_c(\Delta)$, with running time $\tilde{O}(n^2\epsilon^{-2})$ (cf. [41]). Our next result provides an FPRAS for $i_k(G)$ and $m_k(G)$ whose running time exceeds the best-known running time for approximating the independence polynomial or matching polynomial only by a lower order term.

Theorem 1.6. *Let $\Delta \geq 3$ and $\delta \in (0, 1)$. There is a randomized algorithm which, on input a graph $G = (V, E)$ on n vertices of maximum degree at most Δ , an integer $1 \leq k \leq (1 - \delta)n\alpha_c(\Delta)$ and an*

error parameter $\epsilon \in (0, 1)$ outputs an ϵ -relative approximation to $i_k(G)$ with probability $3/4$ in time

$$T + O_{\Delta, \delta}(n^{3/2} \log n \log(n/\epsilon)\epsilon^{-2}),$$

where T is the time to find an $\epsilon/2$ -relative approximation to $Z_G(\lambda)$ for a given $\lambda \in [0, (1-\delta)\lambda_c(\Delta)]$.

Moreover, there exists a constant $C_{\Delta, \delta} > 0$ such that the same conclusion holds for $m_k(G)$ for all $1 \leq k \leq (1-\delta)m^*(G)$, where T is the time to find an $\epsilon/2$ -relative approximation to the matching polynomial $Z_G(\lambda)$ for a given $\lambda \in [0, C_{\Delta, \delta}]$.

Remark. For $1 \leq k \leq c_{\Delta}\sqrt{n}$ (where c_{Δ} is a sufficiently small constant), the running time may be improved to

$$O(k \log n \log \log n) + O_{\Delta}(k\epsilon^{-2} \log n).$$

Moreover, the term $n^{3/2}$ in [Theorem 1.6](#) may be replaced by $\tilde{O}(n)$ by using similar ideas as in the proof of [Theorem 1.3](#). However, since the current bounds on T are $\Omega(n^2)$, we have not pursued this improvement.

Finally, as an intermediate step in obtaining deterministic approximate counting algorithms, we need to deterministically approximate the mean and variance of the size of an independent set or matching drawn from the hard-core or monomer-dimer model respectively. While such algorithms are obtainable by applying algorithms based on the correlation decay method of Weitz [[5](#), [44](#)] to approximate marginals and joint marginals, we instead provide faster deterministic algorithms by adapting the method of Barvinok [[4](#)] and Patel–Regts [[34](#)] to approximate the k th cumulant of the size of the random independent set or matching.

The k th cumulant of a random variable Y is defined in terms of the coefficients of the cumulant generating function $K_Y(t) = \log \mathbb{E}e^{tY}$ (when this expectation exists in a neighborhood of 0). In particular, the k th cumulant is

$$\kappa_k(Y) = K_Y^{(k)}(0).$$

The first and second cumulants are the mean and variance respectively.

Theorem 1.7. *Fix $\Delta \geq 3$, $k \geq 1$, and $\delta \in (0, 1)$. There is an algorithm which, on input a graph $G = (V, E)$ on n vertices of maximum degree at most Δ , $0 < \lambda \leq (1-\delta)\lambda_c(\Delta)$, and an error parameter $\epsilon \in (0, 1)$ outputs an $\epsilon\lambda n$ additive approximation to $\kappa_k(Y)$, where Y is the size of an independent set drawn from the hard-core model on G at fugacity λ . The algorithm runs in time $O_{\Delta, \delta, k}(n(1/\epsilon)^{O_{\Delta, \delta}(1)})$. In particular, this provides an FPTAS for $\mathbb{E}_{\lambda}Y$ and $\text{Var}_{\lambda}Y$ running in time linear in n .*

For claw-free graphs (hence for the size of a random matching Y drawn from the monomer-dimer model), the same holds for all bounded $\lambda > 0$.

1.3. Outline. We prove [Theorems 1.1](#) and [1.2](#) in essentially the same way, given as an input the existence of a zero-free region for the corresponding partition function in the complex plane. For matchings, this is provided by the classical Heilmann–Lieb theorem [[22](#)], while for independent sets this is provided by the theorems of Shearer [[38](#)] and Peters and Regts [[35](#)]. To avoid excessive repetition (and to gain a slight amount of generality) we work with the larger class of independent sets in claw-free graphs (instead of matchings, which are independent sets in line graphs). The generalization of the Heilmann–Lieb theorem to claw-free graphs is due to Chudnovsky and Seymour [[12](#)].

We record these zero-freeness results in [Section 2](#), along with some results from the geometry of polynomials on the consequences of zero-freeness, namely a central limit theorem of Michelen and Sahasrabudhe [[33](#)] and a deterministic approximation algorithm for $Z_G(\lambda)$ (with λ possibly complex) due to Barvinok [[4](#)] and Patel and Regts [[34](#)]. We also record a recent result of Chen, Liu, and Vigoda [[11](#)] on the optimal mixing of Glauber dynamics for bounded marginal spin-systems on bounded degree graphs which is used in our randomized algorithms.

In [Section 3](#), we prove the local central limit theorem for the hard-core model ([Theorem 1.5](#)). Our proof uses both the aforementioned central limit theorem and a variant of the technique of Dobrushin and Tirozzi [\[16\]](#) who proved local central limit theorems for spin models on the integer lattice \mathbb{Z}^d .

In [Section 4](#), we prove our deterministic algorithmic results [Theorems 1.1](#) and [1.2](#).

In [Section 5](#), we prove our randomized algorithmic results [Theorems 1.3](#) and [1.6](#).

In [Section 6](#), we provide a proof of [Theorem 1.5](#) when λ is sufficiently small as a function of Δ using the cluster expansion, a tool from classical statistical physics. While not necessary for the main results, we include this since the proof is simpler and perhaps more intuitive.

Finally, in [Section 7](#), we prove [Theorem 1.7](#).

1.4. Notation. Throughout, we reserve the random variable Y for the size of a random independent set drawn from the hard-core model on a graph G . We use subscripts to indicate the fugacity, e.g. $\mathbb{E}_\lambda Y$ and $\text{Var}_\lambda Y$. We let $\alpha_G(\lambda)$ denote the occupancy fraction of the hard-core model on G at fugacity λ ; that is, $\alpha_G(\lambda) = \frac{\mathbb{E}_\lambda Y}{|V(G)|}$. We let \mathcal{Z} be a standard normal random variable and $\mathcal{N}(x) = e^{-x^2/2}/\sqrt{2\pi}$ denote its density.

Dependence of various constants on input parameters will often be important; we will write $f(n) = O_\Delta(1)$, for instance, to mean that f is bounded by a constant that depends only on Δ .

2. PRELIMINARIES

Definition 2.1. Let $z_1, z_2 \in \mathbb{C}$. We say that z_1 is a δ -additive, ϵ -relative approximation of z_2 if $z_1 = re^{i\theta}z_2 + z_3$ for some $e^{-\epsilon} \leq r \leq e^\epsilon$, $\theta \in \mathbb{R}$, $|\theta| \leq \epsilon$ and $z_3 \in \mathbb{C}$, $|z_3| \leq \delta$. When $\delta = 0$, we simply say that z_1 is an ϵ -relative approximation of z_2 .

The following theorem combines results of Shearer [\[38\]](#) on the non-vanishing of the independence polynomial in a complex disk, Peters and Regts [\[35\]](#) on the non-vanishing of the independence polynomial in a complex neighborhood of $(0, \lambda_c(\Delta))$, and Chudnovsky and Seymour [\[12\]](#) on the real-rootedness of the independence polynomial of claw-free graphs (extending the Heilmann–Lieb theorem on the real-rootedness of the matching polynomial [\[22\]](#)).

Theorem 2.2 ([\[12, 22, 35, 38\]](#)). *Let $\Delta \geq 3$ and $\delta \in (0, 1)$. There exists $c_{2.2} = c_{\delta, \Delta} > 0$ such that for any graph $G = (V, E)$ with maximum degree at most Δ , the partition function $Z_G(\lambda)$ of the hard-core model does not vanish on the region*

$$\mathcal{R}_{\delta, \Delta} := \{z \in \mathbb{C} : 0 \leq \Re(z) \leq (1 - \delta)\lambda_c(\Delta), |\Im(z)| \leq c_{\delta, \Delta}\} \cup \{z \in \mathbb{C} : |z| < (\Delta - 1)^{\Delta-1}/\Delta^\Delta\}.$$

Moreover, if G is claw-free, then all of the roots of $Z_G(\lambda)$ are on the negative real axis below $-e/(\Delta + 1)$. In particular, $Z_G(\lambda)$ does not vanish on the region

$$\mathcal{C}_\Delta := \mathbb{C} \setminus \{z \in \mathbb{C} : \Re(z) \leq -e/(\Delta + 1)\}.$$

Next, we record two consequences of the above zero-free regions of the partition function. The first is an FPTAS for $Z_G(\lambda)$, provided that $\lambda \in \mathbb{C}$ lies in the zero-free region, and follows by using Barvinok’s method [\[4\]](#) and the work of Patel and Regts [\[34\]](#).

Theorem 2.3 ([\[4, 34\]](#)). *Let $\Delta \geq 3$ and $\delta \in (0, 1)$. There exists a deterministic algorithm which, on input a graph $G = (V, E)$ on n vertices with maximum degree at most Δ , a (possibly complex) fugacity $\lambda \in \mathcal{R}_{\delta, \Delta}$, and an approximation parameter $\epsilon \in (0, 1)$, returns an ϵ -relative approximation of $Z_G(\lambda)$ in time $(n/\epsilon)^{O_{\delta, \Delta}(1)}$.*

Moreover, if G is claw-free, then the same conclusion holds for any $\lambda \in \mathcal{C}_\Delta$ with running time $(n/\epsilon)^{O_{\Delta, |\lambda|}(1)}$.

The second is a result of Michelen and Sahasrabudhe [33] on converting zero-free regions of probability generating functions into central limit theorems. We note that in order to establish our results with slightly worse quantitative dependencies, earlier results of Lebowitz, Pittel, Ruelle, and Speer [30] are sufficient.

Theorem 2.4 ([33, Theorem 1.2]). *Let X be a random variable taking values in $\{0, 1, \dots, n\}$ with mean μ and variance σ^2 and let $f_X(z) = \sum_{k=0}^n \mathbb{P}[X = k]z^k$ denote its probability generating function. Let $\delta = \min_{\zeta} |\zeta - 1|$, where ζ ranges over the (complex) roots of f_X . Then,*

$$\sup_{t \in \mathbb{R}} |\mathbb{P}[(X - \mu)/\sigma \leq t] - \mathbb{P}[\mathcal{Z} \leq t]| = O\left(\frac{\log n}{\delta\sigma}\right).$$

For our proof of the LCLT for the hard-core model, we will need the following simple lemma. The precise version we state here appears in work of Berkowitz [6] on quantitative local central limit theorems for triangle counts in $\mathbb{G}(n, p)$ and has been used, for instance, in further work on local central limit theorems for general subgraph counts in random graphs [37]; we include the short proof for the reader's convenience.

Lemma 2.5 ([6, Lemma 3]). *Let X be a random variable supported on the lattice $\mathcal{L} = \alpha + \beta\mathbb{Z}$ and let $\mathcal{N}(x) = e^{-x^2/2}/\sqrt{2\pi}$ denote the density of the standard normal distribution. Then*

$$\sup_{x \in \mathcal{L}} |\beta\mathcal{N}(x) - \mathbb{P}[X = x]| \leq \beta \int_{-\pi/\beta}^{\pi/\beta} |\mathbb{E}[e^{itX}] - \mathbb{E}[e^{it\mathcal{Z}}]| dt + e^{-\pi^2/(2\beta^2)}.$$

Proof. Let $\varphi_X(t) = \mathbb{E}e^{itX}$ and $\varphi(t) = \mathbb{E}e^{it\mathcal{Z}}$. By Fourier inversion on lattices and by Fourier inversion on \mathbb{R} , we have for $x \in \mathcal{L}$ that

$$\mathbb{P}[X = x] = \frac{\beta}{2\pi} \int_{-\pi/\beta}^{\pi/\beta} \varphi_X(t)e^{-itx} dt, \quad \mathcal{N}(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \varphi(t)e^{-itx} dt.$$

Therefore

$$\begin{aligned} |\beta\mathcal{N}(x) - \mathbb{P}[X = x]| &= \frac{\beta}{2\pi} \left| \int_{-\infty}^{\infty} \varphi(t)e^{-itx} dt - \int_{-\pi/\beta}^{\pi/\beta} \varphi_X(t)e^{-itx} dt \right| \\ &\leq \frac{\beta}{2\pi} \int_{-\pi/\beta}^{\pi/\beta} |\varphi(t) - \varphi_X(t)| dt + \frac{\beta}{2\pi} \left| \int_{|t| > \pi/\beta} e^{-itx} \varphi(t) dt \right| \\ &\leq \beta \int_{-\pi/\beta}^{\pi/\beta} |\varphi(t) - \varphi_X(t)| dt + e^{-\pi^2/(2\beta^2)}, \end{aligned}$$

where we used a standard tail bound on Gaussian integrals in the last line. Taking the supremum over $x \in \mathcal{L}$ completes the proof. \square

For our randomized algorithms, we will make use of recent results of Chen, Liu, and Vigoda [11] establishing optimal mixing of the Glauber dynamics for bounded-degree spin systems.

Theorem 2.6 ([11]). *Let $\Delta \geq 3$ and $\delta \in (0, 1)$. For every graph $G = (V, E)$ on n vertices with maximum degree at most Δ and for every $0 < \lambda \leq (1 - \delta)\lambda_c(\Delta)$, the ϵ -mixing time of the Glauber dynamics for the hard-core model on G at fugacity λ is $O_{\Delta, \delta}(n \log(n/\epsilon))$.*

Moreover, for any $C > 0$, the same conclusion holds for all line graphs $G = (V, E)$ on n vertices with maximum degree at most Δ and any $0 < \lambda \leq C$, with the implicit constant depending on Δ and C .

3. PROOF OF THE LOCAL CENTRAL LIMIT THEOREM

The goal of this section is to prove the following quantitative version of [Theorem 1.5](#).

Theorem 3.1. *Fix $\Delta \geq 3$ and $\delta > 0$. Let $\lambda \in (0, \lambda_c(\Delta) - \delta)$. Given a graph G on n vertices with maximum degree at most Δ , draw a random independent set I from the hard core model μ with fugacity λ and let $Y = |I|$. Let $\mu = \mathbb{E}_\lambda Y$ and $\sigma^2 = \text{Var}_\lambda Y$. Then*

$$\sup_{t \in \mathbb{Z}} |\sigma^{-1} \mathcal{N}((t - \mu)/\sigma) - \mathbb{P}[Y = t]| = O_{\Delta, \delta} \left(\min \left(\frac{(\log n)^{5/2}}{\sigma^2}, \frac{1}{\sigma^2} + \frac{\sigma^6 (\log n)^2}{n} \right) \right).$$

Moreover, if G is claw-free, then the same conclusion holds for any $\lambda \in (0, C)$, with the implicit constant depending on Δ and C .

Remark. [Lemma 3.2](#) below shows that $\sigma^2 = \Theta_\Delta(\lambda n)$ (or $\sigma^2 = \Theta_{\Delta, C}(\lambda n)$ in the case of claw-free graphs). Thus, for fixed Δ, δ, C and for a sequence of λ_n in the appropriate region for which $\lambda_n n \rightarrow \infty$, we see that the right hand goes to 0, thereby establishing the qualitative claim of [Theorem 1.5](#).

The proof of [Theorem 3.1](#) requires a few steps. First, we use the zero-free regions given by [Theorem 2.2](#) in order to bound σ .

Lemma 3.2. *Let $\Delta \geq 3$ and $\delta \in (0, 1)$. For any graph $G = (V, E)$ on n vertices with maximum degree at most Δ and for all $\lambda \in (0, \lambda_c(\Delta) - \delta)$, we have*

$$\frac{\lambda n}{(\Delta + 1)(1 + \lambda)^{2+\Delta}} \leq \text{Var}_\lambda Y \leq C_{\delta, \Delta} \lambda n.$$

Moreover, if G is claw-free, then for all $\lambda > 0$,

$$\frac{\lambda n}{(\Delta + 1)(1 + \lambda)^{2+\Delta}} \leq \text{Var}_\lambda Y \leq C_\Delta \lambda n.$$

Proof. The lower bounds follow from [[14](#), Lemma 9]. For the upper bounds, let $N = \alpha(G) \leq n$ be the independence number of G (which is also the degree of $Z_G(\lambda)$ as a polynomial in λ). Since the constant term of the polynomial $Z_G(\lambda)$ is 1, we can write

$$Z_G(\lambda) = \prod_{j=1}^N (1 - \lambda r_j)$$

where r_1, \dots, r_N are the inverses of the complex roots of $Z_G(\lambda)$. Then

$$\begin{aligned} \text{Var}_\lambda Y &= \lambda^2 (\log Z_G(\lambda))'' + \lambda (\log Z_G(\lambda))' \\ &= - \sum_{j=1}^N \left(\frac{\lambda^2 r_j^2}{(1 - \lambda r_j)^2} + \frac{\lambda r_j}{1 - \lambda r_j} \right) \\ &= -\lambda \sum_{j=1}^N \frac{r_j}{(1 - \lambda r_j)^2}. \end{aligned}$$

By [Theorem 2.2](#), $|r_j| = O(\Delta)$ and $|1/(1 - \lambda r_j)^2| = O_{\delta, \Delta}(1)$ when $\lambda \in (0, \lambda_c(\Delta) - \delta)$. In the case of claw-free graphs, $|1/(1 - \lambda r_j)^2| = O_\Delta(1)$ for all $\lambda \geq 0$ by [Theorem 2.2](#). Combining these estimates completes the proof. \square

Next, we control the low Fourier phases of our random variable.

Lemma 3.3. *Let $\Delta \geq 3$ and $\delta \in (0, 1)$. For any graph $G = (V, E)$ on n vertices with maximum degree at most Δ and for any $\lambda \in (0, \lambda_c(\Delta) - \delta)$, we have for all $t \in \mathbb{R}$ that*

$$|\mathbb{E}[e^{itX}] - \mathbb{E}[e^{itZ}]| = O_{\Delta, \delta} \left(\frac{(|t|(\log n)^{3/2} + \log n)}{\sigma} \right),$$

where $X = (Y - \mu)/\sigma$.

Moreover, if G is claw-free, then the same conclusion holds for all $\lambda > 0$ with the implicit constant depending only on Δ .

Proof. We prove the statement for general graphs; the proof for claw-free graphs is completely analogous. By [Theorem 2.4](#) combined with the zero-free region from [Theorem 2.2](#) we have for all $t \in \mathbb{R}$ that

$$|\mathbb{P}[X \leq t] - \mathbb{P}[\mathcal{Z} \leq t]| = O_{\Delta, \delta} \left(\frac{\log n}{\sigma} \right). \quad (3.1)$$

Let X' be X convolved with a centered Gaussian of infinitesimally small variance so that X' has density with respect to the Lebesgue measure on \mathbb{R} ; it suffices to prove the claim for X' and then pass to the limit. For this, we note that

$$\begin{aligned} \mathbb{E}[e^{itX'}] &= \int_{-\infty}^{\infty} e^{itz} p_{X'}(z) dz \\ &= \int_{|z| \leq \tau} e^{itz} p_{X'}(z) dz \pm \mathbb{P}[|X'| \geq \tau] \\ &= \left[e^{itz} \left(\int_{-\tau}^z p_{X'}(z') dz' \right) \right] \Big|_{z=-\tau}^{z=\tau} - \int_{-\tau}^{\tau} ite^{itz} \left(\int_{-\tau}^z p_{X'}(z') dz' \right) dz \pm \mathbb{P}[|X'| \geq \tau] \\ &= e^{it\tau} - \int_{-\tau}^{\tau} ite^{itz} \left(\int_{-\tau}^z p_{X'}(z') dz' \right) dz \pm \mathbb{P}[|X'| \geq \tau] - e^{it\tau} \mathbb{P}[|X'| \geq \tau] \\ &= e^{it\tau} - \int_{-\tau}^{\tau} ite^{itz} \mathbb{P}[X' \in [-\tau, z]] dz \pm \mathbb{P}[|X'| \geq \tau] - e^{it\tau} \mathbb{P}[|X'| \geq \tau] \\ &= e^{it\tau} - \int_{-\tau}^{\tau} ite^{itz} \mathbb{P}[X' \in [-\tau, z]] dz + e^{i\theta} \cdot O_{\Delta, \delta} \left(\frac{\log n}{\sigma} + e^{-\tau^2/4} \right), \end{aligned}$$

for some $\theta \in [0, 2\pi)$, where in the last line, we have used [\(3.1\)](#) along with a standard Gaussian tail bound. Applying the same calculation to \mathcal{Z} instead of X' and taking the difference, we find that

$$\begin{aligned} |\mathbb{E}[e^{itX'}] - \mathbb{E}[e^{itZ}]| &\leq |t| \left| \int_{-\tau}^{\tau} |\mathbb{P}[X' \in [-\tau, z]] - \mathbb{P}[\mathcal{Z} \in [-\tau, z]]| dz \right| + O_{\Delta, \delta} \left(\frac{\log n}{\sigma} + e^{-\tau^2/4} \right) \\ &\leq O_{\Delta, \delta} \left(\frac{(|\tau t| + 1) \log n}{\sigma} + e^{-\tau^2/4} \right). \end{aligned}$$

Since $\sigma \leq n$, setting $\tau = \sqrt{8 \log n}$ gives the desired conclusion. \square

Finally, we control the high Fourier phases of our random variable, following a similar strategy to that of Dobrushin and Tirozzi [\[16\]](#). We need the following elementary lemma.

Lemma 3.4. *Let $G = (V, E)$ be a graph on n vertices with maximum degree at most Δ . Then, there exists a subset $S \subseteq V$ of size $\Omega(n/\Delta^3)$ such that all vertices in S are pairwise distance at least 4 with respect to the graph distance. Moreover, there is an algorithm to find such a subset S in time $O_{\Delta}(n)$.*

Proof. Let v_1, \dots, v_n denote an arbitrary enumeration of the vertices. Initialize $S = \emptyset$. Consider the greedy algorithm which, at each time step, adds the first available vertex to the set S and removes

all vertices within distance 3 of this vertex from consideration. The algorithm stops when there are no more available vertices. The algorithm runs in time $O_\Delta(n)$ and outputs a set S such that any two vertices in S have graph distance at least 4. Moreover, since at each time, $O(\Delta^3)$ vertices are removed, it follows that $|S| = \Omega(n/\Delta^3)$. \square

Lemma 3.5. *Let $\Delta \geq 3$ and $C \geq 1$. There exists $c_{3.5}(\Delta, C) > 0$ satisfying the following. For any graph $G = (V, E)$ on n vertices with maximum degree at most Δ and for any $\lambda \in (0, C]$, let $X = (Y - \mu)/\sigma$. Then, for all $t \in [-\pi\sigma, \pi\sigma]$, we have*

$$|\mathbb{E}[e^{-itX}]| \leq \exp(-c_{3.5}(\Delta, C)\lambda nt^2/\sigma^2).$$

Proof. Rewriting the claim, it suffices to prove that for all $t \in \mathbb{R}$, $|t| \leq \pi$,

$$|\mathbb{E}[e^{-itY}]| \leq \exp(-c_{\Delta, C}\lambda nt^2).$$

Let S be a 4-separated set of vertices of G of size $s = \Omega(n/\Delta^3)$ coming from Lemma 3.4. Let T be the set of vertices that are at distance at least 2 from S in G and let $G[T]$ denote the graph on T induced by G . Let ν denote the distribution on $\mathcal{I}(G[T])$ induced by the hard-core model. We sample I by first sampling $J \sim \nu$ and then sampling from the conditional distribution (induced by the hard-core model and J) on $\mathcal{I}(G[v \cup N(v)])$ for each $v \in S$. The key observation is that these conditional distributions are mutually independent. In particular, given J , we can write

$$Y \stackrel{d.}{=} |J| + X_1 + \dots + X_s$$

where each X_j is an independent random variable with support in $\{0, \dots, \Delta + 1\}$, a probability mass at 0 of $\Omega_{\Delta, C}(1)$, and a probability mass of $\Omega_{\Delta, C}(\lambda)$ at 1. Note that the implicit constant in Ω does not depend on the specific realisation of J .

We claim that for all $|t| \leq \pi$ and all $j \in [s]$, for any realisation of J ,

$$|\mathbb{E}e^{-itX_j}| \leq 1 - c\lambda t^2$$

for some absolute $c = c_{\Delta, C} > 0$. Indeed, for any realisation of J , letting X'_j denote an independent copy of X_j , we have

$$\begin{aligned} |\mathbb{E}e^{-itX_j}|^2 &= \mathbb{E}e^{it(X_j - X'_j)} \\ &= \mathbb{P}[X_j = X'_j] + \sum_{k=1}^{\Delta+1} (\mathbb{P}[X_j - X'_j = k] + \mathbb{P}[X'_j - X_j = k]) \cos(kt) \\ &\leq \mathbb{P}[X_j = X'_j] + \sum_{k=2}^{\Delta+1} (\mathbb{P}[X_j - X'_j = k] + \mathbb{P}[X'_j - X_j = k]) + 2\mathbb{P}[X_j - X'_j = 1] \cos(t) \\ &= 1 - 2\mathbb{P}[X_j - X'_j = 1] + 2\mathbb{P}[X_j - X'_j = 1] \cos(t) \\ &= 1 - 2\mathbb{P}[X_j - X'_j = 1](1 - \cos(t)) \\ &\leq 1 - \frac{1}{4}\mathbb{P}[X_j - X'_j = 1]t^2 \leq 1 - \frac{1}{4}\mathbb{P}[X_j = 1]\mathbb{P}[X'_j = 0]t^2 \leq 1 - c_{\Delta, C}\lambda t^2, \end{aligned}$$

as claimed.

Finally, we have that for any $t \in [-\pi, \pi]$,

$$\begin{aligned} |\mathbb{E}[e^{-itY}]| &\leq \max_J |\mathbb{E}[e^{-itY} \mid J]| \\ &= \max_J \prod_{j=1}^s |\mathbb{E}e^{-itX_j}| \\ &\leq (1 - c_{\Delta, C}\lambda t^2)^{s/2} \end{aligned}$$

$$\leq \exp(-c'n\lambda t^2),$$

for an appropriate $c' = c'_{\Delta, C} > 0$ and the result follows. \square

3.1. Finishing the proof. We now prove [Theorem 3.1](#).

Proof of [Theorem 3.1](#). We prove the result for general graphs; the proof for claw-free graphs is essentially identical. Applying [Lemma 2.5](#) to $X = (Y - \mu)/\sigma \in \alpha + \beta\mathbb{Z}$, where $\alpha = -\mu/\sigma$ and $\beta = 1/\sigma$, and using [Lemmas 3.2](#), [3.3](#), and [3.5](#) we see that for $\sigma \geq 2$,

$$\begin{aligned} \sup_{x \in \mathcal{L}} |\beta \mathcal{N}(x) - \mathbb{P}[X = x]| &\leq \frac{1}{\sigma} \int_{-\pi\sigma}^{\pi\sigma} |\mathbb{E}[e^{itX}] - \mathbb{E}[e^{itZ}]| dt + e^{-\pi^2\sigma^2/2} \\ &\lesssim_{\Delta, \delta} \frac{1}{\sigma} \int_{-\pi\sigma}^{\pi\sigma} \min\left(\frac{|t|(\log n)^{3/2} + \log n}{\sigma}, e^{-c_{3.5}\lambda nt^2/\sigma^2} + e^{-t^2/2}\right) dt + e^{-\pi^2\sigma^2/2} \\ &\lesssim_{\Delta, \delta} \frac{1}{\sigma} \int_{-\pi\sigma}^{\pi\sigma} \min\left(\frac{|t|(\log n)^{3/2} + \log n}{\sigma}, e^{-c_{\Delta, \delta}t^2} + e^{-t^2/2}\right) dt + e^{-\pi^2\sigma^2/2} \\ &\lesssim_{\Delta, \delta} \frac{1}{\sigma} \int_{-C_{\Delta, \delta}\sqrt{\log \sigma}}^{C_{\Delta, \delta}\sqrt{\log \sigma}} \frac{|t|(\log n)^{3/2} + \log n}{\sigma} dt + \frac{1}{\sigma^2} \\ &\lesssim_{\Delta, \delta} \frac{(\log n)^{5/2}}{\sigma^2}. \end{aligned}$$

This gives the first term in the minimum in the statement of [Theorem 3.1](#). For the second term, we may assume that $1 \leq \sigma < \log n$. Let $\lambda' = \lambda/(1+\lambda)$ and observe that the hard-core distribution at fugacity λ is identical to the product distribution $\text{Ber}(\lambda')^{\otimes V}$ conditioned on the configuration being an independent set. Here, $\text{Ber}(\lambda')$ is the random variable which is 1 (or occupied) with probability λ' and 0 (or unoccupied) otherwise. A trivial union bound argument shows that a random sample from $\text{Ber}(\lambda')^{\otimes V}$ is an independent set with probability at least $1 - \lambda'^2 \Delta n = 1 - O_{\Delta}(\lambda^2 \Delta n) = 1 - O_{\Delta}(\sigma^4/n)$, where we have used [Lemma 3.2](#). Therefore, the probability of any configuration under the hard-core model is within a factor of $1 \pm O_{\Delta}(\sigma^4/n)$ of the probability of the same configuration under $\text{Ber}(\lambda')^{\otimes V}$.

Let Y' denote the random variable counting the number of 1s in a random sample from $\text{Ber}(\mu)^{\otimes V}$, and let μ' and σ' denote the mean and standard deviation of Y' . Then, by the classical DeMoivre-Laplace central limit theorem (see [\[36, Chap. VII, Theorem \(4\)\]](#) for the quantitative version used here), we get that for any integer k ,

$$\left| \mathbb{P}[Y' = k] - \frac{1}{\sigma'} \mathcal{N}\left(\frac{k - \mu'}{\sigma'}\right) \right| = O\left(\frac{1}{\sigma'^2}\right).$$

Moreover, from the comparison between the hard-core model and $\text{Ber}(\lambda')^{\otimes V}$ mentioned above, as well as the Chernoff bound for the Binomial distribution, we see that

$$\begin{aligned} \mathbb{P}[Y = k] &= \mathbb{P}[Y' = k] \pm O_{\Delta}(\sigma^4/n) \\ \mu &= \mu'(1 \pm O_{\Delta}(\sigma^4 \log n/n)) \\ \sigma^2 &= \sigma'^2(1 \pm O_{\Delta}(\sigma^6(\log n)^2/n)). \end{aligned}$$

Substituting this in the above gives the desired conclusion. \square

4. DETERMINISTIC ALGORITHMS

For our deterministic algorithms, we will need the following preliminary lemma, which allows us to find a ‘good’ fugacity at which to either apply the LCLT directly or algorithmize its proof.

Lemma 4.1. *Let $\Delta \geq 3$ and $\delta \in (0, 1/2)$. There exists a constant $C_{4.1} = C_{4.1}(\Delta, \delta) \geq 1$ for which the following holds.*

For any $\alpha \leq (1 - \delta)\alpha_c(\Delta)$, there exists a unique $\lambda_ < \lambda_c(\Delta)$ so that $\alpha_{K_{\Delta+1}}(\lambda_*) = \alpha$. Further, for any graph $G = (V, E)$ on n vertices of maximum degree at most Δ and for any $1 \leq k \leq \alpha n$, there exists an integer $t \in \{0, \dots, \lceil C_{4.1} n \lambda_* \rceil\}$ such that*

$$|n\alpha_G(t/(C_{4.1}n)) - k| \leq 1/2.$$

Moreover, if G is claw-free, then for any $1 \leq k \leq (1 - \delta)i^(G)$, there exists an integer $t \in \{0, \dots, \lceil n \cdot 8^{(\Delta+1)/\delta^2} \rceil\}$ such that the same conclusion holds.*

In either case, such an integer t can be found deterministically in time $n^{O_{\delta, \Delta}(1)}$.

Proof. For general graphs, the existence of such an integer t and constant $C_{4.1}$ follows from [14, Lemma 5] (upon replacing [14, Lemma 9] by the optimal variance bound Lemma 3.2). For claw-free graphs, the existence of such an integer t follows similarly as a simple consequence of the following observations: (i) $\alpha_G(\lambda)$ is monotonically increasing in λ ; (ii) since $i^*(G) \geq n/(\Delta + 1)$, we have that for all $\lambda \geq 8^{(\Delta+1)/\delta^2}$

$$\begin{aligned} \mathbb{P}_\lambda[Y \leq (1 - \delta/2)i^*(G)] &= \frac{1}{Z_G(\lambda)} \sum_{k=0}^{(1-\delta/2)i^*(G)} \binom{n}{k} \lambda^k \leq \frac{1}{Z_G(\lambda)} 2^n \lambda^{(1-\delta/2)i^*(G)} \\ &\leq \frac{1}{Z_G(\lambda)} \lambda^{i^*(G)} \cdot 2^n \lambda^{-\delta n/(2\Delta+2)} \leq \delta/4 \end{aligned}$$

from which we see that $n\alpha_G(8^{(\Delta+1)/\delta^2}) \geq (1 - \delta)i^*(G)$; and (iii) for all $\lambda > 0$

$$\frac{d}{d\lambda} \alpha_G(\lambda) = \frac{1}{n\lambda} \text{Var}_\lambda(Y) \leq C_{\delta, \Delta},$$

where the inequality follows from Lemma 3.2.

As for the algorithmic claim, note that for each t in the range, we can deterministically approximate $n\alpha_G(t/(C_{4.1}n))$ to within an additive error of $1/4$ in time $n^{O_{\delta, \Delta}(1)}$ by Theorem 1.7, so that we may find a t with the desired property in the stated time. \square

4.1. EPTAS from the LCLT. Recall that an EPTAS (efficient polynomial-time approximation scheme) is a PTAS (polynomial-time approximation scheme) with a running time of the form $f(\epsilon)n^{O(1)}$ i.e. the degree of the polynomial in n is independent of the error parameter ϵ . By combining our LCLT with Lemma 4.1 and Theorems 1.7 and 2.3, we immediately obtain an EPTAS for $i_k(G)$, for all k bounded away from the relevant barrier.

Theorem 4.2. *Let $\Delta \geq 3$ and $\delta \in (0, 1/2)$. There exists a deterministic algorithm which, on input a graph $G = (V, E)$ on n vertices of maximum degree at most Δ , an integer $1 \leq k \leq (1 - \delta)\alpha_c(\Delta)n$, and an error parameter $\epsilon \in (0, 1)$ outputs an ϵ -relative approximation to $i_k(G)$ in time $n^{O_{\delta, \Delta}(1)} \exp(\tilde{O}_{\delta, \Delta}(\epsilon^{-1}/\sqrt{n}))$.*

Moreover, if G is claw-free, then the same conclusion holds for any $1 \leq k \leq (1 - \delta)i^(G)$.*

Proof. If $k < e^{-5}n/(\Delta + 1)$, then an FPTAS for $i_k(G)$ is already implicit in [13]. Therefore, we may restrict our attention to $k \geq e^{-5}n/(\Delta + 1)$. We may also assume that $\epsilon^{-1} \leq c_{\delta, \Delta} \sqrt{n}/(\log n)^3$ for a sufficiently small constant $c_{\delta, \Delta} > 0$; otherwise $\exp(\tilde{O}_{\delta, \Delta}(\epsilon^{-1}/\sqrt{n})) \geq 4^n$ so that exhaustive enumeration runs in the claimed time.

Fix $k \geq e^{-5}n/(\Delta + 1)$ which also lies in the specified range, let t_k denote the corresponding value of t given by Lemma 4.1, and let $\lambda_k = t_k/C_{4.1}$ denote the corresponding fugacity. The upper bound

on $\alpha'_G(\lambda)$ in the proof of [Lemma 4.1](#) shows that $\lambda_k = \Omega_{\delta,\Delta}(1)$. In particular, $\mu_k := \mu_{\lambda_k}, \sigma_k := \sigma_{\lambda_k}$ satisfy $\mu_k, \sigma_k^2 = \Theta_{\delta,\Delta}(n)$. Moreover, since $|k - \mu_k| \leq 1/2$, it follows by [Theorem 3.1](#) that

$$\begin{aligned} \frac{i_k(G)\lambda_k^k}{Z_G(\lambda_k)} &= \mathbb{P}[Y = k] \\ &= \sigma_k^{-1} \mathcal{N}((k - \mu_k)/\sigma_k) \pm O_{\Delta,\delta}((\log n)^{5/2}/n) \\ &= (1 \pm \epsilon/1000) \cdot \sigma_k^{-1} \mathcal{N}((k - \mu_k)/\sigma_k) \\ &= (1 \pm \epsilon/500) \cdot (\sqrt{2\pi}\sigma_k)^{-1}. \end{aligned}$$

Therefore, letting $\widehat{Z}_G(\lambda_k)$ and $\widehat{\sigma}_k$ denote $\epsilon/1000$ -relative approximations to $Z_G(\lambda_k)$ and σ_k , both of which can be computed deterministically in time $(n/\epsilon)^{O_{\delta,\Delta}(1)}$ by [Theorems 1.7](#) and [2.3](#), it follows that

$$\widehat{i}_k(G) = \lambda_k^{-k} \cdot \widehat{Z}_G(\lambda_k) \cdot (\sqrt{2\pi}\widehat{\sigma}_k)^{-1}$$

is an ϵ -relative approximation to $i_k(G)$, as desired. \square

4.2. FPTAS from the proof of the LCLT. In the previous subsection, we saw how our LCLT gives an EPTAS for $i_k(G)$, for all k bounded away from the relevant barrier. Extending this to an FPTAS requires bypassing the error term $\tilde{O}(1/\sigma^2)$ in [Theorem 3.1](#). For this, instead of approximating $\mathbb{P}[Y = k]$ by $\sigma^{-1}\mathcal{N}((k - \mu)/\sigma)$, we will directly approximate $\mathbb{P}[Y = k]$ to the desired accuracy. The key ingredient required for this is the following.

Lemma 4.3. *Fix $\Delta \geq 3$, $\delta \in (0, 1/2)$, and a parameter $C \geq 1$. There exists a deterministic algorithm which, on input a graph $G = (V, E)$ on n vertices of maximum degree at most Δ , $n^{-2} \leq \lambda \leq (1 - \delta)\lambda_c(\Delta)$, an error parameter $\epsilon \in (0, 1/\sqrt{n})$, and $t \in [-C\sqrt{\log 1/\epsilon}, C\sqrt{\log 1/\epsilon}]$ outputs an ϵ^{50} -relative, ϵ^{50} -additive approximation to $\mathbb{E}_\lambda[e^{itX}]$ in time $(n/\epsilon)^{O_{\delta,\Delta,C}(1)}$, where $X = (Y - \mu_\lambda)/\sigma_\lambda$.*

Moreover, if G is claw-free, the same conclusion holds for all $\lambda \geq n^{-2}$ with running time $(n/\epsilon)^{O_{\delta,C,|\lambda|}(1)}$.

Proof. We provide the proof for general graphs; the proof for claw-free graphs follows by straightforward modifications. For convenience of notation, we denote μ_λ and σ_λ by μ and σ . By [Theorems 1.7](#) and [2.3](#), we can compute ϵ^{100} -relative approximations \widehat{Z} , $\widehat{\mu}$, and $\widehat{\sigma}$ to $Z_G(\lambda)$, μ , and σ in time $(n/\epsilon)^{O_{\delta,\Delta}(1)}$. Note that

$$\begin{aligned} \mathbb{E}_\lambda[e^{itX}] &= e^{-it\mu/\sigma} \cdot \mathbb{E}_\lambda[e^{itY/\sigma}] \\ &= e^{-it\mu/\sigma} \cdot \frac{Z_G(\lambda e^{it/\sigma})}{Z_G(\lambda)}. \end{aligned}$$

We have the following two cases.

Case I: $\lambda \leq \frac{1}{10\Delta}$. In this case, for any $t \in \mathbb{R}$, $\lambda e^{it/\widehat{\sigma}} \in \mathcal{R}_{\delta,\Delta}$, where $\mathcal{R}_{\delta,\Delta}$ is the zero-free region in [Theorem 2.2](#).

Case II: $\lambda > \frac{1}{10\Delta}$. In this case, we may assume that $\lambda e^{it/\widehat{\sigma}} \in \mathcal{R}_{\delta,\Delta}$ for all $t \in \mathbb{R}$, $|t| \leq C\sqrt{\log 1/\epsilon}$. Otherwise, since $\widehat{\sigma}^2 = \Omega_{\delta,\Delta}(n)$ by [Lemma 3.2](#), it follows that $\epsilon^{-1} = \exp(\Omega_{\delta,\Delta,C}(n))$, so that exhaustive enumeration runs in the claimed time.

Hence, in either case, it follows from [Theorem 2.3](#) that an ϵ^{100} -relative approximation to $Z_G(\lambda e^{it/\widehat{\sigma}})$, which we denote by \widehat{Z}_t , can be computed in time $(n/\epsilon)^{O_{\delta,\Delta}(1)}$.

We claim that the output

$$e^{-it\widehat{\mu}/\widehat{\sigma}} \cdot \frac{\widehat{Z}_t}{\widehat{Z}}$$

is an approximation to $\mathbb{E}_\lambda[e^{itX}]$ of the desired quality. Indeed, by the assumed upper bounds on $|t|$ and ϵ and the assumed lower bound on λ , we have for all n sufficiently large (depending on δ, Δ, C) that

$$\frac{e^{it/\sigma}}{e^{it/\hat{\sigma}}} = e^{it(\hat{\sigma}-\sigma)/(\sigma\hat{\sigma})} = e^{i\theta}, \quad \frac{e^{-it\mu}}{e^{-it\hat{\mu}}} = e^{i\theta'}$$

for some $\theta, \theta' \in \mathbb{R}$ satisfying $|\theta|, |\theta'| \leq \epsilon^{97}$. Therefore, $\frac{e^{it\hat{\mu}/\hat{\sigma}}}{e^{it\mu/\sigma}} = e^{i\theta''}$ for some $\theta'' \in \mathbb{R}$, $|\theta''| \leq \epsilon^{90}$ and

$$\begin{aligned} |Z_G(\lambda e^{it/\hat{\sigma}}) - Z_G(\lambda e^{it/\sigma})| &= \left| \sum_{k=0}^n i_k(G) \lambda^k (e^{itk/\hat{\sigma}} - e^{itk/\sigma}) \right| \leq \sum_{k=0}^n i_k(G) \lambda^k |e^{ik\theta} - 1| \\ &\leq \epsilon^{94} \sum_{k=0}^n i_k(G) \lambda^k = \epsilon^{94} Z_G(\lambda). \end{aligned}$$

Hence,

$$\begin{aligned} \mathbb{E}_\lambda[e^{itX}] &= \left(e^{i\theta''} e^{-it\hat{\mu}/\hat{\sigma}} \right) \cdot \frac{Z_G(\lambda e^{it/\hat{\sigma}})}{Z_G(\lambda)} + e^{-it\mu/\sigma} \cdot \frac{Z_G(\lambda e^{it/\sigma}) - Z_G(\lambda e^{it/\hat{\sigma}})}{Z_G(\lambda)} \\ &= \left(e^{i\theta''} e^{-it\hat{\mu}/\hat{\sigma}} \right) \cdot e^{\pm 2\epsilon^{100}} e^{\pm i\epsilon^{100}} \frac{\hat{Z}_t}{\hat{Z}} + w_t \end{aligned}$$

where $w_t \in \mathbb{C}$ with $|w| \leq \epsilon^{94}$, as desired. \square

Given the preceding lemma, [Theorems 1.1](#) and [1.2](#) follow from the proof of [Theorem 3.1](#).

Proof of [Theorems 1.1](#) and [1.2](#). Again, we provide the details only for general graphs i.e. [Theorem 1.2](#); the argument for claw-free graphs is essentially the same. Let $k \geq 1$ be as in the statement of the theorem. Moreover, we may assume that $\epsilon \in (0, 1/n)$, since the statement for larger values of ϵ follows from the statement for $\epsilon = 1/n$. Let t_k denote the integer returned by [Lemma 4.1](#) and let $\lambda := \lambda_k = t_k/C_{4.1}$ denote the corresponding fugacity. Since $\mu := \mu_{\lambda_k} \geq 1/2$, it follows as in the proof of [Lemma 4.1](#) that $\lambda = \Omega_{\delta, \Delta}(1/n)$.

Let $\sigma := \sigma_{\lambda_k}$, $X = (Y - \mu)/\sigma$, and $\gamma = \min(\pi\sigma, C\sqrt{\log 1/\epsilon})$, where C is a sufficiently large constant depending on Δ, δ . Let $\hat{Z}, \hat{\mu}, \hat{\sigma}$ be ϵ^{100} -relative approximations to $Z_G(\lambda), \mu, \sigma$, which can be found deterministically in time $(n/\epsilon)^{O_{\delta, \Delta}(1)}$ by [Theorems 1.7](#) and [2.3](#).

Let $x = (k - \mu)/\sigma$ and $\hat{x} = (k - \hat{\mu})/\hat{\sigma}$. Then, as in the proof of [Theorem 3.1](#), we have

$$\begin{aligned} \mathbb{P}_\lambda[X = x] &= \frac{1}{2\pi\sigma} \int_{-\pi\sigma}^{\pi\sigma} \mathbb{E}_\lambda[e^{itX}] e^{-itx} dt \\ &= \pm\epsilon^{100} + \Re \left(\frac{1}{2\pi\sigma} \int_{-\gamma}^{\gamma} \mathbb{E}_\lambda[e^{itX}] e^{-itx} dt \right) \\ &= \pm 2\epsilon^{75} + \Re \left(\frac{\epsilon^{100}}{2\pi\sigma} \sum_{\ell=-\gamma\epsilon^{-100}}^{\gamma\epsilon^{-100}} \mathbb{E}_\lambda[e^{i\epsilon^{100}\ell X}] e^{-i\epsilon^{100}\ell x} \right) \\ &= \pm 3\epsilon^{75} + \Re \left(\frac{\epsilon^{100}}{2\pi\sigma} \sum_{\ell=-\gamma\epsilon^{-100}}^{\gamma\epsilon^{-100}} \mathbb{E}_\lambda[e^{i\epsilon^{100}\ell X}] e^{-i\epsilon^{100}\ell \hat{x}} \right), \end{aligned}$$

where the first line uses the Fourier inversion formula for lattices, the second line uses the definition of γ , [Lemma 3.5](#), and [Lemma 3.2](#), the third line uses the upper bound on ϵ and the lower bound on λ , and the last line uses the lower bound on λ .

Next, for each integer $\ell \in [-\gamma\epsilon^{-100}, \gamma\epsilon^{-100}]$, let \widehat{Z}_ℓ denote an ϵ^{100} -additive, ϵ^{100} -relative approximation to $\mathbb{E}_\lambda[e^{i\epsilon^{100}\ell X}]$. By [Lemma 4.3](#), these approximations may be found deterministically in time $(n/\epsilon)^{O_{\delta,\Delta}(1)}$. Then,

$$\begin{aligned} (1 \pm \epsilon^{75}) \frac{i_k(G)\lambda^k}{\widehat{Z}} &= \frac{i_k(G)\lambda^k}{Z_G(\lambda)} = \mathbb{P}_\lambda[X = x] \\ &= \pm 3\epsilon^{75} + \Re \left(\frac{\epsilon^{100}}{2\pi\sigma} \sum_{\ell=-\gamma\epsilon^{-100}}^{\gamma\epsilon^{-100}} \mathbb{E}_\lambda[e^{i\epsilon^{100}\ell X}] e^{-i\epsilon^{100}\ell \widehat{x}} \right) \\ &= \pm 4\epsilon^{75} + (1 \pm \epsilon^{50}) \Re \left(\frac{\epsilon^{100}}{2\pi\widehat{\sigma}} \sum_{\ell=-\gamma\epsilon^{-100}}^{\gamma\epsilon^{-100}} \widehat{Z}_\ell e^{-i\epsilon^{100}\ell \widehat{x}} \right). \end{aligned}$$

Finally, since $\mathbb{P}_\lambda[X = x] = \mathbb{P}_\lambda[Y = k] = \Omega_{\delta,\Delta}(1/\sigma)$, we have that

$$(1 \pm \epsilon^{50}) \frac{i_k(G)\lambda^k}{\widehat{Z}} = (1 \pm \epsilon^{50}) \Re \left(\frac{\epsilon^{100}}{2\pi\widehat{\sigma}} \sum_{\ell=-\gamma\epsilon^{-100}}^{\gamma\epsilon^{-100}} \widehat{Z}_\ell e^{-i\epsilon^{100}\ell \widehat{x}} \right),$$

so that the quantity

$$\lambda^{-k} \cdot \widehat{Z} \cdot \Re \left(\frac{\epsilon^{100}}{2\pi\widehat{\sigma}} \sum_{\ell=-\gamma\epsilon^{-100}}^{\gamma\epsilon^{-100}} \widehat{Z}_\ell e^{-i\epsilon^{100}\ell \widehat{x}} \right),$$

which can be computed deterministically in time $(n/\epsilon)^{O_{\delta,\Delta}(1)}$, is an ϵ -relative approximation to $i_k(G)$. \square

5. RANDOMIZED ALGORITHMS

5.1. A quasi-linear time sampling algorithm. We proceed to the proof of [Theorem 1.3](#). We will prove the result for independent sets, noting that the proof for matchings follows identically.

We will need the following preliminary lemma, which is an efficient randomized version of [Lemma 4.1](#).

Lemma 5.1. *Let $\Delta \geq 3$ and $\delta \in (0, 1/2)$. There exists a constant $c_{5.1}(\delta, \Delta) > 0$ and a randomized algorithm with the following property: for any graph $G = (V, E)$ on n vertices of maximum degree at most Δ , for any $1 \leq k \leq (1 - \delta)n\alpha_c(\Delta)$, and for any $\epsilon \in (0, 1)$, the algorithm outputs $\lambda \in (0, (1 - c_{5.1})\lambda_c(\Delta)]$ satisfying*

$$|\mathbb{E}_\lambda Y - k| \leq \sqrt{\text{Var}_\lambda Y}$$

with probability $1 - \epsilon$. The running time of the algorithm is $O_{\Delta,\delta}(n \log(n/\epsilon)(\log n)^3)$.

Remark. If $k \geq n/(3\Delta)$, then it is easily seen that $\lambda \geq 1/(100\Delta)$.

Proof. If $1 \leq k \leq \sqrt{n}$, then it follows from [Corollary 6.3](#) that the deterministic choice $\lambda = k/n$ satisfies the desired conclusion. Therefore, it suffices to consider the case $\sqrt{n} \leq k \leq (1 - \delta)n\alpha_c(\Delta)$. Let $\Lambda = [n^{-1/3}, (1 - c_{5.1})\lambda_c(\Delta)] \cap (\mathbb{Z}/n^2)$. By [Lemma 4.1](#), there exists $\lambda \in \Lambda$ satisfying

$$|\mathbb{E}_\lambda Y - k| \leq 1/2.$$

Further, for any $\lambda \in \Lambda$, it follows from [Theorem 2.4](#) and [Lemma 3.2](#) that

$$|\mathbb{E}_\lambda Y - \text{med}_\lambda Y| = \widetilde{O}_{\delta,\Delta}(1),$$

where med denotes the median. Since $\text{Var}_\lambda Y = \Theta_{\delta,\Delta}(n\lambda)$ by [Lemma 3.2](#), it follows that there exists $\lambda \in \Lambda$ satisfying

$$|\text{med}_\lambda Y - k| \leq \frac{1}{2} \sqrt{\text{Var}_\lambda Y}$$

and it suffices to output such a $\lambda \in \Lambda$.

For any $\lambda \in \Lambda$, there is a randomized algorithm to estimate $\text{med}_\lambda Y$ to within $\sqrt{\text{Var}_\lambda Y}/2$ additive error which succeeds with probability $1 - (\epsilon/n^3)$ and runs in time $O(n \log(n/\epsilon)(\log n)^2)$. Indeed, by [Theorems 2.4](#) and [2.6](#), [Lemma 3.2](#) and the Chernoff bound, this may be accomplished by taking the median of $O_{\delta,\Delta}(\log(n/\epsilon))$ independent runs of the Glauber dynamics, each for $\Theta_{\Delta,\delta}(n \log n)$ steps. Therefore, running binary search with the above primitive takes time $O(n \log(n/\epsilon)(\log n)^3)$ and except with probability $1 - \epsilon$, outputs $\lambda \in \Lambda$ satisfying the desired conclusion. \square

We now present our sampling algorithm. For simplicity of notation, we will restrict attention to the case $k > n/(3\Delta)$; the case $k \leq n/(3\Delta + 1)$ is already handled by the down-up walk in [\[8\]](#) with asymptotically optimal running time. Below, $c_{\Delta,\delta}, C_{\Delta,\delta}, c'_{\Delta,\delta}, C'_{\Delta,\delta}$ are constants depending on Δ, δ whose values can be determined using *a priori* analysis. We will assume that $\epsilon \geq \exp(-n/C'_{\Delta,\delta})$; for smaller ϵ , it follows from [Theorems 2.6](#) and [3.1](#) that rejection sampling for the hard-core model at the fugacity λ given by [Lemma 5.1](#) outputs a distribution within ϵ -TV distance of the uniform distribution on $\mathcal{I}_k[G]$ in time

$$O_{\delta,\Delta}(n \log(n/\epsilon) \cdot \log(n) \cdot \sqrt{n} \log(1/\epsilon)) = O_{\delta,\Delta}(n \log(1/\epsilon)^{5/2} \log n).$$

For $1/100 \geq \epsilon \geq \exp(-n/C'_{\Delta,\delta})$, we use the following algorithm.

- (Preprocessing Step 1) Using [Lemma 5.1](#), find $\lambda \in \Lambda$ such that

$$|\mathbb{E}_\lambda Y - k| \leq \sqrt{\text{Var}_\lambda Y}.$$

- (Preprocessing Step 2) Using [Lemma 3.4](#), find $S \subseteq V$ such that $|S| = \Omega(n/\Delta^3)$ and such that the vertices in S are distance at least 4 apart. Let $T = \{v : v \text{ such that } \text{dist}(v, S) \geq 2\}$.
- (Preprocessing Step 3) Find an independent set I_0 of size k using a similar greedy algorithm as in [Lemma 3.4](#).
- (Initialize Core) Run Glauber dynamics for the hard-core model on G at fugacity λ for $\Theta_{\Delta,\delta}(n \log(n/\epsilon))$ -steps to obtain an independent set I' . Let $J = I' \cap T$.
- (Set parameters) Fix $p = c_{\Delta,\delta} \sqrt{n/\log(1/\epsilon)}$, where $c_{\Delta,\delta} > 0$ is sufficiently small. For each $v \in S$ and for each $K \in \mathcal{I}(G[v \cup N(v)])$, use exhaustive enumeration to compute

$$p_{v,K} = \mathbb{P}_\lambda[I \cap N(v) = K \mid I \cap T = J].$$

Let

$$q_v = \min(p_{v,\emptyset}, p_{v,K_v}),$$

where $K_v \in \mathcal{I}[N(v)]$ denotes the independent set with v occupied and all other vertices unoccupied. Let

$$q = \min_{v \in S} q_v.$$

For each $v \in S$, let W_v be a random variable taking values in $\mathcal{I}[N(v)] \cup \{?\}$ which takes on the value $?$ with probability $2q$, \emptyset with probability $p_{v,\emptyset} - q$, K_v with probability $p_{v,K_v} - q$, and $K \notin \{\emptyset, K_v\}$ with probability $p_{v,K}$.

- (Resample Neighborhoods Step 1) For each $v \in S$, independently, sample W_v . If $W_v \neq ?$, then set $I \cap N(v) = W_v$. Let $S^* = \{v : W_v = ?\}$. If $|S^*| \leq c'_{\Delta,\delta} n$, where $c'_{\Delta,\delta} > 0$ is sufficiently small, then proceed to the final step. Else, proceed to the next step.

- (Resample Neighborhoods Step 2) Let ℓ be the current number of vertices chosen and let $k - \ell = k'$. With probability $p \binom{|S^*|}{k'} 2^{-|S^*|}$ (here, we use the convention that the binomial coefficient is 0 if $k' \notin [0, |S^*|]$) sample a random subset of S^* of size k' and set $I \cap N(v) = K_v$ for the vertices in the subset and $I \cap N(v) = \emptyset$ for the vertices not in the subset; with the remaining probability, proceed to the next step.
- Repeat all steps after preprocessing at most at most $C_{\Delta, \delta} \log(1/\epsilon)^{3/2}$ times for sufficiently large $C_{\Delta, \delta} > 0$; if no valid sample after has been reached, output I_0 .

We now show that the above algorithm satisfies the assertion of [Theorem 1.3](#) for k and ϵ in the specified range.

Proof of [Theorem 1.3](#). Before analyzing the correctness of the algorithm, let us quickly bound its running time. By [Lemmas 3.4](#) and [4.1](#), the preprocessing steps take time $O_{\Delta, \delta}(n \log(n/\epsilon)(\log n)^3)$. Note that the Preprocessing Step 3 can indeed be accomplished using the greedy algorithm in [Lemma 3.4](#) since $\alpha_c(\Delta) \leq 1/(\Delta + 1)$. Each run of Initialize Core takes time $O_{\Delta, \delta}(n \log(n/\epsilon) \log n)$ since each step of Glauber dynamics takes time $O_{\Delta, \delta}(\log n)$ to implement. The step Set parameters takes time $O_{\Delta, \delta}(n)$. Resample Neighborhoods Step 1 takes time $O_{\Delta, \delta}(n)$. In Resample Neighborhoods Step 2, we can compute the probability in time $O(n(\log n)^2)$ by [\[17\]](#) and then sample from the hypergeometric distribution using sampling without replacement in time $O(n \log n)$. Thus, we see that the running time of the algorithm is

$$O_{\delta, \Delta}(n \log(n/\epsilon)(\log n)^3 + n \log(n/\epsilon) \log n \log(1/\epsilon)^{3/2}).$$

We now proceed to the proof of correctness. The idea is that the algorithm may be viewed as implementing rejection sampling where the base sampler outputs I according to the distribution $\mu_{G, \lambda}(\cdot \mid |I| \equiv k \pmod p)$ in time $\tilde{O}_{\delta, \Delta}(n)$. This leads to a $\tilde{O}_{\delta, \Delta}(n)$ time algorithm for approximately sampling from the uniform distribution on $\mathcal{I}_k[G]$ since by [Lemma 5.1](#) and [Theorem 3.1](#), $\mathbb{P}_\lambda[|I| = k \mid |I| \equiv k \pmod p] = \tilde{\Omega}_{\delta, \Delta}(1)$.

To formalize this, we begin by noting that for all $t \in \{0, \dots, p-1\}$ and for all $J \in \mathcal{I}[G[T]]$, it follows from the calculations and notation in the proof of [Lemma 3.5](#) that

$$\begin{aligned} \left| \mathbb{P}_\lambda[|I| \equiv t \pmod p \mid I \cap T = J] - \frac{1}{p} \right| &= \left| \mathbb{P}_\lambda[X_1 + \dots + X_s \equiv t - |J| \pmod p] - \frac{1}{p} \right| \\ &\leq \frac{1}{p} \sum_{\ell=1}^{p-1} \exp(-\Omega_{\delta, \Delta}(n\ell^2/p^2)) \\ &\leq \frac{\epsilon}{100p}, \end{aligned}$$

provided that $c_{\Delta, \delta} > 0$ is chosen to be sufficiently small. Hence, for any J

$$\mathbb{P}_\lambda[I \cap T = J \mid |I| \equiv k \pmod p] = \mathbb{P}_\lambda[I \cap T = J](1 \pm \epsilon/50),$$

so that up to an $\epsilon/49$ -TV distance, the distribution of the set J in Initialize Core is $\mu_{G, \lambda}[I \cap T = \cdot \mid |I| \equiv k \pmod p]$.

Next, for any realisation $\vec{w} = (w_v)_{v \in S}$ of $\vec{W} = (W_v)_{v \in S}$, let $S^*(\vec{w})$ denote the corresponding subset and let $\mu_{\vec{w}}$ denote the uniform distribution on $\{0, 1\}^{S^*(\vec{w})}$ with 0 denoting unoccupied and 1 denoting occupied. Then, sampling from the conditional distribution $\mu_{G, \lambda}(\cdot \mid I \cap T = J)$ is equivalent to first sampling $\vec{W} = \vec{w}$ according to the distribution specified in Set Parameters and then resampling the \cdot according to $\mu_{\vec{w}}$. For any realisation \vec{w} of \vec{W} with $|S^*(\vec{w})| \geq c'_{\Delta, \delta} n$, it follows as above that for any $t \in \{0, \dots, p-1\}$,

$$\left| \mathbb{P}_\lambda[|I| \equiv t \pmod p \mid I \cap T = J, \vec{W} = \vec{w}] - \frac{1}{p} \right| \leq \max_{t'} \left| \mathbb{P}[\text{Binomial}(|S^*(\vec{w})|, 1/2) = t' \pmod p] - \frac{1}{p} \right|$$

$$\leq \frac{\epsilon}{100p},$$

provided that $c_{\Delta,\delta} > 0$ is chosen to be sufficiently small compared to $c'_{\Delta,\delta}$. Since $q = \Omega_{\delta,\Delta}(1)$, it follows by the Chernoff bound that

$$\mathbb{P}_\lambda[|S^*(\vec{w})| \leq c'_{\Delta,\delta}n \mid I \cap T = J] \leq \exp(-\Omega_{\Delta,\delta}(n)) \leq \epsilon/n^3$$

for all sufficiently large n . This shows that

- (P1) The probability of moving directly to the final step from Resample Neighborhoods Step 1 is at most ϵ/n^3 , and
- (P2) Up to an $\epsilon/49$ -TV distance, the $(W_v)_{v \in S}$ in Resample Neighborhoods Step 1 follow the distribution $(W_v)_{v \in S} \mid \{|I| \equiv k \pmod p, I \cap T = J\}$.

Now, we analyze Resample Neighborhoods Step 2. Observe that for any realisation $\vec{W} = \vec{w}$ with $|S^*(\vec{w})| \geq c'_{\Delta,\delta}n$,

$$\begin{aligned} \mathbb{P}_\lambda[|I| = k \mid I \cap T = J, \vec{W} = \vec{w}, |I| \equiv k \pmod p] &= \frac{\mathbb{P}_\lambda[|I| = k \mid I \cap T = J, \vec{W} = \vec{w}]}{\mathbb{P}_\lambda[|I| \equiv k \pmod p \mid I \cap T = J, \vec{W} = \vec{w}]} \\ &= \frac{\mathbb{P}[\text{Binomial}(|S^*(\vec{w})|, 1/2) = k']}{(1 \pm \epsilon/100)/p} \\ &= (1 \pm \epsilon/50)p \binom{|S^*(\vec{w})|}{k'} 2^{-|S^*(\vec{w})|}. \end{aligned}$$

Therefore, Resample Neighborhoods Step 2 samples from a distribution within $\epsilon/49$ -TV distance to the conditional distribution $\mu_{\lambda,G}[\cdot \mid I \cap T = J, \vec{W} = \vec{w}, |I| \equiv k \pmod p]$ and rejects the sample if $|I| \neq k$.

So far, we have not used the property of λ guaranteed by [Lemma 5.1](#). This will now be used to show that the probability of the steps Initialize Core through Resample Neighborhoods Step 2 outputting an independent set of size k is $\Omega_{\Delta,\delta}(1/\sqrt{\log(1/\epsilon)})$. To see this, note that by taking the expectation over J and \vec{W} on both sides in the display equation above and using (P1) and [Theorem 3.1](#), we have that

$$\begin{aligned} \mathbb{E}_{J,\vec{W}} \left[\binom{|S^*(\vec{W})|}{k'} p 2^{-|S^*(\vec{W})|} \mathbb{1}_{|S^*(\vec{W})| \geq c'_{\Delta,\delta}n} \right] &= (1 \pm \epsilon/25) \mathbb{P}_\lambda[|I| = k \mid |I| \equiv k \pmod p] \\ &= \Omega_{\delta,\Delta}(p/\sqrt{n}) \\ &= \Omega_{\delta,\Delta}(1/\sqrt{\log(1/\epsilon)}). \end{aligned}$$

If $|S^*(\vec{W})| \geq c'_{\Delta,\delta}n$, then the quantity inside the expectation is bounded by $O_{\delta,\Delta}(1/\sqrt{\log(1/\epsilon)})$. Hence, by the reverse Markov inequality, with probability $\Omega_{\delta,\Delta}(1)$ (over the choice of J, \vec{W}), the quantity inside the expectation is $\Omega_{\delta,\Delta}(1/\sqrt{\log(1/\epsilon)})$.

To summarize, we have shown the following: a single run of Initialize Core through Resample Neighborhoods Step 2 produces an output with probability $\Omega_{\delta,\Delta}(1/\sqrt{\log(1/\epsilon)})$ and the distribution of this output is within $\epsilon/5$ in TV-distance from the uniform distribution on $\mathcal{I}_k(G)$. Therefore, by the Chernoff bound, repeating this procedure independently $C_{\Delta,\delta} \log(1/\epsilon)^{3/2}$ times for $C_{\Delta,\delta}$ produces an output from a distribution on $\mathcal{I}_k(G)$ which is within ϵ in TV-distance of the uniform distribution on $\mathcal{I}_k(G)$. \square

5.2. A faster FPRAS. The FPRAS for $i_k(G)$ and $m_k(G)$ is substantially simpler than the sampling algorithm above. As before, we will present the proof only for $i_k(G)$ with the proof for $m_k(G)$ being similar.

Proof of Theorem 1.6. We have the following two cases:

Case I: $1 \leq k \leq c_\Delta \sqrt{n}$, where $c_\Delta > 0$ is a sufficiently small constant which can be determined *a priori*. Let

$$p_k := \mathbb{P}[J \in \mathcal{I}_k],$$

where J is a uniformly random subset of V of size exactly k . By the union bound, it follows that

$$p_k = 1 - O\left(n\Delta \cdot \frac{k^2}{n^2}\right) \geq \frac{1}{2},$$

provided that c_Δ is sufficiently small. Since

$$i_k(G) = \binom{n}{k} p_k,$$

it suffices to obtain an ϵ -relative approximation of p_k . Let S_1, \dots, S_ℓ denote independent samples from the uniform distribution on size k subsets of V . Then, by the Chernoff bound,

$$\frac{\mathbb{1}[S_1 \in \mathcal{I}_k(G)] + \dots + \mathbb{1}[S_\ell \in \mathcal{I}_k(G)]}{\ell} = (1 \pm \epsilon)p_k$$

with probability at least $3/4$, provided that $\ell > C/\epsilon^2$ for a sufficiently large constant C .

For the running time, note that sampling a uniformly random subset of size k takes time $O(k \log n)$, checking whether it is an independent set takes time $O_\Delta(k)$, and computing the binomial coefficient $\binom{n}{k}$ takes time $O(k \log n \log \log n)$, so that the total running time is

$$O(k \log n \log \log n) + O_\Delta(k\epsilon^{-2} \log n).$$

Case II: $c_{\Delta,\delta} \sqrt{n} \leq k \leq (1 - \delta)\alpha_c(\Delta)n$. In this case, we first use [Lemma 4.1](#) to find a suitable λ . By [Theorem 3.1](#),

$$p_k := \frac{i_k(G)\lambda^k}{Z_G(\lambda)} = \mathbb{P}_\lambda[|I| = k] = \Omega_{\Delta,\delta}\left(\frac{1}{\sqrt{n\lambda}}\right).$$

Let I_1, \dots, I_ℓ denote independent samples obtained by running the Glauber dynamics for the hard-core model at fugacity λ for $O_{\Delta,\delta}(n \log n)$ steps. Then, by the Chernoff bound,

$$\frac{\mathbb{1}[I_1 \in \mathcal{I}_k(G)] + \dots + \mathbb{1}[I_\ell \in \mathcal{I}_k(G)]}{\ell} = (1 \pm \epsilon/4)p_k$$

with probability at least $3/4$, provided that $\ell > C_{\Delta,\delta}\epsilon^{-2}\sqrt{n\lambda}$, for a sufficiently large constant $C_{\Delta,\delta}$.

For the running time, sampling each I_i takes time $O_{\Delta,\delta}(n(\log n) \log(n/\epsilon))$, finding its size takes time $O_\Delta(n)$, computing λ^k takes time $O(k \log n \log \log n)$ and approximating $Z_G(\lambda)$ to within an $\epsilon/2$ -relative approximation takes time T , which gives the desired conclusion. \square

6. CLUSTER EXPANSION

In this section, we treat the case of small activities λ using the cluster expansion, a classical tool from statistical physics. The cluster expansion (or Mayer series [\[32\]](#)) is a formal infinite series for $\log Z_G(\lambda)$. For an introduction to the cluster expansion, see [\[18, Chapter 5\]](#).

We introduce the cluster expansion in the special case of the hard-core model on a graph G . A *cluster* Γ is an ordered tuple of vertices from G . The size of Γ , denoted $|\Gamma|$, is the length of the tuple. The incompatibility graph of $\Gamma = (v_1, \dots, v_k)$, $H(\Gamma)$, is the graph with vertex set $\{v_1, \dots, v_k\}$ and an edge between v_i, v_j , $i \neq j$, if $v_i \in N(v_j) \cup \{v_j\}$ in G . The Ursell function of a graph H is the function

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

The cluster expansion is the formal infinite power series

$$\log Z_G(\lambda) = \sum_{\Gamma} \phi(H(\Gamma)) \lambda^{|\Gamma|},$$

where the sum is over all clusters of vertices from G . In fact, in this setting the cluster expansion is simply the Taylor series for $\log Z_G(\lambda)$ around 0, with terms organized by clusters instead of grouping all terms of order k together.

To use the cluster expansion as an enumeration tool, it is essential to bound its rate of convergence. We will use the convergence criteria of Kotecký and Preiss [29] (though the zero-freeness result of Shearer [38] along with the lemma of Barvinok [4] on truncating Taylor series would also work). This lemma bounds the additive error of truncating the cluster expansion after a given number of terms.

Lemma 6.1. *Let G be a graph of maximum degree at most Δ on n vertices and suppose $0 < \lambda < \frac{e}{\Delta+1}$. Then*

$$\sum_{\Gamma: |\Gamma| \geq k} \left| \phi(H(\Gamma)) \lambda^{|\Gamma|} \right| \leq n (\lambda e (\Delta + 1))^k. \quad (6.1)$$

Proof. This is a consequence of the main result of [29]. We can express the hard-core model as a polymer model in the setting of [29] by defining each vertex to be a polymer with weight λ . Taking $a(v) = 1$ for all $v \in V(G)$ and $\exp(d(v)) = \frac{1}{\lambda e (\Delta + 1)}$, we have for all $v \in V(G)$,

$$\sum_{u \in N(v) \cup \{v\}} \lambda e^{a(u)+d(u)} \leq (\Delta + 1) \lambda e \frac{1}{\lambda e (\Delta + 1)} = 1.$$

Then by the main theorem in [29], for all $v \in V(G)$,

$$\sum_{\Gamma \ni v} \left| \phi(H(\Gamma)) \lambda^{|\Gamma|} \left(\frac{1}{\lambda e (\Delta + 1)} \right)^{|\Gamma|} \right| \leq 1.$$

Restricting the sum to clusters of size at least k and summing over all $v \in V(G)$ gives (6.1). \square

The cluster expansion is a very convenient tool for studying the cumulants of the random variable $Y = |I|$ (see e.g. [10, 15, 24]). In particular, when the cluster expansion converges, we have the formula

$$\kappa_k(Y) = \sum_{\Gamma} |\Gamma|^k \phi(H(\Gamma)) \lambda^{|\Gamma|}. \quad (6.2)$$

Lemma 6.2. *Fix $\Delta \geq 2, \delta > 0$ and suppose $\lambda \leq \frac{1-\delta}{e(\Delta+1)}$. Then for all fixed $k \geq 1$, and all graphs G of maximum degree Δ on n vertices,*

$$\sum_{\Gamma} |\Gamma|^k \lambda^k \phi(H(\Gamma)) = n\lambda + O_{k,\delta}(n\lambda^2 \Delta^2).$$

Proof. Since the contribution to the left-hand side from clusters of size 1 is $n\lambda$, it suffices to show that

$$\sum_{\Gamma: |\Gamma| \geq 2} |\Gamma|^k \lambda^{|\Gamma|} |\phi(H(\Gamma))| = O_{k,\delta}(n\lambda^2 \Delta^2).$$

Applying Lemma 6.1, we have

$$\sum_{\Gamma: |\Gamma| \geq 2} |\Gamma|^k \lambda^{|\Gamma|} |\phi(H(\Gamma))| \leq n \sum_{t \geq 2} t^k (\lambda e (\Delta + 1))^t \leq C_{k,\delta} n \lambda^2 (\Delta + 1)^2. \quad \square$$

As in [24, Corollary 23], we can immediately deduce bounds on the mean and variance and a central limit theorem from Lemma 6.2.

Corollary 6.3. *Let G be a graph of maximum degree Δ . If $\lambda \leq \frac{1-\delta}{e(\Delta+1)}$, then the following hold:*

- (1) $\mathbb{E}_\lambda Y = n\lambda + O_\delta(n\lambda^2\Delta^2)$.
- (2) $\text{Var}_\lambda Y = n\lambda + O_\delta(n\lambda^2\Delta^2)$.

If in addition $\lambda n \rightarrow \infty$ as $n \rightarrow \infty$, then

- (3) *The random variable Y satisfies a central limit theorem.*

Proof. We prove these statements via the cumulants of Y .

$$\mathbb{E}_\lambda Y = \kappa_1(Y) = \sum_{\Gamma} |\Gamma| \lambda^{|\Gamma|} \phi(H(\Gamma)) = n\lambda + O_\delta(n\lambda^2\Delta^2).$$

$$\text{Var}_\lambda Y = \kappa_2(Y) = \sum_{\Gamma} |\Gamma|^2 \lambda^{|\Gamma|} \phi(H(\Gamma)) = n\lambda + O_\delta(n\lambda^2\Delta^2).$$

Now let $X = (Y - \mathbb{E}_\lambda Y) / \sqrt{\text{Var}_\lambda Y}$. By definition $\kappa_1(X) = 0$ and $\kappa_2(X) = 1$. All the cumulants of a standard Gaussian random variable are 0 except for the second which is 1, and so to prove a central limit theorem it suffices to show that for any fixed $k \geq 3$, $\kappa_k(X) \rightarrow 0$ as $n \rightarrow \infty$. We have

$$\begin{aligned} |\kappa_k(X)| &= \left| \sum_{\Gamma} \frac{|\Gamma|^k}{\text{Var}(Y)^{k/2}} \lambda^{|\Gamma|} \phi(H(\Gamma)) \right| \\ &\leq \frac{1}{\text{Var}(Y)^{3/2}} \sum_{\Gamma} |\Gamma|^k \lambda^{|\Gamma|} |\phi(H(\Gamma))| \\ &= (1 + o(1))(n\lambda)^{-1/2} = o(1). \end{aligned} \quad \square$$

Note that if $\lambda n \rightarrow \rho > 0$ as $n \rightarrow \infty$, then [Lemma 6.2](#) shows that $\kappa_k(Y) \rightarrow \rho$ for each fixed k , which implies that Y converges in distribution to a Poisson random variable of mean ρ .

Theorem 6.4. *Fix $\Delta \geq 3$. If G_n is a sequence of graphs of maximum degree Δ on n vertices and $n^{-1} \ll \lambda \leq \frac{1}{100\Delta^2}$, then the random variable Y satisfies a local central limit theorem as $n \rightarrow \infty$.*

Proof. We follow the proof strategy of [\[25, Theorem 19\]](#) that proves a local central limit theorem in the setting of polymer models satisfying the equivalent of [Lemma 6.2](#).

Let $X = (Y - \mathbb{E}_\lambda Y) / \sqrt{\text{Var}_\lambda Y}$, and let $\phi_X(t) = \mathbb{E}_\lambda e^{itX}$ be the characteristic function of X (and ϕ_Y the characteristic function of Y).

First we prove that there exists $c > 0$ so that for all $t \in [-\pi, \pi]$, $|\phi_X(t)| \leq e^{-ct^2}$. Using the cluster expansion we write

$$\log \phi_Y(t) = \sum_{\Gamma} \left(e^{it|\Gamma|} - 1 \right) \phi(H(\Gamma)) \lambda^{|\Gamma|},$$

and so

$$\begin{aligned} \text{Re} \log \phi_Y(t) &= \sum_{\Gamma} (\cos(t|\Gamma|) - 1) \phi(H(\Gamma)) \lambda^{|\Gamma|} \\ &= n\lambda(\cos(t) - 1) + \sum_{\Gamma: |\Gamma| \geq 2} (\cos(t|\Gamma|) - 1) \phi(H(\Gamma)) \lambda^{|\Gamma|} \\ &\leq -\frac{t^2 n \lambda}{5} + \sum_{\Gamma: |\Gamma| \geq 2} t^2 |\Gamma|^2 \phi(H(\Gamma)) \lambda^{|\Gamma|} \\ &\leq -\frac{t^2 n \lambda}{5} + -\frac{t^2 n \lambda}{10} \leq -\frac{t^2 n \lambda}{10}, \end{aligned}$$

where we have used the bounds $\cos(t) - 1 \leq -t^2/5$ and $1 - \cos(tx) \leq (tx)^2$. Exponentiating then gives $|\phi_Y(t)| \leq e^{-t^2 n \lambda / 10}$, which, along with [Corollary 6.3](#), implies that $|\phi_X(t)| \leq e^{-ct^2}$ for some $c > 0$.

To finish the proof we will apply [Lemma 2.5](#) with $\alpha = \mathbb{E}Y$ and $\beta = \sqrt{\text{Var}(Y)}$, which says

$$\begin{aligned} \sup_{x \in \mathcal{L}} |\beta \mathcal{N}(x) - \mathbb{P}[X = x]| &\leq \beta \int_{-\pi/\beta}^{\pi/\beta} |\phi_X(t) - \phi_Z(t)| dt + e^{-\pi^2/(2\beta^2)} \\ &\leq o(\beta) + \beta \int_{-\infty}^{\infty} |\phi_X(t) - \phi_Z(t)| dt, \end{aligned}$$

and so to prove the LCLT it suffices to show that

$$\int_{-\infty}^{\infty} |\phi_X(t) - \phi_Z(t)| dt = o(1).$$

By the central limit theorem of [Corollary 6.3](#) we have that $\phi_X(t) \rightarrow \phi_Z(t)$ as $n \rightarrow \infty$. Moreover $|\phi_X(t) - \phi_Z(t)|$ is an integrable function since it is bounded by $e^{-ct^2} + e^{-t^2/2}$ from the bound above. Applying dominated convergence completes the proof. \square

7. DETERMINISTIC APPROXIMATION OF CUMULANTS IN LINEAR TIME

In this section we prove [Theorem 1.7](#). We prove the theorem first in the regime of cluster expansion convergence, then extend to more general zero-free regions. We prove the theorem for general graphs, noting that the proof for claw-free graphs is similar.

Lemma 7.1. *For all graphs G of maximum degree Δ , all $0 < \lambda \leq \frac{1-\delta}{e(\Delta+1)}$, and all fixed $k \geq 1$ there is a deterministic algorithm to give an $\epsilon \lambda n$ additive approximation to $\kappa_k(Y)$. The algorithm runs in time $O_{\Delta, \delta, k}(n \cdot (1/\epsilon)^{O_{\Delta, \delta}(1)})$.*

Proof. The algorithm will be to compute a truncation of the cluster expansion for $\kappa_k(Y)$. Recall that in the regime of cluster expansion convergence, we have

$$\kappa_k(Y) = \sum_{\Gamma} |\Gamma|^k \phi(H(\Gamma)) \lambda^{|\Gamma|}.$$

Now let $T_t^{(k)} = \sum_{|\Gamma| < t} |\Gamma|^k \phi(H(\Gamma)) \lambda^{|\Gamma|}$ be the truncation keeping only clusters of size less than t . By [\(6.2\)](#) and [\(6.1\)](#) we have

$$\left| \kappa_k(Y) - T_t^{(k)} \right| \leq n \sum_{j \geq t} j^k (\lambda e(\Delta + 1))^j.$$

By taking $t = \Omega\left(\frac{\log(\epsilon)}{\log(\lambda e(\Delta + 1))} + \frac{k^2}{\delta^2}\right)$, we have that $\left| \kappa_k(Y) - T_t^{(k)} \right| \leq n \lambda \epsilon$. This truncated cluster expansion can be computed in time $n \cdot \exp(O(t \log \Delta)) = O_{\Delta, \delta, k}(n \cdot (1/\epsilon)^{O_{\Delta, \delta}(1)})$ using the algorithm of [\[23, 34\]](#). \square

Next we give a general algorithm when λ is not necessarily in the regime of cluster expansion convergence.

Proof of Theorem 1.7. Since [Lemma 7.1](#) covers the case of $\lambda \leq 1/(2e(\Delta + 1))$, we will assume here that $\lambda \geq 1/(2e(\Delta + 1))$. The general algorithm is an adaptation of the approximate counting algorithm of Barvinok and Patel and Regts via an expression for $\kappa_k(Y)$ in terms of derivatives of $\log Z_G(\lambda)$. The k th cumulant of Y can be written as a linear combination of the first k derivatives of $\log Z_G(\lambda)$ in λ , where (for λ in the considered range) the size of the coefficients in the linear combination can be bounded in terms of only k and Δ . Hence, it suffices to give an ϵn additive approximation to $\frac{d^k}{d\lambda^k} \log Z_G(\lambda)$ in the stated running time.

Let $\delta > 0$ be small enough so that $\lambda \in \mathcal{R}_{\delta, \Delta}$. Following Barvinok [4] and Peters and Regts [35], there is a polynomial f of degree $D = D(\delta, \Delta)$ that maps the unit circle in the complex plane into the region $\mathcal{R}_{\delta, \Delta}$, sending 0 to 0 and 1 to λ . Let

$$\hat{Z}(y) = Z_G(f(y))$$

so that $\hat{Z}(1) = Z_G(\lambda)$. In particular, \hat{Z} is a polynomial in y of degree $N \leq Dn$. Let r_1, \dots, r_N denote the inverses of the roots of \hat{Z} so that $\hat{Z}(y) = \prod_{i=1}^N (1 - r_i y)$. By Theorem 2.2, there is an $\eta = \eta(\delta, \Delta) \in (0, 1)$ so that $|r_i| \leq \eta$ for $i = 1, \dots, N$.

The first k derivatives of $\log Z_G(\lambda)$ with respect to λ can be written in terms of the first k derivatives of $\log \hat{Z}(y)$ with respect to y and those of f with respect to y . Using the chain rule we obtain

$$\frac{d^k \log Z_G(\lambda)}{d\lambda^k} = \sum_{j=1}^k b_j \frac{d^j \log \hat{Z}(y)}{dy^j}$$

where the coefficients b_j depend only on the first j derivatives of the bounded-degree polynomial f and thus are bounded. For instance, we have

$$\frac{d \log Z_G(\lambda)}{d\lambda} = \frac{\frac{d \log \hat{Z}(y)}{dy}}{\frac{df(y)}{dy}}$$

and

$$\frac{d^2 \log Z_G(\lambda)}{d\lambda^2} = \frac{\frac{d^2 \log \hat{Z}(y)}{dy^2}}{\frac{df(y)}{dy}} - \frac{\frac{d \log \hat{Z}(y)}{dy} \cdot \frac{d^2 f(y)}{dy^2}}{\left(\frac{df(y)}{dy}\right)^2}$$

In particular, it now suffices to compute an ϵn additive approximation to $\frac{d^j \log \hat{Z}(y)}{dy^j}$ for $j = 1, \dots, k$. We can write

$$\begin{aligned} \frac{d^k \log \hat{Z}(y)}{dy^k} &= \frac{d^k}{dy^k} \sum_{i=1}^N \log(1 - r_i y) = -(k-1)! \sum_{i=1}^N \frac{r_i^k}{(1 - r_i y)^k} \\ &= -(k-1)! \sum_{i=1}^N r_i^k \sum_{s=0}^{\infty} \binom{k-1+s}{k-1} (r_i y)^s. \end{aligned}$$

Now setting

$$T_t^{(k)} = -(k-1)! \sum_{i=1}^N r_i^k \sum_{s=0}^t \binom{k-1+s}{k-1} r_i^s$$

we have

$$\left| T_t^{(k)} - \frac{d^k \log \hat{Z}(y)}{dy^k} \right| \leq (k-1)! N \sum_{s=t+1}^{\infty} (s+k)^k \eta^s = O(n\eta^t),$$

and so for $t = \Omega_{\Delta, \delta}(\log(1/\epsilon) + k^2)$, the truncation error can be made at most ϵn as desired. Moreover, using the algorithm of Patel–Regts [34], $T_t^{(k)}$ can be computed in time $ne^{O_{\Delta, \delta}(t)} = O_{\Delta, \delta, k}(n(1/\epsilon)^{O_{\Delta, \delta}(1)})$ for this choice of t .

The FPTAS for the mean and variance follow from the additive approximations and Lemma 3.2. \square

REFERENCES

- [1] Vedat Levi Alev and Lap Chi Lau, *Improved analysis of higher order random walks and applications*, Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC), 2020, pp. 1198–1211. [3](#)
- [2] Yeganeh Alimohammadi, Nima Anari, Kirankumar Shiragur, and Thuy-Duong Vuong, *Fractionally log-concave and sector-stable polynomials: counting planar matchings and more*, Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing (STOC), 2021, pp. 433–446. [1](#)
- [3] Nima Anari, Kuikui Liu, and Shayan Oveis Gharan, *Spectral independence in high-dimensional expanders and applications to the hardcore model*, 2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS), IEEE, 2020, pp. 1319–1330. [1](#)
- [4] Alexander Barvinok, *Combinatorics and complexity of partition functions*, vol. 9, Springer, 2016. [4](#), [5](#), [6](#), [7](#), [21](#), [24](#)
- [5] Mohsen Bayati, David Gamarnik, Dimitriy Katz, Chandra Nair, and Prasad Tetali, *Simple deterministic approximation algorithms for counting matchings*, Proceedings of the thirty-ninth annual ACM Symposium on Theory of Computing (STOC), 2007, pp. 122–127. [1](#), [6](#)
- [6] Ross Berkowitz, *A quantitative local limit theorem for triangles in random graphs*, arXiv preprint arXiv:1610.01281 (2016). [8](#)
- [7] Ivona Bezáková, Andreas Galanis, Leslie Ann Goldberg, and Daniel Štefankovič, *Inapproximability of the independent set polynomial in the complex plane*, SIAM Journal on Computing **49** (2019), STOC18–395–STOC18–448. [5](#)
- [8] Russ Bubley and Martin Dyer, *Path coupling: A technique for proving rapid mixing in Markov chains*, Proceedings 38th Annual Symposium on Foundations of Computer Science, IEEE, 1997, pp. 223–231. [3](#), [17](#)
- [9] Pjotr Buys, Andreas Galanis, Viresh Patel, and Guus Regts, *Lee-Yang zeros and the complexity of the ferromagnetic Ising model on bounded-degree graphs*, Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM, 2021, pp. 1508–1519. [5](#)
- [10] Sarah Cannon and Will Perkins, *Counting independent sets in unbalanced bipartite graphs*, Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM, 2020, pp. 1456–1466. [21](#)
- [11] Zongchen Chen, Kuikui Liu, and Eric Vigoda, *Optimal mixing of Glauber dynamics: Entropy factorization via high-dimensional expansion*, Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing (STOC), 2021, pp. 1537–1550. [3](#), [5](#), [6](#), [8](#)
- [12] Maria Chudnovsky and Paul Seymour, *The roots of the independence polynomial of a clawfree graph*, Journal of Combinatorial Theory, Series B **97** (2007), 350–357. [6](#), [7](#)
- [13] Ewan Davies, Matthew Jenssen, and Will Perkins, *A proof of the Upper Matching Conjecture for large graphs*, Journal of Combinatorial Theory, Series B (to appear). [2](#), [13](#)
- [14] Ewan Davies and Will Perkins, *Approximately counting independent sets of a given size in bounded-degree graphs*, 48th International Colloquium on Automata, Languages, and Programming (ICALP), vol. 198, 2021, pp. 62:1–62:18. [1](#), [2](#), [3](#), [4](#), [5](#), [9](#), [13](#)
- [15] RL Dobrushin, *Estimates of semi-invariants for the Ising model at low temperatures*, Translations of the American Mathematical Society-Series 2 **177** (1996), 59–82. [21](#)
- [16] RL Dobrushin and Brunello Tirozzi, *The central limit theorem and the problem of equivalence of ensembles*, Communications in Mathematical Physics **54** (1977), 173–192. [7](#), [10](#)
- [17] Martín Farach-Colton and Meng-Tsung Tsai, *On the complexity of computing prime tables*, International Symposium on Algorithms and Computation, Springer, 2015, pp. 677–688. [18](#)
- [18] Sacha Friedli and Yvan Velenik, *Statistical mechanics of lattice systems: a concrete mathematical introduction*, Cambridge University Press, 2017. [20](#)
- [19] Andreas Galanis, Daniel Štefankovič, and Eric Vigoda, *Inapproximability of the partition function for the antiferromagnetic Ising and hard-core models*, Combinatorics, Probability and Computing **25** (2016), 500–559. [2](#)
- [20] Chris D Godsil, *Matching behaviour is asymptotically normal*, Combinatorica **1** (1981), 369–376. [4](#)
- [21] Nicholas JA Harvey, Piyush Srivastava, and Jan Vondrák, *Computing the independence polynomial: from the tree threshold down to the roots*, Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM, 2018, pp. 1557–1576. [5](#)
- [22] Ole J Heilmann and Elliott H Lieb, *Theory of monomer-dimer systems*, Communications in Mathematical Physics **25** (1972), 190–232. [4](#), [6](#), [7](#)
- [23] Tyler Helmuth, Will Perkins, and Guus Regts, *Algorithmic Pirogov–Sinai theory*, Probability Theory and Related Fields **176** (2020), 851–895. [23](#)
- [24] Matthew Jenssen and Will Perkins, *Independent sets in the hypercube revisited*, Journal of the London Mathematical Society **102** (2020), 645–669. [21](#)

- [25] Matthew Jenssen, Will Perkins, and Aditya Potukuchi, *Independent sets of a given size and structure in the hypercube*, arXiv preprint arXiv:2106.09709 (2021). [22](#)
- [26] Mark Jerrum and Alistair Sinclair, *Approximating the permanent*, SIAM Journal on Computing **18** (1989), 1149–1178. [1](#), [3](#)
- [27] Mark Jerrum, Alistair Sinclair, and Eric Vigoda, *A polynomial-time approximation algorithm for the permanent of a matrix with nonnegative entries*, Journal of the ACM (JACM) **51** (2004), 671–697. [1](#)
- [28] Pieter Kasteleyn, *Graph theory and crystal physics*, Graph Theory and Theoretical Physics (1967), 43–110. [1](#)
- [29] Roman Kotecký and David Preiss, *Cluster expansion for abstract polymer models*, Communications in Mathematical Physics **103** (1986), 491–498. [21](#)
- [30] J. L. Lebowitz, B. Pittel, D. Ruelle, and E. R. Speer, *Central limit theorems, Lee-Yang zeros, and graph-counting polynomials*, J. Combin. Theory Ser. A **141** (2016), 147–183. [8](#)
- [31] Jingcheng Liu, Alistair Sinclair, and Piyush Srivastava, *The Ising partition function: Zeros and deterministic approximation*, Journal of Statistical Physics **174** (2019), 287–315. [5](#)
- [32] Joseph E Mayer and Elliott Montroll, *Molecular distribution*, The Journal of Chemical Physics **9** (1941), 2–16. [20](#)
- [33] Marcus Michelen and Julian Sahasrabudhe, *Central limit theorems and the geometry of polynomials*. [6](#), [8](#)
- [34] Viresh Patel and Guus Regts, *Deterministic polynomial-time approximation algorithms for partition functions and graph polynomials*, SIAM Journal on Computing **46** (2017), 1893–1919. [4](#), [5](#), [6](#), [7](#), [23](#), [24](#)
- [35] Han Peters and Guus Regts, *On a conjecture of Sokal concerning roots of the independence polynomial*, Michigan Math. J. **68** (2019), 33–55. [4](#), [5](#), [6](#), [7](#), [24](#)
- [36] V. V. Petrov, *Sums of independent random variables*, Ergebnisse der Mathematik und ihrer Grenzgebiete, Band 82, Springer-Verlag, New York-Heidelberg, 1975, Translated from the Russian by A. A. Brown. [12](#)
- [37] Ashwin Sah and Mehtaab Sawhney, *Local limit theorems for subgraph counts*. [8](#)
- [38] James B. Shearer, *On a problem of Spencer*, Combinatorica **5** (1985), 241–245. [6](#), [7](#), [21](#)
- [39] Allan Sly, *Computational transition at the uniqueness threshold*, 2010 IEEE 51st Annual Symposium on Foundations of Computer Science (FOCS), IEEE, 2010, pp. 287–296. [2](#)
- [40] Allan Sly and Nike Sun, *Counting in two-spin models on d -regular graphs*, Annals of Probability **42** (2014), 2383–2416. [2](#)
- [41] Daniel Štefankovič, Santosh Vempala, and Eric Vigoda, *Adaptive simulated annealing: A near-optimal connection between sampling and counting*, Journal of the ACM (JACM) **56** (2009), 1–36. [5](#)
- [42] Daniel Štefankovič, Eric Vigoda, and John Wilmes, *On counting perfect matchings in general graphs*, Latin American Symposium on Theoretical Informatics, Springer, 2018, pp. 873–885. [1](#)
- [43] Leslie G Valiant, *The complexity of computing the permanent*, Theoretical Computer Science **8** (1979), 189–201. [1](#)
- [44] Dror Weitz, *Counting independent sets up to the tree threshold*, Proceedings of the thirty-eighth annual ACM Symposium on Theory of Computing (STOC), 2006, pp. 140–149. [1](#), [2](#), [6](#)

DEPARTMENT OF STATISTICS, STANFORD UNIVERSITY

Email address: visheshj@stanford.edu

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

Email address: math@willperkins.org

DEPARTMENT OF MATHEMATICS, MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Email address: {asah,msawhney}@mit.edu