

Sampling, counting, and large deviations for triangle-free graphs near the critical density

Matthew Jenssen

Department of Mathematics
King's College London
London, UK
matthew.jenssen@kcl.ac.uk

Will Perkins

School of Computer Science
Georgia Tech
Atlanta, USA
wperkins3@gatech.edu

Aditya Potukuchi

Department of EECS
York University
Toronto, Canada
apotu@yorku.ca

Michael Simkin

Department of Mathematics
MIT
Cambridge, USA
msimkin@mit.edu

Abstract—We study the following combinatorial counting and sampling problems: can we sample from the Erdős-Rényi random graph $G(n, p)$ conditioned on triangle-freeness? Can we approximate (either algorithmically or with a formula) the probability that $G(n, p)$ is triangle-free? These are prototypical instances of forbidden substructure problems ubiquitous in combinatorics. The algorithmic questions are instances of approximate sampling and counting for a hypergraph hard-core model.

Estimating the probability that $G(n, p)$ has no triangles is a fundamental question in probabilistic combinatorics and one that has led to the development of many important tools in the field. Through the work of several authors, the asymptotics of the logarithm of this probability are known if $p = o(n^{-1/2})$ or if $p = \omega(n^{-1/2})$. The regime $p = \Theta(n^{-1/2})$ is more mysterious, as this range witnesses a dramatic change in the typical structural properties of $G(n, p)$ conditioned on triangle-freeness. As we show, this change in structure has a profound impact on the performance of sampling algorithms.

We give two different efficient sampling algorithms for this problem (and complementary approximate counting algorithms), one that is efficient when $p < c/\sqrt{n}$ and one that is efficient when $p > C/\sqrt{n}$ for constants $c, C > 0$. The latter algorithm involves a new approach for dealing with large defects in the setting of sampling from low-temperature spin models.

Our algorithmic results can be used to give an asymptotic formula for the logarithm of the probability $G(n, p)$ is triangle-free when $p < c/\sqrt{n}$. This algorithmic approach to large deviation problems in random graphs is very different than the known approaches in the subcritical regime $p = o(n^{-1/2})$ (based on the Poisson paradigm) and in the supercritical regime $p = \omega(n^{-1/2})$ (based on regularity lemmas or hypergraph containers); in fact, to the best of our knowledge, no asymptotic formula for the log probability in the regime $p = \Theta(n^{-1/2})$ was even conjectured previously.

Index Terms—Approximate counting and sampling, triangle-free graphs, random graphs, large deviations

I. INTRODUCTION

Let $\mathcal{T} = \mathcal{T}(n)$ denote the set of (labeled) triangle-free graphs on n vertices. Understanding properties of this set is a central topic in extremal and probabilistic combinatorics. Mantel's Theorem [1], one of the earliest results in extremal graph theory, solves an optimization problem: which $G \in \mathcal{T}(n)$ has the most edges? Erdős, Kleitman, and Rothschild [2] gave an asymptotic formula for $|\mathcal{T}(n)|$ and proved an important structural property of the uniform distribution on $\mathcal{T}(n)$: almost all triangle-free graphs are bipartite. More generally, one can

study $\mu_p(\mathcal{T})$, the probability that the Erdős-Rényi random graph $G(n, p)$ is triangle-free, a question considered already in the papers of Erdős and Rényi initiating the study of random graphs [3].

Understanding such non-existence probabilities and more general large deviation probabilities for subgraphs in random graphs often involves deep mathematical tools from combinatorics and probability theory, including regularity lemmas, nonlinear large deviations, container methods, and entropic methods (e.g. [4]–[11]). These techniques often connect large deviation probabilities to structural information about graphs drawn from the appropriate conditional distribution.

Here we will be interested in the connection between structural properties of the conditional distribution, algorithms to sample from the conditional distribution, and approximations of the non-existence probability (either via algorithms or formulas).

In our case, the distribution of interest is that of $G(n, p)$ conditioned on triangle-freeness. We denote this distribution by $\mu_{\mathcal{T}, p}$. The result of Erdős, Kleitman, and Rothschild not only gives an asymptotic formula for $\mu_{1/2}(\mathcal{T})$, but also gives an efficient algorithm to sample approximately from $\mu_{\mathcal{T}, 1/2}$: sample a nearly balanced partition of a set of n vertices and include each crossing edge independently with probability $1/2$. This gives a random bipartite graph on n vertices, whose distribution is very close to $\mu_{\mathcal{T}, 1/2}$. Osthus, Prömel, and Taraz [12] later showed that an analogous result holds for much smaller p : for any fixed $\varepsilon > 0$, if $p \geq (1+\varepsilon)\sqrt{3 \log n}/\sqrt{n}$ then $G \sim \mu_{\mathcal{T}, p}$ is bipartite whp. In the same way as above, this immediately yields an efficient approximate sampling algorithm.

What happens for smaller p ? T. Łuczak [13] had previously showed that if $p \geq C/\sqrt{n}$ for large C , then whp G has a cut containing most of the edges; that is, it is *approximately* bipartite. On the other hand, when $p = o(n^{-1/2})$, no such large cut or global structure exists, and instead behavior is captured by the ‘Poisson paradigm’ and Janson’s Inequality [14].

Here we will be interested in the regime $p = \Theta(n^{-1/2})$, in which this global structure emerges. Beyond Łuczak’s result and consequences of Janson’s Inequality, very little is known in this regime about either the properties of $\mu_{\mathcal{T}, p}$, the value of $\mu_p(\mathcal{T})$, or the existence of efficient sampling algorithms. In this paper we will show that these three questions are closely

related.

In different settings, very similar questions about approximation of probabilities and understanding conditional probability measures are studied in computer science in the area of approximate counting and sampling. Here a canonical object is the partition function of the hard-core model or the generating function for independent sets of a graph or hypergraph G :

$$Z_G(\lambda) = \sum_{I \in \mathcal{I}(G)} \lambda^{|I|}, \quad (1)$$

where $\mathcal{I}(G)$ is the set of independent sets of G . The related sampling question is about sampling from the hard-core measure, the probability measure on $\mathcal{I}(G)$ defined by

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

The complexity of approximate sampling and counting for the hard-core model on graphs of maximum degree Δ is now very well understood: when λ is below a threshold $\lambda_c(\Delta) \approx \frac{e}{\Delta}$, there are efficient sampling and counting algorithms [15]–[19], and when λ is above this threshold no such algorithms exist unless $\text{NP}=\text{RP}$ [20]–[22]. The critical value λ_c marks a phase transition of the hard-core model on the infinite Δ -regular tree. The phase transition behavior is also reflected in finite graphs: when $\lambda < \lambda_c$, the hard-core model on any graph of maximum degree Δ exhibits several ‘nice’ behaviors including decay of correlations and rapid convergence of local dynamics (Markov chains), while there exist graphs for which these properties fail for $\lambda > \lambda_c$ (in particular, for the random Δ -regular bipartite graph [23]).

The probability that $G(n, p)$ is triangle-free and its conditional distribution given this event can be written in terms of a hard-core partition function and measure. Given a graph G on vertex set $[n]$, we will often identify G with its edge set and let $|G|$ denote the number of edges in G . Let μ_p denote the distribution of the Erdős-Rényi random graph $G(n, p)$, that is, $\mu_p(G) = p^{|G|}(1-p)^{\binom{n}{2}-|G|}$. Let

$$Z(\lambda) = \sum_{G \in \mathcal{T}} \lambda^{|G|}.$$

With $\lambda = \frac{p}{1-p}$, we have the identities

$$\mu_p(\mathcal{T}) = \frac{Z(\lambda)}{(1+\lambda)^{\binom{n}{2}}}, \quad (2)$$

and

$$\mu_{\mathcal{T},p}(G) = \frac{\lambda^{|G|}}{Z(\lambda)}, \quad G \in \mathcal{T}.$$

From this perspective, $\mu_{\mathcal{T},p}$ may be viewed as a hypergraph hard-core measure. Let $H = (V, E)$ be the hypergraph with $V = \binom{[n]}{2}$ where $\{x, y, z\} \in E$ if and only if x, y, z form a triangle (when viewed as edges of a graph on vertex set $[n]$). Then $\mu_{\mathcal{T},p} = \mu_{H,\lambda}$ where $\lambda = p/(1-p)$.

The algorithmic and probabilistic behavior of hypergraph hard-core models on bounded-degree, uniform hypergraphs is more complex and less well understood than that of graphs.

While for graphs $\lambda_c(\Delta)$ marks a computational threshold, a phase transition on a particular infinite graph, a threshold for fast mixing of local dynamics, and a threshold for a strong form of correlation decay (strong spatial mixing), in hypergraphs these thresholds need not coincide (see, e.g., the results and discussion in [24]–[27]).

Our main results combine these points of view and give some understanding of triangle-free graphs in this critical regime $p = \Theta(n^{-1/2})$. We will give efficient approximate counting and sampling algorithms when $p \leq cn^{-1/2}$ and $p \geq Cn^{-1/2}$ for constants $c, C > 0$. We will then use the algorithmic results to prove an asymptotic formula for $\log \mu_p(\mathcal{T})$ when $p \leq cn^{-1/2}$.

To state our algorithmic results we need a couple of definitions. The *Glauber dynamics* is a Markov chain defined by a probability distribution μ on vectors (in our case indicator functions of edges of a graph) that proceeds by choosing a random coordinate and updating this coordinate according to the conditional distribution of μ given all other coordinates. The mixing time of a Markov chain is the number of steps required to guarantee convergence to the stationary distribution μ (to within $1/4$ total variation distance) from a worst-case starting point. See Section II for full definitions. We say that $\hat{Z} > 0$ is an ε -relative approximation to $Z > 0$ if $e^{-\varepsilon} \hat{Z} \leq Z \leq e^{\varepsilon} \hat{Z}$.

Theorem I.1. *Let $c < 1/\sqrt{2}$ be fixed and suppose $p \leq cn^{-1/2}$. The Glauber dynamics for sampling from $\mu_{\mathcal{T},p}$ has mixing time $O(n^2 \log n)$.*

In particular, this gives randomized algorithms running in time polynomial in n and $1/\varepsilon$ that

- 1) *Output M so that with probability at least $2/3$, M is an ε -relative approximation to $\mu_p(\mathcal{T})$.*
- 2) *Output $G \in \mathcal{T}$ with distribution $\hat{\mu}$ so that $\|\hat{\mu} - \mu_{\mathcal{T},p}\|_{TV} < \varepsilon$.*

In the terminology of approximate counting and sampling, these are a Fully Polynomial-time Randomized Approximation Scheme (FPRAS) and an efficient sampling scheme.

Up to the constant $1/\sqrt{2}$, Theorem I.1 is sharp. When $p \geq Cn^{-1/2}$ for C a large constant, the Glauber dynamics mix slowly due to the emergence of global structure (large max-cuts) in triangle-free graphs.

Theorem I.2. *There exists $C > 0$ so that for $p \geq Cn^{-1/2}$, the Glauber dynamics for sampling from $\mu_{\mathcal{T},p}$ has mixing time $\Omega(e^{\sqrt{n}})$.*

On the other hand, even in this higher density regime we can still sample from $\mu_{\mathcal{T},p}$ and approximate $\mu_p(\mathcal{T})$ efficiently.

Theorem I.3. *There exist $C, C' > 0$, so that for $p \geq Cn^{-1/2}$ and $\varepsilon \geq C'e^{-\sqrt{n}}$ there are randomized algorithms running in time polynomial in n and $1/\varepsilon$ that*

- 1) *Output M so that with probability at least $2/3$, M is an ε -relative approximation to $\mu_p(\mathcal{T})$.*
- 2) *Output $G \in \mathcal{T}$ with distribution $\hat{\mu}$ so that $\|\hat{\mu} - \mu_{\mathcal{T},p}\|_{TV} \leq \varepsilon$.*

We next describe our results on the asymptotics of $\log \mu_p(\mathcal{T})$. Through the work of Erdős, Kleitman, and Rothschild [2], Janson, T. Łuczak, and Ruciński [14], Prömel and Steger [28], and T. Łuczak [13], the asymptotics of $\log \mu_p(\mathcal{T})$ are known when p is either much smaller or much larger than $n^{-1/2}$.

Theorem I.4. [[2], [13], [14], [28]] For $p = o(n^{-1/2})$,

$$\log \mu_p(\mathcal{T}) = -(1 + o(1)) \binom{n}{3} p^3. \quad (3)$$

For $p = \omega(n^{-1/2})$,

$$\log \mu_p(\mathcal{T}) = (1 + o(1)) \frac{n^2 \log(1 - p)}{4}. \quad (4)$$

It is worth noting that the methods used for small and large p in Theorem I.4 are very different. The case when $p = o(1/\sqrt{n})$ is proved using Janson's Inequality (in fact this powerful tool was first introduced in [14] to study this problem), and the case when $p = \omega(1/\sqrt{n})$ is proved using (a sparse version of) Szemerédi's regularity lemma, by showing that almost all triangle-free graphs with $\omega(n^{3/2})$ edges are nearly bipartite.

Sharper results (including first-order asymptotics of $\mu_p(\mathcal{T})$) are known when p is sufficiently smaller or sufficiently larger than $n^{-1/2}$. For $p \leq n^{-1/2-\varepsilon}$ with $\varepsilon > 0$ constant, asymptotics of $\mu_p(\mathcal{T})$ are known via Janson's Inequality [14] and extensions [29]–[31]. For $p \geq c\sqrt{\log n} \cdot n^{-1/2}$, asymptotics of $\mu_p(\mathcal{T})$ are known by quantifying how close $G \sim \mu_{\mathcal{T},p}$ is to being bipartite [2], [12], [32].

When $p = \Theta(n^{-1/2})$ it is known via [14] that $-\log \mu_p(\mathcal{T}) = \Theta(n^{3/2})$ but the leading constant is not known. Our algorithmic approach allows us to determine the asymptotics of the logarithm when $p \leq cn^{-1/2}$ for a constant $c > 0$.

Let $W(\cdot)$ denote the Lambert-W function, the inverse of the function xe^x .

Theorem I.5. Fix $0 < c < 1/\sqrt{e}$. If $p = (1 + o(1))cn^{-1/2}$, then

$$\lim_{n \rightarrow \infty} \frac{\log \mu_p(\mathcal{T})}{n^{3/2}} = \frac{1}{2} \left[\frac{W(2c^2)^{3/2} + 3W(2c^2)^{1/2}}{3\sqrt{2}} - c \right]. \quad (5)$$

In fact we can deduce from Theorem I.5 the existence of a phase transition (in the sense of statistical physics and non-analyticities of thermodynamic functions, see e.g. [33]) in the regime $p = \Theta(n^{-1/2})$.

Corollary I.6. A phase transition for $G(n, p)$ conditioned on triangle-freeness occurs at $p = c^*/\sqrt{n}$ for some $c^* \in [1/\sqrt{e}, 4.342)$ in the sense that the function

$$f(c) := \liminf_{n \rightarrow \infty} \frac{1}{n^{3/2}} \log \mu_{c/\sqrt{n}}(\mathcal{T}) \quad (6)$$

is non-analytic at $c = c^*$.

The proofs of Theorem I.5 and Corollary I.6 are not included in this extended abstract and appear in [34] along with additional results on more general lower tail probabilities for triangles in $G(n, p)$.

A. Techniques

The techniques we use to prove the main three results (algorithms at low densities, algorithms at high densities, log asymptotics at low densities) are related and combine ideas from computer science, statistical physics, and combinatorics.

The first result is the most straightforward: using the path coupling technique of Bubley and Dyer [35] along with an initial ‘burn-in’ period [36] to control the max degree of a graph, we show that the Glauber dynamics for $\mu_{\mathcal{T},p}$ are rapidly mixing.

For the second result, efficient sampling at high densities, a different algorithm is needed. In fact, using structural results from [32] we prove that the Glauber dynamics for $\mu_{\mathcal{T},p}$ mix slowly when $p \geq C/\sqrt{n}$ (Theorem I.2). These structural results (following those in [13]) say that almost all $G \sim \mu_{\mathcal{T},p}$ have a unique max-cut (A, B) that contains almost all of its edges, and moreover, the graphs induced by A and by B have controlled maximum degree. In the language of statistical physics we can say almost all $G \sim \mu_{\mathcal{T},p}$ are close to some ‘ground state’ collection of bipartite graphs with bipartition (A, B) . We call the edges within the parts of a max-cut (A, B) ‘defect edges’. While the existence of a large cut is a bottleneck for the Glauber dynamics, we use the same structural result as an algorithmic tool in proving Theorem I.3, reducing the sampling problem to that of sampling from $\mu_{\mathcal{T},p}$ conditioned on a max cut (A, B) and a small max degree among the defect edges.

With this approach, the main technical step is to show that one can efficiently sample defect edges from (approximately) the correct distribution. There is a rough duality to the problem in that defect edges approximately behave like a *disordered* triangle-free graph $H \sim \mu_{\mathcal{T},p}$ where $p = c/\sqrt{n}$ and c is small. However in reality the distribution is significantly more complex. The probability of seeing a set of defect edges H is proportional to a certain hard-core model partition function determined by H . Using the machinery of the *cluster expansion* (see Section II-C), the defect distribution can be viewed as an exponential random graph with an unbounded number of subgraph counts in the exponent, conditioned on a max-degree bound and on triangle-freeness. In a similar spirit to Theorem I.1, we sample defect edges via edge-update Glauber dynamics. However, the implementation is more complicated here. A key step is to use the cluster expansion to accurately estimate edge marginals needed in the implementation of the dynamics. This use of cluster expansion can be viewed as a way to study an exponential random graph as a perturbation of an Erdős-Rényi random graph.

There is a line of previous algorithmic results (e.g. [37]–[44]) on efficient approximate counting and sampling for spin models (hard-core, Potts, coloring, etc.) on structured instances (expanders, lattices, random graphs) in low-temperature (i.e. strong interaction strength) or high-density regimes that are computationally hard (more precisely #BIS-hard [45], [46]) in the worst-case. These algorithms are based on the framework of abstract polymer models in which defects from a ground

state configuration or set of configurations are modeled as an auxiliary spin model and then analyzed or sampled from using cluster expansion or Markov chain algorithms. Other low-temperature algorithms include suitably initialized Markov chains [47]–[49]. In both cases, it is crucial to the algorithms or to the analysis that defects from the appropriate ground state in the spin model are small: at most logarithmic size. See the discussion in [44, Section 2.2] and the percolation conditions in [48].

What is new in our approach to proving Theorem I.3 is that we can efficiently sample even though the defect edges form connected graphs with large degrees. At a high level, the novelty of our approach is that in previous works, the algorithms or analyses were based on perturbations around the empty set of defects; here one can interpret our analysis of Glauber dynamics on defect edges as a perturbation around a measure of independent (Erdős-Rényi) defect edges, and instead of needing to bound the size of their connected components, we need to bound their distance from independence, measured in terms of the mixing time of Glauber dynamics.

The third result, on the log asymptotics of $\mu_p(\mathcal{T})$, requires a new technique which combines several different ideas. We omit the full proof in this extended abstract, but it appears in [34]. Here we highlight the role that Markov chain mixing results (like Theorem I.1) can play in proving an asymptotic formula for $\log \mu_p(\mathcal{T})$.

We first observe that to approximate $\log \mu_p(\mathcal{T})$ (or equivalently, $\log Z(\lambda)$), it suffices to approximate the expected number of edges in a sample from $\mu_{\mathcal{T},p}$. Indeed with $p = \lambda/(1+\lambda)$ we have the identity

$$\mathbb{E}_{\mathcal{T},p}(|G|) = \sum_{G \in \mathcal{T}} |G| \frac{\lambda^{|G|}}{Z(\lambda)} = \lambda (\log Z(\lambda))', \quad (7)$$

where $\mathbb{E}_{\mathcal{T},p}$ denotes expectation with respect to the measure $\mu_{\mathcal{T},p}$. In particular,

$$\log Z(\lambda) = \int_0^\lambda \frac{\mathbb{E}_{\mathcal{T},\theta/(1+\theta)}(|G|)}{\theta} d\theta. \quad (8)$$

Such a relation between the expected size of an independent set and the hard-core partition function in the case of graphs has been used in extremal combinatorics in, e.g. [50], [51].

Second, we estimate the expectation $\mathbb{E}_{\mathcal{T},p}(|G|)$ by conditioning on the graph $G - v$ for a uniformly chosen v , and writing the expected degree of v in terms of a graph hard-core model. This allows us to reduce the study of $\mu_{\mathcal{T},p}$, a hypergraph hard-core measure, to the study of graph hard-core measures where more tools are available. In particular we are able to use the cluster expansion for the hard-core model on graphs whose current known analogues in the hypergraph setting are too weak for our purposes. In particular, we use the cluster expansion to accurately estimate the expected size of an independent set drawn from the hard-core model on a graph, knowing only that the graph is approximately regular and has relatively few short cycles, and that the activity parameter λ is sufficiently small as a function of the maximum degree. A key ingredient in this estimation is Theorem I.1, fast mixing

of the Glauber dynamics for $\mu_{\mathcal{T},p}$. This allows us to exploit a connection between mixing and the concentration inequalities (established in [52]) to show that a typical graph sampled from $\mu_{\mathcal{T},p}$ is approximately regular.

B. Open questions, future directions, and related work

This paper serves as a proof-of-concept for two new techniques, one for efficient sampling of low-temperature (or high density) spin models with large defects and the other for establishing log asymptotics of large deviation probabilities for subgraph counts in random graphs in near-critical regimes. The main future directions are to explore the power of these two methods in other settings. We give a couple of concrete questions here.

Question 1. *Can one obtain the asymptotics of the logarithm of more general lower-tail probabilities for triangles in $G(n, p)$ in the critical regime $p = \Theta(n^{-1/2})$ using variants of these techniques based on an anti-ferromagnetic Ising model (penalizing instead of forbidding triangles as in the hard-core model)?*

For some positive answers to this question, see [34].

It is also natural to consider forbidden subgraphs H other than triangles. If H is non-bipartite, the ‘critical density’ is determined by the 2-density of H , $m_2(H) = \max_{F \subseteq H; |E(F)| \geq 2} \frac{|E(F)|-1}{|V(F)|-2}$.

Question 2. *For non-bipartite graph H , can one sample efficiently from $G(n, p)$ conditioned on the event $\{X_H = 0\}$ or determine asymptotics of $\log \mu_p(X_H = 0)$ in the critical regime $p = \Theta(n^{-1/m_2(H)})$?*

See the discussion on non-existence and lower tail probabilities for $p = \omega(n^{-1/m_2(H)})$ in [6, Section 1.3] and in [11]; and for $p = o(n^{-1/m_2(H)})$ in [14], [30], [31], [53], [54].

We anticipate that our approach to high-density sampling will be useful more generally in the setting of low-temperature or high-density spin systems mentioned above. Concretely, we ask if this approach can be used to find efficient algorithms for the hard-core model on random bipartite graphs.

Question 3. *Can the high-level approach of the algorithm of Theorem I.3 be used to sample from the hard-core model on random Δ -regular bipartite graphs for all $\lambda > \lambda_c(\Delta)$?*

Currently efficient algorithms are known for $\lambda = \Omega(\log \Delta / \Delta)$ from [44] where the authors point out the barrier of polymers of polynomial size at smaller values of λ .

C. Organization

In Section II we provide some preliminaries on Markov chain mixing and the cluster expansion. In Section III we prove Theorem I.1 via a path coupling argument with burn-in. In Section IV we prove Theorem I.3, showing the existence of efficient sampling and counting algorithms at high densities. In Section VII we prove Theorem VII, showing that the Glauber dynamics for sampling from $\mu_{\mathcal{T},p}$ mixes slowly for large p .

II. PRELIMINARIES

A. Notation

For a graph $G = (V, E)$ we denote the number of edges by $|G|$ and its maximum degree by $\Delta(G)$. All graphs in this paper will have n vertices unless specified otherwise. For $v \in V$ we let $d_G(v)$ denote the degree of the vertex v and we let $d_G = 2|G|/n$ denote the average degree. We use μ_p to refer to the distribution of the Erdős-Rényi random graph $G(n, p)$; that is, $\mu_p(G) = p^{|G|}(1-p)^{\binom{n}{2}-|G|}$. We let $\mathcal{T}$ denote the set of triangle-free graphs on n vertices or the event that a random graph is triangle-free. The conditional distribution $\mu_p(\cdot|\mathcal{T})$ is denoted $\mu_{\mathcal{T},p}$. The partition function $Z(\lambda) = \sum_{G \in \mathcal{T}} \lambda^{|G|}$ will always be used with $\lambda = p/(1-p)$ giving the identity $\mu_p(\mathcal{T}) = (1-p)^{\binom{n}{2}} Z(\lambda)$.

B. Markov chains and mixing times

Our sampling algorithms for both the low density and high density cases will use Markov chains. In the low density case, the Markov chain approach will directly give an efficient sampling algorithm, while in the high density case, Markov chains will be one part of a more complicated algorithm.

The basic idea to approximately sample from a target distribution μ is to design a Markov chain with μ as the stationary distribution so that a single step of the chain can be implemented efficiently, and so that the chain converges quickly. The convergence time is often quantified by the *mixing time*. With $\mu_t^{X_0}$ denoting the t -step distribution of the Markov chain starting from the state X_0 , the mixing time is

$$\tau_{\text{mix}} = \max_{X_0} \min \left\{ t : \left\| \mu_t^{X_0} - \mu \right\|_{TV} < 1/4 \right\}. \quad (9)$$

Here $\|\cdot - \cdot\|_{TV}$ denotes the total variation distance between probability measures on the same space, and the choice of $1/4$ is arbitrary and can be reduced to any $\varepsilon > 0$ by running the chain a factor $O(\log(1/\varepsilon))$ more than τ_{mix} . For background on Markov chains and mixing times see [55], [56].

We will consider Markov chains on graphs; these will have the form of the *Glauber dynamics*: at each step we choose a potential edge $e \in \binom{V}{2}$ uniformly at random and then resample it (whether it is in the graph or not) conditioned on the status of all other edges. More generally, the Glauber dynamics can be applied to any high-dimensional probability distribution by choosing a random coordinate and resampling the value of that coordinate conditioned on the rest of the vector.

For the Erdős-Rényi distribution $G(n, p)$ the Glauber dynamics are particularly simple: at each step one of the $\binom{n}{2}$ possible edges is chosen uniformly; with probability p the edge is included in G and with probability $1 - p$ it is not included.

For $\mu_{\mathcal{T},p}$, the Glauber dynamics are similar: at each step one of the $\binom{n}{2}$ possible edges is chosen uniformly; with probability p the edge is included in G , but only if its inclusion would not form a triangle with other edges already present.

C. Cluster expansion and the hard-core model

Recall $Z_G(\lambda)$, the partition function of the hard-core model on a graph G from (1). The cluster expansion is a formal power series for $\log Z_G(\lambda)$; in fact, it is the Taylor series around $\lambda = 0$. Conveniently, the terms of the cluster expansion have a nice combinatorial interpretation (see e.g. [57], [58]). A *cluster* $\Gamma = (v_1, \dots, v_k)$ is a tuple of vertices from G such that the induced graph $G[\{v_1, \dots, v_k\}]$ is connected. We let $\mathcal{C}(G)$ denote the set of all clusters of G . We call k the size of the cluster and denote it by $|\Gamma|$. Given a cluster Γ , the *incompatibility graph* G_Γ is the graph on vertex set Γ (considered as a multiset) with an edge between v_i, v_j if either v_i, v_j are adjacent on G or $i \neq j$ and v_i, v_j correspond to the same vertex in G . In particular, by the definition of a cluster, the incompatibility graph G_Γ is connected.

As a formal power series, the cluster expansion is the infinite series

$$\log Z_G(\lambda) = \sum_{\Gamma \in \mathcal{C}(G)} \phi_G(\Gamma) \lambda^{|\Gamma|}, \quad (10)$$

where

$$\phi_G(\Gamma) = \frac{1}{|\Gamma|!} \sum_{\substack{A \subseteq E(G_\Gamma) \\ \text{spanning, connected}}} (-1)^{|A|}. \quad (11)$$

If the graph G is clear from the context we will often write $\phi(\Gamma)$ in place of $\phi_G(\Gamma)$.

The cluster expansion converges absolutely if λ lies inside a disk $D \subset \mathbb{C}$ so that $Z_G(\xi) \neq 0$ for all $\xi \in D$. We will use the following lemma, [32, Lemma 4.1], which gives a sufficient condition for convergence and bounds the error in truncating the cluster expansion.

Lemma II.1. *Suppose G is a graph on n vertices with maximum degree Δ , and suppose $|\lambda| \leq \frac{1}{2e(\Delta+1)}$. Then the cluster expansion converges absolutely. Moreover, for any set $U \subseteq V(G)$ such that $1 \leq |U| \leq \min\{2, k\}$ we have*

$$\sum_{\substack{\Gamma: \Gamma \supseteq U, \\ |\Gamma| \geq k}} |\phi(\Gamma)| |\lambda|^{|\Gamma|} \leq (2e)^k \Delta^{k-|U|} |\lambda|^k.$$

We remark that the conclusion of [32, Lemma 4.1] is slightly weaker than the statement above (the absolute value signs appear outside of the sum), however, it is readily checked that the proof from [32] gives the statement of Lemma II.1.

III. SAMPLING AT LOW DENSITY

In this section we prove Theorem I.1. Our proof will proceed by a path coupling argument where we use an initial burn-in period to make sure that the graphs in the coupling have reasonable maximum degree. To this end it will be useful to observe that $\mu_{\mathcal{T},p}$ is stochastically dominated by μ_p .

Lemma III.1. *For every $p \in [0, 1]$, the measure μ_p stochastically dominates $\mu_{\mathcal{T},p}$. That is, there is a coupling of the two distributions so that with probability 1, $G \subseteq G'$, where $G \sim \mu_{\mathcal{T},p}$ and $G' \sim \mu_p$.*

Proof. Sample $G \sim \mu_{\mathcal{T},p}$ by sampling one edge at a time given the previous history. At any step, the conditional probability of any edge is at most p . We can therefore couple the sampling with an edge-by-edge sampling of $G(n,p)$ so that an edge is present in the sample from $\mu_{\mathcal{T},p}$ only if it is present in the sample from μ_p . $\square$

This fact combined with Chernoff's inequality and a union bound over vertices yields the following corollary which will allow us to condition on the event $\Delta(G) \leq (1+\varepsilon)np$.

Corollary III.2. *Let $p, \varepsilon \in [0, 1]$ and let $G \sim \mu_{\mathcal{T},p}$. Then with $\mu = (n-1)p$*

$$\mathbb{P}(\Delta(G) \geq (1+\varepsilon)\mu) \leq ne^{-\varepsilon^2\mu/3}.$$

It will also be useful to observe that we can couple the triangle-free Glauber dynamics for $\mu_{\mathcal{T},p}$ with the Glauber dynamics for μ_p , in which one of the $\binom{n}{2}$ potential edges is chosen at random and resampled at each step.

Lemma III.3. *There exists a coupling $(X_t, Y_t)_{t \geq 0}$ of the Glauber dynamics for $\mu_{\mathcal{T},p}$ with the Glauber dynamics for μ_p such that if X_0 is a subgraph of Y_0 , then X_t is a subgraph of Y_t for all $t \geq 0$.*

Proof. We make the same edge updates in both chains, unless a pair $\{u, v\} \in \binom{[n]}{2}$ is chosen such that u, v have a common neighbour in X_t , in which case we set $X_{t+1} = X_t$. $\square$

A coupon collector argument gives us that the mixing time for Glauber dynamics for μ_p is $O(n^2 \log n)$. As a result, we have the following corollary.

Corollary III.4. *There exists $C > 0$ so that the following holds for any $c > 0$. Let $p \leq c/\sqrt{n}$. Let $(X_t)_{t \geq 0}$ be a run of the Glauber dynamics for $\mu_{\mathcal{T},p}$ on n vertices with X_0 arbitrary. If $t \geq Cn^2 \log(n/\varepsilon)$,*

$$\mathbb{P}(\Delta(X_t) \geq np + n^{1/3}) \leq ne^{-n^{2/3}/(3np)} + \varepsilon.$$

Proof. Consider the coupling $(X_t, Y_t)_{t \geq 0}$ from Lemma III.3 with $X_0 = Y_0$. Let Y be sampled from μ_p . For $t \geq Cn^2 \log(n/\varepsilon)$ we have

$$\|Y - Y_t\|_{TV} \leq \varepsilon.$$

By the Chernoff bound and the union bound, we have $\mathbb{P}(\Delta(Y) \geq np + n^{1/3}) \leq ne^{-n^{2/3}/(3np)}$. So it follows that $\mathbb{P}(\Delta(Y_t) \geq np + n^{1/3}) \leq ne^{-n^{2/3}/(3np)} + \varepsilon$. Since X_t is a subgraph of Y_t the result follows. $\square$

Now let $\Omega(p) \subseteq \mathcal{T}$ be the set of graphs in $\mathcal{T}$ whose maximum degree is at most $np + n^{1/3}$.

Lemma III.5. *Let $p = (1 + o(1))c/\sqrt{n}$ where $c < 1/\sqrt{2}$ is fixed. Let $(X_t, Y_t) \in \Omega(p) \times \Omega(p)$ and let (X_{t+1}, Y_{t+1}) be copies of the Glauber dynamics for $\mu_{\mathcal{T},p}$ coupled to attempt the same updates. Then*

$$\mathbb{E}[d(X_{t+1}, Y_{t+1})] \leq \left(1 - \frac{\delta}{\binom{n}{2}}\right) d(X_t, Y_t),$$

where $\delta = 1 - 2(np + n^{1/3})p = 1 - 2c^2 - o(1)$ and $d(\cdot, \cdot)$ is the Hamming distance between graphs.

Proof. By the path coupling technique [35], it suffices to consider the case where $d(X_t, Y_t) = 1$ i.e. X_t, Y_t differ in one edge. Suppose WLOG that $\{i, j\}$ is an edge of X_t , but not Y_t . Let d_i, d_j denote the degree of i, j in Y_t respectively and note that $d_i, d_j \leq np + n^{1/3}$. If the update was performed on $\{i, j\}$, then $X_{t+1} = Y_{t+1}$. If the update was performed on one of the $d_i + d_j - 2$ potential edges that form a triangle together with $\{i, j\}$, then this edge would be added to Y_{t+1} but not to X_{t+1} with probability at most p . If the update is performed on any other potential edge, then $d(X_{t+1}, Y_{t+1}) = d(X_t, Y_t)$. Thus

$$\mathbb{E}[d(X_{t+1}, Y_{t+1})] \leq 1 + \frac{d_i + d_j}{\binom{n}{2}} \cdot p - \frac{1}{\binom{n}{2}} \leq 1 - \frac{\delta}{\binom{n}{2}},$$

as claimed. $\square$

We are now ready to prove Theorem I.1.

Proof of Theorem I.1. Let $(X_t, Y_t)_{t \geq 0}$ be a coupling of the Glauber dynamics for $\mu_{\mathcal{T},p}$ to attempt the same updates with X_0, Y_0 chosen arbitrarily. Now let $T = 6Cn^2 \log n$ where C is as in Corollary III.4.

For $t \geq T$, let $\mathcal{E}_t$ denote the event $\{X_t, Y_t \in \Omega(p)\}$ and note that by Corollary III.4 (with $\varepsilon = n^{-5}$) and a union bound we have $\mathbb{P}(\mathcal{E}_t^c) \leq 2(ne^{-n^{2/3}/(3np)} + n^{-5}) = O(n^{-5})$. We then have,

$$\begin{aligned} \mathbb{P}(X_{2T} \neq Y_{2T}) &\leq \mathbb{E}[d(X_{2T}, Y_{2T})] \\ &\leq \mathbb{E}[d(X_{2T}, Y_{2T}) \mid \mathcal{E}_{2T-1}] + O(n^{-3}) \\ &\leq \left(1 - \frac{\delta}{\binom{n}{2}}\right) \mathbb{E}[d(X_{2T-1}, Y_{2T-1})] + O(n^{-3}). \end{aligned}$$

For the first inequality we used Markov. For the second inequality we used that $d(X_{2T}, Y_{2T}) \leq n^2$ and $\mathbb{P}(\mathcal{E}_{2T-1}^c) = O(n^{-5})$. For the third inequality we used Lemma III.5. Iterating the above $T-1$ more times, we conclude that

$$\begin{aligned} \mathbb{P}(X_{2T} \neq Y_{2T}) &\leq \left(1 - \frac{\delta}{\binom{n}{2}}\right)^T \mathbb{E}[d(X_T, Y_T)] + O(Tn^{-3}) \\ &= o(1), \end{aligned}$$

where for the final inequality we used that $\mathbb{E}[d(X_T, Y_T)] \leq n^2$.

It follows that

$$\tau_{\text{mix}} \leq 2T.$$

The existence of an efficient sampling scheme follows immediately, since a single step of the Glauber dynamics can be implemented efficiently.

The existence of an FPRAS for approximating $\mu_p(\mathcal{T})$ follows from a standard reduction of approximate counting to sampling. Using the adaptive simulated annealing technique from [59], which requires an efficient approximate sampling algorithm for $\mu_{\mathcal{T},p'}$ for all $0 \leq p' \leq p$, one obtains an FPRAS for $\mu_p(\mathcal{T})$ running in time $O(n^3 \varepsilon^{-2} \log^6 n)$. The

algorithm proceeds by using samples from $\mu_{\mathcal{T},p'}$ to estimate the ratio $\mu_{p''}(\mathcal{T})/\mu_{p'}(\mathcal{T})$ for adaptively chosen pairs p', p'' along a ‘cooling schedule’, then multiplying these estimates in a telescoping product. $\square$

IV. OVERVIEW FOR SAMPLING AT HIGH DENSITIES

In this section we give an overview of the proof of Theorem I.3, with the remaining proof details to follow in Section V. The proof relies on a structural result for graphs drawn from $\mu_{\mathcal{T},p}$ established in [32] in the regime where $p = C/\sqrt{n}$ and C is large. To state the result we introduce some notation and definitions from [32].

Definition IV.1. *Call a partition (A, B) of $[n]$ weakly balanced if $||A| - |B|| \leq n/10$. Let Π_{weak} denote the set of weakly balanced partitions of $[n]$.*

We set $\alpha = \frac{1}{96e^8}$. This constant is taken from [32]. Its precise form does not matter, only that it is sufficiently small. For a partition (A, B) of $[n]$ let

$$\mathcal{T}_{A,B,\lambda}^w = \{G \in \mathcal{T} : \Delta(G[A] \cup G[B]) \leq \alpha/\lambda\}, \quad (12)$$

where $G[A], G[B]$ denote the induced subgraphs of G on vertex sets A, B respectively.

We think of the set $\mathcal{T}_{A,B,\lambda}^w$ as those graphs $G \in \mathcal{T}$ which ‘align well’ with the bipartition (A, B) (i.e. (A, B) is a large cut for G).

Let

$$Z_{A,B}^w(\lambda) := \sum_{G \in \mathcal{T}_{A,B,\lambda}^w} \lambda^{|G|}. \quad (13)$$

In other words, $Z_{A,B}^w(\lambda)$ is the contribution to the partition function $Z(\lambda) = \sum_{G \in \mathcal{T}} \lambda^{|G|}$ from graphs that align with (A, B) .

Now define a distribution $\mu_{\text{weak},\lambda}$ on $\mathcal{T}$ as follows.

- 1) Sample $(A, B) \in \Pi_{\text{weak}}$ with probability proportional to $Z_{A,B}^w(\lambda)$.
- 2) Sample $G \in \mathcal{T}_{A,B,\lambda}^w$ with probability proportional to $\lambda^{|G|}$.
- 3) Output G .

Let

$$Z_{\text{weak}}(\lambda) = \sum_{(A,B) \in \Pi_{\text{weak}}} Z_{A,B}^w(\lambda).$$

Throughout this section we set $C > 0$ to be a sufficiently large constant.

Theorem IV.2 ([32, Theorem 2.9]). *For $\lambda \geq Cn^{-1/2}$,*

$$\|\mu_{\mathcal{T},p} - \mu_{\text{weak},\lambda}\|_{TV} = o(1),$$

where $p = \frac{\lambda}{1+\lambda}$.

Here we make the error estimate explicit.

Proposition IV.3. *For $\lambda \geq Cn^{-1/2}$,*

$$\|\mu_{\mathcal{T},p} - \mu_{\text{weak},\lambda}\|_{TV} = O(e^{-\sqrt{n}}),$$

where $p = \frac{\lambda}{1+\lambda}$. Moreover

$$Z(\lambda) = (1 + O(e^{-\sqrt{n}}))Z_{\text{weak}}(\lambda).$$

We defer the proof of Proposition IV.3 to Section VI.

With Proposition IV.3 in hand, we turn our attention to sampling from the measure $\mu_{\text{weak},\lambda}$. For this, we split Step 2 in the definition of $\mu_{\text{weak},\lambda}$ into two substeps: first we sample the ‘defect edges’ within parts A and B and then we sample edges crossing the partition (A, B) . The second of these substeps will be relatively simple. As we will see, the crossing edges have the distribution of a (graph) hard-core model, and the condition in (12) ensures that this model is subcritical in the sense that λ is small enough as a function of the max degree that the cluster expansion converges and efficient sampling algorithms exist. Most of the work in proving Theorem I.3 will come from implementing the first step and showing we can sample defect edges efficiently. To make this more precise, we make some further definitions.

Given a partition (A, B) of $[n]$ and a pair of graphs $S \subseteq \binom{A}{2}, T \subseteq \binom{B}{2}$, we write $S \square T$ as a shorthand for the Cartesian product of the graphs $(A, S), (B, T)$, i.e., the graph with vertex set $V(S \square T) = A \times B$ and edge set $E(S \square T)$ equal to $\{(a, b), (a, b')\} : \{b, b'\} \in T\} \cup \{(a, b), (a', b)\} : \{a, a'\} \in S\}$.

The significance of the Cartesian product for us comes from the following lemma.

Lemma IV.4. *Let (A, B) be a partition of $[n]$ and suppose $S \subseteq \binom{A}{2}, T \subseteq \binom{B}{2}$ such that $S \cup T$ is triangle-free. Let $\mathcal{G}(S, T)$ be the set of triangle-free graphs G so that $G[A] = S$ and $G[B] = T$. Then*

$$\sum_{G \in \mathcal{G}(S, T)} \lambda^{|G|} = \lambda^{|S|+|T|} Z_{S \square T}(\lambda),$$

where $Z_{S \square T}(\lambda)$ is the hard-core partition function on the graph $S \square T$.

Proof. In what follows we identify the vertex $(u, v) \in V(S \square T) = A \times B$ with the edge $\{u, v\} \in \binom{[n]}{2}$. The proof follows from the observation that if I is an independent set in the graph $S \square T$, then the graph on $[n]$ with edge set $S \cup T \cup I$ is a triangle-free graph in $\mathcal{G}(S, T)$, and likewise for any $G \in \mathcal{G}(S, T)$, $E(G) \cap (A \times B)$ forms an independent set in $S \square T$, giving a one-to-one correspondence. $\square$

Let $\mathcal{D}_{A,B,\lambda}^w$ denote the set of pairs (S, T) such that $S \subseteq \binom{A}{2}, T \subseteq \binom{B}{2}, \Delta(S \cup T) \leq \alpha/\lambda$ and $S \cup T$ is triangle free. In other words,

$$\mathcal{D}_{A,B,\lambda}^w = \{(G[A], G[B]) : G \in \mathcal{T}_{A,B,\lambda}^w\}, \quad (14)$$

the set of possible defect graphs with respect to (A, B) . By Lemma IV.4 we see that

$$Z_{A,B}^w(\lambda) = \sum_{(S,T) \in \mathcal{D}_{A,B,\lambda}^w} \lambda^{|S|+|T|} Z_{S \square T}(\lambda),$$

and we may rewrite the definition of $\mu_{\text{weak},\lambda}$ as

- 1) Sample $(A, B) \in \Pi_{\text{weak}}$ with probability proportional to $Z_{A,B}^w(\lambda)$.
- 2) Sample $(S, T) \in \mathcal{D}_{A,B,\lambda}^w$ with probability proportional to $\lambda^{|S|+|T|} Z_{S \square T}(\lambda)$.

3) Sample $E_{\text{cr}} \subseteq A \times B$ from the hard-core model on $S \sqcup T$ at activity λ .

4) Let $E = S \cup T \cup E_{\text{cr}}$ and output the graph $G = ([n], E)$.

As mentioned above, Step 2 will be our main focus. Given a weakly balanced partition (A, B) let $\nu_{A,B,\lambda}$ be the distribution on $\mathcal{D}_{A,B,\lambda}^w$ defined by

$$\nu_{A,B,\lambda}(S, T) = \frac{\lambda^{|S|+|T|} Z_{S \sqcup T}(\lambda)}{Z_{A,B}^w(\lambda)}. \quad (15)$$

Our main task is to show that we can sample from $\nu_{A,B,\lambda}$ efficiently. First we show that for $\lambda \geq Cn^{-1/2}$, the edge-update Glauber dynamics for $\nu_{A,B,\lambda}$ mixes rapidly.

Proposition IV.5. *For $\lambda \geq Cn^{-1/2}$, the mixing time of the edge-update Glauber dynamics for $\nu_{A,B,\lambda}$ is $O(n^2 \log n)$, uniformly over weakly balanced (A, B) .*

We prove Proposition IV.5 in the next section. We use path coupling in a similar spirit to the proof of Theorem I.1, however now the task is significantly more difficult due to the complex interactions between edges under the measure $\nu_{A,B,\lambda}$. In particular, if we condition on the entire graph outside of the edge e , computing the conditional marginal of e to sufficient accuracy is a non-trivial task (whereas this is trivial for $\mu_{\mathcal{T},p}$). Our proof of Proposition IV.5 leans heavily on the cluster expansion to estimate these edge marginals.

Once we have proved Proposition IV.5, we show that we can approximately implement the steps of the Glauber dynamics efficiently to obtain a sampling algorithm for $\nu_{A,B,\lambda}$.

Proposition IV.6. *If $\lambda > C/\sqrt{n}$ and (A, B) is weakly balanced then there are randomized algorithms running in time polynomial in n , $1/\varepsilon$ and $\log(1/\delta)$ (uniformly over all (A, B)) that*

- 1) *Output M so that with probability at least $1 - \delta$, M is an ε -relative approximation to $Z_{A,B,\lambda}^w(\lambda)$.*
- 2) *Output $G \in \mathcal{T}$ with distribution $\hat{\nu}$ so that $\|\hat{\nu} - \nu_{A,B,\lambda}\|_{TV} \leq \varepsilon$.*

We prove Proposition IV.6 in the next section.

V. PROOFS FOR SAMPLING AT HIGH DENSITY

In this section we prove Propositions IV.5 and IV.6 and deduce Theorem I.3. We begin with Proposition IV.5.

A. Glauber dynamics for the defect distribution

A key step in the proof of Proposition IV.5 is to show that $\nu_{A,B,\lambda}$ is stochastically dominated by two independent copies of relatively sparse Erdős-Rényi random graphs on A, B respectively. Given a set V and $q \in [0, 1]$, we let $\mu(V, q)$ denote the measure associated to the Erdős-Rényi random graph on vertex set V and edge probability q .

Lemma V.1. *For $\lambda \geq Cn^{-1/2}$, the measure $\nu_{A,B,\lambda}$ is stochastically dominated by $\mu(A, q) \times \mu(B, q)$ where $q = \lambda e^{-n\lambda^2/5}$.*

Lemma V.1 shows that as λ gets larger samples from the defect measure $\nu_{A,B,\lambda}$ become increasingly sparse. We will

prove Lemma V.1 in the following subsection (Section V-B). We record the following corollary.

Corollary V.2. *Let $\lambda \geq Cn^{-1/2}$. There exists $C' > 0$ so that the following holds. Let $(X_t)_{t \geq 0}$ be a run of the Glauber dynamics for $\nu_{A,B,\lambda}$ with X_0 arbitrary. If $t \geq C'n^2 \log(n/\varepsilon)$,*

$$\mathbb{P}(\Delta(X_t) \geq 2nq) \leq ne^{-nq/3} + \varepsilon$$

where $q = \lambda e^{-n\lambda^2/5}$.

We will prove Proposition IV.5 via path coupling. For $(\mathbf{S}, \mathbf{T}) \sim \nu_{A,B,\lambda}$, $S \subseteq \binom{A}{2}$, $T \subseteq \binom{B}{2}$ and $e \in \binom{A}{2} \cup \binom{B}{2}$, define the conditional edge marginal

$$p(e|S, T) := \mathbb{P}(e \in \mathbf{S} \cup \mathbf{T} \mid (\mathbf{S} \cup \mathbf{T}) \setminus e = S \cup T).$$

If $\mathbb{P}((\mathbf{S} \cup \mathbf{T}) \setminus e = S \cup T) = 0$ then we define $p(e|S, T)$ to be 0.

Let $(X_t, Y_t)_{t \geq 0}$, denote two runs of the Glauber dynamics of $\nu_{A,B,\lambda}$ coupled so that given $t \geq 0$ and (X_t, Y_t) we perform the following update:

- Pick $e \in \binom{A}{2} \cup \binom{B}{2}$ uniformly at random and let $U \sim U[0, 1]$.
- If $U \leq p(e|X_t)$ then set $X_{t+1} = X_t \cup e$. Otherwise set $X_{t+1} = X_t \setminus e$.
- Similarly, if $U \leq p(e|Y_t)$ then set $Y_{t+1} = Y_t \cup e$. Otherwise set $Y_{t+1} = Y_t \setminus e$.

We refer to this as the *optimal* coupling. Let

$$\Omega(\lambda) = \left\{ (S, T) \in \mathcal{D}_{A,B,\lambda}^w : \Delta(S \cup T) \leq 2n\lambda e^{-n\lambda^2/5} \right\}. \quad (16)$$

In analogy to Lemma III.5, our goal is to prove the following path coupling result.

Lemma V.3. *Let $\lambda \geq Cn^{-1/2}$. Let $(X_t, Y_t) \in \Omega(\lambda) \times \Omega(\lambda)$ and let (X_{t+1}, Y_{t+1}) be copies of the optimally coupled Glauber dynamics for $\nu_{A,B,\lambda}$. Then*

$$\mathbb{E}[d(X_{t+1}, Y_{t+1})] \leq \left(1 - \frac{1}{n^2}\right) d(X_t, Y_t),$$

where $d(\cdot, \cdot)$ is the Hamming distance between graphs.

We note that Proposition IV.5 follows from Lemma V.3 via the same burn-in argument that was used to deduce Theorem I.1 from Lemma III.5 (only now we use Corollary V.2 in place of Corollary III.4).

To prove Lemma V.3 we need a good understanding of the edge marginals that appear in the definition of the optimal coupling above. This will be a goal of the next subsection.

B. Stochastic domination and computing edge marginals in $\nu_{A,B,\lambda}$

Throughout this section we fix a weakly balanced partition (A, B) . For the proofs to come, it will be convenient to use the following shorthand. Given a graph $G \subseteq \binom{A}{2} \cup \binom{B}{2}$ we let $G_\square$ denote the graph $G[A] \sqcup G[B]$. We will identify G

with the pair $(G[A], G[B])$ and write $\nu_{A,B,\lambda}(G)$ in place of $\nu_{A,B,\lambda}(G[A], G[B])$ and $p(e|G)$ in place of $p(e|G[A], G[B])$.

Lemma V.1 follows immediately from the following.

Lemma V.4. *Let $\lambda \geq Cn^{-1/2}$. For $e \in \binom{A}{2} \cup \binom{B}{2}$, $G \subseteq \binom{A}{2} \cup \binom{B}{2}$,*

$$p(e|G) \leq \lambda e^{-n\lambda^2/5}.$$

Proof. We may assume that $\nu_{A,B,\lambda}(G \cup e) > 0$ and $e \notin G$ else the result is trivial. In this case,

$$p(e|G) = \frac{\nu_{A,B,\lambda}(G \cup e)}{\nu_{A,B,\lambda}(G) + \nu_{A,B,\lambda}(G \cup e)} \quad (17)$$

$$= \frac{\lambda Z_{(G \cup e)_\square}(\lambda) / Z_{G_\square}(\lambda)}{1 + \lambda Z_{(G \cup e)_\square}(\lambda) / Z_{G_\square}(\lambda)}. \quad (18)$$

Let us estimate the numerator in the expression above. First note that $\Delta((G \cup e)_\square) \leq 2\Delta(G \cup e) \leq 2\alpha/\lambda$ since $G \cup e \in \mathcal{D}_{A,B,\lambda}^w$ by assumption. Since $\lambda \leq (2e(1 + \alpha/\lambda))^{-1}$, we may apply Lemma II.1 (the cluster expansion).

Let $\mathcal{C}'(G_\square)$ denote the set of non-constant clusters of $G_\square$ and let

$$\mathcal{C}'' = \mathcal{C}'((G \cup e)_\square) \setminus \mathcal{C}'(G_\square). \quad (19)$$

Assume WLOG that $e \in \binom{A}{2}$, then

$$\begin{aligned} \log \left(\frac{Z_{(G \cup e)_\square}(\lambda)}{Z_{G_\square}(\lambda)} \right) &= \sum_{\Gamma \in \mathcal{C}''} \phi(\Gamma) \lambda^{|\Gamma|} \\ &= -\lambda^2 |B| + \sum_{\Gamma \in \mathcal{C}'' : |\Gamma| \geq 3} \phi(\Gamma) \lambda^{|\Gamma|}. \end{aligned} \quad (20)$$

For the second equality we used that there are no clusters of size one in $\mathcal{C}''$ and the clusters of size two are of the form (γ_1, γ_2) where $\{\gamma_1, \gamma_2\}$ is an edge of $(G \cup e)_\square$ but not of $G_\square$, and there are $|B|$ such edges. Call such an edge $\{\gamma_1, \gamma_2\}$ an *e-edge*. Now note that if $\Gamma \in \mathcal{C}''$, then Γ must contain an *e-edge* $\{\gamma_1, \gamma_2\}$. By Lemma II.1 (applied with $k = 3$ and $U = \{\gamma_1, \gamma_2\}$ an *e-edge*) we have

$$\begin{aligned} \left| \sum_{\Gamma \in \mathcal{C}'' : |\Gamma| \geq 3} \phi(\Gamma) \lambda^{|\Gamma|} \right| &\leq |B| (2e)^3 \Delta((G \cup e)_\square) \lambda^3 \\ &\leq |B| 2^4 e^3 \alpha \lambda^2 \\ &\leq \lambda^2 |B| / 2, \end{aligned}$$

where for the second inequality we used that $\Delta((G \cup e)_\square) \leq 2\Delta(G \cup e) \leq 2\alpha/\lambda$ since $G \cup e \in \mathcal{D}_{A,B,\lambda}^w$ by assumption. Returning to (20) we conclude that

$$\log \left(\frac{Z_{(G \cup e)_\square}(\lambda)}{Z_{G_\square}(\lambda)} \right) \leq -\lambda^2 |B| / 2 \quad (21)$$

and so by (17)

$$p(e|G) \leq \lambda e^{-|B|\lambda^2/2} \leq \lambda e^{-n\lambda^2/5},$$

where for the last inequality we used that (A, B) is weakly balanced. $\square$

With Lemma V.3 in mind, we now prove a result that compares the marginals of an edge e in two configurations

that differ in exactly one edge f . First we introduce some notation.

For $e \in \binom{A}{2}$, we call a pair in $A \times B$ of the form $\{(u, x), (v, x)\}$ an *e-pair* (and similarly if $e \in \binom{B}{2}$). For $e, f \in \binom{A}{2} \cup \binom{B}{2}$, let $\mathcal{C}_{e,f}(G_\square)$ denote the set

$$\{\Gamma \in \mathcal{C}(G_\square) : \Gamma \text{ contains both an } e\text{-pair and an } f\text{-pair}\}.$$

Recall the definition of $\Omega(\lambda)$ from (16).

Lemma V.5. *Let $\lambda \geq Cn^{-1/2}$. Let $e, f \in \binom{A}{2} \cup \binom{B}{2}$, $G \in \Omega(\lambda)$. If $p(e|G \cup f) > 0$, then*

$$|p(e|G) - p(e|G \cup f)| \leq 2\lambda e^{-\lambda^2 n/5} R$$

where

$$R := \sum_{\Gamma \in \mathcal{C}_{e,f}((G \cup e)_\square)} |\phi(\Gamma)| \lambda^{|\Gamma|} + \sum_{\Gamma \in \mathcal{C}_{e,f}((G \cup e \cup f)_\square)} |\phi(\Gamma)| \lambda^{|\Gamma|}.$$

Proof. We continue with the notation in the proof of Lemma V.4, with the additional assumption that $f \notin G$. Let $\mathcal{C}''$ be as in (19) and let

$$\tilde{\mathcal{C}} := \mathcal{C}'((G \cup e \cup f)_\square) \setminus \mathcal{C}'((G \cup f)_\square)$$

By (20) we have

$$\begin{aligned} \log \left(\frac{Z_{(G \cup e)_\square}(\lambda)}{Z_{G_\square}(\lambda)} \right) - \log \left(\frac{Z_{(G \cup e \cup f)_\square}(\lambda)}{Z_{(G \cup f)_\square}(\lambda)} \right) \\ = \sum_{\Gamma \in \mathcal{C}''} \phi_1(\Gamma) \lambda^{|\Gamma|} - \sum_{\Gamma \in \tilde{\mathcal{C}}} \phi_2(\Gamma) \lambda^{|\Gamma|}, \end{aligned} \quad (22)$$

where $\phi_1 = \phi_{(G \cup e)_\square}$ and $\phi_2 = \phi_{(G \cup e \cup f)_\square}$. We note that if $\Gamma \in \mathcal{C}'' \cup \tilde{\mathcal{C}}$ then Γ must contain an *e-pair*. Note also that if $\Gamma \in \tilde{\mathcal{C}}$ does not contain an *f-pair*, then $\Gamma \in \mathcal{C}''$ and $\phi_1(\Gamma) = \phi_2(\Gamma)$. Similarly if $\Gamma \in \mathcal{C}''$ does not contain an *f-pair*, then $\Gamma \in \tilde{\mathcal{C}}$ and $\phi_1(\Gamma) = \phi_2(\Gamma)$. We conclude that for a cluster $\Gamma \in \mathcal{C}'' \cup \tilde{\mathcal{C}}$ to contribute to the difference of sums in (22), Γ must contain both an *e-pair* and an *f-pair*. Thus (dropping the subscripts from ϕ_1, ϕ_2)

$$\left| \log \left(\frac{Z_{(G \cup e)_\square}(\lambda)}{Z_{G_\square}(\lambda)} \right) - \log \left(\frac{Z_{(G \cup e \cup f)_\square}(\lambda)}{Z_{(G \cup f)_\square}(\lambda)} \right) \right| \leq R. \quad (23)$$

We note that if Γ contains both an *e-pair* and an *f-pair* then then $|\Gamma| \geq 3$ (we may assume that $e \neq f$ else the lemma is trivial). Since there are at most n *e-pairs* an application of Lemma II.1 (applied with $k = 3$ and U an *e-pair*) shows that

$$\begin{aligned} \sum_{\Gamma \in \mathcal{C}_{e,f}((G \cup e \cup f)_\square)} |\phi(\Gamma)| \lambda^{|\Gamma|} &= O(n \Delta((G \cup e \cup f)_\square) \lambda^3) \\ &= O(n^2 q \lambda^3) \\ &= O(C^4 e^{-C^2/5}) \end{aligned}$$

where for the penultimate inequality we used that $\Delta((G \cup e \cup f)_\square) = O(nq)$ since $G \in \Omega(\lambda)$. The same bound also holds for $\sum_{\Gamma \in \mathcal{C}_{e,f}((G \cup e)_\square)} |\phi(\Gamma)| \lambda^{|\Gamma|}$. It follows that

$R = O(C^4 e^{-C^2/5})$ and so $R \leq 1$ for C sufficiently large. As a consequence

$$\left| \frac{\lambda Z_{(G \cup e)_{\square}}(\lambda)}{Z_{G_{\square}}(\lambda)} - \frac{\lambda Z_{(G \cup e \cup f)_{\square}}(\lambda)}{Z_{(G \cup f)_{\square}}(\lambda)} \right| \leq \frac{\lambda Z_{(G \cup e)_{\square}}(\lambda)}{Z_{G_{\square}}(\lambda)} |e^R - 1| \leq 2\lambda e^{-\lambda^2 n/5} R$$

where for the second inequality we used (21) and the fact that $R \leq 1$. The result now follows from (17) and the fact that $|x/(1+x) - y/(1+y)| < |x-y|$ for $x, y > 0$. $\square$

C. Bounding cluster sums

As in the previous section, we fix a weakly balanced partition (A, B) .

Our goal in this section is to effectively bound the sums over clusters appearing in Lemma V.5.

Given a graph H , let $\mathcal{Q}(H)$ denote the set of labeled spanning trees of H . The following is the tree-graph bound of Penrose which will be convenient when bounding sums of Ursell functions.

Lemma V.6 ([60]). *For a graph $H = (V, E)$,*

$$\left| \sum_{\substack{A \subseteq E: \\ (V, A) \text{ connected}}} (-1)^{|A|} \right| \leq |\mathcal{Q}(H)|.$$

Given a graph G and cluster $\Gamma \in \mathcal{C}(G)$ such that $|\Gamma| = k$, we identify the vertex set of G_{Γ} (the incompatibility graph of Γ) with $[k]$.

Lemma V.7. *Let $k \geq 1$ and let Q be a tree on vertex set $[k]$. Let $G \subseteq \binom{A}{2} \cup \binom{B}{2}$, $f \in \binom{A}{2} \cup \binom{B}{2}$.*

$$\sum_{\substack{e \in \binom{A}{2} \cup \binom{B}{2}, \\ e \neq f}} \sum_{\substack{\Gamma \in \mathcal{C}_{e,f}((G \cup e \cup f)_{\square}): \\ |\Gamma| = k}} \mathbf{1}_{Q \in \mathcal{Q}(H_{e,\Gamma})} \leq 2k^5 n^2 (\Delta + 1)^{k-3}, \quad (24)$$

where $\Delta = \Delta((G \cup f)_{\square})$ and $H_{e,\Gamma}$ is the incompatibility graph of Γ with respect to the graph $(G \cup e \cup f)_{\square}$.

Proof. Recall that a cluster Γ of size k is a k -tuple of elements of $A \times B$, i.e., Γ is a map $[k] \rightarrow A \times B$.

We pick an f -pair $\tilde{f} = \{u, v\}$ (there are at most n choices for $\tilde{f}$) and $z, w \in [k]$, $\{x, y\} \in E(Q)$ (there are at most $2k^3$ choices for (x, y, w, z)) and construct a cluster Γ such that for some $e \in \binom{A}{2} \cup \binom{B}{2}$ where $e \neq f$:

- 1) $Q \in \mathcal{Q}(H_{e,\Gamma})$.
- 2) $\Gamma \in \mathcal{C}_{e,f}((G \cup e \cup f)_{\square})$.
- 3) $\Gamma(z) = u, \Gamma(w) = v$.
- 4) Either
 - a) $\Gamma(a) = \Gamma(b)$ or $\Gamma(\{a, b\})$ is an edge of $(G \cup f)_{\square}$ for all $\{a, b\} \in E(Q)$ (in other words, Q is a labeled spanning tree of the incompatibility graph of Γ with respect to $(G \cup f)_{\square}$), or
 - b) $\Gamma(\{x, y\})$ is not an edge of $(G \cup f)_{\square}$ and $\Gamma(x) \neq \Gamma(y)$ and $\{x, y\}$ is the closest edge of Q to $\{w, z\}$

(in terms of graph distance in Q) with these properties.

Suppose without loss of generality that $\delta_Q(\{x, y\}, \{w, z\}) = \delta_Q(x, w)$ where δ_Q denotes graph distance in Q . Let $\{z', z\}$ be an edge of Q such that z is in a separate component to x, y, w in the graph Q with the edge $\{z, z'\}$ removed. Consider the graph obtained by deleting edges $\{z, z'\}$ and $\{x, y\}$ from Q and call its three components Q_1, Q_2, Q_3 . Assume without loss of generality that Q_1 contains w and x , Q_2 contains z , and Q_3 contains y . Let P be the path in Q_1 from w to x . By (4), $\Gamma(a) = \Gamma(b)$ or $\Gamma(\{a, b\})$ is an edge of $(G \cup f)_{\square}$ for all $\{a, b\} \in E(P)$. Since $\Gamma(w)$ is fixed, there are at most $(\Delta + 1)^{|P|}$ choices for $(\Gamma(v) : v \in V(P))$. Given $(\Gamma(v) : v \in V(P))$ (and so in particular $\Gamma(x)$) there are at most n choices for $\Gamma(y)$ since $\Gamma(\{x, y\})$ need not be an edge of $(G \cup f)_{\square}$ but does need to be an edge of $(G \cup e \cup f)_{\square}$ for some e . If $\Gamma(\{x, y\})$ is not an edge of $(G \cup f)_{\square}$ then $\Gamma(\{x, y\})$ specifies the pair e for which (1) and (2) hold.

Observe that since $Q \in \mathcal{Q}(H_{e,\Gamma})$ we have that either $\Gamma(a) = \Gamma(b)$ or $\Gamma(\{a, b\})$ is an edge of $(G \cup e \cup f)_{\square}$ for all $\{a, b\} \in E(Q)$. We conclude that there are at most $(\Delta + 2)^{|Q_1| - |P|}$ choices for $(\Gamma(v) : v \in V(Q_1) \setminus V(P))$ since $(G \cup e \cup f)_{\square}$ has maximum degree at most $\Delta + 1$. Similarly there are at most $(\Delta + 2)^{|Q_2|}$ choices for $(\Gamma(v) : v \in V(Q_2))$, and at most $(\Delta + 2)^{|Q_3|}$ choices for $(\Gamma(v) : v \in V(Q_3))$ (since $\Gamma(y)$ has been specified). In total there are at most

$$2k^3 n^2 (\Delta + 2)^{|Q_1| + |Q_2| + |Q_3|} = 2k^3 n^2 (\Delta + 2)^{k-3}$$

choices for Γ satisfying (1)-(4) for some f -pair $\{u, v\}$, some $\{x, y\} \in E(Q)$ and some $e \in \binom{A}{2} \cup \binom{B}{2}$, $e \neq f$. Finally we note that every such cluster is counted at most k^2 times in the sum (24). $\square$

Let Q_k denote the set of all labeled trees on vertex set $[k]$.

Lemma V.8. *If $G \subseteq \binom{A}{2} \cup \binom{B}{2}$, $f \in \binom{A}{2} \cup \binom{B}{2}$, then*

$$\sum_{\substack{e \in \binom{A}{2} \cup \binom{B}{2}, \\ e \neq f}} \sum_{\substack{\Gamma \in \mathcal{C}_{e,f}((G \cup e \cup f)_{\square}): \\ |\Gamma| = k}} |\phi(\Gamma)| \leq 20^k n^2 \Delta^{k-3}, \quad (25)$$

where $\Delta = \Delta((G \cup f)_{\square})$.

Proof. By Lemma V.6, the LHS of 25 is at most

$$\begin{aligned} & \frac{1}{k!} \sum_{\substack{e \in \binom{A}{2} \cup \binom{B}{2}, \\ e \neq f}} \sum_{\substack{\Gamma \in \mathcal{C}_{e,f}((G \cup e \cup f)_{\square}): \\ |\Gamma| = k}} \left| \sum_{\substack{A \subseteq E(H_{e,\Gamma}): \\ ([k], A) \text{ connected}}} (-1)^{|A|} \right| \\ & \leq \frac{1}{k!} \sum_{\substack{e \in \binom{A}{2} \cup \binom{B}{2}, \\ e \neq f}} \sum_{\substack{\Gamma \in \mathcal{C}_{e,f}((G \cup e \cup f)_{\square}): \\ |\Gamma| = k}} \sum_{Q \in Q_k} \mathbf{1}_{Q \in \mathcal{Q}(H_{e,\Gamma})} \\ & \leq \frac{1}{k!} \sum_{Q \in Q_k} \sum_{\substack{e \in \binom{A}{2} \cup \binom{B}{2}, \\ e \neq f}} \sum_{\substack{\Gamma \in \mathcal{C}_{e,f}((G \cup e \cup f)_{\square}): \\ |\Gamma| = k}} \mathbf{1}_{Q \in \mathcal{Q}(H_{e,\Gamma})} \\ & \leq 20^k n^2 \Delta^{k-3} \end{aligned}$$

where for the final inequality we used Lemma V.7 and Cayley's formula: $|Q_k| = k^{k-2}$. $\square$

In the next subsection, we prove Lemma V.3 and hence also Proposition IV.5.

D. Proof of Lemma V.3

As before, it suffices to consider the case where $d(X_t, Y_t) = 1$ i.e. X_t and Y_t differ in one edge. Suppose WLOG that f is an edge of X_t , but not Y_t . Consider a randomly chosen $e \in \binom{A}{2} \cup \binom{B}{2}$ chosen for update. There are three cases to consider:

Case 1: $e = f$. In this case

$$\mathbb{E}[d(X_{t+1}, Y_{t+1}) - d(X_t, Y_t) \mid e] = -1.$$

Case 2: $X_t \cup e$ contains a triangle: In this case

$$\mathbb{E}[d(X_{t+1}, Y_{t+1}) - d(X_t, Y_t) \mid e] \leq q$$

since the distance increases only if $Y_{t+1} = Y_t \cup e$ which happens with probability at most q by Lemma V.4. We note that the number of pairs e such that $X_t \cup \{e, f\}$ contains a triangle is at most $2qn$ (it is here that we use the assumption $X_t \in \Omega(\lambda)$).

Case 3: $X_t \cup e$ does not contain a triangle and $e \neq f$. Given such an e we have

$$\mathbb{E}[d(X_{t+1}, Y_{t+1}) - d(X_t, Y_t) \mid e] \leq |p(e|X_t) - p(e|Y_t)|$$

since the only way for the distance to increase is if $X_{t+1} = X_t \cup e$ and $Y_{t+1} = Y_t$ which occurs with probability at most $|p(e|X_t) - p(e|Y_t)|$ by the definition of the optimal coupling. Here we have used that $X_t \in \Omega(\lambda)$ so that $\Delta(X_t) \leq 2n\lambda e^{-n\lambda^2/5} \leq \alpha/\lambda - 1$ and so $X_t \cup e$ does not violate the maximum degree condition in the definition of $\mathcal{D}_{A,B,\lambda}^w$. Let H denote the set of all $e \neq f$ such that $X_t \cup e$ does not contain a triangle. Let $q := \lambda e^{-\lambda^2 n/5}$ and for $e \in H$ let

$$R_e := \sum_{\Gamma \in \mathcal{C}_{e,f}((Y_t \cup e)_\square)} |\phi(\Gamma)| \lambda^{|\Gamma|} + \sum_{\Gamma \in \mathcal{C}_{e,f}((X_t \cup e)_\square)} |\phi(\Gamma)| \lambda^{|\Gamma|}.$$

By Lemmas V.5 and V.8 (recalling that $X_t = Y_t \cup f$) we conclude that

$$\begin{aligned} & \sum_{e \in H} \mathbb{E}[d(X_{t+1}, Y_{t+1}) - d(X_t, Y_t) \mid e] \\ & \leq 2q \sum_{e \in H} R_e \\ & \leq 4qn^2 \sum_{k \geq 3} 20^k (2nq)^{k-3} \lambda^k \\ & \leq 10^5 q n^2 \lambda^3 \end{aligned}$$

where for the final inequality we used that $40\lambda nq \leq 1/2$. Letting $N = \binom{|A|}{2} + \binom{|B|}{2}$, it follows that,

$$\begin{aligned} \mathbb{E}[d(X_{t+1}, Y_{t+1})] - 1 & \leq \frac{1}{N} (-1 + 2q^2 n + 10^5 q n^2 \lambda^3) \\ & \leq -\frac{1}{2N} \\ & \leq -\frac{1}{n^2}, \end{aligned}$$

for C sufficiently large. $\square$

E. Sampling from the defect distribution

With Proposition IV.5 in hand, we now prove Proposition IV.6 and show that we can approximately sample from the defect distributions $\nu_{A,B,\lambda}$. We will show that we can implement each step of the Glauber dynamics efficiently and up to small error. For this, we need the following lemma that will allow us to approximate edge marginals in $\nu_{A,B,\lambda}$.

Lemma V.9. [[19], [61]] *Let G be a graph on n vertices and let $\lambda \leq 1/\Delta(G)$. Given $\varepsilon, \delta > 0$ there exist randomized algorithms that runs in time polynomial in $n, \log(1/\varepsilon), \log(1/\delta)$ that*

- 1) *Output M so that with probability at least $1 - \delta$, M is an ε -relative approximation to $Z_G(\lambda)$.*
- 2) *Output G with distribution $\hat{\mu}$ so that $\|\hat{\mu} - \mu\|_{TV} \leq \varepsilon$ where μ is the hard-core measure on G at activity λ .*

Proof of Proposition IV.6. Part (1) of the Proposition IV.6 follows from part (2) via a standard reduction of approximate counting to sampling. We now prove part (2). Let $(X_t)_{t \geq 0}$ be a run of the Glauber dynamics for $\nu_{A,B,\lambda}$ with X_0 arbitrary. By Proposition IV.5 there exists $K > 0$ such that if $T := Kn^2 \log(n/\varepsilon)$ then $\|\mathcal{L}(X_T) - \nu_{A,B,\lambda}\|_{TV} \leq \varepsilon$. We will define a dynamic (Y_t) that will serve as an approximation to (X_t) .

Recall from (17) that if $\nu_{A,B,\lambda}(G \cup e) > 0$ then

$$p(e|G) = \frac{\lambda Z_{(G \cup e)_\square}(\lambda)}{Z_{G_\square}(\lambda) + \lambda Z_{(G \cup e)_\square}(\lambda)}, \quad (26)$$

and $p(e|G) = 0$ otherwise. Let $\delta = 1/T$ and let $\varepsilon' = \varepsilon/T$. By Lemma V.9, in $\text{poly}(n, 1/\varepsilon)$ time we can compute M_1, M_2 in so that with probability $\geq 1 - \delta$, M_1, M_2 are ε' -relative approximations to $Z_{(G \cup e)_\square}(\lambda), Z_{G_\square}(\lambda)$ respectively. Let

$$\hat{p}(e|G) = \frac{\lambda M_1}{M_2 + \lambda M_1},$$

and note that with probability $\geq 1 - \delta$, $\hat{p}(e|G)$ is a $2\varepsilon'$ -relative approximation to $p(e|G)$ and so $|p(e|G) - \hat{p}(e|G)| \leq e^{2\varepsilon'} - 1 \leq 3\varepsilon'$ since $p(e|G) \leq 1$.

Define (Y_t) as follows. Given Y_t perform the following update:

- Pick $e \in \binom{A}{2} \cup \binom{B}{2}$ uniformly at random and let $U \sim U[0, 1]$.
- Compute $\hat{p}(e|Y_t)$. If $U \leq \hat{p}(e|Y_t)$ then set $Y_{t+1} = Y_t \cup e$. Otherwise set $Y_{t+1} = Y_t \setminus e$.

We note that we can compute Y_T in $\text{poly}(n, 1/\varepsilon)$ time and by coupling $(Y_t), (X_t)$ in the obvious way we see that

$$\mathbb{P}(X_T = Y_T) \geq [(1 - \delta)(1 - 3\varepsilon)]^T \geq 1 - T(\delta + 3\varepsilon) \geq 1 - 4\varepsilon.$$

By the coupling inequality we conclude that

$$\|\mathcal{L}(Y_T) - \mathcal{L}(X_T)\|_{TV} \leq \mathbb{P}(X_T \neq Y_T) \leq 4\varepsilon,$$

and so by the triangle inequality

$$\|\mathcal{L}(Y_T) - \nu_{A,B,\lambda}\|_{TV} \leq 5\varepsilon. \quad \square$$

F. Proof of Theorem I.3

First we prove part (1) of the theorem. First observe that by symmetry, $Z_{A,B}^w(\lambda)$ depends only on the level of imbalance $||A| - |B||$ of the partition (A, B) . If $k = ||A| - |B||$ for some weakly balanced partition (A, B) we set $Z_k(\lambda) := Z_{A,B}^w(\lambda)$ (so in particular $k \leq n/10$). By Proposition IV.6, for each $k \in [0, n/10]$ we can compute M_k in $\text{poly}(n, 1/\varepsilon)$ time such that with probability $\geq 1 - 1/n$, M_k is an ε -relative approximation to $Z_k(\lambda)$. We can therefore compute

$$M := \sum_{k=0}^{n/10} \binom{n}{n/2 + k/2} M_k$$

in $\text{poly}(n, 1/\varepsilon)$ time and, by a union bound, with probability $\geq 9/10$, M is an ε -relative approximation to

$$Z_{\text{weak}}(\lambda) = \sum_{k=0}^{n/10} \binom{n}{n/2 + k/2} Z_k(\lambda).$$

Part (1) of the theorem now follows from Proposition IV.3.

We now turn to part (2). Let $\varepsilon > 0$. By Proposition IV.3 it suffices to show that we can sample from $\mu_{\text{weak}, \lambda}$ up to total variation distance ε in time polynomial in n and $1/\varepsilon$. For this, we show that we can implement each step in the definition of $\mu_{\text{weak}, \lambda}$ up to total variation distance ε in time polynomial in n and $1/\varepsilon$.

Sampling a partition. Step 1 in the definition of $\mu_{\text{weak}, \lambda}$ can be rewritten as:

- 1) Sample an integer k with probability proportional to $Z_k(\lambda)$.
- 2) Sample a partition (A, B) of $[n]$ such that $||A| - |B|| = k$ uniformly at random.

As above it suffices to observe that for each $k \in [0, n/10]$ we can compute M_k in $\text{poly}(n, 1/\varepsilon)$ time such that with probability $\geq 1 - 1/n$, M_k is an ε -relative approximation to $Z_k(\lambda)$.

Sampling defects. Proposition IV.6 shows that we can implement Step 2 in the definition of $\mu_{\text{weak}, \lambda}$, i.e., we can sample from $\nu_{A,B,\lambda}$ up to ε total variation distance in $\text{poly}(n, 1/\varepsilon)$ time.

Sampling crossing edges. Given (A, B) and $(S, T) \in \mathcal{D}_{A,B,\lambda}^w$ from Steps 1 and 2, we note that $\lambda \leq 1/\Delta(S \sqcap T)$ since $\Delta(S \sqcap T) \leq 2\Delta(S \cup T) \leq 2\alpha/\lambda$ (by the definition of $\mathcal{D}_{A,B,\lambda}^w$). We may therefore apply Lemma V.9 to output sample from the hard-core model on $S \sqcap T$ at activity λ up to total variation distance ε in $\text{poly}(n, 1/\varepsilon)$ time.

VI. PROOF OF PROPOSITION IV.3

We begin with a special case of [6, Theorem 1.7].

Theorem VI.1. *For all $\delta > 0$ there exist $C, c > 0$ so that if $m \geq Cn^{3/2}$, then all but a e^{-cm} proportion of graphs in $\mathcal{T}(n, m)$ admit a cut of size at least $(1 - \delta)m$.*

We note that in [6] the result is stated with $o(1)$ in place of the effective error e^{-cm} . However, as noted by in previous

works (see e.g. the comment following [62, Theorem 1.2]) the proof in [6] gives the error stated above. We remark that Theorem VI.1 was first proved (with $o(1)$ in place of e^{-cm}) by T. Łuczak [13].

We will in fact need a slight refinement of Theorem VI.1 (Theorem VI.3 below) which follows from combining [62, Proposition 6.1] and [62, Claim 6.2]. To state the result we need a definition.

Definition VI.2. *Let G be a graph and (A, B) a partition of its vertex set. We say (A, B) is a dominating cut of G if*

$$d_G(v, B) \geq d_G(v, A) \text{ for all } v \in A$$

and similarly with A, B swapped.

Recall also the definition of a weakly balanced partition from Definition IV.1.

Theorem VI.3. *For all $\delta > 0$ sufficiently small, there exist $C, c > 0$ so that if $m \geq Cn^{3/2}$, then all but a e^{-cm} proportion of graphs in $\mathcal{T}(n, m)$ admit a dominating, weakly balanced cut of size at least $(1 - \delta)m$.*

The results of [62] are also stated with a $o(1)$ error in place of e^{-cm} . However one readily checks that the short proof of Theorem VI.3 from Theorem VI.1 presented in [62] yields the error stated here.

For $\lambda, \delta > 0$, let $\mathcal{L}(\lambda, \delta)$ denote the set of all $G \in \mathcal{T}$ such that G admits a weakly balanced, dominating cut of size $\geq |G| - 2\delta\lambda n^2$.

For a subset $\mathcal{R} \subseteq \mathcal{T}$ and $\lambda > 0$ we define the measure on $\mathcal{R}$ given by

$$\mu_{\mathcal{R}, \lambda}(G) = \frac{\lambda^{|G|}}{Z(\mathcal{R}, \lambda)},$$

where

$$Z(\mathcal{R}, \lambda) = \sum_{G \in \mathcal{R}} \lambda^{|G|}.$$

We use the following result from [32] (Proposition 3.4). Recall that we let $Z(\lambda) = Z(\mathcal{T}, \lambda)$.

Proposition VI.4. *For all $\delta > 0$ sufficiently small, there exists $C, c > 0$ such that if $\lambda \geq C/\sqrt{n}$ then*

$$\|\mu_{\mathcal{T}, p} - \mu_{\mathcal{L}, \lambda}\|_{TV} \leq e^{-c\lambda n^2},$$

and

$$Z(\lambda) = (1 + O(e^{-c\lambda n^2}))Z(\mathcal{L}, \lambda),$$

where $\mathcal{L} = \mathcal{L}(\lambda, \delta)$ and $p = \lambda/(1 + \lambda)$

The proposition is an easy consequence of Theorem VI.3. Here we have included the effective error that the proof in [32] gives.

Finally we note the following result from [32] (Proposition 3.5).

Proposition VI.5. *For all $\delta > 0$ sufficiently small, there exists $C > 0$ such that if $\lambda \geq C/\sqrt{n}$ then*

$$\|\mu_{\mathcal{L}, \lambda} - \mu_{\text{weak}, \lambda}\|_{TV} = O(e^{-\sqrt{n}}),$$

and

$$Z(\mathcal{L}, \lambda) = \left(1 + O(e^{-\sqrt{n}})\right) Z_{\text{weak}}(\lambda),$$

where $\mathcal{L} = \mathcal{L}(\lambda, \delta)$.

Combining the above two propositions, we immediately arrive at a proof of Proposition IV.3.

VII. SLOW MIXING AT HIGH DENSITY

In this section we prove Theorem I.2. We import some structural results from [32]. Roughly speaking, these results will show that almost all graphs in $\mathcal{T}(n)$ have a unique max cut in a rather robust sense.

First we need a definition. Given a graph G and subsets $X, Y \subseteq V(G)$, we let $E(X, Y)$ denote the set of edges with one endpoint in X and one endpoint in Y .

Definition VII.1. Given a graph G and a partition (A, B) , we call G an (A, B) - λ -expander if

$$d_G(v, B) \geq \lambda n/30 \text{ for all } v \in A,$$

and

$$|E(X, Y)| \geq \lambda |X||Y|/10$$

whenever $X \subseteq A$, $|X| \geq \lambda n/100$, and $Y \subseteq B$, $|Y| \geq n/6$, and both statements hold also with A, B swapped.

We use Lemma 5.3 from [32].

Lemma VII.2. There exists $C > 0$ so that if $\lambda \geq C/\sqrt{n}$ the following holds. Let (A, B) be a weakly balanced partition, and let $(S, T) \in \mathcal{D}_{A, B, \lambda}^w$. Sample $E_{\text{cr}} \subseteq A \times B$ according to the hard-core model on $S \square T$ at activity λ . Let G be the graph $([n], E_{\text{cr}})$. Then

$$\mathbb{P}(G \text{ is an } (A, B)\text{-}\lambda\text{-expander}) \geq 1 - e^{-\sqrt{n}}.$$

For a weakly balanced partition (A, B) , let $\mathcal{E}_{A, B, \lambda}$ denote the set of $G \in \mathcal{T}(n)$ such that $(G[A], G[B]) \in \mathcal{D}_{A, B, \lambda}^w$ and G is an (A, B) - λ -expander. Let

$$\mathcal{E}_0 := \mathcal{T}(n) \setminus \bigcup_{\substack{(A, B) \\ \text{weakly balanced}}} \mathcal{E}_{A, B, \lambda}.$$

Let $p = \lambda/(1 + \lambda)$. We will use a conductance argument using $\mathcal{E}_0$ as a ‘bottleneck’ to lower bound the mixing time of the Glauber dynamics of $\mu_{\mathcal{T}, p}$.

If $G \in \mathcal{E}_0$ then either

- 1) $(G[A], G[B]) \notin \mathcal{D}_{A, B, \lambda}^w$ for all weakly balanced (A, B) , or
- 2) $(G[A], G[B]) \in \mathcal{D}_{A, B, \lambda}^w$ for some weakly balanced (A, B) but G is not an (A, B) - λ -expander.

Let $\mathcal{E}_0^1, \mathcal{E}_0^2$ denote the set of $G \in \mathcal{E}_0$ satisfying (1), (2) respectively. Proposition IV.3 shows that $\mu_{\mathcal{T}, p}(\mathcal{E}_0^1) = O(e^{-\sqrt{n}})$ and Proposition IV.3 taken together with Lemma VII.2 shows that $\mu_{\mathcal{T}, p}(\mathcal{E}_0^2) = O(e^{-\sqrt{n}})$. We conclude that

$$\mu_{\mathcal{T}, p}(\mathcal{E}_0) = O(e^{-\sqrt{n}}).$$

Note that by symmetry, if $(A, B), (A', B')$ are two partitions such that $\|A\| - \|B\| = \|A'\| - \|B'\|$ then $\mu_{\mathcal{T}, p}(\mathcal{E}_{A, B, \lambda}) =$

$\mu_{\mathcal{T}, p}(\mathcal{E}_{A', B', \lambda})$. Moreover, if (A, B) is weakly balanced then there are at least $\binom{n}{n/2 - n/10} = e^{\Omega(n)}$ other partitions with the same imbalance. It follows that

$$\mu_{\mathcal{T}, p}(\mathcal{E}_{A, B, \lambda}) \leq e^{-\Omega(n)}$$

for all weakly balanced (A, B) . It follows that we can take a union of sets $\mathcal{E}_{A, B, \lambda}$, call it $\mathcal{E}_1$, so that

$$\frac{1}{4} \leq \mu_{\mathcal{T}, p}(\mathcal{E}_1) \leq \frac{1}{2}.$$

We will show that for the Glauber dynamics to leave the set $\mathcal{E}_1$ it must pass through $\mathcal{E}_0$. By a standard conductance argument (see [63, Claim 2.3]) this implies that the mixing time is at least

$$\frac{\mu_{\mathcal{T}, p}(\mathcal{E}_1)}{8\mu_{\mathcal{T}, p}(\mathcal{E}_0)} \geq \Omega(e^{\sqrt{n}}).$$

It remains to show that for the Glauber dynamics to leave the set $\mathcal{E}_1$ it must pass through $\mathcal{E}_0$. Since the Glauber dynamics updates a single edge at a time, it suffices to show that if $G_1 \in \mathcal{E}_{A, B, \lambda}$ and $G_2 \in \mathcal{E}_{A', B', \lambda}$ for distinct weakly balanced partitions $(A, B), (A', B')$ then $|E(G_1) \Delta E(G_2)| \geq 2$. This is what we show now.

Since the partitions $(A, B), (A', B')$ are distinct, either $A \cap B' \neq \emptyset$ or $B \cap A' \neq \emptyset$. Assume WLOG that $B \cap A' \neq \emptyset$.

Suppose first that $|A \cap B'| < \lambda n/100$. By assumption there exists $v \in B \cap A'$. Since G_1 is an (A, B) - λ -expander we then have

$$\begin{aligned} d_{G_1}(v, A') &\geq d_{G_1}(v, A) - d_{G_1}(v, A \cap B') \\ &\geq \lambda n/30 - |A \cap B'| \\ &\geq \lambda n/50. \end{aligned}$$

On the other hand, since $G_2 \in \mathcal{D}_{A', B', \lambda}^w$, $d_{G_2}(v, A') \leq \alpha/\lambda$ and so

$$|E(G_1) \Delta E(G_2)| \geq \lambda n/50 - \alpha/\lambda \geq 2.$$

We may therefore assume that $|A \cap B'| \geq \lambda n/100$. In particular $A \cap B' \neq \emptyset$, and so by a similar argument to the above, we may assume that $|A' \cap B| \geq \lambda n/100$ also. By symmetry (swapping the roles of A and B) we may also assume that $|A \cap A'| \geq \lambda n/100$ and $|B \cap B'| \geq \lambda n/100$.

Since (A, B) is weakly balanced we have $|B| \geq n/3$. Suppose WLOG that $|B \cap B'| \geq |B \cap A'|$ so that in particular $|B \cap B'| \geq n/6$. Since G_1 is an (A, B) - λ -expander we then have

$$|E_{G_1}(A \cap B', B \cap B')| \geq \lambda |A \cap B'| |B \cap B'|/10, \quad (27)$$

and so there exists $v \in A \cap B'$ such that $d_{G_1}(v, B \cap B') \geq \lambda |B \cap B'|/10 \geq \lambda n/60$. On the other hand, since $G_2 \in \mathcal{D}_{A', B', \lambda}^w$, $d_{G_2}(v, B') \leq \alpha/\lambda$ and so

$$|E(G_1) \Delta E(G_2)| \geq \lambda n/60 - \alpha/\lambda \geq 2.$$

This concludes the proof. $\square$

ACKNOWLEDGMENTS

The authors thank Wojciech Samotij for very helpful conversations on this topic. MJ is supported by a UK Research and Innovation Future Leaders Fellowship MR/W007320/2. WP supported in part by NSF grant CCF-2309708. AP is supported by an NSERC Discovery grant.

REFERENCES

- [1] W. Mantel, “Problem 28,” *Wiskundige Opgaven*, vol. 10, no. 60-61, p. 320, 1907.
- [2] P. Erdős, D. Kleitman, and B. Rothschild, “Asymptotic enumeration of K_n -free graphs,” *Colloquio Internazionale sulle Teorie Combinatorie (Rome, 1973)*, no. 17, pp. 19–27, 1973.
- [3] P. Erdős and A. Rényi, “On the evolution of random graphs,” *Publ. Math. Inst. Hung. Acad. Sci.*, vol. 5, no. 1, pp. 17–60, 1960.
- [4] S. Chatterjee and S. S. Varadhan, “The large deviation principle for the Erdős-Rényi random graph,” *European Journal of Combinatorics*, vol. 32, no. 7, pp. 1000–1017, 2011.
- [5] S. Chatterjee, “Large deviations for random graphs,” *Lecture Notes in Mathematics*, vol. 2197, 2017.
- [6] J. Balogh, R. Morris, and W. Samotij, “Independent sets in hypergraphs,” *Journal of the American Mathematical Society*, vol. 28, no. 3, pp. 669–709, 2015.
- [7] D. Saxton and A. Thomason, “Hypergraph containers,” *Inventiones mathematicae*, vol. 201, no. 3, pp. 925–992, 2015.
- [8] S. Chatterjee and A. Dembo, “Nonlinear large deviations,” *Advances in Mathematics*, vol. 299, pp. 396–450, 2016.
- [9] R. Eldan, “Gaussian-width gradient complexity, reverse log-Sobolev inequalities and nonlinear large deviations,” *Geometric and Functional Analysis*, vol. 28, no. 6, pp. 1548–1596, 2018.
- [10] N. Cook and A. Dembo, “Large deviations of subgraph counts for sparse Erdős-Rényi graphs,” *Advances in Mathematics*, vol. 373, p. 107289, 2020.
- [11] G. Kozma and W. Samotij, “Lower tails via relative entropy,” *The Annals of Probability*, vol. 51, no. 2, pp. 665–698, 2023.
- [12] D. Osthus, H. J. Prömel, and A. Taraz, “For which densities are random triangle-free graphs almost surely bipartite?” *Combinatorica*, vol. 23, no. 1, pp. 105–150, 2003.
- [13] T. Łuczak, “On triangle-free random graphs,” *Random Structures & Algorithms*, vol. 16, no. 3, pp. 260–276, 2000.
- [14] S. Janson, T. Łuczak, and A. Ruciński, “An exponential bound for the probability of nonexistence of a specified subgraph in a random graph,” in *Random graphs 87, Poznań*, J. J. M. Karonński and A. Ruciński, Eds., 1987, pp. 73–87.
- [15] D. Weitz, “Counting independent sets up to the tree threshold,” in *Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing*, ser. STOC 2006. ACM, 2006, pp. 140–149.
- [16] N. Anari, K. Liu, and S. O. Gharan, “Spectral independence in high-dimensional expanders and applications to the hardcore model,” *SIAM Journal on Computing*, no. 0, pp. FOCSS20–1, 2021.
- [17] Z. Chen, K. Liu, and E. Vigoda, “Optimal mixing of Glauber dynamics: Entropy factorization via high-dimensional expansion,” in *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, 2021, pp. 1537–1550.
- [18] N. Anari, V. Jain, F. Koehler, H. T. Pham, and T.-D. Vuong, “Entropic independence: optimal mixing of down-up random walks,” in *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, 2022, pp. 1418–1430.
- [19] Y. Chen and R. Eldan, “Localization schemes: A framework for proving mixing bounds for Markov chains,” in *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2022, pp. 110–122.
- [20] A. Sly, “Computational transition at the uniqueness threshold,” in *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*. Las Vegas, NV, USA: IEEE, Oct. 2010, pp. 287–296.
- [21] A. Sly and N. Sun, “Counting in two-spin models on d -regular graphs,” *Annals of Probability*, vol. 42, no. 6, pp. 2383–2416, 2014.
- [22] A. Galanis, D. Štefankovič, and E. Vigoda, “Inapproximability of the partition function for the antiferromagnetic Ising and hard-core models,” *Combinatorics, Probability and Computing*, vol. 25, no. 4, pp. 500–559, Jul. 2016.
- [23] E. Mossel, D. Weitz, and N. Wormald, “On the hardness of sampling independent sets beyond the tree threshold,” *Probability Theory and Related Fields*, vol. 143, no. 3-4, pp. 401–439, 2009.
- [24] J. Hermon, A. Sly, and Y. Zhang, “Rapid mixing of hypergraph independent sets,” *Random Structures & Algorithms*, vol. 54, no. 4, pp. 730–767, 2019.
- [25] I. Bezáková, A. Galanis, L. A. Goldberg, H. Guo, and D. Štefankovič, “Approximation via correlation decay when strong spatial mixing fails,” *SIAM Journal on Computing*, vol. 48, no. 2, pp. 279–349, 2019.
- [26] D. Galvin, G. McKinley, W. Perkins, M. Sarantis, and P. Tetali, “On the zeroes of hypergraph independence polynomials,” *Combinatorics, Probability and Computing*, vol. 33, no. 1, pp. 65–84, 2024.
- [27] F. Bencs and P. Buys, “Optimal zero-free regions for the independence polynomial of bounded degree hypergraphs,” *arXiv preprint arXiv:2306.00122*, 2023.
- [28] H. J. Prömel and A. Steger, “On the asymptotic structure of sparse triangle free graphs,” *Journal of Graph Theory*, vol. 21, no. 2, pp. 137–151, 1996.
- [29] N. C. Wormald, “The perturbation method and triangle-free random graphs,” *Random Structures & Algorithms*, vol. 9, no. 1-2, pp. 253–270, 1996.
- [30] D. Stark and N. Wormald, “The probability of non-existence of a subgraph in a moderately sparse random graph,” *Combinatorics, Probability and Computing*, vol. 27, no. 4, pp. 672–715, 2018.
- [31] F. Mousset, A. Noever, K. Panagiotou, and W. Samotij, “On the probability of nonexistence in binomial subsets,” *Annals of Probability*, vol. 48, no. 1, pp. 493–525, 2020.
- [32] M. Jenssen, W. Perkins, and A. Potukuchi, “On the evolution of structure in triangle-free graphs,” *arXiv preprint arXiv:2312.09202*, 2023.
- [33] S. Friedli and Y. Velenik, *Statistical Mechanics of Lattice Systems: a Concrete Mathematical Introduction*. Cambridge University Press, 2017.
- [34] M. Jenssen, W. Perkins, A. Potukuchi, and M. Simkin, “Lower tails for triangles inside the critical window,” *arXiv preprint arXiv:2411.18563*, 2024.
- [35] R. Bubbly and M. Dyer, “Path coupling: A technique for proving rapid mixing in Markov chains,” in *Proceedings 38th Annual Symposium on Foundations of Computer Science*. IEEE, 1997, pp. 223–231.
- [36] M. Dyer and A. Frieze, “Randomly coloring graphs with lower bounds on girth and maximum degree,” *Random Structures & Algorithms*, vol. 23, no. 2, pp. 167–179, 2003.
- [37] T. Helmuth, W. Perkins, and G. Regts, “Algorithmic Pirogov-Sinai theory,” *Probability Theory and Related Fields*, vol. 176, pp. 851–895, 2020.
- [38] M. Jenssen, P. Keevash, and W. Perkins, “Algorithms for #BIS-hard problems on expander graphs,” *SIAM Journal on Computing*, vol. 49, no. 4, pp. 681–710, 2020.
- [39] Z. Chen, A. Galanis, L. A. Goldberg, W. Perkins, J. Stewart, and E. Vigoda, “Fast algorithms at low temperatures via Markov chains,” *Random Structures & Algorithms*, vol. 58, no. 2, pp. 294–321, 2021.
- [40] A. Galanis, L. A. Goldberg, and J. Stewart, “Fast algorithms for general spin systems on bipartite expanders,” *ACM Transactions on Computation Theory (TOCT)*, vol. 13, no. 4, pp. 1–18, 2021.
- [41] —, “Fast mixing via polymers for random graphs with unbounded degree,” *Information and Computation*, vol. 285, p. 104894, 2022.
- [42] T. Helmuth, M. Jenssen, and W. Perkins, “Finite-size scaling, phase coexistence, and algorithms for the random cluster model on random graphs,” in *Annales de l’Institut Henri Poincaré (B) Probabilités et statistiques*, vol. 59, no. 2. Institut Henri Poincaré, 2023, pp. 817–848.
- [43] C. Carlson, E. Davies, N. Fraiman, A. Kolla, A. Potukuchi, and C. Yap, “Algorithms for the ferromagnetic Potts model on expanders,” in *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2022, pp. 344–355.
- [44] Z. Chen, A. Galanis, D. Štefankovič, and E. Vigoda, “Sampling colorings and independent sets of random regular bipartite graphs in the non-uniqueness region,” in *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2022, pp. 2198–2207.
- [45] M. Dyer, L. A. Goldberg, C. Greenhill, and M. Jerrum, “The relative complexity of approximate counting problems,” *Algorithmica*, vol. 38, no. 3, pp. 471–500, 2004.
- [46] J.-Y. Cai, A. Galanis, L. A. Goldberg, H. Guo, M. Jerrum, D. Štefankovič, and E. Vigoda, “# BIS-hardness for 2-spin systems

- on bipartite bounded degree graphs in the tree non-uniqueness region,” *Journal of Computer and System Sciences*, vol. 82, no. 5, pp. 690–711, 2016.
- [47] R. Gheissari and A. Sinclair, “Low-temperature Ising dynamics with random initializations,” in *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, 2022, pp. 1445–1458.
 - [48] A. Blanca and R. Gheissari, “Sampling from the Potts model at low temperatures via Swendsen–Wang dynamics,” in *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2023, pp. 2006–2020.
 - [49] A. Galanis, L. A. Goldberg, and P. Smolarova, “Sampling from the random cluster model on random regular graphs at all temperatures via Glauber dynamics,” in *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2023)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2023.
 - [50] E. Davies, M. Jenssen, W. Perkins, and B. Roberts, “Independent sets, matchings, and occupancy fractions,” *Journal of the London Mathematical Society*, vol. 96, no. 1, pp. 47–66, 2017.
 - [51] —, “On the average size of independent sets in triangle-free graphs,” *Proceedings of the American Mathematical Society*, vol. 146, no. 1, pp. 111–124, 2018.
 - [52] M. Luczak, “Concentration of measure and mixing for Markov chains,” in *Discrete Mathematics & Theoretical Computer Science*, 2008, pp. 95–120.
 - [53] S. Janson, “Poisson approximation for large deviations,” *Random Structures & Algorithms*, vol. 1, no. 2, pp. 221–229, 1990.
 - [54] S. Janson and L. Warnke, “The lower tail: Poisson approximation revisited,” *Random Structures & Algorithms*, vol. 48, no. 2, pp. 219–246, 2016.
 - [55] R. Montenegro and P. Tetali, “Mathematical aspects of mixing times in Markov chains,” *Foundations and Trends in Theoretical Computer Science*, vol. 1, no. 3, pp. 237–354, 2006.
 - [56] D. A. Levin and Y. Peres, *Markov chains and mixing times*. American Mathematical Soc., 2017, vol. 107.
 - [57] A. D. Scott and A. D. Sokal, “The repulsive lattice gas, the independent-set polynomial, and the Lovász local lemma,” *Journal of Statistical Physics*, vol. 118, no. 5-6, pp. 1151–1261, 2005.
 - [58] W. G. Faris, “Combinatorics and cluster expansions,” *Probability Surveys*, vol. 7, pp. 157–206, 2010.
 - [59] D. Štefankovič, S. Vempala, and E. Vigoda, “Adaptive simulated annealing: A near-optimal connection between sampling and counting,” *Journal of the ACM (JACM)*, vol. 56, no. 3, pp. 1–36, 2009.
 - [60] O. Penrose, “Convergence of fugacity expansions for classical systems,” *Statistical Mechanics: Foundations and Applications*, p. 101, 1967.
 - [61] M. R. Jerrum, L. G. Valiant, and V. V. Vazirani, “Random generation of combinatorial structures from a uniform distribution,” *Theoretical Computer Science*, vol. 43, pp. 169–188, 1986.
 - [62] J. Balogh, R. Morris, W. Samotij, and L. Warnke, “The typical structure of sparse K_{r+1} -free graphs,” *Transactions of the American Mathematical Society*, vol. 368, no. 9, pp. 6439–6485, 2016.
 - [63] M. Dyer, A. Frieze, and M. Jerrum, “On counting independent sets in sparse graphs,” *SIAM Journal on Computing*, vol. 31, no. 5, pp. 1527–1541, 2002.