

On the average size of independent sets in triangle-free graphs

Davies, Ewan; Jenssen, Matthew; Perkins, Will; Roberts, Barnaby

DOI:

[10.1090/proc/13728](https://doi.org/10.1090/proc/13728)

License:

Creative Commons: Attribution-NonCommercial-NoDerivs (CC BY-NC-ND)

Document Version

Peer reviewed version

Citation for published version (Harvard):

Davies, E, Jenssen, M, Perkins, W & Roberts, B 2018, 'On the average size of independent sets in triangle-free graphs', *Proceedings of the American Mathematical Society*, vol. 146, no. 1, pp. 111-124.
<https://doi.org/10.1090/proc/13728>

[Link to publication on Research at Birmingham portal](#)

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

ON THE AVERAGE SIZE OF INDEPENDENT SETS IN TRIANGLE-FREE GRAPHS

EWAN DAVIES, MATTHEW JENSSEN, WILL PERKINS, AND BARNABY ROBERTS

ABSTRACT. We prove an asymptotically tight lower bound on the average size of independent sets in a triangle-free graph on n vertices with maximum degree d , showing that an independent set drawn uniformly at random from such a graph has expected size at least $(1 + o_d(1))\frac{\log d}{d}n$. This gives an alternative proof of Shearer's upper bound on the Ramsey number $R(3, k)$. We then prove that the total number of independent sets in a triangle-free graph with maximum degree d is at least $\exp\left[\left(\frac{1}{2} + o_d(1)\right)\frac{\log^2 d}{d}n\right]$. The constant $1/2$ in the exponent is best possible. In both cases, tightness is exhibited by a random d -regular graph.

Both results come from considering the hard-core model from statistical physics: a random independent set I drawn from a graph with probability proportional to $\lambda^{|I|}$, for a fugacity parameter $\lambda > 0$. We prove a general lower bound on the occupancy fraction (normalized expected size of the random independent set) of the hard-core model on triangle-free graphs of maximum degree d . The bound is asymptotically tight in d for all $\lambda = O_d(1)$.

We conclude by stating several conjectures on the relationship between the average and maximum size of an independent set in a triangle-free graph and give some consequences of these conjectures in Ramsey theory.

1. INTRODUCTION

1.1. Independent sets in triangle-free graphs. Ajtai, Komlós, and Szemerédi [1] proved that any triangle-free graph G on n vertices with average degree d has an independent set of size at least $c\frac{\log d}{d}n$ for a small constant c . Shearer [23] later improved the constant to 1, asymptotically as $d \rightarrow \infty$, showing that such a graph has an independent set of size at least $f(d) \cdot n$ where $f(d) = \frac{d \log d - d + 1}{(d-1)^2} = (1 + o_d(1))\frac{\log d}{d}$. Here, and in what follows, logarithms will always be to base e . We use standard asymptotic notation, with subscripts indicating which parameters the implied functions depend on. For example we write $o_d(1)$ for a quantity that tends to zero as d tends to infinity.

The off-diagonal Ramsey number $R(3, k)$ is the least integer n such that any graph on n vertices contains either a triangle or an independent set of size k . The above result of Ajtai, Komlós, and Szemerédi and a result of Kim [17] show that $R(3, k) = \Theta(k^2/\log k)$. Shearer's result gives the current best upper bound, showing that $R(3, k) \leq (1 + o(1))k^2/\log k$. Independent work of Bohman and Keevash [6] and Fiz Pontiveros, Griffiths, and Morris [12] shows that $R(3, k) \geq (1/4 + o(1))k^2/\log k$. Reducing the factor 4 gap between these bounds is a major open problem in Ramsey theory.

We prove a lower bound on the *average size* of an independent set in a triangle-free graph of maximum degree d , matching the asymptotic form of Shearer's result, and in turn giving an alternative proof of the above upper bound on $R(3, k)$.

Date: March 6, 2017.

Theorem 1. *Let G be a triangle-free graph on n vertices with maximum degree d . Let $\mathcal{I}(G)$ be the set of all independent sets of G . Then*

$$\frac{1}{|\mathcal{I}(G)|} \sum_{I \in \mathcal{I}(G)} |I| \geq (1 + o_d(1)) \frac{\log d}{d} n.$$

Moreover, the constant ‘1’ is best possible.

This result is weaker than Shearer’s [23] in that instead of average degree d we require maximum degree d . Our result is stronger in that we show that the *average size* of an independent set from such a graph is of size at least $(1 + o_d(1)) \frac{\log d}{d} n$, while Shearer shows the largest independent set is of at least this size (by analyzing a randomized greedy algorithm).

The proof of Theorem 1 is an adaptation of the authors’ method used in [11] to give a strengthening of results of Kahn [15] and Zhao [26] that a d -regular graph has at most as many independent sets as a disjoint union of $K_{d,d}$ ’s on the same number of vertices. We make this explicit in Section 2, giving a unified method for proving these classical results of Shearer and Kahn. After sharing a draft of this paper with colleagues, we discovered that James Shearer also knew the proof of the lower bound in Theorem 1 and presented a sketch of it at the SIAM Conference on Discrete Mathematics in 1998, but never published it [22].

To see that Theorem 1 directly implies the upper bound $(1 + o(1))k^2 / \log k$ on $R(3, k)$, suppose that G is triangle free with no independent set of size k . Then G must have maximum degree less than k . Applying Theorem 1 we see the independence number is at least $(1 + o_k(1)) \frac{\log k}{k} n$ but less than k , and so $n < (1 + o_k(1)) \frac{k^2}{\log k}$ as required. Of course this reasoning simply uses the average size of an independent set as a lower bound for the maximum size. In Section 5 we consider whether the discrepancy between the maximum and average size can be exploited to improve the upper bound on $R(3, k)$.

After bounding the average size of an independent set from G , we next give lower bounds on the total number of independent sets in G , $|\mathcal{I}(G)|$.

Theorem 2. *Let G be a triangle-free graph on n vertices with maximum degree d . Then*

$$|\mathcal{I}(G)| \geq e^{\left(\frac{1}{2} + o_d(1)\right) \frac{\log^2 d}{d} n}.$$

Moreover, the constant 1/2 in the exponent is best possible.

In comparison, improving on previous results of Cooper and Mubayi [9], Cooper, Dutta, and Mubayi [8] proved that any triangle-free graph of average degree d has at least $e^{\left(\frac{1}{4} + o_d(1)\right) \frac{\log^2 d}{d} n}$ independent sets.

As a simple corollary we get the following lower bound without degree restrictions.

Corollary 1. *Let G be a triangle-free graph on n vertices. Then*

$$|\mathcal{I}(G)| \geq e^{\left(\frac{\sqrt{2 \log 2}}{4} + o(1)\right) \sqrt{n} \log n}.$$

This improves on a result of Cooper, Dutta, and Mubayi [8] by a factor of $\sqrt{2}$ in the exponent; they also provide a construction based on the analysis of the triangle-free process in [6, 12] showing that the optimal constant is at most $1 + \log 2 \approx 1.693$ (as compared to the constant $\frac{\sqrt{2 \log 2}}{4} \approx .294$ in Corollary 1).

In Section 4 we show that the constants ‘1’ in Theorem 1 and ‘1/2’ in Theorem 2 are best possible, using the random d -regular graph (with d fixed as the number of vertices tends to infinity) as an example. It remains an interesting open question to determine whether these constants are tight when the degree of the graph grows with the number of vertices; for example, if $d = n^c$ for some $c \in (0, 1)$.

1.2. The hard-core model. Theorems 1 and 2 are special cases of results on a more general model from statistical physics; the *hard-core* distribution on the independent sets \mathcal{I} of a graph G . For more on the hard-core model and its links with statistical physics and combinatorics, see [7]. The distribution depends on a fugacity parameter $\lambda > 0$ and is given by

$$\Pr[I] = \frac{\lambda^{|I|}}{\sum_{J \in \mathcal{I}} \lambda^{|J|}}.$$

The denominator, $P_G(\lambda) = \sum_{J \in \mathcal{I}} \lambda^{|J|}$, is the *partition function* of the hard-core model on G (also called the *independence polynomial* of G). If $\lambda = 1$, the partition function $P_G(1)$ is the total number of independent sets and the hard-core distribution is simply the uniform distribution over all independent sets of G .

In what follows G will always be a graph on n vertices. We write $\alpha(G)$ for the size of the largest independent set in G . The expected size of an independent set drawn from the hard-core model on G at fugacity λ is

$$(1) \quad \bar{\alpha}_G(\lambda) := \sum_{I \in \mathcal{I}} |I| \cdot \Pr[I] = \frac{\sum_{I \in \mathcal{I}} |I| \lambda^{|I|}}{P_G(\lambda)} = \frac{\lambda P'_G(\lambda)}{P_G(\lambda)} = \lambda \cdot (\log P_G(\lambda))'.$$

The key result of this paper is the following general lower bound on $\bar{\alpha}_G(\lambda)$ for triangle-free graphs. The lower bound is written naturally in terms of the Lambert W function, $W(z)$: for $z > 0$, $W(z)$ denotes the unique positive real satisfying the relation $W(z)e^{W(z)} = z$. It will be useful to note that for $z \geq e$ we have $W(z) \geq \log z - \log \log z$.

Theorem 3. *Let G be a triangle-free graph with maximum degree d . Then for any $\lambda > 0$,*

$$(2) \quad \frac{1}{n} \bar{\alpha}_G(\lambda) \geq \frac{\lambda}{1 + \lambda} \frac{W(d \log(1 + \lambda))}{d \log(1 + \lambda)}.$$

The quantity $\frac{1}{n} \bar{\alpha}_G(\lambda)$, the expected fraction of vertices of G in the random independent set, is known as the *occupancy fraction* of G at fugacity λ .

It is interesting to note that the lower bound in Theorem 3 is not monotone in λ whereas for any graph G , $\bar{\alpha}_G(\lambda)$ is monotone increasing (see Proposition 1). Simply substituting $\lambda = 1$ into (3), does not quite suffice to prove Theorem 1. Surprisingly it turns out to be better to use a smaller λ and then appeal to the monotonicity of $\bar{\alpha}_G(\lambda)$. As an example, using $\lambda = 1/\log d$ in (3) is enough to prove Theorem 1. This shows that Theorem 1 holds even when we replace the average size of an independent set with a weighted average biased toward small sets. In fact we can afford to bias using any λ of the form $d^{-o(1)}$. For smaller λ the lower bound in Theorem 3 becomes asymptotically weaker. For example taking $\lambda = d^{-s}$ where $s \in (0, 1)$ is fixed, (3) gives $\bar{\alpha}_G(\lambda) \geq (1 - s + o_d(1)) \frac{\log d}{d} n$. This lower bound is tight and we in fact show that Theorem 3 (extended by monotonicity) is asymptotically tight for all $\lambda = O_d(1)$ in Section 4.

Turning to Theorem 2 and the problem of counting independent sets, we again give a more general statement about the hard-core model. We bound the partition function of a

triangle-free graph at any fugacity λ using the fact that the occupancy fraction is the scaled logarithmic derivative of the partition function. Indeed from (1.2) it follows that

$$(3) \quad \frac{1}{n} \log P_G(\lambda) = \frac{1}{n} \int_0^\lambda \frac{\bar{\alpha}_G(t)}{t} dt.$$

Using the lower bound on the occupancy fraction from Theorem 3 in the integral above gives the lower bound on the partition function which we state below, noting that Theorem 2 is a direct consequence.

Theorem 4. *Let G be a triangle-free graph on n vertices with maximum degree d . Then for all $\lambda > 0$,*

$$P_G(\lambda) \geq \exp\left(\left[W(d \log(1 + \lambda))^2 + 2W(d \log(1 + \lambda))\right] \frac{n}{2d}\right).$$

In particular, taking $\lambda = 1$, we see that G has at least $e^{(\frac{1}{2} + o_d(1)) \frac{\log^2 d}{d} n}$ independent sets.

2. LOWER BOUNDS ON THE OCCUPANCY FRACTION

The proof of Theorem 3 begins in the same way as the triangle-free case of the proof in [11] that the occupancy fraction of the hard-core model on any d -regular graph is at most that of $K_{d,d}$. The main difference is that in the last step we ask for the minimum instead of the maximum of the same constrained optimization problem. The method is similar to that of Shearer [24] in lower bounding the independence number of K_r -free graphs and that of Alon [3] in lower bounding the independence number of graphs of average degree d in which neighborhoods are r -colorable. Alon and Spencer [4] used this technique to give a result like Theorem 1 but with a worse constant. See also chapters 21 and 22 of [19] for an exposition of some of these results, and [18] for a recent application to hypergraphs. Unlike in previous works, in this paper we use the full power of choosing the fugacity parameter λ , rather than considering only the uniform distribution.

There are two key steps: First, we define a random variable that depends on two sources of randomness: the random independent set drawn from the hard-core model on G and a uniformly chosen random vertex $v \in V(G)$. We then express the occupancy fraction in terms of two different expectations involving this random variable. This gives a constraint on the distribution of the random variable. We then optimize over all random variables that satisfy the constraint, and deduce a bound on the occupancy fraction.

Recall that $\bar{\alpha}_G(\lambda)$ is the expected size of an independent set drawn from the hard-core model on G at fugacity λ , and $\frac{1}{n} \bar{\alpha}_G(\lambda)$ is the occupancy fraction. In places where no confusion should arise we omit G from the notation.

Proof of Theorem 3. Let I be an independent set of G drawn from the hard-core model at fugacity λ . For any vertex v in G , consider $I_0 = I \setminus N(v)$ (noting that I_0 may include v itself). We say a vertex v is *occupied* if $v \in I$ and *unoccupied* otherwise. Furthermore we say v is *uncovered* if $N(v) \cap I = \emptyset$ and *covered* otherwise. Note that if v is covered, it must be unoccupied. We begin by stating some basic properties of the hard-core model.

Fact 1: $\Pr[v \in I | v \text{ uncovered}] = \frac{\lambda}{1 + \lambda}$.

This follows from the observation that for any realization of I_0 under the event that v is uncovered, there are exactly two possibilities for I : $I = I_0$ and $I = I_0 \cup \{v\}$, since v is uncovered and so may be added.

Fact 2: $\Pr[v \text{ uncovered} | v \text{ has } j \text{ uncovered neighbors}] = (1 + \lambda)^{-j}$.

First note that the event that v is uncovered is the same as the event that its uncovered neighbors are all unoccupied. The probability that an uncovered neighbor of v is unoccupied is $\frac{1}{1+\lambda}$ by Fact 1. Moreover, since G is triangle free the uncovered neighbors of v form an independent set and so they are in fact unoccupied independently with probability $\frac{1}{1+\lambda}$.

We now write the occupancy fraction as:

$$\begin{aligned} \frac{1}{n} \bar{\alpha}(\lambda) &= \frac{1}{n} \sum_{v \in G} \Pr[v \in I] \\ (4) \quad &= \frac{\lambda}{1 + \lambda} \cdot \frac{1}{n} \sum_{v \in G} \Pr[v \text{ uncovered}] \end{aligned}$$

$$(5) \quad = \frac{\lambda}{1 + \lambda} \cdot \frac{1}{n} \sum_{v \in G} \sum_{j=0}^d \Pr[v \text{ has } j \text{ uncovered neighbors}] \cdot (1 + \lambda)^{-j}$$

where (2) follows from Fact 1 and (2) from Fact 2. Now let Z be the random variable that counts the number of uncovered neighbors of a uniformly chosen vertex v . Z has two layers of randomness, that of I drawn from the hard-core measure, and that of selecting v at random. Interpreting the RHS of (2) in terms of Z , we obtain

$$(6) \quad \frac{1}{n} \bar{\alpha}(\lambda) = \frac{\lambda}{1 + \lambda} \mathbb{E}[(1 + \lambda)^{-Z}].$$

An alternative way to relate the occupancy fraction to Z is to observe that in the sum $\sum_v \sum_{u \sim v} \Pr[u \in I]$ each vertex u appears $\deg(u)$ times. Any uncovered neighbor of v is occupied with probability $\frac{\lambda}{1+\lambda}$ by Fact 1 and any covered neighbor is occupied with probability zero. Then since G has maximum degree d we have

$$(7) \quad \frac{1}{n} \bar{\alpha}(\lambda) = \frac{1}{n} \sum_{v \in G} \Pr[v \in I] \geq \frac{1}{dn} \sum_v \sum_{u \sim v} \Pr[u \in I] = \frac{\lambda}{1 + \lambda} \frac{\mathbb{E} Z}{d}.$$

We aim to minimize the occupancy fraction subject to the constraints on the distribution of Z given by (2) and (2). In fact we relax the optimization problem to optimize over all distributions of random variables Z that satisfy these constraints, not only those that arise from the hard-core model on a graph.

In the calculation below we show that Jensen's inequality applied to (2) implies that the minimizer is achieved by the unique *constant* random variable Z that satisfies the constraints with equality. From (2) and (2) we have

$$\frac{1 + \lambda}{\lambda} \cdot \frac{1}{n} \bar{\alpha}(\lambda) \geq \max \left\{ \frac{\mathbb{E} Z}{d}, (1 + \lambda)^{-\mathbb{E} Z} \right\} \geq \min_{x \in \mathbb{R}^+} \max \left\{ \frac{x}{d}, (1 + \lambda)^{-x} \right\}.$$

To compute the minimum observe that the first of our lower bounds is increasing in x , and the second is decreasing. Then the minimum occurs at the value of x which makes the two bounds equal i.e. the x that satisfies

$$x e^{\log(1+\lambda)x} = d$$

and hence

$$\log(1 + \lambda)x = W(d \log(1 + \lambda)).$$

The result follows. \square

A simple observation is the fact that lower bounds at a small fugacity λ imply the same bounds at higher fugacities.

Proposition 1. *For any graph G , the expected size of an independent set $\bar{\alpha}_G(\lambda)$ is monotone increasing in λ .*

Proof. We will show that the derivative of $\bar{\alpha}(\lambda)$ with respect to λ is positive. We write (using P for $P_G(\lambda)$)

$$\begin{aligned}\bar{\alpha}'(\lambda) &= \left(\frac{\lambda P'}{P}\right)' = \frac{P'}{P} + \frac{\lambda P P'' - \lambda(P')^2}{P^2} \\ &= \frac{P'}{P} + \frac{1}{\lambda} \left(\frac{\lambda^2 P''}{P} - \left(\frac{\lambda P'}{P}\right)^2\right) \\ &= \frac{\mathbb{E}(|I|) + \mathbb{E}(|I|^2) - \mathbb{E}(|I|) - (\mathbb{E}(|I|))^2}{\lambda} \\ &= \frac{\text{Var}(|I|)}{\lambda} \geq 0\end{aligned}$$

where I is a random independent set drawn from the hard-core model at fugacity λ . \square

Now using Theorem 3 and Proposition 1 we prove the main statement of Theorem 1. We leave the proof of tightness to Section 4.

Proof of Theorem 1. Substituting $\lambda = 1/\log d$ in (3) and recalling the bound $W(z) \geq \log z - \log \log z$ for $z \geq e$ we obtain

$$\frac{1}{n} \bar{\alpha}_G(\lambda) \geq (1 + o(1)) \frac{\log d}{d}.$$

The result then follows from the monotonicity given by Proposition 1. \square

So far we have been concerned with lower bounds on the occupancy fraction of a triangle-free graph of maximum degree d . In [11] the current authors considered upper bounds, showing that the occupancy fraction of a d -regular graph is at most that of $K_{d,d}$. The proof method is essentially identical to the proof of Theorem 3. To make this connection explicit, we give a proof of the upper bound result here, restricted to the triangle-free case.

Theorem 5 ([11]). *Let G be a d -regular triangle-free graph, then for any $\lambda > 0$*

$$\frac{1}{n} \bar{\alpha}_G(\lambda) \leq \frac{1}{2d} \bar{\alpha}_{K_{d,d}}(\lambda)$$

with equality only if G is a disjoint union of copies of $K_{d,d}$.

Proof. Note that since G is triangle free, equation (2) holds for G and since G is d -regular, (2) holds with equality throughout. That is we have

$$(8) \quad \frac{1}{n} \bar{\alpha}_G(\lambda) = \frac{\lambda}{1+\lambda} \frac{\mathbb{E} Z}{d} = \frac{\lambda}{1+\lambda} \mathbb{E}[(1+\lambda)^{-Z}],$$

where we recall that Z is a random variable bounded between 0 and d . Now instead of asking for the minimum value of $\frac{1}{n} \bar{\alpha}_G(\lambda)$ over all distributions of Z as we did in the proof of Theorem 3, we ask instead for the maximum. Note that since $0 \leq Z/d \leq 1$ convexity of the function $x \mapsto (1+\lambda)^{-x}$ implies that

$$(9) \quad (1+\lambda)^{-Z} \leq \frac{Z}{d} (1+\lambda)^{-d} + 1 - \frac{Z}{d}.$$

Substituting this into (2) and using linearity of expectation yields

$$(10) \quad \frac{1}{n} \bar{\alpha}_G(\lambda) = \frac{\lambda}{1+\lambda} \frac{\mathbb{E} Z}{d} \leq \frac{\lambda(1+\lambda)^{d-1}}{2(1+\lambda)^d - 1},$$

where one can check that the right hand side is the occupancy fraction of $K_{d,d}$. For uniqueness, note that to have equality in (2) we must have had equality in (2) which is only possible if Z takes only the values 0 and d . This distribution of Z can only occur in a disjoint union of copies of $K_{d,d}$. To see this recall that Z is the number of uncovered neighbors of a randomly selected vertex v . The only way every vertex v can always have either 0 or d uncovered neighbors is for all the neighbors of v to have the same neighborhood. For d -regular graphs this property holds only in graphs consisting disjoint unions of $K_{d,d}$. \square

Note that by the integral in (1.2), Theorem 5 immediately implies that for a triangle-free graph G , $\frac{1}{n} \log P_G(\lambda) \leq \frac{1}{2d} \log P_{K_{d,d}}(\lambda)$. Taking $\lambda = 1$ implies that G has at most as many independent sets as a disjoint union of copies of $K_{d,d}$ on the same number of vertices. This last statement was famously proved by Kahn [15] (in the case where G is bipartite) using the entropy method.

3. COUNTING INDEPENDENT SETS

In this section we prove the main statement of Theorem 4 (and hence Theorem 2) by integrating the lower bound on the occupancy fraction of a triangle-free graph given in Theorem 3. Again, the proof of tightness is deferred to Section 4. We also prove Corollary 1.

Proof of Theorem 4. By (1.2) and Theorem 3 we have

$$(11) \quad \begin{aligned} \log P_G(\lambda) &\geq \frac{n}{d} \int_0^\lambda \frac{W(d \log(1+t))}{(1+t) \log(1+t)} dt \\ &= \frac{n}{d} \int_0^{W(d \log(1+\lambda))} \frac{1}{(1+u)} du \\ &= \frac{n}{2d} [W(d \log(1+\lambda))^2 + 2W(d \log(1+\lambda))]. \end{aligned}$$

where for the first equality we used the substitution $u = W(d \log(1+t))$. In particular when $\lambda = 1$, using the inequality $W(z) \geq \log z - \log \log z$ for $z \geq e$, we have

$$\log P_G(\lambda) \geq \left(\frac{1}{2} + o_d(1) \right) \frac{\log^2 d}{d} n. \quad \square$$

Proof of Corollary 1. In a triangle-free graph the neighborhood of any vertex forms an independent set. Let d be the largest degree of a vertex in G , then we have the bound

$$P_G(\lambda) \geq \max \left\{ (1+\lambda)^d, \exp \left[\frac{n}{2d} W(d \log(1+\lambda))^2 \right] \right\},$$

by considering the neighborhood of a vertex of maximum degree and by inequality (3). The first expression is increasing in d while the second is decreasing, and at

$$d = \frac{1}{2} \sqrt{\frac{n}{2 \log(1+\lambda)}} \log \left(\frac{n \log(1+\lambda)}{2} \right)$$

they are equal. It follows that for $\lambda > 0$,

$$(12) \quad P_G(\lambda) \geq \exp \left[\frac{1}{2} \sqrt{\frac{n \log(1 + \lambda)}{2}} \log \left(\frac{n \log(1 + \lambda)}{2} \right) \right].$$

Take $\lambda = 1$ to complete the proof. \square

Inequality (3) may be of independent interest, giving a general lower bound for the independence polynomial of a triangle-free graph on n vertices.

4. RANDOM REGULAR GRAPHS

Here we will show that the constants in Theorems 1 to 4 are tight, as evidenced by the random d -regular graph (conditioned on having no triangles).

To do this, we first leave the world of finite graphs and consider the hard-core model on the infinite d -regular tree T_d . The hard-core model can be defined on a infinite graph by taking the limit of hard-core models on a growing sequence of finite graphs (in this case d -regular trees of increasing depth) with some boundary conditions specified. There is a unique *translation invariant* hard-core measure on the infinite d -regular tree (see e.g. [16]), and under this measure the probability that any given vertex is in the independent set is $\alpha_{T_d}(\lambda)$, where $\alpha_{T_d}(\lambda)$ is the solution of the equation

$$(13) \quad \lambda = \frac{\alpha}{(1 - \alpha)} \left(\frac{1 - \alpha}{1 - 2\alpha} \right)^d.$$

Using α_{T_d} to denote the occupancy fraction of the infinite d -regular tree is standard, however we note that our notation for the analogous quantity $\frac{1}{n} \bar{\alpha}_G(\lambda)$ in a finite graph G graph on n vertices includes the scaling $1/n$. We compare the lower bound in Theorem 3 to $\alpha_{T_d}(\lambda)$:

Proposition 2. *For every triangle-free graph G of maximum degree d and any $\lambda = O_d(1)$,*

$$\frac{1}{n} \bar{\alpha}_G(\lambda) \geq (1 + o_d(1)) \alpha_{T_d}(\lambda).$$

In the language of statistical physics, $\alpha_{T_d}(\lambda)$ is the *replica symmetric* occupancy fraction of a d -regular graph of large girth. Extremality of a replica symmetric solution has been proved in several contexts including [10, 21, 25] but usually requires a condition such as an attractive potential function (e.g. the ferromagnetic Ising model) or bipartiteness in the case of the hard-core or monomer-dimer models. Proposition 2 states that in fact under very minimal conditions (triangle free), the replica symmetric solution is a lower bound to the true occupancy fraction for all $\lambda = O_d(1)$, at least asymptotically in d . This result can be compared with [11, Theorem 2] where we used stronger assumptions, showing that a d -regular, vertex-transitive bipartite graph has occupancy fraction strictly greater than $\alpha_{T_d}(\lambda)$, though we only believe that result is best possible for λ below the *uniqueness threshold* on the tree, $\lambda \leq \frac{(d-1)^{d-1}}{(d-2)^d} = \Theta\left(\frac{1}{d}\right)$, in the sense that for any $\epsilon > 0$, there exists some finite graph G with $\alpha_G(\lambda) < \alpha_{T_d}(\lambda) + \epsilon$.

Proof of Proposition 2. In what follows all asymptotic notation refers to the limit $d \rightarrow \infty$. We first show that $\alpha_{T_d}(\lambda) = (1 + o(1)) \frac{W(\lambda d)}{d}$. We then show that for $\lambda = o(1)$, our bound on $\bar{\alpha}_G(\lambda)$ for triangle-free graphs given by Theorem 3 is asymptotically $(1 + o(1)) \frac{W(\lambda d)}{d}$. We

then appeal to monotonicity (Proposition 1) to extend this lower bound to all $\lambda = O(1)$. Letting $z = \frac{\alpha_{T_d}(\lambda)}{1-2\alpha_{T_d}(\lambda)}d$, equation (4) defining $\alpha_{T_d}(\lambda)$ becomes

$$(14) \quad z \left(1 + \frac{z}{d}\right)^{d-1} = \lambda d.$$

First note that from (4) it follows easily that $z = O(\log d)$ and so $\left(1 + \frac{z}{d}\right)^{d-1} = (1 + o(1))e^z$. It then follows that $z = W(\lambda d + o(\lambda d))$. Using the fact that $\frac{dW}{dx} = \frac{W(x)}{x(1+W(x))} \leq \frac{W(x)}{x}$ along with the Mean Value Theorem, we deduce

$$z = (1 + o(1))W(\lambda d).$$

By the definition of z we then have

$$(15) \quad \alpha_{T_d}(\lambda) = \frac{z}{d} \left(1 + 2\frac{z}{d}\right)^{-1} = \frac{z}{d} + O\left(\frac{z^2}{d^2}\right) = (1 + o(1))\frac{W(d\lambda)}{d}.$$

Let us now suppose that $\lambda = o(1)$. By Theorem 3 we then have

$$\frac{1}{n}\bar{\alpha}_G(\lambda) \geq \frac{\lambda}{1+\lambda} \frac{W(d\log(1+\lambda))}{d\log(1+\lambda)} = (1 + o(1))\frac{W(d\lambda + O(d\lambda^2))}{d} = (1 + o(1))\frac{W(d\lambda)}{d},$$

where for the last equality we use the same argument via the derivative of W as above. In particular if $\lambda = 1/\log d$ say, we have that

$$\frac{1}{n}\bar{\alpha}_G(\lambda) \geq (1 + o(1))\frac{\log d}{d}$$

and by Proposition 1, the same holds when $\lambda = \Theta(1)$. Now when $\lambda = \Theta(1)$, $\alpha_{T_d}(\lambda) = (1 + o(1))\frac{\log d}{d}$ by (4) and so the result follows. \square

In the context of the proof technique of Section 2, we can understand the derivation of (4) by imposing two local constraints on the hard-core model: that every vertex has the same probability of being occupied, and that conditioned on a vertex v not being occupied, the events that each of its neighbors are uncovered are *independent* events. Contrast this with our lower bound, which essentially says that the worst case for the occupancy fraction is if the number of these events that hold is always a constant; that is, the corresponding events are as negatively correlated as possible. Now if d is large, then the sum of d independent indicator random variables is highly concentrated so it is plausible that the most negatively correlated case is not far from the independent case. Proposition 2 makes this precise.

Returning to the world of finite graphs, we note that in fact $\alpha_{T_d}(\lambda)$ is approximately achieved by finite graphs for a very large range of λ , and so Theorem 3 is asymptotically tight. Bhatnagar, Sly, and Tetali [5] give a precise description of the local distribution of the hard-core model on a random d -regular graph $G_{d,n}$ for λ below a specified value $\lambda_{\text{upper}} = d^{1+o(1)}$ (just below the condensation threshold of the model). They show that the local distribution converges to that of the unique translation invariant hard-core measure on the infinite d -regular tree described above, and so in particular for $\lambda < \lambda_{\text{upper}}$,

$$\frac{1}{n}\bar{\alpha}_{G_{d,n}}(\lambda) = \alpha_{T_d}(\lambda) + o_n(1),$$

whp as $n \rightarrow \infty$ (note that this is a much stronger notion of approximation than that of Proposition 2, with the error term tending to 0 with n instead of d). Then provided we condition on $G_{d,n}$ being triangle free, which occurs with positive probability (depending only

on d), a random d -regular graph indeed exhibits the tightness of Theorem 3. Let us quickly remark that when $\lambda = o_d(1)$ one can get better asymptotic agreement than that given in Proposition 2. For example, when $\lambda = d^{-s}$ for some fixed $0 < s < 1$, one can show that both the lower bound in Theorem 3 and $\alpha_{T_d}(\lambda)$ can be written as $\frac{W(\lambda d)}{d} + O_d\left(\frac{\log d}{d^{1+s}}\right)$ and so Theorem 3 is tight in a rather strong sense in this range.

An easy consequence of the results of [5] (either by their Sections 3.2 and 3.3 or by integrating the occupancy fraction) is that Theorem 4 is tight asymptotically in d in the exponent. In particular, whp,

$$P_{G_{d,n}}(1) = e^{\left(\frac{1}{2} + o_d(1)\right) \frac{\log^2 d}{d} n},$$

showing that the constant in the exponent of Theorem 2 is tight.

Finally let us note that for any fixed positive integer k , $G_{d,n}$ in fact has girth at least k with positive probability (depending only on d). It follows that even if a stronger lower bound on girth is assumed in Theorems 1 to 4, the results cannot be improved asymptotically.

5. ON THE RATIO OF THE MAXIMUM AND AVERAGE INDEPENDENT SET SIZE

In light of the above result showing that the average size of an independent set in a triangle-free graph with maximum degree d is at least $(1 + o_d(1)) \frac{\log d}{d} n$, we now raise the question of whether the largest independent set should be significantly larger. This gives a new way to pursue an upper bound on $R(3, k)$. There is always a gap between the maximum and average size of an independent set (since the empty set is an independent set), but in general the ratio of maximum to average size can be arbitrarily close to 1. For example, the complete graph K_n has maximum independent set size 1 with average size $n/(n+1)$.

We conjecture that such a narrow gap cannot occur in triangle-free graphs. The following three conjectures make this claim precise in different ways.

Conjecture 1. *For every triangle-free graph G ,*

$$\frac{\alpha(G)}{\bar{\alpha}_G(1)} \geq 4/3.$$

Replacing $4/3$ with any number strictly greater than 1 would give an improvement to the $R(3, k)$ bound. The graph with the smallest ratio $\alpha/\bar{\alpha}(1)$ we have found is the triangle-free cyclic graph that exhibits the bound $R(3, 9) \geq 36$ [14]. For this graph $\alpha/\bar{\alpha}(1) = \frac{197136}{137585} = 1.43283\dots$. We choose $4/3$ since it is a nice fraction less than 1.43 and since it is the ratio of maximum to average size in a triangle. One might wonder if the extremal $R(3, k)$ graphs are good candidates for pushing the ratio $\alpha/\bar{\alpha}(1)$ down to 1. However, for large k it may be the case that graphs arising from the triangle-free process are asymptotically extremal, as is conjectured in [12]. We believe that for such graphs the ratio $\alpha/\bar{\alpha}(1)$ in fact converges to 2. This motivates the following conjectures.

Conjecture 2. *For every triangle-free graph G of minimum degree d ,*

$$\frac{\alpha(G)}{\bar{\alpha}_G(1)} \geq 2 - o_d(1).$$

Conjecture 3. *For every $\epsilon > 0$, there exists $\lambda > 0$ so that for all triangle-free graphs G ,*

$$\frac{\alpha(G)}{\bar{\alpha}_G(\lambda)} \geq 2 - \epsilon.$$

Lemma 1. *The following improvements to Shearer's upper bound on $R(3, k)$ would follow from the above conjectures and Theorems 1 and 3.*

- (1) *Conjecture 1 implies $R(3, k) \leq (3/4 + o(1))k^2 / \log k$.*
- (2) *Conjecture 2 implies $R(3, k) \leq (1/2 + o(1))k^2 / \log k$.*
- (3) *Conjecture 3 implies $R(3, k) \leq (1/2 + o(1))k^2 / \log k$.*

Proof. (1) and (3) are immediate as Theorem 3 implies that for any fixed $\lambda > 0$, $\bar{\alpha}_G(\lambda) \geq (1 + o_d(1)) \frac{\log d}{d} n$ in any maximum degree d , triangle-free graph G on n vertices.

To show (2) we take any triangle-free graph of maximum degree d , greedily removing any vertices of degree at most $\frac{d}{2 \log d}$ and their neighbors. \square

One possible approach to the above conjectures is via the following simple consequence of the proof of Proposition 1. For any graph G

$$\alpha(G) = \lim_{\lambda \rightarrow \infty} \bar{\alpha}_G(\lambda) = \bar{\alpha}_G(1) + \int_1^\infty \frac{\text{Var}_\lambda(|I|)}{\lambda} d\lambda.$$

In this paper we gave a lower bound for the expected size of an independent set drawn from a triangle-free graph according to the hard-core model. The above equation shows that one approach to the above conjectures would be do to the same for the variance.

It is not clear that there is anything special about excluding triangles as opposed to other graphs, and so we make the following more general conjecture.

Conjecture 4. *For every K_r -free graph G ,*

$$\frac{\alpha(G)}{\bar{\alpha}_G(1)} \geq 1 + \frac{1}{r}.$$

For every K_r -free graph G of minimum degree d ,

$$\frac{\alpha(G)}{\bar{\alpha}_G(1)} \geq 2 - o_d(1),$$

with r fixed as $d \rightarrow \infty$.

To conclude, we give a lower bound on $\alpha/\bar{\alpha}(\lambda)$ that is tight for disjoint unions of r -cliques. The result also shows that $\alpha/\bar{\alpha}(1)$ is bounded away from 1 for all graphs containing an independent set of linear size.

Theorem 6. *For any graph G ,*

$$\frac{\alpha(G)}{\bar{\alpha}_G(\lambda)} \geq 1 + \frac{\alpha(G)}{\lambda n},$$

or equivalently,

$$(16) \quad P(\lambda) \geq \left(\frac{\lambda}{\alpha} + \frac{1}{n} \right) P'(\lambda).$$

This is tight when G is a union of r -cliques, for any r .

Proof. Let $P(\lambda)$ be the partition function of G . Define

$$Q(\lambda) = P(\lambda) - \frac{\lambda}{\alpha} P'(\lambda) - \frac{1}{n} P'(\lambda).$$

Then

$$\frac{Q(\lambda)}{P(\lambda)} = 1 - \frac{\lambda P'(\lambda)}{\alpha P(\lambda)} - \frac{1 P'(\lambda)}{n P(\lambda)} = 1 - \frac{\bar{\alpha}(\lambda)}{\alpha} - \frac{\bar{\alpha}(\lambda)}{\lambda n},$$

and so after rearranging, it suffices to show that $Q(\lambda) \geq 0$ for all λ . In fact, we will show that all coefficients of $Q(\lambda)$ as a polynomial in λ are non-negative. Let $Q[k]$, denote the coefficient of λ^k in $Q(\lambda)$, and let $i_k = P[k]$ be the number of independent sets of size k in G . Then

$$Q[k-1] = i_{k-1} - \frac{k-1}{\alpha} i_{k-1} - \frac{k}{n} i_k.$$

So it is enough to show for $k = 1, \dots, \alpha$ that

$$(17) \quad k \frac{i_k}{i_{k-1}} \leq n - \frac{n(k-1)}{\alpha}.$$

The following inequality is due to Moon and Moser [20] (though it is usually stated for cliques instead of independent sets):

$$(18) \quad \frac{1}{k^2 - 1} \left(k^2 \frac{i_k(G)}{i_{k-1}(G)} - n \right) \leq \frac{i_{k+1}(G)}{i_k(G)}.$$

We proceed by induction on k , from $k = \alpha$ down to $k = 1$. If $k = \alpha$, then (5) gives:

$$\frac{i_k(G)}{i_{k-1}(G)} \leq \frac{n}{k^2}$$

which is exactly (5) with $\alpha = k$.

Now inductively assume $\frac{i_{k+1}}{i_k} \leq n(1 - \frac{k}{\alpha}) \frac{1}{k+1}$, and plug this into the RHS of (5). Rearranging, this gives

$$\frac{i_k}{i_{k-1}} \leq n \left(\frac{1}{k} - \frac{1}{\alpha} + \frac{1}{k\alpha} \right),$$

which is (5). □

The integrated form of (6) is a simple bound on the partition function in terms of the independence number and number of vertices: for any graph G on n vertices with independence number α ,

$$P_G(\lambda) \leq \left(1 + \frac{\lambda n}{\alpha} \right)^\alpha.$$

This bound is tight for unions of cliques of the same size and is implicit in the work of Moon and Moser. It appears explicitly in [13], generalizing the $\lambda = 1$ case from [2].

ACKNOWLEDGEMENTS

We thank Rob Morris for his helpful comments on this paper, Noga Alon and Benny Sudakov for recalling James Shearer's 1998 SIAM talk, and James Shearer for sharing his transparencies from that talk with us.

REFERENCES

- [1] Miklós Ajtai, János Komlós, and Endre Szemerédi, *A note on Ramsey numbers*, Journal of Combinatorial Theory, Series A **29** (1980), no. 3, 354–360.
- [2] Vladimir E Alekseev, *On the number of maximal independent sets in graphs from hereditary classes*, Combinatorial-algebraic methods in discrete optimization, University of Nizhny Novgorod (1991), 5–8.
- [3] Noga Alon, *Independence numbers of locally sparse graphs and a Ramsey type problem*, Random Structures and Algorithms **9** (1996), no. 3, 271–278.
- [4] Noga Alon and Joel H Spencer, *The probabilistic method*, 3rd ed., John Wiley & Sons, 2011.
- [5] Nayantara Bhatnagar, Allan Sly, and Prasad Tetali, *Decay of correlations for the hardcore model on the d -regular random graph*, Electronic Journal of Probability **21** (2016).
- [6] Tom Bohman and Peter Keevash, *Dynamic concentration of the triangle-free process*, arXiv preprint arXiv:1302.5963 (2013).
- [7] Graham Brightwell and Peter Winkler, *Hard constraints and the Bethe lattice: adventures at the interface of combinatorics and statistical physics*, Proc. Int'l. Congress of Mathematicians, vol. III, 2002, pp. 605–624.
- [8] Jeff Cooper, Kunal Dutta, and Dhruv Mubayi, *Counting independent sets in hypergraphs*, Combinatorics, Probability and Computing **23** (2014), no. 04, 539–550.
- [9] Jeff Cooper and Dhruv Mubayi, *Counting independent sets in triangle-free graphs*, Proceedings of the American Mathematical Society **142** (2014), no. 10, 3325–3334.
- [10] Péter Csikvári, *Lower matching conjecture, and a new proof of Schrijver’s and Gurvits’s theorems*, arXiv preprint arXiv:1406.0766 (2014).
- [11] Ewan Davies, Matthew Jenssen, Will Perkins, and Barnaby Roberts, *Independent sets, matchings, and occupancy fractions*, arXiv preprint arXiv:1508.04675 (2015).
- [12] Gonzalo Fiz Pontiveros, Simon Griffiths, and Robert Morris, *The triangle-free process and $R(3, k)$* , arXiv preprint arXiv:1302.6279 (2013).
- [13] David Galvin, *An upper bound for the number of independent sets in regular graphs*, Discrete Mathematics **309** (2009), no. 23, 6635–6640.
- [14] Charles M Grinstead and Sam M Roberts, *On the Ramsey numbers $R(3, 8)$ and $R(3, 9)$* , Journal of Combinatorial Theory, Series B **33** (1982), no. 1, 27–51.
- [15] Jeff Kahn, *An entropy approach to the hard-core model on bipartite graphs*, Combinatorics, Probability and Computing **10** (2001), no. 03, 219–237.
- [16] Frank P Kelly, *Stochastic models of computer communication systems*, Journal of the Royal Statistical Society. Series B (Methodological) (1985), 379–395.
- [17] Jeong Han Kim, *The Ramsey number $R(3, t)$ has order of magnitude $t^2/\log t$* , Random Structures & Algorithms **7** (1995), no. 3, 173–207.
- [18] Alexandr Kostochka, Dhruv Mubayi, and Jacques Verstraëte, *On independent sets in hypergraphs*, Random Structures & Algorithms **44** (2014), no. 2, 224–239.
- [19] Michael Molloy and Bruce Reed, *Graph colouring and the probabilistic method*, vol. 23, Springer Science & Business Media, 2013.
- [20] JW Moon and Leo Moser, *On a problem of Turán*, Magyar Tud. Akad. Mat. Kutató Int. Közl **7** (1962), 283–286.
- [21] Nicholas Ruoizzi, *The Bethe partition function of log-supermodular graphical models*, Advances in Neural Information Processing Systems, 2012, pp. 117–125.
- [22] James B Shearer, personal communication.
- [23] ———, *A note on the independence number of triangle-free graphs*, Discrete Mathematics **46** (1983), no. 1, 83–87.
- [24] ———, *On the independence number of sparse graphs*, Random Structures & Algorithms **7** (1995), no. 3, 269–271.
- [25] Alan S Willsky, Erik B Sudderth, and Martin J Wainwright, *Loop series and Bethe variational bounds in attractive graphical models*, Advances in neural information processing systems, 2008, pp. 1425–1432.
- [26] Yufei Zhao, *The number of independent sets in a regular graph*, Combinatorics, Probability and Computing **19** (2010), no. 02, 315–320.

LONDON SCHOOL OF ECONOMICS

E-mail address: {`e.s.davies,m.o.jenssen,b.j.roberts`}@lse.ac.uk

UNIVERSITY OF BIRMINGHAM

E-mail address: `math@willperkins.org`