

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

Ewan Davies

London School of Economics and Political Science

E.S.Davies@lse.ac.uk

August 3, 2016

Joint work with W. Perkins (Birmingham), M. Jenssen, B. Roberts (LSE)

Introduction

- The **Ramsey number** $R(3, k)$ is the least integer n such that every graph on n vertices contains either a triangle or an independent set of size k .
- It is known that

$$\left(\frac{1}{4} + o(1)\right) \frac{k^2}{\log k} \leq R(3, k) \leq (1 + o(1)) \frac{k^2}{\log k}$$

Bohman, Keevash 2013; Fiz Pontiveros, Griffiths, Morris 2013
Shearer 1983

- Both these bounds have been thought correct by different people at different times and reducing the gap is a major open problem in Ramsey theory.

New method

- The focus of this talk is a new method which leads to an alternative proof of Shearer's upper bound.
- The method is of independent interest, applicable to other problems such as bounding the number of matchings in d -regular graphs.
- Our method suggests a new approach to improving the upper bound on $R(3, k)$.

Theorem 1 (D., Jensen, Perkins, Roberts 2016)

The *average* size of an independent set in a *triangle-free* graph on n vertices with *maximum degree* d is at least

$$(1 + o_d(1)) \frac{\log d}{d} n.$$

The constant 1 is best possible.

In a triangle-free graph the *neighbourhood* of a vertex is an *independent set*, hence for any triangle-free graph

$$\alpha(G) \geq \max \left\{ d, (1 + o(1)) \frac{\log d}{d} n \right\},$$

which immediately implies $R(3, k) \leq (1 + o(1))k^2 / \log k$.

Results

Theorem 2 (D., Jensen, Perkins, Roberts 2016)

The *number* of independent sets in a *triangle-free* graph G on n vertices with *maximum degree* d is at least

$$e^{\left(\frac{1}{2} + o_d(1)\right) \frac{\log^2 d}{d} n}.$$

The constant $\frac{1}{2}$ is best possible.

Then for any triangle-free graph we have at least

$$\max \left\{ 2^d, e^{\left(\frac{1}{2} + o(1)\right) \frac{\log^2 d}{d} n} \right\} \geq e^{\left(\frac{\sqrt{2 \log 2}}{4} + o(1)\right) \sqrt{n} \log n},$$

independent sets, improving a result of Cooper, Dutta and Mubayi by a factor $\sqrt{2}$ in the exponent.

The hard-core model

The **hard-core model** on a graph G is simply an independent set I chosen at random from G .

The distribution is defined by

$$\Pr(I) = \frac{\lambda^{|I|}}{P_G(\lambda)} \text{ where } P_G(\lambda) = \sum_I \lambda^{|I|}$$

is the **partition function** or independence polynomial of G . The parameter $\lambda > 0$ is the **fugacity** and controls whether small or large independent sets are preferred.

Our method resembles previous work on bounding the size of independent sets by Shearer, Alon, and Alon–Spencer. The main difference is that we use the full power of the hard-core model at general fugacity λ .

The partition function

The partition function contains important information about the graph and the hard-core model, for instance

- At fugacity 1 the distribution is uniform so $P_G(1)$ is the number of independent sets in G .
- The expected size of the random independent set I is

$$\mathbb{E}|I| = \sum_I |I| \frac{\lambda^{|I|}}{P_G(\lambda)} = \frac{\lambda P'_G(\lambda)}{P_G(\lambda)} = \lambda \left(\log P_G(\lambda) \right)'.$$

- The variance of the size of I can be written in terms of $P_G(\lambda)$.
- As $\lambda \rightarrow \infty$, $\mathbb{E}|I| \rightarrow \alpha(G)$, the maximum size of an independent set in G .

Main result

Theorem 3 (D., Jansen, Perkins, Roberts 2016)

Let G be a *triangle-free graph* on n vertices with *maximum degree d* . The expected size of an independent set drawn according to the hard-core model at fugacity λ on G is at least

$$\frac{\lambda}{1 + \lambda} \frac{W(d \log(1 + \lambda))}{d \log(1 + \lambda)} n.$$

Where $W(x)$ is the Lambert W function satisfying $W(xe^x) = x$.

Theorems 1 and 2 are simple consequences of Theorem 3, showing the power of bounding the partition function of the hard-core model.

Proof of Theorem 1

The **average** size of an independent set in a **triangle-free graph** on n vertices with **maximum degree d** is at least $(1 + o(1)) \frac{\log d}{d} n$.

Choose $\lambda = 1/\log d$ so that $\lambda = o(1)$ and $\log \lambda = o(\log d)$.

Using the fact that $W(x) \geq \log x - \log \log x$ for $x \geq e$ we have

$$\begin{aligned}\mathbb{E}|I| &\geq \frac{\lambda}{1 + \lambda} \frac{W(d \log(1 + \lambda))}{d \log(1 + \lambda)} n \\ &\geq (1 + o(1)) \frac{\log d}{d} n.\end{aligned}$$

Theorem 1 now follows from the fact that $\mathbb{E}|I|$ is increasing in λ ,

$$\frac{d}{d\lambda} \mathbb{E}|I| = \frac{d}{d\lambda} \frac{\lambda P'_G}{P_G} = \dots = \frac{\text{Var}|I|}{\lambda} > 0,$$

a formula which will have significance later. □

Proof of Theorem 2

Since we have $\mathbb{E}|I| = \frac{\lambda P_G'}{P_G} = \lambda(\log P_G)'$ and $P_G(0) = 1$, by integrating $\mathbb{E}|I|/\lambda$ we obtain $\log P_G$.


$$\begin{aligned}\log P_G(\lambda) &= \int_0^\lambda \frac{\mathbb{E}_t|I|}{t} dt \\ &\geq \frac{n}{d} \int_0^\lambda \frac{W(d \log(1+t))}{\log(1+t)} dt \\ &= \frac{n}{d} \int_0^{W(d \log(1+\lambda))} (1+u) du \\ &= \frac{n}{2d} \left[W(d \log(1+\lambda))^2 + 2W(d \log(1+\lambda)) \right].\end{aligned}$$

For $\lambda = 1$ we have $\log P_G(1) \geq \left(\frac{1}{2} + o(1)\right) \frac{\log^2 d}{d} n$. □

Proof of Theorem 3

Recall that Theorem 3 is a lower bound on $\mathbb{E}|I|$ in a triangle-free graph of maximum degree d .

We define an experiment, picking I from the hard core model on G and a vertex v uniformly at random. We then consider which of the neighbours of v are **uncovered**, which means none of their neighbours are in I .



vertex in I : ●

uncovered vertex: ●

covered vertex: ●

Proof of Theorem 3

There are two important facts.

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

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

These facts rely on an important property of the hard-core model: conditioned on a boundary, what happens inside and outside are independent.

Proof of Theorem 3

We now write $\mathbb{E}|I|$ as

$$\begin{aligned}\mathbb{E}|I| &= \sum_{v \in G} \Pr[v \in I] \\ &= \frac{\lambda}{1 + \lambda} \cdot \sum_{v \in G} \Pr[v \text{ uncovered}] && \text{by Fact 1} \\ &= \frac{\lambda}{1 + \lambda} \cdot \sum_{v \in G} \sum_{j=0}^d \Pr[v \text{ has } j \text{ uncovered neighbours}] \cdot (1 + \lambda)^{-j} \\ &&& \text{by Fact 2}\end{aligned}$$

Let Z be a random variable which counts the number of uncovered neighbours of a uniformly random vertex. Then the last line gives

$$\mathbb{E}|I| = \frac{n\lambda}{1 + \lambda} \mathbb{E} \left[(1 + \lambda)^{-Z} \right] \geq \frac{n\lambda}{1 + \lambda} (1 + \lambda)^{-\mathbb{E}Z}$$

Proof of Theorem 3

But we can also write

$$\begin{aligned}\mathbb{E}|I| &= \sum_{v \in G} \Pr[v \in I] \\ &\geq \frac{1}{d} \sum_{v \in G} \sum_{u \sim v} \Pr[u \in I]\end{aligned}$$

since each vertex u appears $\deg u \leq d$ times in the double sum,

$$= \frac{n\lambda}{1 + \lambda} \frac{\mathbb{E}Z}{d}$$

by **Fact 1**.

Proof of Theorem 3

Then

$$\begin{aligned}\mathbb{E}|I| &\geq \frac{n\lambda}{1+\lambda} \max \left\{ (1+\lambda)^{-\mathbb{E}Z}, \frac{\mathbb{E}Z}{d} \right\} \\ &\geq \frac{n\lambda}{1+\lambda} \max_{z \in \mathbb{R}} \left\{ (1+\lambda)^{-z}, \frac{z}{d} \right\}\end{aligned}$$

to minimise, set the two functions of z equal so that $z \log(1+\lambda)e^{z \log(1+\lambda)} = d \log(1+\lambda)$ and use $W(xe^x) = x$,

$$\geq \frac{\lambda}{1+\lambda} \frac{W(d \log(1+\lambda))}{d \log(1+\lambda)} n,$$

as required. □

Tightness

- Our method shows that $\mathbb{E}|I|$ is minimized when Z is a constant.
- The translation-invariant hard-core measure on the infinite d -regular tree asymptotically achieves the bound for $\lambda = O(1)$. In this case the indicators of $u \sim v$ being uncovered are iid Bernoullis, so their sum Z is concentrated.
- A random d -regular graph is triangle free with positive probability and Bhatnagar, Sly and Tetali showed that the hard-core model on a random d -regular graph approaches that of the infinite tree for λ in a suitable range.
- These arguments show that Theorem 3, and hence Theorems 1 and 2 are asymptotically tight for certain ranges of λ .

Ramsey Theory

In deducing the upper bound on $R(3, k)$ from Theorem 1 we used the **average size** as a lower bound for the **maximum size** of the independent set. We cannot improve Theorem 1, but one could improve the upper bound by proving the maximum is somewhat larger than the average:

Conjecture 1

For every triangle free graph, the maximum independent set is at least $4/3$ times the average size.

Conjecture 2

For every triangle free graph of minimum degree d , the maximum independent set is at least $2 - o(1)$ times the average size.

By Theorem 1, each of these immediately imply an improvement to the upper bound on $R(3, k)$ which has stood since 1983.

Ramsey Theory

We proved a lower bound on the **expected size** of an independent set from the hard-core model to bound $R(3, k)$. To improve the bound we want to understand the **maximum size**.

Recall that

$$\frac{d}{d\lambda} \mathbb{E}_\lambda |I| = \frac{\text{Var}_\lambda |I|}{\lambda},$$

and that $\mathbb{E}_\lambda |I| \rightarrow \alpha(G)$ as $\lambda \rightarrow \infty$. Then

$$\alpha(G) = \mathbb{E}_\lambda |I| + \int_\lambda^\infty \frac{\text{Var}_t |I|}{t} dt,$$

which shows that one approach to the above conjectures would be to do the same for **variance**.

Thank you

Bonus slide: maximisation

The method also proves the triangle-free case of a strengthening of a result due to Kahn. For a d -regular triangle-free graph we actually showed the equality

$$\mathbb{E}|I| = \frac{n\lambda}{1+\lambda} \frac{\mathbb{E}Z}{d} = \frac{n\lambda}{1+\lambda} \mathbb{E} \left[(1+\lambda)^{-Z} \right].$$

By $0 \leq Z \leq d$ and convexity we have $(1+\lambda)^{-Z} \leq \frac{Z}{d}(1+\lambda)^{-d} + 1 - \frac{Z}{d}$, and hence

$$\mathbb{E}|I| \leq \frac{\lambda(1+\lambda)^{d-1}}{2(1+\lambda)^d - 1} n.$$

A quick calculation reveals this upper bound to be $\mathbb{E}|I|$ when G is a disjoint union of $\frac{n}{2d}$ $K_{d,d}$'s. By integrating we have $P_G(\lambda)^{1/n} \leq P_{K_{d,d}}(\lambda)^{1/2d}$, giving a unified probabilistic proof of the results of Shearer and Kahn.