Structure and Randomness in Extremal Combinatorics

Barnaby Roberts

A thesis submitted for the degree of

Doctor of Philosophy

Department of Mathematics

London School of Economics

and Political Science

21/07/2017

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is my own work, with the following exceptions.

Chapter 2 is based on the paper 'Triangle-free subgraphs of random graphs' written with Peter Allen, Julia Böttcher and Yoshiharu Kohayakawa.

Chapter 4 is based on the paper 'The size-Ramsey number of powers of paths' with Dennis Clemens, Matthew Jenssen, Yoshiharu Kohayakawa, Natasha Morrison, Guilherme Oliveira Mota and Damian Reding.

Chapter 5 is based on the papers 'Independent sets matchings and occupancy fractions' and 'On the average size of independent sets in triangle-free graphs' both written with Ewan Davies, Matthew Jenssen and Will Perkins.

I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that this thesis consists of 39646 words.

Abstract

In this thesis we prove several results in extremal combinatorics from areas including Ramsey theory, random graphs and graph saturation. We give a random graph analogue of the classical Andrásfai, Erdős and Sós theorem showing that in some ways subgraphs of sparse random graphs typically behave in a somewhat similar way to dense graphs. In graph saturation we explore a 'partite' version of the standard graph saturation question, determining the minimum number of edges in H-saturated graphs that in some way resemble H themselves. We determine these values for K_4, paths, and stars and determine the order of magnitude for all graphs. In Ramsey theory we give a construction from a modified random graph to solve a question of Conlon, determining the order of magnitude of the size-Ramsey numbers of powers of paths. We show that these numbers are linear. Using models from statistical physics we study the expected size of random matchings and independent sets in d-regular graphs. From this we give a new proof of a result of Kahn determining which d-regular graphs have the most independent sets. We also give the equivalent result for matchings which was previously unknown and use this to prove the Asymptotic Upper Matching Conjecture of Friedland, Krop, Lundow and Markström. Using these methods we give an alternative proof of Shearer's upper bound on off-diagonal Ramsey numbers.

Acknowledgments

I would like to thank my supervisors, Peter Allen and Jozef Skokan, for their amazing generosity, providing wisdom, time, and beer in equal abundance. For the interesting problems you made me aware of, the poorly written drafts you patiently edited, and always being willing to help and chat, I am incredibly grateful. Having supervisors who both shared several of their own project suggestions, whilst also giving me freedom to choose projects for myself, is something I certainly do not take for granted.

Also from the LSE mathematics department I would like to thank Julia Böttcher for being as a third supervisor, giving detailed feedback and advice on many occasions. Also to Rebecca Lumb and Kate Barker for their remarkable efficiency at dealing with everything administrative and for creating such a friendly and fun atmosphere in the department.

As well as my supervisors and Julia, I have also had the privilege to co-author with, Dennis Clemens, Ewan Davies, Matthew Jenssen, Yoshiharu Kohayakawa, Natasha Morrison, Guilherme Mota, Will Perkins, and Damian Reding. It was joyful and humbling to work with such amazingly talented mathematicians. Particular thanks I would like to give to Ewan and Matthew, with whom I have written over half my papers. It has been great fun doing my PhD alongside both of you and I am grateful for the way you patiently listened to so many of my bad ideas in exchange for just a few good ones.

Two of my school teachers in particular were instrumental in taking maths from being something I did well at, to something I thoroughly enjoyed. I am grateful to Mr Moran and Mr McEwan for their infectious enthusiasm for maths and problem solving in general.

For putting up with me for longer than anyone I would like to thank my parents and siblings for their constant support throughout my life. Your pushing me academically may not have always been appreciated in the way it should have been at the time but without it there is no way I would have had the chance to enjoy the studies I have done over the past eight years.

Hana Tsuruhara has been an ever reliable source of friendship, humour and fun throughout my studies going above and beyond the girlfriend duties by not only listening as I attempt to explain my work, but also pretending to understand and be interested. I am very grateful for your unfailing support.

Finally I would like to thank my God and Saviour who provides a peace and joy that prevail over even the most tedious aspects of thesis writing and editing.

Contents

1

Introduction

Since Erdős' [32] introduction of the probabilistic method in the late 1940's, using ideas from probability has become widespread and fruitful in combinatorics. Whilst it may at first seem counter-intuitive that a randomised method could prove a deterministic result, it turns out that often a randomly generated graph will have properties that it is very hard to explicitly construct a graph with. For example, in 1959 Erdős [33] showed there there are graphs with both arbitrarily high girth and chromatic number. Whilst there are now non-random proofs of this statement, Erdős' proof is arguably still the simplest and most elegant. For this example it is not simply enough to take a randomly generated graph. Although randomly generated graphs with enough edges tend to have high chromatic number they also typically have low girth. However they tend to only have a small number short cycles and so, after removing vertices from these short cycles, the graph that remains has both high chromatic number and high girth.

This concept has similarities to ideas in other fields. For example it is more difficult to demonstrate a transcendental number than it is to prove that such numbers exist. The first number to be proved to be transcendental was

$$\sum_{n=1}^{\infty} 10^{-n!},$$

by Liouville [67]. However it is easy to see that there are just countably many algebraic numbers, and so uncountably many transcendental numbers. In this way, 'almost all numbers' are transcendental, even though it can be difficult to find them.

One could attribute this to the idea that if we construct a graph, or define a number, ourselves, then since we will likely define it in a relatively simple and finite way, such a graph or number will be in some sense *simple* and *structured*. If the graph property we desire requires the graph to be in some sense *complex* or *unstructured* then displaying it explicitly may prove substantially harder than arguing non-constructively that such a graph (or number) exists.

Of course we have not yet defined what we mean by a graph being *structured*, and in fact we will avoid doing so for the reason that there will be two competing notions of structure. Under several natural (informal) definitions of *structure* for graphs, randomly generated graphs are with high probability very unstructured. And yet, there is a sense, that we will look at shortly, in which randomly generated graphs typically exhibit a lot of structure.

At first we would like to think of a graph as being structured if the vertex set can be partitioned into a small number of parts so that for any given pair of vertices we can determine whether they define an edge simply by knowing which parts the two vertices are in. A very simple example is that of complete bipartite graphs which consist of vertices in two parts with edges just between vertices from distinct parts. This definition seems particularly suited to problems involving counting particular subgraphs. For example it is easy to count the number of copies of C_4, the cycle on four vertices, in a graph defined by a small number of parts as above. One downside of this definition is that other simple graphs that we might consider structured do not fit into this category. Cycles C_n, grid graphs $P_n \times P_m$ and hypercubes H_{2^d} are examples of such graphs. We could instead consider a graph to be structured if it can be described easily or in few words. By both definitions it is straightforward to show that randomly generating a graph will almost certainly generate an unstructured graph. In particular if we generate a graph on n vertices by flipping unbiased coins independently for each pair of vertices to choose which pairs have an edge then it requires an average of $\binom{n}{2}$ *bits* of information to describe the graph.

This model of generating a random graph is called $G(n,p)$. In general for a positive integer n and a real number $p \in [0,1]$ we let $G(n,p)$ denote the random graph model on n vertices where each pair of vertices contains an edge independently with probability p. Of course if $p = 0$ or 1, this gives the empty graph or complete graph respectively, but if p is not near either 0 or 1, the result is random and unpredictable. Yet, whilst the exact graph outputted by this model may be unpredictable, it turns out there are many properties of the random graph that we can be highly confident of, even before seeing the output. For example, if $n \geqslant 20$ and $p = 1/2$ then we can be more than 99% sure that the graph generated will be connected, in the sense that we can walk from any vertex to any other along edges. We can also accurately estimate many parameters of the random graph with a high degree of confidence. For large values of n, with probability very near to 1, $G(n, 1/2)$ has roughly $\frac{1}{2}\binom{n}{2}$ edges. We will make this more precise later in the thesis.

Erdős' idea of using random graphs, or modifications of random graphs, to display properties that it is hard to exhibit constructively, has been used many times since 1959. The results presented in Chapter 4 are of this flavour. We use modified random graphs to prove the existence of graphs without too many edges that satisfy a certain property.

Over the last few decades, the random graph model $G(n,p)$ itself has become an important object of study, rather than merely a useful tool. Questions about $G(n,p)$ typically choose a graph property $P(n)$ and ask for the asymptotic behaviour of the probability that $G(n,p)$ satisfies $P(n)$ as n tends to infinity. We say that $G(n,p)$ has $P(n)$ *with high probability* if the probability that $G(n,p)$ satisfies $P(n)$ tends to 1 as n tends to infinity.

One particular theme is taking properties that hold deterministically for the complete graph K_n, and studying related problems for the random graph $G(n,p)$. For example, a well known result due to Mantel [70] is that, for any n, the graph on n vertices with the most edges whilst not containing a triangle, is also the largest bipartite subgraph of K_n. That is to say, the complete bipartite graph $K_{\lceil n/2\rceil,\lfloor n/2\rfloor}$ has the most edges of any triangle-free graph on n vertices. We can ask whether the same phenomena happens for other graphs. Given a graph G, we can ask if the largest bipartite subgraph is the largest triangle-free subgraph, where by largest we mean the subgraph with the most edges. Babai, Simonovits and Spencer [9] showed that in $G(n, 1/2)$ the probability that the largest bipartite subgraph is also the largest triangle-free subgraph tends to 1 as n tends to infinity. In Chapter 2 we prove a result of this type giving a random graph analogue of a theorem of Andrásfai, Erdős and Sós [8]. This theorem states that triangle-free graphs on n vertices with minimum degree greater than $2n/5$ are bipartite. We show that it is almost always the case that triangle-free subgraphs of random graphs with a corresponding minimum degree condition are nearly bipartite. The particular emphasis of our result is determining sharp bounds for how near to bipartite such subgraphs must be.

Curiously, results of this type can be seen as saying that some structural properties of K_n are typically shared by random graphs. In particular, it is in general, the fact that random graphs are highly connected with edges distributed all across the graph, that tends to result in such properties holding. This is the reason we avoided defining 'structure' in some way that would exclude random graphs.

For some problems random constructions are far from best and deterministic techniques are needed. We study one such set of problems in Chapter 3 when we study graph saturation. One property of random graphs is that typically different parts of a random graph all look quite similar. One could say that in a randomly generated graph "no vertex is special". For example, usually all vertices of the random graph have a very similar number of neighbours. In the extremal graphs that arise in the saturation problems of Chapter 3 there tend to be a small number of 'special' vertices with very high degree whilst all other vertices have very few neighbours.

Chapter 5 contains results that are both deterministic and probabilistic. The probabilistic results of this chapter however, do not concern creating a graph randomly, but choosing a random independent set or matching from a fixed graph. We look at the expected size of a random independent set or matching drawn from a d-regular graph. We show that the d-regular graph that maximises the expected fraction of vertices in a random independent set, is the complete bipartite graph $K_{d,d}$. We show that $K_{d,d}$ also maximise the expected fraction of vertices in a random matching. We prove both these results using very little probability theory, instead primarily relying on linear programming. We prove these results when the random independent set or matching is drawn uniformly and also when it is drawn from a more general distribution. We then show how these results about random independent sets and matchings imply the deterministic results that the graph consisting of multiple copies of $K_{d,d}$, has the most independent sets and matchings of any d-regular graph on a fixed number of vertices. In Chapter 5 we also show that, with the right probability distribution, a random independent set drawn from a triangle-free graph on n vertices with maximum degree d, contains on average at least $\frac{\log d}{d}n$ vertices. This gives a new proof of Shearer's upper bound [81] on the off-diagonal Ramsey numbers $R(3,k)$. Intriguingly, and counter-intuitively, our method works best when we randomly choose the independent set in a way that is biased towards selecting smaller independent sets.

1.1 Thesis overview

1.1.1 Chapter 2

This chapter is based on the paper 'Triangle-free subgraphs of random graphs' [5] and looks at taking extremal graph theory problems and asking to what extent analogues of these hold for random graphs. For example; a theorem of Andrásfai, Erdős and Sós [8] states that triangle-free graphs on n vertices with minimum degree greater than $2n/5$ are bipartite. We show that for any $\varepsilon > 0$, with high probability, all triangle-free subgraphs of $G(n,p)$ with minimum degree at least $(\frac{2}{5}+\varepsilon)pn$ are 'close' to bipartite. By close to bipartite we mean that such triangle-free subgraph can be made bipartite by removing just $O(n/p)$ edges.

We use our methods to give the same treatment to a theorem of Thomassen [83]. This theorem states that triangle-free graphs on n vertices with minimum degree at least $(\frac{1}{3}+\varepsilon)n$ for any ε have bounded chromatic number. The bound on the chromatic number is a function of ε that does not depend on n. We show that for all $\varepsilon > 0$ there exists r_ε so that, with high probability, all triangle-free subgraphs of $G(n,p)$ with minimum degree at least $(\frac{1}{3}+\varepsilon)pn$ are 'close' to r_ε-partite. Again this means that such subgraphs can be made r_ε-partite by removing at most $O(n/p)$ edges.

It turns out that this caveat that we have to remove $O(n/p)$ edges is necessary. For some values of p there exist triangle-free subgraphs of $G(n,p)$ with minimum degree very close to $\frac{1}{2}pn$ which even after removing $\Omega(n/p)$ edges have arbitrarily high chromatic number. We construct such graphs to show that our main results are sharp.

1.1.2 Chapter 3

This chapter is based on the paper 'Partite saturation problems' [77]. Saturation problems look at graphs that avoid a certain substructure but in some sense nearly contain that structure. For example we say a graph is *triangle-saturated* if it has no triangles but adding any edge would create a triangle. In this chapter we look at a set of problems in what we call *partite saturation* in which the saturated graphs must bear some specific resemblance to the structure they avoid. For a graph H we let $H[n]$ denote the blow-up of H where each vertex of H is replaced by an independent set of size n and each edge is replaced by a complete bipartite graph between the corresponding independent sets. For example $K_r[n]$ is the complete balanced r-partite graph on rn vertices. If G is a subgraph of $H[n]$ we refer to a copy of H in G as *partite* if it has one vertex in each of the parts of $H[n]$. We say $G \subseteq H[n]$ is $(H, H[n])$-partite-saturated if G contains no partite copy of H but the addition of any extra edge from $H[n]$ would create one. We study the minimum number of edges in such graphs which we call the *partite saturation number*. This type of question was first approached by Ferrara, Jacobson, Pfender, and Wenger [41] who showed (among other things) that all $(K_3, K_3[n])$-partite saturated contain at least $6n-6$ edges. We show that for large enough n it is the case that $(K_4, K_4[n])$-partite-saturated graphs always have at least $18n-21$ edges. We determine the unique graph that attains this bound. We also determine the partite saturation numbers of paths and stars. Finally, we discover an interesting connection between the 2-connectivity of a graph H, and its partite saturation number. A graph is *2-connected* if it requires at least two vertices to be removed, in order to break it into multiple components. Equivalently, it has no *cut vertex*, the removal of which separates the graph. If H is not 2-connected all $(H, H[n])$-partite-saturated graphs have quadratically many edges whilst if H is 2-connected there exist $(H, H[n])$-partite-saturated graphs with linearly many edges.

1.1.3 Chapter 4

This chapter is based on the paper 'The size-Ramsey number of powers of paths' [21]. This is a type of problem in Ramsey theory which concerns partitioning the edge set of graphs in such a way that each part avoids some particular structure. Conversely it also asks for graphs which cannot be partitioned in such a way. We use colours to refer to the parts, and so by a q-colouring of a graph G, we simply mean a partition of the edge set of G into q parts. Given graphs G and H and a positive integer q we say that *G is q-Ramsey for H*, denoted $G \to (H)_q$, if every q-colouring of the edges of G contains a monochromatic copy of H. If $q = 2$ then we will just say G is Ramsey for H.

The classic Ramsey theory question asks for the smallest number of vertices of a graph which is Ramsey for H. In this case we need only consider G which are complete graphs as extra edges can only help. Size Ramsey theory instead asks for the smallest number of edges in a graph G which is Ramsey for H. Here it is not enough to only consider complete graphs. In fact it is the range of possible graphs that require consideration, that makes these questions so interesting. Determining the size-Ramsey numbers of paths was a problem asked by Erdős [34]. Beck [11] showed that these numbers are linear as exhibited by a random graph. Conlon [22] asked whether the same is true for powers of paths. Here random graphs are not sufficient, however by using a modification of a random graph, we showed that the size-Ramsey numbers of powers of paths are indeed linear. Using a random graph here is something of a choice in that it is really just the property that edges are distributed somewhat evenly across the graph that we require. Graphs with this property can be constructed explicitly, but their existence is easier to prove probabilistically.

1.1.4 Chapter 5

This chapter is an amalgamation of the papers 'Independent sets, matchings and occupancy fractions' [26] and 'On the average size of independent sets in triangle-free graphs' [27]. In this chapter we look primarily at counting independent sets and matchings in d-regular graphs, determining which such graphs have the most independent sets and matchings. For independent sets this was already known due to work by Kahn [56] for bipartite graphs via the entropy method, and by Zhao [87] using the bipartite swapping trick for all d-regular graphs. The answer to both these questions is that so long as the number of vertices of the graph, n, is a multiple of $2d$ the optimal graph is the disjoint union on $n/2d$ copies of the complete bipartite graph $K_{d,d}$. We use probabilistic models from statistical physics to turn these problems into linear optimisation problems. Perhaps the most striking element of our method is that we exploit a connection between the average size of independent sets (and respectively matchings) when drawn according to a particular family of probability distributions, and the total number of independent sets (or matchings) in the graph. This family of distributions includes the uniform distribution and so our results also show that on a fixed number of vertices n, so long as $2d|n$ the graph on n vertices with the largest average independent set size is $n/2d$ copies of $K_{d,d}$. The same holds for matchings. We then use our results to solve the asymptotic upper matching conjecture of Friedland, Krop, Lundow and Markström [43] and the equivalent result for independent sets. This conjecture, roughly speaking, required estimating the upper bound on the number of independent sets, of fixed size, in d-regular graphs to within a multiplicative factor that is sub-exponential. We give a bound that is off by a factor of $\sqrt{n}$. We use the same techniques to give an alternative proof of a result of Shearer on the upper bound for off-diagonal Ramsey numbers $R(3,k)$. This problem asks how large an independent set there must be in triangle-free graphs. Curiously our method approaches these very different sounding problems in a remarkably similar way. We are again looking at the average size of an independent set in a graph; this time specifically triangle-free graphs. We show that the average size of an independent set in a triangle-free graph on n vertices with maximum degree d is at least $(1+o(1))\frac{\log d}{d}n$. Shearer's result was the same except with the words *average* and *maximum* switched. He showed that the largest independent sets in triangle-free graphs with average degree d contain at least $(1+o(1))\frac{\log d}{d}n$ vertices.

1.2 Notation

We write $[n]$ for the set $\{1,...,n\}$, and the notation $x = (1 \pm \varepsilon)$ is used to mean $x \in [1-\varepsilon, 1+\varepsilon]$.

For disjoint sets of vertices X and Y in G we will use $E_G(X,Y)$ to denote the set of edges between X and Y in G and $E_G(X)$ to denote the set of edges of G with both ends in X. We denote the sizes of these sets by $e_G(X, Y)$ and $e_G(X)$ respectively. We will use $N_G(v, X)$ to denote the set of vertices in X which are adjacent to a vertex v of G and $\deg_G(v, X)$ for the number of vertices in $N_G(v, X)$. In a graph G we say a vertex is a *common neighbour* of a pair of vertices if it is adjacent to both of them. For two vertices u, v we will write $N_G(u, v, X)$ for the common neighbourhood $N_G(u, X) \cap N_G(v, X)$ of u and v in X, and $\deg_G(u, v, X)$ for its size. For $X = V(G)$ we will simply use $N_G(v)$, $\deg_G(v)$ and $N_G(u, v)$. When it is clear which graph is being referred to, we omit the subscripts. For a graph G and a vertex set S we will use $G[S]$ to denote the induced subgraph of G on the set S. We let $V(G)$ denote the vertex set of a graph G and let $v(G) = |V(G)|$ be the number of vertices in G. We use $\delta(G)$ to denote the minimum degree of G and $\Delta(G)$ the maximum degree.

We use $o(.)$, $\omega(.)$, $O(.)$, $\Omega(.)$ notation in the standard way, where $f(n) = o(g(n))$ if $\lim_{n\to\infty} f(n)/g(n) = 0$ and $f(n) = \omega(g(n))$ if the same limit tends to ∞. We say $f(n) = O(g(n))$ if there is a constant C such that, for all n, $f(n) \leqslant g(n)$ and say that $f(n) = \Omega(g(n))$ if there is a constant $c > 0$ such that for all n we have $f(n) \geqslant c \cdot g(n)$. We use $\ll$ only in informal discussions where $a \ll b$ will mean that our argument holds so long as a is small enough compared to b (depending just on b). Throughout this thesis we shall omit floor and ceiling symbols when they do not affect our argument.

The next few sections will give an overview of some of the themes of this thesis including some history, related work as well as some methods we will make direct use of.

1.3 Chromatic threshold and minimum degree conditions in triangle-free graphs

One of the earliest results in extremal graph theory is Mantel's Theorem [70] which states that any graph on n vertices with more than $\lfloor n/2 \rfloor \cdot \lceil n/2 \rceil$ edges contains a triangle. An even simpler (and weaker) statement says that if a graph on n vertices has minimum degree greater than $n/2$ it must contain a triangle. This can be seen by noting that for any edge, by the pigeon-hole principle, the neighbourhoods of the endpoints must overlap, giving a triangle. In fact this shows further that in a graph with minimum degree greater than $n/2$ *every* edge is in a triangle. The minimum degree condition here is tight as shown by a complete balanced bipartite graph.

But what if we did reduce the minimum degree condition? What would we be able to say about graphs that satisfy some smaller minimum degree condition but still happen to be triangle free? Andrásfai, Erdős and Sós [8] approached this question, proving that so long as the minimum degree of an n-vertex, triangle-free graph is greater than $2n/5$, the graph is bipartite. We will give a short proof of this result here.

Theorem 1.1 (Andrásfai, Erdős and Sós). *If G is a triangle-free graph on n vertices with $\delta(G) > 2n/5$ then G is bipartite.*

Proof. Suppose for contradiction that G is a counter-example on n vertices and furthermore assume G has the most edges of any counter-example on n vertices. As G is not bipartite there must be an odd cycle. Choose a shortest odd cycle with vertices $v_1, ..., v_k$. We claim that the shortest cycle is a C_5. If not we could add the edge v_1v_4 and the graph would still be triangle free, as if there was a triangle v_1v_4x then we would have had a C_5 in G on the vertices $v_1v_2v_3v_4x$. This is a contradiction to the choice of G as the counter-example with the most edges. So now the vertices $v_1, ..., v_5$ form a C_5 and we know that each of these five vertices has more than $2n/5$ neighbours. From this we can see that there are at least $2n - 4$ edges leaving this C_5. By the pigeon-hole principle one of the $n - 5$ vertices not in the C_5 is adjacent to at least $\frac{2n-4}{n-5} > 2$ vertices of the C_5. As this vertex is adjacent to an integer number of vertices of the C_5 it is adjacent to at least three such vertices, two of these must be adjacent and this gives a triangle. This contradiction completes the proof. □

It is easy to extend the proof above to see that the only non bipartite triangle-free graphs with minimum degree *at least* $2n/5$ (rather than strictly greater than $2n/5$) are complete balanced blow-ups of C_5. That is, graphs which can be partitioned into 5 equal sizes vertex sets $V_1, ..., V_5$ with edges just between V_i and V_{i+1} modulo 5.

There followed a series of results showing that triangle-free graphs satisfying lower minimum degree conditions still have small chromatic number. Häggkvist [48] showed that triangle-free graphs with minimum degree greater than $\frac{3n}{8}$ are 3-colourable. Jin [54] reduced the minimum degree condition of Häggkvist's result to $\frac{10n}{29}$, matching a construction of Häggkvist [48]. Thomassen [83] showed that for any $\varepsilon > 0$ there exists r_ε such that if H is triangle-free and $\delta(H) > (\frac{1}{3} + \varepsilon)n$ then H is r_ε-partite. Finally, in 2006, Brandt and Thomassé [18] proved that all triangle-free graphs with minimum degree strictly greater than $\frac{n}{3}$ are 4-partite. No similar result with a minimum degree condition lower than $\frac{n}{3}$ holds as shown by Hajnal (see [38]) who exhibited triangle-free graphs with minimum degree $(\frac{1}{3} - \varepsilon)n$ and arbitrarily high chromatic number.

The fact that minimum degree conditions above $\frac{1}{3}n$ lead to triangle-free graphs having small chromatic number whilst triangle-free graphs with smaller minimum degree can have arbitrarily large chromatic number leads to $\frac{1}{3}$ being known as the *Chromatic Threshold* of the triangle.

For a graph H, the Chromatic Threshold of H, $\delta_\chi(H)$, is defined as the infimum over all $d > 0$ such that there exists a constant $C = C(H, d)$ such that every H-free graph G with minimum degree at least $d|G|$ satisfies $\chi(G) \leqslant C$. Thus for the triangle K_3 we have $\delta_\chi(K_3) = \frac{1}{3}$. In [3] the chromatic threshold for all non-bipartite graphs was determined. For $r \geqslant 3$ it was shown that every graph with chromatic number r has chromatic threshold either $\frac{r-3}{r-2}$, or $\frac{2r-5}{2r-3}$, or $\frac{r-2}{r-1}$.

1.4 The random graph model $G(n,p)$

The random graph model $G(n,p)$ where n is a natural number and $p \in [0,1]$ is a way of randomly generating a graph on n vertices. For each pair of vertices x, y a biased coin is flipped. If the coin lands on heads (which happens with probability p) the edge xy is put in the graph. Otherwise it is not. All coin flips are independent of each other.

Of course as the output of this procedure is random we cannot guarantee for sure that the graph created satisfies any interesting properties but we will be able to say that some properties 'almost always hold'.

For a property $\mathcal{P}$ we say that $\mathcal{P}$ holds *asymptotically almost surely* (a.a.s) or *with high probability* (w.h.p) if

$$\lim_{n\to\infty} \mathbb{P}(G(n,p) \text{ satisfies } \mathcal{P}) = 1\,.$$

As an example we will prove that for any $\varepsilon > 0$ with $p = \omega(n^{-2})$ the random graph $G(n,p)$ a.a.s has $(1 \pm \varepsilon)p\binom{n}{2}$ edges and if $p = \omega(n^{-2/3})$ then a.a.s every vertex is in $(1 \pm \varepsilon)p^3n^2/2$ triangles. To do so we first introduce the Chernoff bound, see for example [53]. We use $\mathrm{Bin}(n,p)$ to denote the binomial distribution with n trials and success probability p for each trial.

Theorem 1.2. *Let X be a random variable with distribution $\mathrm{Bin}(n,p)$ and $0 < \delta < \frac{3}{2}$. Then*

$$\mathbb{P}(X < (1-\delta)\mathbb{E}X) < \exp\left(\tfrac{-\delta^2}{3}\mathbb{E}X\right) \quad \textit{and} \quad \mathbb{P}(X > (1+\delta)\mathbb{E}X) < \exp\left(\tfrac{-\delta^2}{3}\mathbb{E}X\right).$$

Since the number of edges in $G(n,p)$ has distribution $\mathrm{Bin}(\binom{n}{2}, p)$ we see that with probability at least $1 - 2\exp(\frac{-\varepsilon^2}{3}p\binom{n}{2})$ the number of edges in $G(n,p)$ is between $(1-\varepsilon)p\binom{n}{2}$

and $(1+\varepsilon)p\binom{n}{2}$. This probability tends to one for $p=\omega(n^{-2})$.

Counting triangles is slightly harder as the number of triangles is not binomially distributed. We proceed by first showing that all vertices have roughly pn neighbours and then that within each neighbourhood there are roughly $p\binom{pn}{2}$ edges.

First note that for each the vertex the number of neighbours has distribution $\mathrm{Bin}(n-1,p)$ and hence each vertex has $(1-\frac{\varepsilon}{3})p(n-1)$ neighbours with probability at least

$$1-2\exp\left(-\frac{\varepsilon^2}{27}p(n-1)\right),$$

and so the probability that all vertices have this many neighbours is at least

$$1-2n\exp\left(-\frac{\varepsilon^2}{27}p(n-1)\right).$$

For the rest of the argument we assume that all vertices have neighbourhoods in this range. We can now use the Chernoff bound again to say that with high probability the number of edges in each vertex neighbourhood is

$$(1\pm\varepsilon/3)p\binom{(1\pm\varepsilon/3)pn}{2},$$

which lies within $(1\pm\varepsilon)p^3n^2/2$. For a vertex v let X denote the number of edges in the neighbourhood of v; equivalently the number of triangles v is in. Conditioned on the value of $\deg(v)$ the distribution of X is $\mathrm{Bin}(\deg(v),p)$ and so with probability at least

$$1-2\exp\left(-\frac{\varepsilon^2}{27}\mathbb{E}[X]\right)\geqslant 1-2\exp\left(-\frac{\varepsilon^2}{27}(1-\varepsilon)p^3n^2/2\right),$$

the vertex v is in $(1\pm\varepsilon)p^3n^2/2$ triangles. Hence, conditioned on all vertices having degrees in $(1\pm\varepsilon/3)p(n-1)$, with probability at least

$$1-2n\exp\left(-\frac{\varepsilon^2}{27}(1-\varepsilon)p^3n^2/2\right),$$

all vertices are in that many triangles. Putting all this together we see that with probability at least

$$1-2\exp\left(-\frac{\varepsilon^2}{27}(1-\varepsilon)p^3n^2/2\right)-2n\exp\left(-\frac{\varepsilon^2}{27}p(n-1)\right),$$

each vertex is in $(1\pm\varepsilon)p^3n^2/2$ triangles. This probability tends to one for $p=\omega(n^{-2/3})$.

1.5 Sparse analogues of extremal graph theory results

Looking at versions of classical extremal graph theory problems in the setting of random graphs has proved to be an interesting, fruitful and popular area of study.

Recall that Mantel's theorem [70] states that all triangle-free graphs on n vertices have at most $\lceil \frac{n}{2} \rceil \lfloor \frac{n}{2} \rfloor$ edges, achieving equality if and only if the graph is a complete bipartite graph which is as balanced as possible. If one rephrases 'triangle-free graphs on n vertices' as 'triangle-free spanning subgraphs of K_n then it seems natural to replace K_n with other graphs and ask similar questions. If we replace K_n with the random graph $G(n, p)$ we can ask how many edges the largest triangle-free spanning subgraph of $G(n, p)$ typically has. We could also ask whether the largest triangle-free subgraphs of $G(n, p)$ are bipartite as is the case for K_n. Babai, Simonovits and Spencer [9] showed that this is indeed true with high probability as n tends to infinity so long as $p \geqslant \frac{1}{2}$. This was improved by Brightwell, Panagiotou and Steger [19] who proved the same conclusion holds for $p \geqslant n^{-c}$ for some positive number c. DeMarco and Kahn [29] showed that there exists a constant C such that this is still the case for any $p \geqslant C(\log n/n)^{\frac{1}{2}}$.

With the aim of giving a systematic method to approach these kind of problems, Kohayakawa [63] and Rödl (unpublished) developed a sparse analogue of Szemerédi's Regularity Lemma and together with Łuczak [60] formulated the KŁR conjecture which asserts the existence of a corresponding 'counting lemma'. The definitions and statements relating to the original Szemerédi's Regularity Lemma [82] can be obtained from the discussion below by taking $p = 1$. We postpone discussion of the KŁR conjecture to Chapter 2.

We use the following definitions of regularity. We define the *density* $d(U, V)$ of a pair of disjoint vertex sets (U, V) to be the value $e(U, V)/|U||V|$. A pair (U, V) is called (ε, d, p)-*lower-regular* if for any sets $U' \subseteq U$, $V' \subseteq V$ satisfying $|U'| \geqslant \varepsilon|U|$, $|V'| \geqslant \varepsilon|V|$ we have $d(U', V') \geqslant (d - \varepsilon)p$. We say a pair (U, V) is (ε, d, p)-*regular* if $d(U, V) \geqslant dp$ and for any sets $U' \subseteq U$, $V' \subseteq V$ satisfying $|U'| \geqslant \varepsilon|U|$, $|V'| \geqslant \varepsilon|V|$ we have $d(U', V') = (d(U, V) \pm \varepsilon p)$. We say (U, V) is (ε, p)-*regular* if it is (ε, d, p)-*regular* for some d.

An (ε, p)-*regular-partition* of a graph H is a vertex partition $V_0 \cup V_1 \cup \cdots \cup V_t$ of $V(G)$ with $|V_0| \leqslant \varepsilon|V|$ and $|V_1| = |V_2| = \cdots = |V_t|$ such that all but at most $\varepsilon\binom{t}{2}$ pairs (V_i, V_j) with $i, j \geqslant 1$ are (ε, p)-regular. The corresponding (ε, d, p)-*reduced graph* R is the graph with vertex set $[t]$ where ij is an edge precisely if (V_i, V_j) is an (ε, d, p)-lower-regular pair in H. The following version of the Sparse Regularity Lemma can be deduced from [2, Lemma 12]. We show how in Chapter 2.

Lemma 1.3 (Sparse regularity lemma, minimum degree version)**.** *For all $\beta \in [0, 1]$, $\varepsilon > 0$ and every integer t_0 there exists $t_1 \geqslant 1$ such that for all $d \in [0, 1]$ the following holds for any $p > 0$. For any graph G on n vertices with minimum degree βpn, such that for any $X, Y \subseteq V(G)$ with $|X|, |Y| \geqslant \frac{\varepsilon n}{t_1}$ we have $e(X, Y) \leqslant (1 + \frac{1}{1000}\varepsilon^2)p|X||Y|$, there is an (ϵ, p)-*

regular-partition of $V(G)$ with (ε, d, p)-reduced graph R satisfying $\delta(R) \geqslant (\beta - d - \varepsilon)|V(R)|$ and $t_0 \leqslant |V(R)| \leqslant t_1$. Furthermore, for each $i \in V(R)$ the number of $j \in V(R)$ such that (V_i, V_j) is not (ε, p)-regular is at most $\varepsilon v(R)$, and for each $i \in V(R)$ and $v \in V_i$, at most $(d + \varepsilon)pn$ neighbours of v lie in $\bigcup_{j:ij \notin R} V_j$.

Note that the regularity lemma above is not specifically for $G(n, p)$ but for graphs in which the density of edges between pairs of large sets is never much greater than p. For $p = \omega(n^{-1})$ the random graph $G(n, p)$ a.a.s. satisfies this.

The primary reason for the power of the Szemerédi regularity lemma is that it comes with a counting lemma. The counting lemma ensures that once a regular partition is given the number of copies of a small fixed graph H can be estimated accurately. For the *sparse* regularity lemma such a counting result is not true. In particular there are triples (V_1, V_2, V_3) (with $|V_i| = n$ for each i) such that each (V_i, V_j) is (ε, d, p)-regular but with *no* triangles rather than approximately $(dpn)^3$ as one might expect. Such examples however were shown by Kohayakawa, Łuczak and Rödl [60] to be in some sense 'rare' for triples. We discuss this in Chapter 2.

1.6 Graph saturation

For a graph H we say another graph G is *H-saturated* if it contains no copy of H, yet adding any new edge would result in a copy of H. For example a complete bipartite graph is K_3-saturated, as it is triangle free but every pair of non-adjacent vertices has a common neighbour and so adding an edge between any non-adjacent pair would result in a triangle. It is then natural to ask what are the greatest and least number of edges an H-saturated graph can have. Asking for the most edges in an H-saturated graph on a fixed number of vertices is identical to the Turán problem asking for the most edges in an H-free graph; if an H-free graph is not saturated there is an edge that can be added to give an H-free graph with more edges. Looking for the minimum number of edges in an H-saturated graph is a problem with a very different feel. We define $\mathrm{sat}(H, n)$ to be the minimum number of edges over all H-saturated graphs on n vertices and call this the saturation number of H. Returning to triangles we can easily show that $\mathrm{sat}(K_3, n) = n - 1$. Firstly note that if G is K_3-saturated, it must be connected, as adding an edge between distinct connected components cannot create a triangle. Therefore all K_3-saturated graphs contain a spanning tree, and hence at least $n - 1$ edges. The star $K_{1,n-1}$ on n vertices has exactly $n - 1$ edges and is K_3-saturated, giving an extremal construction. The study of these saturation numbers was initiated by Erdős, Hajnal and Moon [36] who proved that $\mathrm{sat}(K_r, n) = (r-2)(n - \frac{1}{2}(r-1))$. Kászonyi and Tuza [57] later showed that cliques have the largest saturation number of any graph on r vertices which in particular implies that for any H the saturation number $\mathrm{sat}(H, n)$ grows linearly in n. This is in stark contrast to the Turán numbers which are quadratic in n for any graph H which is non-bipartite.

1.7 Ramsey theory

Ramsey theory deals with results on partitioned structures, claiming that in any partition of some sufficiently large structure one of the parts will have some nice property. It is often convenient to use colours to label the different parts of the partition. We begin by stating Ramsey's theorem.

Theorem 1.4. *For every natural number k there exists $N \in \mathbb{N}$ such that in any colouring of the edges of K_N with red and blue there exists a monochromatic copy of K_k.*

The Ramsey number $R(k)$ is defined to be the least N for which the above conclusion holds. Thus there must be a way to colour the edges of $K_{R(k)-1}$ with red and blue to avoid a monochromatic K_k but no such way to colour $K_{R(k)}$.

A result of Erdős and Szekeres [40] states that $R(k) \leqslant \binom{2k-2}{k-1} \leqslant 4^k$. We will give a short proof of the upper bound 4^k.

Proof. We do so by proving the stronger statement that the size of the largest red and largest blue cliques add up to at least $2k$. We use induction on k noting that for $k = 1$ the result is clear. Now suppose $k \geqslant 2$ and let G be a complete graph on 4^k vertices in which the edges have each been coloured either red or blue. Choose an arbitrary vertex x. Either x has at least $\frac{1}{2}4^k$ neighbours in blue or in red. Without loss of generality assume x has at least this many neighbours in blue and let A denote the set of vertices adjacent to x in blue. Now choose an arbitrary vertex y in A. The vertex y either has at least 4^{k-1} blue neighbours in A or that many red neighbours in A. Let B be the set of at least 4^{k-1} vertices that are adjacent to y in the most common colour. By the induction hypothesis there is a red clique and a blue clique in B such that their combined size is at least $2(k-1)$. We can add the vertex x to the blue clique and the vertex y to the clique of the same colour that y is adjacent to B in. This gives a red clique and a blue clique covering at least $2k$ vertices in total. $\square$

Erdős [32] also gave a lower bound on the Ramsey number $R(k)$ using the probabilistic method to show that $R(k) \geqslant 2^{(k-1)/2}$. We give the proof here.

Proof. Letting $N = \lceil 2^{(k-1)/2} \rceil$ we need to show the existence of a colouring of the edges of K_N with two colours such that there is no monochromatic copy of K_k. Rather than constructing such a colouring explicitly, we will colour K_N randomly and show that there is a positive probability that the colouring avoids containing a monochromatic K_k. If we colour each edge red or blue independently with 50% probability of each colour, then the expected number of monochromatic copies of K_k is

$$\binom{N}{k} 2^{1-\binom{k}{2}} .$$

If we can show this is less than one then there must have been a colouring with no monochromatic copies of K_k. Recalling that $N = \lceil 2^{(k-1)/2} \rceil$ we see that

$$\binom{N}{k} < \frac{N^k}{k!} \leqslant \frac{2^{\binom{k}{2}}}{k!} ,$$

and so the expected number of monochromatic copies of K_k is less than $2/k! < 1$. □

We can generalise these ideas in a number of ways. We could ask to find monochromatic copies of graphs other than cliques and this has been a large area of research over many years with Ramsey numbers of paths, cycles, trees and bounded degree graphs in particular receiving a lot of attention. We could also change the graph that is being coloured. We say that a graph G is *Ramsey for* H if any 2-colouring of the edges of G contains a monochromatic copy of H. We denote this by $G \to H$. In particular $K_{R(k)} \to K_k$. This sets us up to introduce the concept of *size-Ramsey* numbers. Rather than looking for how small a complete graph can be such that however it is coloured it contains a particular monochromatic subgraph, the study of size-Ramsey asks for how few edges a graph can have whilst being Ramsey for H. We define the size-Ramsey number $\hat{r}(H)$ to be this minimum.

$$\hat{r}(H) := \min\{e(G) : G \to H\} ,$$

where the minimum is taken over all graphs. It is clear that for any H we have $\hat{r}(H) \leqslant \binom{R(H)}{2}$ as the clique on $R(H)$ vertices is Ramsey for H by definition and has this many edges. The study of these numbers was introduced by Erdős, Faudree, Rousseau and Schelp [35] who in particular were interested in graphs where the trivial bound could be substantially improved.

Erdős [34] posed the problem of determining the order of magnitude of $\hat{r}(P_n)$. This was resolved by Beck [11], who showed that the size-Ramsey numbers of paths are linear in the length of the path.

1.8 Partition functions

In Chapter 5 we use the monomer-dimer and hard-core models of statistical physics to study problems in d-regular graphs. These models give a probability distribution for selecting a random matching and independent set respectively from a graph.

1.8.1 The monomer-dimer model

The matching polynomial (or matching partition function) of a graph G is defined to be

$$M_G(\lambda) = \sum_{M \in \mathcal{M}} \lambda^{|M|} \, ,$$

where $\mathcal{M}$ is the set of all matchings in G, including the empty matching. For a matching M the size of M, denoted by $|M|$, is the number of *edges* in M. In particular we see that $M_G(1)$ is the total number of matchings in G. Letting λ tend to infinity we see that

$$pm(G) = \lim_{\lambda \to \infty} \frac{M_G(\lambda)}{\lambda^{|G|/2}} \, ,$$

where $pm(G)$ denotes the number of perfect matchings of G. For each $\lambda > 0$ we can define a probability distribution on $\mathcal{M}$ by choosing each matching M with probability proportional to $\lambda^{|M|}$. Thus, a particular matching M is chosen with probability

$$\frac{\lambda^{|M|}}{M_G(\lambda)} \, .$$

This distribution is known as the monomer-dimer model. Setting $\lambda = 1$ gives the uniform distribution on the set of matchings. The edges of the matching are the 'dimers' whilst the vertices that are not incident to any dimers are referred to as 'monomers'. This notion comes from chemistry, originating with the study of Roberts [78] on the adsorption of oxygen and hydrogen on a flat tungsten surface. The λ parameter is known as the *fugacity.*

For more on the monomer-dimer model we recommend reading [51] which in particular shows there is never a phase transition in the monomer-dimer model.

1.8.2 The hard-core model

The independence polynomial (or independent set partition function) of a graph G is

$$P_G(\lambda) = \sum_{I \in \mathcal{I}} \lambda^{|I|},$$

where $\mathcal{I}$ denotes the set of independent sets in G including the empty set. For $\lambda = 1$ this gives the total number of independent sets in G. Similarly to matchings, we can use this to define a probability distribution on the set of independent sets of a graph. This will be called the hard-core model. Each independent set I will be chosen with probability

$$\frac{\lambda^{|I|}}{P_G(\lambda)},$$

and so again setting $\lambda = 1$ gives the uniform distribution.

The origin of the hard-core model lies with the hard-sphere model. The hard-sphere model is a probability distribution of unit radius spheres placed randomly in a large volume such that no two overlap. This gives a way of modelling the distribution of gases. The hard-core model is a discretisation of the hard-sphere model. The vertices of the independent set can be thought of as the centres, or *cores*, of the spheres whilst requiring the random set be independent ensures the centres are not too close.

1.8.3 Spatial Markov property

An important feature of the two models above, which we use heavily in Chapter 5, is that they have what is known as the spatial Markov property. Specifically, with the hard-core model, if we choose a set A of vertices that separates the graph G into multiple connected components $G_1, G_2, ...$ then conditioned on which vertices of A are in the independent set I the intersection of I with each G_i is independent. In particular, if we choose a vertex v and a natural number r we can look at the set of vertices at distance exactly r from v. Call this set $N^r(v)$. Let $B^r(v)$ denote the set of all vertices at distance *at most* r from v. If we condition on $I \cap N^r(v)$ then we can determine the distribution of $I \cap B^{r-1}(v)$ without knowing anything about the graph outside $N^r(v)$. In particular, conditioned on $I \cap N^r(v)$ the sets $I \cap B^{r-1}$ and $I \setminus B^r(v)$ are independent. Furthermore, if we let $\tilde{B}^{r-1}(v)$ denote the vertices in $B^{r-1}(v)$ that are not adjacent to a vertex of $I \cap N^r(v)$ then the distribution of $I \cap \tilde{B}^{r-1}(v)$ is precisely the hard-core model run on $G[\tilde{B}^{r-1}(v)]$ with the same value of λ.

1.9 Linear programming and duality

Linear programming deals with optimising the value of a linear expression under linear inequality constraints. For example we may wish to maximise the expression

$$2x + y$$

under the constraints

$$x \geqslant 0, \quad y \geqslant 0, \quad x + y \leqslant 1, \quad 4x + y \leqslant 2\,.$$

It is helpful to write a linear program in its *canonical form*, that is:

$$\begin{array}{ll} \text{maximise} & \mathbf{c}^T\mathbf{x} \\ \text{subject to} & A\mathbf{x} \leqslant \mathbf{b} \\ \text{and} & \mathbf{x} \geqslant 0\,. \end{array}$$

In the example above we have

$$\mathbf{c} = \begin{pmatrix} 2 \\ 1 \end{pmatrix}, \quad \mathbf{x} = \begin{pmatrix} x \\ y \end{pmatrix}, \quad A = \begin{pmatrix} 1 & 1 \\ 4 & 1 \end{pmatrix}, \quad \mathbf{b} = \begin{pmatrix} 1 \\ 2 \end{pmatrix}.$$

Such a problem as above is referred to as the *primal* problem. We can convert the primal into a *dual* problem. The dual gives an upper bound on the solution to the primal. For a primal in the canonical form above the dual program is:

$$\begin{array}{ll} \text{minimise} & \mathbf{b}^T\mathbf{y} \\ \text{subject to} & A^T\mathbf{y} \geqslant \mathbf{c} \\ \text{and} & \mathbf{y} \geqslant 0\,. \end{array}$$

It is possible for the primal to be *unbounded* (with $\mathbf{c}^T\mathbf{x}$ able to be arbitrarily large) or *infeasible* (meaning no value of $\mathbf{x}$ satisfies all constraints). Similarly the dual may also be unbounded or infeasible.

The 'weak duality theorem' tells us that for any $\mathbf{x}$ that satisfies the primal constraints and any $\mathbf{y}$ that satisfies the dual constraints the value of $\mathbf{c}^T\mathbf{x}$ is always at most that of $\mathbf{b}^T\mathbf{y}$. The 'strong duality theorem' states that if the primal has an optimal solution $\mathbf{x}^*$ then the dual also has an optimal $\mathbf{y}^*$ and furthermore $\mathbf{c}^T\mathbf{x}^* = \mathbf{b}^T\mathbf{y}^*$. That is to say that when an optimum value exists for the primal, the dual also has the same optimal value. By the weak duality theorem if we have a value that we believe is optimal for the primal it is enough to find some $\mathbf{y}$ satisfying the dual constraints such that $\mathbf{b}^T\mathbf{y}^*$ attains this value.

After a little thought we may suspect that the optimal value of the example primal above is $\frac{4}{3}$ attained at $(\frac{1}{3}, \frac{2}{3})$. With $\mathbf{y} = (\frac{2}{3}, \frac{1}{3})$ we see that $\mathbf{b}^T\mathbf{y}$ also attains the value $\frac{4}{3}$. This confirms that $\frac{4}{3}$ is indeed the optimal value of the primal.

In our applications of linear programming we will in fact use equality constraints rather than inequality constraints. Equality constraints can be created by using two inequality constraints. The main change is that the extra constraints in the primal give extra freedom in the dual which means we do not have to require the dual variables to be positive.

2

Triangle-free subgraphs of random graphs

2.1 Introduction

In a 1948 edition of the recreational maths journal Eureka, Blanche Descartes [30] proved that triangle-free graphs can have arbitrarily large chromatic number, and thus be complex in structure. This motivates the question of which additional restrictions on the class of triangle-free graphs allow for a bound on the chromatic number. By Mantel's theorem [70], the densest triangle-free graphs are balanced complete bipartite graphs. So we may first ask whether triangle-free graphs H with minimum degree somewhat below $\frac{1}{2}v(H)$ are still necessarily bipartite. This is true, as Andrásfai, Erdős and Sós showed in 1974.

Theorem 2.1 (Andrásfai, Erdős, Sós [8]). *All triangle-free graphs H with $\delta(H) > \frac{2}{5}v(H)$ are bipartite.*

Triangle-free graphs of smaller minimum degree do not need to be bipartite, as blow-ups of a 5-cycle illustrate. But one may still ask whether their chromatic number is bounded (questions of this type were first addressed by Erdős and Simonovits in [38]). In 2002 Thomassen [83] proved that this is the case for triangle-free graphs of minimum degree at least $(\frac{1}{3} + \varepsilon)n$.

Theorem 2.2 (Thomassen [83]). *For any $\varepsilon > 0$ there exists r_ε such that if H is triangle-free and $\delta(H) > (\frac{1}{3} + \varepsilon)v(H)$ then H is r_ε-colourable.*

A construction of Hajnal (see [38]) shows that the minimum degree bound in this theorem cannot be replaced by $(\frac{1}{3} - \varepsilon)n$. A much stronger result was established by Brandt and

Thomassé [18], who showed that triangle-free graphs H with $\delta(H) > \frac{1}{3}n$ are 4-colourable.

In this chapter we are interested in random graph analogues of Theorem 2.1 and Theorem 2.2. Establishing such analogues for prominent results in extremal graph theory has been a particularly fruitful area of study in the last few years. A good overview can be found in Conlon's survey paper [23].

In order to study these kinds of questions systematically, Kohayakawa [63] and Rödl (unpublished) developed a sparse analogue of Szemerédi's Regularity Lemma, and, together with Łuczak [60] formulated the KŁR conjecture which asserts the existence of a corresponding 'counting lemma'. Recently Conlon, Samotij, Schacht and Gowers [25] proved this conjecture (see also [10, 79]). It is easy (as observed in [25]) to use these results to prove 'approximate' random versions of Theorems 2.1 and 2.2, as well as to re-prove Mantel's theorem for random graphs. Thus if $p \gg n^{-1/2}$ then asymptotically almost surely (a.a.s.) the random graph $G(n,p)$ has the property that all subgraphs with minimum degree a little larger than $\frac{2}{5}pn$ can be made bipartite by deleting $o(pn^2)$ edges. Similarly, the sparse random version of Mantel's theorem obtained states that any subgraph with a little more than half the edges of $G(n,p)$ contains a triangle.

One might expect that all subgraphs of $G(n,p)$ with minimum degree a little larger than $\frac{2}{5}pn$ are bipartite. Indeed, an alternative sparse random version of Mantel's theorem, proved by DeMarco and Kahn [29], states that a largest triangle-free subgraph of $G(n,p)$ coincides exactly with a largest bipartite subgraph for $p \gg (\log n/n)^{1/2}$. However, subgraphs of $G(n,p)$ with minimum degree larger than $\frac{2}{5}pn$ which are not bipartite do exist (see Theorem 2.5 below). We determine for all p how far from bipartite such graphs can be.

Theorem 2.3. *For any $\gamma > 0$, there exists C such that for any $p(n)$ the random graph $\Gamma = G(n,p)$ a.a.s. has the property that all triangle-free spanning subgraphs $H \subseteq \Gamma$ with $\delta(H) \geqslant (\frac{2}{5}+\gamma)pn$ can be made bipartite by removing at most $\min\left(Cp^{-1}n, (\frac{1}{4}+\gamma)pn^2\right)$ edges.*

In addition we derive an analogous random graph version of Theorem 2.2.

Theorem 2.4. *For any $\gamma > 0$, there exist C and r such that for any $p(n)$ the random graph $\Gamma = G(n,p)$ a.a.s. has the property that all triangle-free spanning subgraphs $H \subseteq \Gamma$ with $\delta(H) \geqslant (\frac{1}{3}+\gamma)pn$ can be made r-partite by removing at most $\min\left(Cp^{-1}n, (\frac{1}{2r}+\gamma)pn^2\right)$ edges.*

Up to the values of C, these theorems are best possible as shown by the theorem below.

Theorem 2.5. *For any $\gamma > 0$ and $r \in \mathbb{N}$, there exist constants $c, c' > 0$ such that if $n^{-1/2}/c' \leqslant p(n) \leqslant c'$ then $\Gamma = G(n,p)$ a.a.s has a triangle-free spanning subgraph H with $\delta(H) \geqslant (\frac{1}{2} - \gamma)pn$ which cannot be made r-partite by removing fewer than $cp^{-1}n$ edges.*

Note that for $p \ll n^{-1/2}$ the minimum in each of Theorems 2.3 and 2.4 is achieved by the second term and that these statements are easy: For such values of p only a tiny fraction of the edges of $G(n,p)$ are in triangles and the question reduces to asking for the largest bipartite (respectively, r-partite) subgraph of $G(n,p)$. For p close to 1, by the original Theorems 2.1 and 2.2, the conclusion of Theorem 2.5 becomes false. For this reason we need a condition of the form $p \leqslant c'$.

It would be interesting to obtain analogous results for K_r-free subgraphs of $G(n,p)$ for $r > 3$. It would also be interesting to know whether Theorem 2.4 could be improved to generalise the result of Brandt and Thomassé. We conjecture that this is the case.

Organisation In Section 2 we will introduce some of the main tools that will be used throughout the chapter. Section 3 of this chapter will give a method of constructing a triangle-free subgraph from a given, randomly generated graph. We will then prove a series of results about this construction which will result in proving Theorem 2.5. In Section 4 we will state and prove some properties that a.a.s. $\Gamma = G(n,p)$ possesses. We will then use these properties in Section 5 to prove Theorem 2.3, and in Section 6 to prove Theorem 2.4.

2.2 Tools

Probability We write $\mathrm{Bin}(n,p)$ for the binomial distribution with n trials and success probability p. Our proofs we will make frequent use of the following Chernoff bound, which is an immediate corollary of [53, Theorem 2.1].

Lemma 2.6 (Chernoff bound)**.** *Let X be a random variable with distribution* $\mathrm{Bin}(n,p)$ *and* $0 < \delta < \frac{3}{2}$. *Then*

$$\mathbb{P}(X < (1-\delta)\mathbb{E}X) < \exp\left(\tfrac{-\delta^2}{3}\mathbb{E}X\right) \quad \textit{and} \quad \mathbb{P}(X > (1+\delta)\mathbb{E}X) < \exp\left(\tfrac{-\delta^2}{3}\mathbb{E}X\right).$$

Sparse regularity We define the *density* $d(U,V)$ of a pair of disjoint vertex sets (U,V) to be the value $e(U,V)/|U||V|$. A pair (U,V) is called (ε,d,p)-*lower-regular* if for any sets $U' \subseteq U$, $V' \subseteq V$ satisfying $|U'| \geqslant \varepsilon|U|$, $|V'| \geqslant \varepsilon|V|$ we have $d(U',V') \geqslant (d-\varepsilon)p$. We say a pair (U,V) is *(ε,d,p)-regular* if $d(U,V) \geqslant dp$ and for any sets $U' \subseteq U$, $V' \subseteq V$ satisfying $|U'| \geqslant \varepsilon|U|$, $|V'| \geqslant \varepsilon|V|$ we have $d(U',V') = (d(U,V) \pm \varepsilon)$. We say (U,V) is *(ε,p)-regular* if it is *(ε,d,p)-regular* for some d.

An *(ε,p)-regular-partition* of a graph H is a vertex partition $V_0 \cup V_1 \cup \cdots \cup V_t$ of $V(G)$ with $|V_0| \leqslant \varepsilon|V|$ and $|V_1| = |V_2| = \cdots = |V_t|$ such that all but at most $\varepsilon\binom{t}{2}$ pairs (V_i,V_j) with $i,j \geqslant 1$ are (ε,p)-regular. The corresponding *(ε,d,p)-reduced graph R* is the graph with vertex set $[t]$ where ij is an edge precisely if (V_i,V_j) is an (ε,d,p)-lower-regular pair in H. The following version of the Sparse Regularity Lemma can be deduced from [2, Lemma 12]

Lemma 2.7 (Sparse Regularity Lemma, Minimum Degree Form Version)**.** *For all $\beta \in [0,1]$, $\varepsilon > 0$ and every integer t_0 there exists $t_1 \geqslant 1$ such that for all $d \in [0,1]$ the following holds for any $p > 0$. For any graph G on n vertices with minimum degree βpn, such that for any $X,Y \subseteq V(G)$ with $|X|,|Y| \geqslant \frac{\varepsilon n}{t_1}$ we have $e(X,Y) \leqslant (1+\frac{1}{1000}\varepsilon^2)p|X||Y|$, there is a regular-partition of $V(G)$ with (ε,d,p)-reduced graph R satisfying $\delta(R) \geqslant (\beta-d-\varepsilon)|V(R)|$ and $t_0 \leqslant |V(R)| \leqslant t_1$. Furthermore, for each $i \in V(R)$ the number of $j \in V(R)$ such that (V_i,V_j) is not (ε,p)-regular is at most $\varepsilon v(R)$, and for each $i \in V(R)$ and $v \in V_i$, at most $(d+\varepsilon)pn$ neighbours of v lie in $\bigcup_{j:ij\notin R} V_j$.*

The statement above is identical to that in [2] except for the final 'Furthermore' conclusion. That we can assume no part is in many irregular pairs follows from the proof there. The final condition can be obtained by applying the statement in [2] with $\varepsilon/100$ replacing ε and removing vertices from $V_1,\ldots,V_{v(R)}$ to V_0, keeping the sizes of the V_i equal, until no vertices failing the condition remain. Initially, by regularity and by the upper bound on densities in G, we remove at most $\frac{\varepsilon}{20}n$ vertices. Thereafter, we remove vertices only because they have at least $\varepsilon pn/2$ neighbours in the current set V_0. If at some point in the process V_0 has $\varepsilon n/10$ vertices, then it contains at least $\varepsilon^2pn^2/40$ edges, so contains a bipartite subgraph with at least $\varepsilon^2pn^2/80$ edges, in contradiction to the density assumption on G. We conclude the process stops before this point, as desired.

Note that the regularity lemma above is not specifically for $G(n,p)$ but for graphs in which the density edges between pairs of large sets is never much greater than p. For $p = \omega(\frac{\log n}{n})$ the random graph $G(n,p)$ a.a.s. satisfies this, see for example Lemma 2.14 part (c).

When applying the Sparse Regularity Lemma we will wish to say that if H is triangle-free then the reduced graph is also triangle-free. In order to do this we use the following regularity inheritance lemma, which is [4, Lemma 1.27] and is based on techniques from [62].

Lemma 2.8 (Regularity Inheritance)**.** *For any $0 < \varepsilon', d$ there exist ε_0 and C' such that for any $0 < \varepsilon < \varepsilon_0$ and any $0 < p = p(n) < 1$ the random graph $\Gamma = G(n,p)$ a.a.s. has the following property. For any $X, Y \subseteq V(\Gamma)$ with $|X|, |Y| \geqslant C' \max\{p^{-2}, p^{-1}\log n\}$ and any subgraph H of $\Gamma[X,Y]$ which is (ε, d, p)-lower-regular, there are at most $C' \max\{p^{-2}, p^{-1}\log n\}$ vertices v of $V(\Gamma)$ such that $(X \cap N_\Gamma(v), Y \cap N_\Gamma(v))$ is not (ε', d, p)-lower-regular in H.*

We shall also want the following consequence of this lemma, stating that for every regular partition of every $H \subseteq G(n,p)$ the neighbourhoods of most vertices induce lower-regular subgraphs on the regular pairs of the partition.

Lemma 2.9. *For any $0 < \varepsilon', d < 1$ there exist ε_0 and C' such that for any $t_1 \in \mathbb{N}$ and any $p > 2C't_1 n^{-1/2}$ the random graph $\Gamma = G(n,p)$ a.a.s. satisfies the following. For any $0 < \varepsilon < \varepsilon_0$, any spanning subgraph H of Γ and any (ε, d, p)-regular-partition $V_0 \cup V_1 \cup \cdots \cup V_t$ of H with $t \leqslant t_1$ and reduced graph R, all but at most $\binom{t_1}{2} C' \max\{p^{-2}, p^{-1}\log n\}$ vertices v of H have the property that for each $ij \in E(R)$ the pair $(N_\Gamma(v) \cap V_i, N_\Gamma(v) \cap V_j)$ is (ε', d, p)-lower-regular in H.*

Proof. By applying Lemma 2.8 with ε' and d we are given ε_0 and C'. Suppose $p \geqslant 2C'tn^{-1/2}$ and that Γ satisfies the probable event of Lemma 2.8. Now let $H \subseteq \Gamma$ and a partition $V_0 \cup V_1 \cup \cdots \cup V_t$ of H with reduced graph R be given. Let $ij \in E(R)$. For large enough n we have $C' \max\{p^{-2}, p^{-1}\log n\} \leqslant C' \max\{\frac{n}{4C'^2 t_1^2}, \frac{\sqrt{n}\log n}{2C't_1}\} \leqslant \frac{n}{2t_1} \leqslant |V_i|, |V_j|$. So we conclude from Lemma 2.8 that for all but at most $C' \max\{p^{-2}, p^{-1}\log n\}$ vertices $v \in V(H)$ the pair $(N_\Gamma(v) \cap V_i, N_\Gamma(v) \cap V_j)$ is (ε', d, p)-lower-regular in H. The lemma follows by summing over all $ij \in E(R)$. □

The following lemma combines Lemma 2.7 with Lemma 2.8 to give a regular partition of a triangle-free subgraph H for which the reduced graph is triangle-free.

Lemma 2.10. *For any $0 < \varepsilon, d, \beta < 1$ and any t_0 there exist c and t_1 such that for $p \geqslant cn^{-1/2}$ in $\Gamma = G(n,p)$ a.a.s. any triangle-free subgraph H with $\delta(H) > \beta pn$ has an (ε, d, p)-regular-partition $V_0 \cup V_1 \cup \cdots \cup V_t$ with $t_0 \leqslant t \leqslant t_1$ such that the corresponding reduced graph R is triangle-free and has minimum degree at least $(\beta - d - \varepsilon)v(R)$.*

Proof. Suppose we are given $\varepsilon, d, \beta, t_0$ as in the lemma statement. Set $\varepsilon' = \frac{d}{3}$ and apply Lemma 2.8 (Regularity Inheritance) to ε' and d to obtain ε_0 and C'. Now apply Lemma 2.7

(Sparse Regularity, Minimum Degree Form) with d, β, t_0 as given and with ε also required to be smaller than ε_0. This gives t_1. Take $c = 6t_1C'$.

Lemma 2.7 has given us an (ε, d, p)-regular-partition of H with reduced graph R that satisfies all the conditions we require except that of R being triangle free. Suppose for a contradiction there is a triangle in R. This corresponds to an (ε, d, p)-lower-regular triple (X, Y, Z). First observe that $|X| = |Y| \geqslant \frac{n}{2t_1}$ and for $p(n) \geqslant cn^{-1/2}$ we have $\frac{n}{4t_1} > C' \max\{p^{-2}, p^{-1}\log n\}$. By lower-regularity of (X, Z) and (Y, Z), at least $\frac{1}{2}|Z|$ vertices z of Z have $\deg_H(z, X) \geqslant \frac{d}{2}p|X|$ and also $\deg_H(z, Y) \geqslant \frac{d}{2}p|Y|$. Furthermore, for all but at most $C' \max\{p^{-2}, p^{-1}\log n\} \leqslant \frac{|Z|}{3}$ vertices z of Z, the pair $\big(N_\Gamma(z, X), N_\Gamma(z, Y)\big)$ is (ε', d, p)-lower-regular. Choosing a vertex $z \in Z$ which satisfies both conditions, by regularity of $\big(N_\Gamma(z, X), N_\Gamma(z, Y)\big)$ the edge density of $\big(N_H(z, X), N_H(z, Y)\big)$ is at least $(d - \varepsilon)p > 0$. This gives a triangle, the desired contradiction. $\square$

Finally, we need the following special case of the Slicing Lemma.

Lemma 2.11 (Slicing Lemma). *Let (V_i, V_j) be (ε, d, p)-lower-regular. For any $X \subseteq V_i$, $Y \subseteq V_j$ such that $|X| \geqslant d|V_i|, |Y| \geqslant d|V_j|$ the pair (X, Y) is $(\frac{\varepsilon}{d}, d, p)$-lower-regular.*

Proof. Let $X' \subseteq X$, $Y' \subseteq Y$ satisfy $|X'| \geqslant \frac{\varepsilon}{d}|X| \geqslant \varepsilon|V_i|$ and $|Y'| \geqslant \frac{\varepsilon}{d}|Y| \geqslant \varepsilon|V_j|$. So $d(X', Y') \geqslant (d - \varepsilon)p \geqslant \big(d - \frac{\varepsilon}{d}\big)p$. $\square$

2.3 Proof of Theorem 2.5

Recall that Theorem 2.5 asserts that for any $\gamma > 0$ and $r \in \mathbb{N}$, there are $c, c' > 0$ such that for any $n^{-1/2}/c' \leqslant p \leqslant c'$ the random graph $G(n, p)$ a.a.s. contains a subgraph which is triangle-free, whose minimum degree is at least $\big(\frac{1}{2} - \gamma\big)pn$, and which cannot be made r-partite by removing any $cp^{-1}n$ edges.

The idea of the proof of this theorem is as follows. Let $\Gamma = G(n, p)$ and partition $[n]$ into sets $B = [n/2]$ and $A = [n] \setminus B$. We remove all edges in A. We further 'sparsify' $\Gamma[B]$, keeping edges with a suitable probability p'. The goal of this 'sparsification' is to obtain a subgraph of $\Gamma[B]$ which is still complex enough for the rest of the argument, but is such that for each vertex a in A the number of edges in $N(a, B)$ is negligible compared to the degree of a (see Lemma 2.12(*b*)). Observe that this subgraph is distributed as the following inhomogeneous random graph model. We define $G(n, p, p')$ to be the random graph on $[n]$ obtained by letting pairs of vertices within $[n/2]$ be edges independently with probability pp', letting pairs in $[n] \setminus [n/2]$ all be non-edges, and letting all other pairs be edges independently with probability p.

We next use the fact, first proved in [30], that there exists a triangle-free graph F which is not r-partite. Let $[\ell]$ be the vertex set of F. We place a 'random blow-up' of F into B as follows: We partition B into ℓ equal sets $B_1, \ldots, B_\ell$ and keep only those edges in B running between B_i and B_j with $ij \in F$. Finally, we remove in B all edges with an endpoint whose degree in B deviates too much from expectation, and then all edges between A and B which are in a triangle with a vertex from B. This last step is the only step in which we delete edges between A and B.

It is easy to check that the resulting graph is triangle-free by construction. Using some properties of $G(n,p,p')$ and the blow-up of F we can also show that it cannot be made r-partite by deleting $cp^{-1}n$ edges. Moreover, using the fact that for each vertex a in A the number of edges in $N(a,B)$ is small and hence in the last step not many edges were deleted at any vertex, we can also conclude that the minimum degree of the resulting graph is at least $\left(\frac{1}{2} - \gamma\right)pn$.

The typical properties of $G(n,p,p')$ we need are the following.

Lemma 2.12. *For any $\varepsilon > 0$ and $K \geqslant 10$, there exists $0 < c < \varepsilon$ such that the following holds. If $Kn^{-1/2} \leqslant p(n) \leqslant \varepsilon^2 c/(10^4 K^2)$ and $p' = cK^2 p^{-2} n^{-1}$, then a.a.s. the random graph $G(n,p,p')$ has the following properties. Let $B = [n/2]$ and $A = [n] \setminus B$.*

(*a*) $\deg(b,A), \deg(a,B) = \left(\frac{1}{2} \pm \varepsilon\right)pn$ *for every $a \in A$ and $b \in B$.*

(*b*) *For each $a \in A$, at most $p'p^3n^2$ edges have both ends in $N(a,B)$.*

(*c*) *For each $b \in B$ with $\deg(b,B) \geqslant \frac{1}{10}p'pn$, the number of vertices $a \in A$ such that there exists $b' \in B$ with abb' a triangle is at most $pn\big(1 - (1-p)^{\deg(b,B)}\big)$.*

(*d*) *At most $cp^{-1}n$ edges in B are incident to some $b \in B$ with $\deg(b,B) \geqslant pp'n$ or $\deg(b,B) \leqslant \frac{1}{10}p'pn$.*

(*e*) *$e(U,V) > 2cp^{-1}n$ for every pair of disjoint sets $U, V \subseteq B$ with $|U|, |V| \geqslant 2n/K$.*

We delay the proof of this lemma to after the proof of Theorem 2.5.

Proof of Theorem 2.5. Given $\gamma > 0$ and $r \in \mathbb{N}$, let F be a triangle-free graph which is not r-partite. Let $\ell = v(F)$. We set $K = 8r\ell$ and

$$\varepsilon = \tfrac{1}{400}\gamma r^{-2}\ell^{-2} \,. \tag{2.1}$$

Now we let $c > 0$ with $c < \varepsilon$ be returned by Lemma 2.12 for input ε and K. We choose $c' = \min\left(\frac{1}{K}, \frac{c}{10^4}\right)$.

Given $n^{-1/2}/c' \leqslant p(n) \leqslant c'$, let $p' = cK^2p^{-2}n^{-1}$. Observe that $p' \leqslant 1$ by choice of p. Let $B = [n/2]$, and $A = [n] \setminus B$. We generate $\Gamma = G(n,p)$, and let G_1 be the subgraph of Γ obtained by sparsifying B, keeping edges independently with probability p' and removing all edges of A. Since G_1 is distributed as $G(n,p,p')$, by Lemma 2.12 it a.a.s. satisfies the properties (a)–(e). We now condition on G_1 satisfying these properties.

Partition B into ℓ equal sets $B_1, \dots, B_\ell$. Let G_2 be the subgraph of G_1 obtained by keeping only edges of the form ab with $a \in A$ and $b \in B$, or of the form bb' with $b \in B_i$ and $b' \in B_j$ for some $ij \in F$. We claim that $G_2[B]$ is far from r-partite.

Claim 2.13. *$G_2[B]$ cannot be made r-partite by deleting any $2cp^{-1}n$ edges.*

Proof. Given a (not necessarily proper) r-colouring $\chi : B \to [r]$, we define a majority r-colouring $\chi' : [\ell] \to [r]$ by setting $\chi'(i)$ equal to the smallest j such that $\left|\chi^{-1}(j) \cap B_i\right| \geqslant |B_i|/r$. Since F is not r-partite, the colouring χ' is not proper, and hence there exists $ij \in F$ such that $\chi'(i) = \chi'(j)$. The subsets B'_i and B'_j of B_i and B_j respectively which are given colour $\chi'(i)$ by χ are by construction disjoint and each of size at least $n/(4r\ell) = 2n/K$. Thus by Lemma 2.12(e) we have $e(B'_i, B'_j) > 2cp^{-1}n$, and the claim follows. □

Now we let G_3 be obtained from G_2 by deleting all edges of $G_2[B]$ which use a vertex $b \in B$ with $\deg(b, B) \geqslant pp'n$ or $\deg(b, B) \leqslant pp'n/10$. By Lemma 2.12$(d)$ the number of edges deleted is at most $cp^{-1}n$.

Finally, we let H be obtained from G_3 by deleting all edges ab of G_3 with $a \in A$ and $b \in B$ such that there exists $b' \in B$ with abb' a triangle of G_3. Observe that since A is independent in H, any triangle of H has at most one vertex in A. By construction of H, there are no triangles with exactly one vertex in A, so any triangle of H has all three vertices in B. But then the three vertices of a triangle in H would lie in sets B_i, B_j and B_k with ijk a triangle in F, and we chose F to be a triangle-free graph. We conclude that H is triangle-free. Furthermore, if H can be made r-partite by deleting $cp^{-1}n$ edges, then certainly $H[B]$ can be made r-partite by deleting $cp^{-1}n$ edges. But since we deleted at most $cp^{-1}n$ edges from $G_2[B]$ in order to obtain $G_3[B]$, and no further edges to obtain $H[B]$, this implies $G_2[B]$ can be made r-partite by deleting at most $2cp^{-1}n$ edges, in contradiction to Claim 2.13.

It remains only to show that $\delta(H) \geqslant \left(\frac{1}{2} - \gamma\right)pn$. First consider any vertex $b \in B$. By

Lemma 2.12(a) we have $\deg_{G_1}(b, A) \geqslant \left(\frac{1}{2} - \varepsilon\right)pn$. By construction, no edge from b to A was deleted in creating G_2 from G_1, or G_3 from G_2. By construction of G_3, either $\deg_{G_3}(b, B) = 0$, in which case no edge from b to A was deleted in creating H, or we have $\frac{1}{10}pp'n \leqslant \deg_{G_1}(b, B) \leqslant pp'n$. By Lemma 2.12($c$) we conclude that the total number of edges deleted from b to A in forming H from G_3 is at most

$$pn\big(1 - (1-p)^{pp'n}\big) \leqslant p^3p'n^2 \leqslant 64r^2\ell^2cpn \overset{(2.1)}{\leqslant} \tfrac{1}{2}\gamma pn\,,$$

because $c < \varepsilon$. Thus we have

$$d_H(b) \geqslant \left(\tfrac{1}{2} - \varepsilon\right)pn - \tfrac{1}{2}\gamma pn \overset{(2.1)}{\geqslant} \left(\tfrac{1}{2} - \gamma\right)pn$$

as desired.

Now consider any $a \in A$. Again by Lemma 2.12(a) we have $\deg_{G_1}(a, B) \geqslant \left(\frac{1}{2} - \varepsilon\right)pn$. Again no edges from a to B are deleted in forming G_2 or G_3. In forming H from G_3, we delete edges from a to each of b and b' in B whenever abb' forms a triangle in G_3. Since $G_3[B]$ is a subgraph of $G_1[B]$, this means that we delete at most $2 \cdot e\big(N_{G_1}(a; B)\big)$ edges from a to B, which by Lemma 2.12(b) is at most $2p'p^3n^2$. Thus we have

$$d_H(a) \geqslant \left(\tfrac{1}{2} - \varepsilon\right)pn - 2p'p^3n^2 \overset{(2.1)}{\geqslant} \left(\tfrac{1}{2} - \tfrac{1}{2}\gamma\right)pn - \tfrac{1}{2}\gamma pn = \left(\tfrac{1}{2} - \gamma\right)pn\,,$$

which completes the proof. $\square$

We now give the proof of Lemma 2.12.

Proof of Lemma 2.12. Choose $c = \min\{\frac{1}{2}\varepsilon, K^{-2}\}$. These properties follow from easy applications of the Chernoff bound, Lemma 2.6. We omit the proof of (a) as it is standard.

(b): By property (a) we may assume that there are at most $(\frac{1}{2}+\varepsilon)pn$ vertices in $N(a, B)$ for each $a \in A$. Now consider an arbitrary set S of $(\frac{1}{2} + \varepsilon)pn$ vertices in B. The expected number of edges in S is $\binom{|S|}{2}p'p \leqslant \frac{1}{2}|S|^2p'p$. By Lemma 2.6 the probability that S has more than $|S|^2p'p \leqslant p'p^3n^2$ edges is less than $\exp(\frac{-1}{6}|S|^2p'p) \leqslant \exp(-\frac{1}{100}p'p^3n^2) = \exp(-\frac{1}{100}K^2cpn)$. Hence the claimed property follows by taking a union bound over all $a \in A$.

(c): Assume that we first only reveal the edges of $G(n, p, p')$ in B and consider a vertex $b \in B$ for which $\deg(b, B) \geqslant \frac{1}{10}p'pn$. Now reveal also the edges between A and B. Then a fixed $a \in A$ forms a triangle with b in which the third vertex is also in B with probability $p \cdot (1 - (1-p)^{\deg(b,B)})$. Therefore the expected number of such $a \in A$ is

$$\frac{1}{2}np(1 - (1-p)^{\deg(b,B)}) \geqslant \frac{1}{2}np \cdot (1 - (1-p)^{p'pn/10}) \geqslant \frac{1}{40}p'p^3n^2\,,$$

where the inequality follows from $1-(1-p)^{p'pn/10} \geqslant \frac{1}{10}p'p^2n - \frac{1}{100}p'^2p^4n^2 \geqslant \frac{1}{20}p'p^2n$, which uses $p' = K^2cp^{-2}n^{-1}$. Hence by Lemma 2.6 the probability that there are more than $pn(1-(1-p)^{\deg(b,B)})$ such $a \in A$ is less than $\exp(-10^{-3}p'p^3n^2) = \exp(-10^{-3}K^2cpn)$. Taking a union bound over vertices in B the claimed property follows.

(d): Two applications of Lemma 2.6 and simple union bounds show that a.a.s. for any $S \subseteq B$ with $|S| = n/(2K^2)$ we have

$$e(S) \leqslant (1+\varepsilon)p'p\binom{|S|}{2} \quad \text{and} \tag{2.2}$$

$$e(S, B\setminus S) = (1\pm\varepsilon)p'p|S||B\setminus S|\,, \tag{2.3}$$

since $p \leqslant \varepsilon^2 c/(10^4K^2)$. This implies that for any $S \subseteq B$ with $|S| \leqslant n/(2K^2)$ the number of edges in B adjacent to S is at most

$$(1+\varepsilon)p'p\binom{n/(2K^2)}{2} + (1+\varepsilon)p'p\frac{n}{2K^2}\Big(\frac{n}{2} - \frac{n}{2K^2}\Big) \leqslant (1+\varepsilon)p'p\frac{n}{2K^2}\cdot\frac{n}{2} \leqslant \frac{1}{2}cp^{-1}n\,.$$

Hence, with $C = \{b \in B\colon \deg(b,B) \leqslant \frac{1}{10}p'pn\}$ and $D = \{b \in B\colon \deg(b,B) \geqslant p'pn\}$, the claimed property follows if $|C| \leqslant n/(2K^2)$ and $|D| \leqslant n/(2K^2)$.

So assume that there is $C' \subseteq C$ with $|C'| = n/(2K^2)$. But then $e(C', B\setminus C') \leqslant |C'|\frac{1}{10}p'pn \leqslant \frac{1}{20K^2}p'pn^2$, contradicting (2.3). Similarly, assuming there is $D' \subseteq D$ with $|D'| = n/(2K^2)$ and using (2.2) we get

$$e(D', B\setminus D') \geqslant |D'|p'pn - 2e(D') \geqslant \frac{n^2p'p}{2K^2} - (1+\varepsilon)p'p\Big(\frac{n}{2K^2}\Big)^2 \geqslant \frac{1}{3K^2}p'pn^2\,,$$

contradicting (2.3).

(e): For any disjoint $U, V \subseteq B$ each with at least $\frac{2n}{K}$ vertices the expected number of edges between U and V is $|U||V|p'p \geqslant \frac{4n^2}{K^2}p'p = 4cp^{-1}n$, so the result follows from another application of Lemma 2.6 and a union bound (using $p \leqslant \varepsilon^2c/(10^4K^2)$). □

2.4 Auxiliary properties of $G(n,p)$

In this section we list some typical properties of $G(n,p)$, which we shall use in the proofs of Theorems 2.3 and 2.4.

Lemma 2.14. *For any $0 < \varepsilon < \frac{3}{2}$ and $M \in \mathbb{N}$ and any $p = \omega\big(\frac{\ln n}{n}\big)$, the graph $\Gamma = G(n,p)$ a.a.s. satisfies the following.*

(a) $\deg_\Gamma(v) = (1\pm\varepsilon)pn$ *for every $v \in V(\Gamma)$.*

(*b*) $e_\Gamma(A) \leqslant \max\{|A|^2 p, 9n\}$ *for every* $A \subseteq V(\Gamma)$.

(*c*) $e_\Gamma(A,B) = (1\pm\varepsilon)p|A||B|$ *for every disjoint* $A, B \subseteq V(\Gamma)$ *with* $|A|, |B| \geqslant \frac{n}{M}$. *If on the other hand* $|A| < M^{-1}n$, *then* $e_\Gamma(A,B) \leqslant (1+\varepsilon)pM^{-1}n^2$.

(*d*) *For any* $A \subseteq V(\Gamma)$ *with* $|A| \geqslant \frac{n}{M}$ *all but at most* $10M\varepsilon^{-2}p^{-1}$ *vertices in* $V(\Gamma)$ *have* $(1\pm\varepsilon)p|A|$ *neighbours in* A.

Proof. These properties follow from standard applications of the Chernoff bound, Lemma 2.6. Here we only show (*b*); the other properties follow similarly.

Suppose that A is an arbitrarily chosen vertex subset. The expected number of edges in A is $\binom{|A|}{2}p \leqslant |A|^2 p$. By Lemma 2.6 the probability that there are more than $|A|^2 p$ edges in A is less than $\exp(\frac{-1}{3}\binom{|A|}{2}p) \leqslant \exp(\frac{-1}{7}|A|^2 p)$. For $|A| \geqslant 3p^{-1/2}n^{1/2}$ this probability is less than $\exp(\frac{-9}{7}n)$ and so taking a union bound over all subsets the probability that Property (*b*) fails for a set of size at least $3p^{-1/2}n^{1/2}$ is less than $2^n \exp(\frac{-9}{7}n)$, which tends to zero. A set A with $|A| < 3p^{-1/2}n^{1/2}$ is less likely to have more than $9n$ edges than a set B with $|B| = 3p^{-1/2}n^{1/2} \leqslant n$. Therefore, since $|B|^2 p = 9n$ and by the previous argument, the probability that a set A of size less than $3p^{-1/2}n^{1/2}$ has more than $9n$ edges tends to zero. □

The next lemma shows that for any partition $V\big(G(n,p)\big) = A \cup B$ with neither A nor B very small, most edges of $G(n,p)$ have 'typical' neighbourhoods in each set.

Lemma 2.15. *For any* $0 < \varepsilon < \frac{1}{2}$, $M \in \mathbb{N}$ *and* $p = \omega\big(\frac{\ln n}{n}\big)$ *in* $\Gamma = G(n,p)$ *a.a.s. for any two subsets* A, B *of* $V(\Gamma)$ *with* $\frac{n}{M} \leqslant |A|, |B|$ *all but at most* $10^3 M\varepsilon^{-2}p^{-1}n$ *edges* uv *in* Γ *satisfy all of the following:*

- $\deg_\Gamma(u,A), \deg_\Gamma(v,A) = (1\pm\varepsilon)p|A|$.
- $\deg_\Gamma(u,B), \deg_\Gamma(v,B) = (1\pm\varepsilon)p|B|$.
- $\deg_\Gamma(u,v,B) \geqslant (1-\varepsilon)p^2|B|$.

Proof. By Lemma 2.14(*d*) we may assume that all but a set S of at most $20M\varepsilon^{-2}p^{-2}$ vertices in Γ have $(1\pm\varepsilon)p|B|$ neighbours in B and $(1\pm\varepsilon)p|A|$ neighbours in A. By Lemma 2.14(*c*) we further may assume that we have

$$e(S,A) \leqslant (1+\varepsilon)p \cdot 20M\varepsilon^{-2}p^{-2}n = 20(1+\varepsilon)M\varepsilon^{-2}p^{-1}n\,. \tag{2.4}$$

We now consider an arbitrary vertex v in $V \setminus S$ and two arbitrary sets $P, Q \subseteq N(v)$ satisfying $|P| \geqslant (1-\frac{1}{2}\varepsilon)p|B|$ and $|Q| \geqslant 100M\varepsilon^{-2}p^{-1}$. The probability that all vertices in Q have fewer than $(1-\varepsilon)p^2|B| \leqslant (1-\frac{1}{2}\varepsilon)p|P|$ neighbours in P is less than

$$\exp\Big(-\frac{\varepsilon^2}{12}p|P||Q|\Big) \leqslant \exp\Big(-\frac{\varepsilon^2}{12}p \cdot \frac{1}{2}p\frac{n}{M} \cdot 100M\varepsilon^{-2}p^{-1}\Big) \leqslant \exp(-3pn)\,.$$

Since $P, Q \subseteq N(v)$ we have $|P|, |Q| \leqslant (1+\varepsilon)pn$. So, taking a union bound, the probability that there exist v, P, Q as above is less than $n2^{(1+\varepsilon)pn}2^{(1+\varepsilon)pn}\exp(-3pn)$ which tends to zero as n tends to infinity for $p = \omega(\log n/n)$. Hence a.a.s. each vertex v in $V \setminus S$ has at most $100M\varepsilon^{-2}p^{-1}$ neighbours u such that $\deg(u, v, B) < (1-\varepsilon)p^2|B|$. Summing over v we obtain at most $100M\varepsilon^{-2}p^{-1}n$ such edges, which along with the edges incident to S by (2.4) gives at most $10^3M\varepsilon^{-2}p^{-1}n$ edges. □

The following lemma is crucial in the proofs of Theorems 2.3 and 2.4. Before stating it we need some definitions. For any $s \in \mathbb{N}$, the *s-star* is the star $K_{1,s}$. The vertex of degree s in the s-star is called its *centre*, all other vertices are its *leaves*. For $A \subseteq V(\Gamma)$ and $0 < q, \varepsilon < 1$ we say that an s-star with centre x is *(q,ε)-bad* for A if there is $S \subseteq N_\Gamma(x, A)$ with $|S| \leqslant qp|A|$ such that each leaf y of the s-star satisfies $\deg_\Gamma(y, S) \geqslant (1+\varepsilon)qp^2|A|$; in other words y has substantially more neighbours in S than expected. We also say that S *witnesses* this badness.

When we use this definition, we will choose a star with centre x and set $S = N_\Gamma(x, A) \setminus N_H(x, A)$, where H is a triangle-free subgraph of Γ with large minimum degree, and we will choose our star such that that $N_\Gamma(y, S)$ is quite large for each leaf y. Now if the star is good it follows that S itself must be quite large, so that the degree of x in H cannot be too large, leading to a contradiction to the minimum degree of H. The following lemma however implies that bad stars cover only $\mathcal{O}(p^{-1}n)$ edges, which is where the sharp bounds in Theorems 2.3 and 2.4 come from.

Lemma 2.16. *For every $0 < \varepsilon < 1$ and every p the random graph $G(n,p)$ a.a.s. satisfies the following. For every $A \subseteq V(\Gamma)$ with $\frac{n}{3} \leqslant |A|$, every q with $\varepsilon < q < 1$, and every $s \geqslant 100q^{-1}\varepsilon^{-2}p^{-1}$ there are fewer than $\frac{1}{2}p^{-1}$ vertex disjoint s-stars in $V(\Gamma) \setminus A$ which are (q,ε)-bad for A.*

Proof. First let A be fixed. Consider an s-star with centre x and a set $S \subseteq N_\Gamma(x, A)$ with $|S| \leqslant qp|A|$. By the Chernoff bound, Lemma 2.6, the probability that S witnesses that this star is (q,ε)-bad for A is less than $\exp\big(\frac{-\varepsilon^2}{3} \cdot qp^2|A|s\big)$. Observe that $|S| \leqslant qp|A| \leqslant pn$ and that we may assume $s \leqslant \deg_\Gamma(x) \leqslant 2pn$ by Lemma 2.14(a). So by taking a union bound over choices of S for a single s-star, and then considering collections of $\frac{1}{2}p^{-1}$ vertex disjoint s-stars, and taking another union bound over all such collections, we obtain that the probability that there are at least $\frac{1}{2}p^{-1}$ disjoint (q,ε)-bad stars for A in $V(\Gamma)\setminus A$ is less than

$$\Big(n\cdot 2^{2pn}\Big)^{\frac{1}{2}p^{-1}}\cdot\Big(2^{pn}\exp\big(\tfrac{-\varepsilon^2}{3}qp^2|A|s\big)\Big)^{\frac{1}{2}p^{-1}}\leqslant\big(2^{4pn}\exp\big(\tfrac{-\varepsilon^2}{9}qp^2ns\big)\big)^{\frac{1}{2}p^{-1}}.$$

By taking a union bound over choices of A we find that the probability that there is A such that $\frac{1}{2}p^{-1}$ stars $K_{1,s}$ outside A are (q,ε)-bad for A is less than

$$2^n\big(2^{4pn}\exp\big(\tfrac{-\varepsilon^2}{9}qp^2ns\big)\big)^{\frac{1}{2}p^{-1}}\leqslant\exp\big(n+2n-\tfrac{\varepsilon^2}{18}qpns\big),$$

which tends to zero for $s\geqslant 100\varepsilon^{-2}q^{-1}p^{-1}$. (Observe that we do not have to take a union bound over s, because for $s'>s$ any s-star which is a subgraph of a (q,ε)-bad s'-star is also (q,ε)-bad.) $\square$

2.5 Proof of Theorem 2.3

Recall that Theorem 2.3 states the following.

Theorem 2.3. *For any $\gamma>0$, there exists C such that for any $p(n)$ the random graph $\Gamma=G(n,p)$ a.a.s. has the property that all triangle-free spanning subgraphs $H\subseteq\Gamma$ with $\delta(H)\geqslant(\frac{2}{5}+\gamma)pn$ can be made bipartite by removing at most $\min\big(Cp^{-1}n,(\frac{1}{4}+\gamma)pn^2\big)$ edges.*

The main strategy of the proof is as follows. We first apply Lemma 2.10 (which is a consequence of the Sparse Regularity Lemma) to H to obtain a dense triangle-free reduced graph R of H with minimum degree above $\frac{2}{5}v(R)$, which by the Andrásfai–Erdős–Sós Theorem, Theorem 2.1, is bipartite. We conclude that H can be made bipartite by removing $o(pn^2)$ edges. Hence in a maximum cut $X\cup Y$ of H we have $e_H(X),e_H(Y)=o(pn^2)$. Our goal will then be to improve this bound on $e_H(X)$ and $e_H(Y)$ by distinguishing between 'typical' and 'atypical' edges in these sets and applying the results established in the previous section to count these, using that $X\cup Y$ is a maximum cut and that H is triangle-free.

Proof of Theorem 2.3. Let

$$\varepsilon=\frac{\gamma^4}{10^7},\quad d=\frac{\gamma^2}{10^3},\quad \eta=d+3\varepsilon,\quad \beta=\frac{2}{5}+\gamma,\quad t_0=\frac{1}{\varepsilon}\tag{2.5}$$

and let c and t_1 be the values attained by applying Lemma 2.10 with inputs ε, d, β and t_0. Let $M = t_1^2$, and let

$$C = \max\left(10^{10}\varepsilon^{-2}, c^2\right). \tag{2.6}$$

We first consider the easy case that p is small. If $p \leqslant n^{-7/4}$, then the expected number of paths with two edges in $G(n,p)$ is at most $p^2n^3 \leqslant n^{-1/2}$. In particular a.a.s there are no such paths, so a.a.s. $G(n,p)$ is bipartite and the statement of Theorem 2.3 holds trivially. We may therefore assume $p \geqslant n^{-7/4}$, so by Lemma 2.6 a.a.s. $G(n,p)$ has at most $\left(\frac{1}{2}+\gamma\right)pn^2$ edges. Now if G is any graph with at most $\left(\frac{1}{2}+2\gamma\right)pn^2$ edges, then we can make G bipartite by removing all the edges of G not in a maximum cut. Since a maximum cut of G contains at least half its edges, we remove at most $\left(\frac{1}{4}+\gamma\right)pn^2$ edges. Again, if $\min\left(Cp^{-1}n, (\frac{1}{4}+\gamma)pn^2\right) = (\frac{1}{4}+\gamma)pn^2$, which occurs when $p \leqslant cn^{-1/2}$, the statement of Theorem 2.3 follows.

It remains to consider the hard case that $p \geqslant cn^{-1/2}$. We now assume $\Gamma = G(n,p)$ satisfies the properties stated in Lemma 2.14 with input ε and M, Lemma 2.15 with input ε and M, Lemma 2.16 with input ε and Lemma 2.10 for the parameters given above.

Consider any triangle-free $H \subseteq \Gamma$ with $\delta(H) \geqslant (\frac{2}{5}+\gamma)pn$ and let $X \cup Y$ be a maximum cut of the vertex set of H. Assume without loss of generality that $e_H(X) \geqslant e_H(Y)$. Our goal is to show $e_H(X) \leqslant \frac{1}{2}Cp^{-1}n$. We start with the following observation.

Claim 2.17. *$e_H(X) \leqslant \eta pn^2$.*

Proof of Claim 2.17. By the property asserted by Lemma 2.10 we obtain an (ε, d, p)-regular partition $V(\Gamma) = V_0 \cup V_1 \cup \cdots \cup V_t$ of H with $t_0 \leqslant t \leqslant t_1$ whose corresponding reduced graph R is triangle-free and has minimum degree at least $(\frac{2}{5}+\gamma-d-\varepsilon)v(R) > \frac{2}{5}v(R)$. Therefore, by the Andrásfai–Erdős–Sós Theorem, Theorem 2.1, R is bipartite.

By Lemma 2.14(a) at most $\varepsilon n(1+\varepsilon)pn$ edges have at least one end in V_0. Moreover, since at most an ε-fraction of all pairs are irregular, by Lemma 2.14(c) at most $\varepsilon(1+\varepsilon)pn^2$ edges are contained in irregular pairs. Finally, at most dpn^2 edges are in pairs with density less than d. We conclude that at most $(d+2(1+\varepsilon)\varepsilon)pn^2 \leqslant \eta pn^2$ edges of H do not lie in pairs corresponding to edges of R, which proves the claim. □

We next bound the sizes of X and Y.

Claim 2.18. *$\left(\frac{2}{5}+\frac{1}{2}\gamma\right)n \leqslant |X|, |Y| \leqslant \left(\frac{3}{5}-\frac{1}{2}\gamma\right)n$.*

Proof of Claim 2.18. Suppose for a contradiction that X satisfies $|X| > (\frac{3}{5}-\frac{1}{2}\gamma)n$ and hence $|Y| < (\frac{2}{5} + \frac{1}{2}\gamma)$. Then by Lemma 2.14($c$) we see that $e_H(X,Y) \leqslant e_\Gamma(X,Y) \leqslant (1+\varepsilon)(\frac{3}{5} - \frac{1}{2}\gamma)(\frac{2}{5} + \frac{1}{2}\gamma)pn^2$.

On the other hand, by our minimum degree condition $2e_H(X)+e_H(X,Y) \geqslant (\frac{2}{5}+\gamma)pn|X|$, and similarly $2e_H(Y) + e_H(X,Y) \geqslant (\frac{2}{5} + \gamma)pn|Y|$. Since $e_H(X), e_H(Y) \leqslant \eta pn^2$ this gives $e_H(X,Y) \geqslant (\frac{2}{5} + \gamma)pn \cdot \max\{|X|,|Y|\} - 2\eta pn^2$. Since $\max\{|X|,|Y|\} \geqslant (\frac{3}{5} - \frac{1}{2}\gamma)n$ we obtain $e_H(X,Y) \geqslant \big((\frac{3}{5} - \frac{1}{2}\gamma)(\frac{2}{5} + \gamma) - 2\eta\big)pn^2$, a contradiction.

So $|X| \leqslant (\frac{3}{5} - \frac{1}{2}\gamma)n$, and analogously $|Y| \leqslant (\frac{3}{5} - \frac{1}{2}\gamma)n$, proving the claim. □

We next define

$$\tilde{X} = \left\{x \in X : \deg_H(x,X) \geqslant \gamma \cdot \deg_H(x)\right\},$$

a set of vertices with high degree in X, which require special treatment later on. The next claim shows that $\tilde{X}$ is small and contains at most half of the edges in X.

Claim 2.19. *$|\tilde{X}| \leqslant \frac{1}{100}\gamma n$, and if $e_H(X) > \frac{1}{2}Cp^{-1}n$ then $e_H(\tilde{X}) \leqslant \frac{1}{2}e_H(X)$.*

Proof of Claim 2.19. By Claim 2.17 and the definition of $\tilde{X}$ we have

$$\eta pn^2 \geqslant e_H(X) \geqslant \frac{1}{2}|\tilde{X}|\gamma\delta(H) \geqslant \frac{\gamma}{2}\Big(\frac{2}{5} + \gamma\Big)pn|\tilde{X}|\,, \tag{2.7}$$

hence $|\tilde{X}| \leqslant \frac{2\eta n}{\gamma(2/5+\gamma)} \leqslant 5\gamma^{-1}\eta n \leqslant \gamma n/100$ by (2.5).

For the second part of the claim assume that $e_H(X) > \frac{1}{2}Cp^{-1}n$. By Lemma 2.14($b$) we have $e_H(\tilde{X}) \leqslant e_\Gamma(\tilde{X}) \leqslant \max\{|\tilde{X}|^2p, 9n\}$. If this maximum is attained by $9n$, then we are done because $9n \leqslant \frac{1}{4}Cp^{-1}n < \frac{1}{2}e_H(X)$. Otherwise $e_H(\tilde{X}) \leqslant |\tilde{X}|^2p$, and since $|\tilde{X}| \leqslant \frac{1}{100}\gamma n$, we have

$$|\tilde{X}|^2p \leqslant \frac{1}{100}\gamma pn|\tilde{X}| \leqslant \frac{\gamma}{4}\Big(\frac{2}{5} + \gamma\Big)pn|\tilde{X}| \overset{(2.7)}{\leqslant} \frac{1}{2}e_H(X)\,,$$

and we are also done. □

We continue by removing ‘atypical’ edges from H. Let H' be the graph obtained from H by removing edges from $E_H(X)$ which do not satisfy the conditions of Lemma 2.15 with respect to the partition $X \cup Y$. We also remove the edges in $E_H(\tilde{X})$. By Lemma 2.15 and Claim 2.19 we have $e_H(X) \leqslant \frac{1}{2}Cp^{-1}n$ or

$$e_H(X) - e_{H'}(X) \leqslant 10^3\varepsilon^{-2}p^{-1}n + \frac{1}{2}e_H(X) \overset{(2.6)}{\leqslant} \frac{1}{10}Cp^{-1}n + \frac{1}{2}e_H(X)\,. \tag{2.8}$$

Our goal in the remainder is to bound the number of H'-edges in X.

Let xz be any H'-edge in X. We have

$$\deg_\Gamma(x,z,Y) \geqslant (1-\varepsilon)p^2|Y| \tag{2.9}$$

by construction of H', so this common neighbourhood constitutes many Γ-triangles xzy, for each of which either xy or zy is not present in H'. We now would like to direct the edges in X according which of these two cases is more common – however, it turns out that we need to favour vertices not in $\tilde{X}$ in this process; so we direct with a bias.

More precisely, for any H'-edge in X, if one of its vertices is in $\tilde{X}$ call it x, otherwise let x be any vertex of the edge. Let x' be the other vertex of the edge. We direct xx' towards x if

$$|N_\Gamma(x, x', Y) \setminus N_{H'}(x, Y)| \geqslant \frac{2}{3} \deg_\Gamma(x, x', Y)\,,$$

that is if many edges from x to $N_\Gamma(x, x', Y)$ were deleted. We direct xx' towards x' otherwise, in which case we have

$$|N_\Gamma(x, x', Y) \setminus N_{H'}(x', Y)| > \frac{1}{3} \deg_\Gamma(x, x', Y)\,,$$

An *s-in-star* in this directed graph is an s-star such that all edges are directed towards the centre. Recall that an s-star with centre x is (q, ε)-bad for Y if there is a witness $S \subseteq N_\Gamma(x, Y)$ with $|S| \leqslant qp|Y|$ such that each leaf z of the s-star satisfies $\deg_\Gamma(z, S) \geqslant (1 + \varepsilon)qp^2|Y|$. The next claim shows that in-stars in $H'[X]$ are bad. We define

$$s = 10^3\varepsilon^{-2}p^{-1}\,, \quad \tilde{q} = (1 - 2\varepsilon)\frac{2}{3}\,, \quad q = (1 - 2\varepsilon)\frac{1}{3}\,.$$

Claim 2.20. *Each s-in-star in $H'[X]$ with centre $x \in \tilde{X}$ is $(\tilde{q}, \varepsilon)$-bad for Y, and each s-in-star in $H'[X]$ with centre $x \notin \tilde{X}$ is (q, ε)-bad for Y.*

Proof of Claim 2.20. First assume F is an s-in-star with centre $x \in \tilde{X}$ which is not $(\tilde{q}, \varepsilon)$-bad. We first show that this implies

$$|N_\Gamma(x, Y) \setminus N_{H'}(x, Y)| > \tilde{q}p|Y|\,. \tag{2.10}$$

Indeed, assume otherwise. Then, since F is not $(\tilde{q}, \varepsilon)$-bad for Y we have for $S = N_\Gamma(x, Y) \setminus N_{H'}(x, Y)$ that there is a leaf z of F such that

$$|N_\Gamma(x, z, Y) \setminus N_{H'}(x, Y)| = \deg_\Gamma(z, S) < (1 + \varepsilon)\tilde{q}p^2|Y| \leqslant \frac{2}{3}(1 - \varepsilon)p^2|Y|\,.$$

This however contradicts the fact that F is an in-star and thus

$$|N_\Gamma(x, z, Y) \setminus N_{H'}(x, Y)| \geqslant \frac{2}{3} \deg_\Gamma(x, z, Y) \overset{(2.9)}{\geqslant} \frac{2}{3}(1 - \varepsilon)p^2|Y|\,.$$

Accordingly (2.10) holds.

Since $\deg_H(x, Y) = \deg_{H'}(x, Y)$ we conclude that

$$\deg_H(x, Y) \leqslant \deg_\Gamma(x, Y) - \tilde{q}p|Y| \leqslant (1 + \varepsilon)p|Y| - (1 - 2\varepsilon)\frac{2}{3}p|Y| \leqslant \Big(\frac{1}{3} + 3\varepsilon\Big)p|Y|\,.$$

Because $X \cup Y$ is a maximum cut this implies by Claim 2.18 that

$$\deg_H(x) \leqslant 2\Big(\frac{1}{3} + 3\varepsilon\Big)p\Big(\frac{3}{5} - \frac{1}{2}\gamma\Big)n < \Big(\frac{2}{5} + \gamma\Big)pn\,,$$

contradicting the minimum degree of H.

For the second part of the claim assume that F is an s-in-star with centre $x \notin \tilde{X}$ which is not (q, ε)-bad. By similar logic to the proof of (2.10), this implies that

$$|N_\Gamma(x, Y) \setminus N_{H'}(x, Y)| > qp|Y|$$

by using that for any leaf z of F we have $|N_\Gamma(x, z, Y) \setminus N_{H'}(x, Y)| > \frac{1}{3}\deg_\Gamma(x, z, Y)$. Also analogously, this implies that $\deg_H(x, Y) \leqslant (\frac{2}{3} + 3\varepsilon)p|Y|$. Recall that $x \notin \tilde{X}$ means that $\deg_H(x, X) < \gamma \deg_H(x)$ and hence $\deg_H(x) \leqslant \frac{1}{1-\gamma}\deg_H(x, Y) \leqslant (1 + 2\gamma)\deg_H(x, Y)$. Thus, by Claim 2.18,

$$\deg_H(x) \leqslant (1+2\gamma)\Big(\frac{2}{3} + 3\varepsilon\Big)p\Big(\frac{3}{5} - \frac{1}{2}\gamma\Big)n \leqslant \Big(\frac{2}{3} + \frac{5}{3}\gamma\Big)p\Big(\frac{3}{5} - \frac{1}{2}\gamma\Big)n < \Big(\frac{2}{5} + \gamma\Big)pn\,,$$

again contradicting the minimum degree of H. □

By Lemma 2.16, however, the number of s-stars in Γ which are either $(\tilde{q}, \varepsilon)$-bad or (q, ε)-bad is less than p^{-1}. So Claim 2.20 implies that the number of vertex disjoint s-in-stars in $H'[X]$ is less than p^{-1}. The following claim shows that this implies that $e_{H'}(X)$ is small.

Claim 2.21. $e_{H'}(X) \leqslant \frac{1}{10}Cp^{-1}n.$

Proof of Claim 2.21. Assume for a contradiction that $e_{H'}(X) > \frac{1}{10}Cp^{-1}n \geqslant 10^4\varepsilon^{-2}p^{-1}n$. Using a greedy argument, we will show that we then can find more than p^{-1} stars in $H'[X]$ which are s-in-stars (with $s = 10^3\varepsilon^{-2}p^{-1}$). Indeed, the average in-degree is at least $10^4\varepsilon^{-2}p^{-1}$, so we can find at least one $(10^3\varepsilon^{-2}p^{-1})$-in-star. If we remove from $H'[X]$ this star and all edges adjacent to it this accounts for at most $(1 + s)(1 + \varepsilon)pn \leqslant 2spn$ edges. So we can repeat this process p^{-1} times, after which at most $2sn = 2 \cdot 10^3\varepsilon^{-2}p^{-1}n$ edges have been deleted from $H'[X]$, hence $H[X]$ still contains more than $10^3\varepsilon^{-2}p^{-1}n$ edges in X, still giving an average in-degree of at least $10^3\varepsilon^{-2}p^{-1}$, and hence we can find another $(10^3\varepsilon^{-2}p^{-1})$-in-star, which is the desired contradiction. □

Now (2.8) and Claim 2.21 imply $e_H(Y) \leqslant e_H(X) \leqslant \frac{1}{2}Cp^{-1}n$, hence H can be made bipartite by removing at most $Cp^{-1}n$ edges as claimed. □

2.6 Proof of Theorem 2.4

The proof of Theorem 2.4 adds the techniques developed for the proof of Theorem 2.3 to ideas used in [3, 69]. Our strategy is as follows. Given a subgraph H of $\Gamma = G(n,p)$ with $\delta(H) \geqslant \big(\frac{1}{3}+\gamma\big)pn$, we will apply the sparse regularity lemma to obtain a regular partition $V(H) = V_0 \cup \cdots \cup V_t$ with (ε, d, p)-reduced graph R. We let W be the set of all vertices whose degree, in Γ, to some set V_i is far from the expected $p|V_i|$, and then for each $I \subseteq [t]$ we let N_I be the subset of vertices in $V(H) \setminus W$ with many H-neighbours in exactly the clusters $\{V_i : i \in I\}$, which gives a partition of $V(H)$ into $2^t + 1$ sets. We will show that there are $O(p^{-1}n)$ edges in W and in each N_I, hence we can remove all such edges to obtain a graph with bounded chromatic number. We do this by showing that W is too small to contain many edges, and that the same is true for any N_I such that $R[I]$ contains an edge. If on the other hand $R[I]$ is independent, we use an argument similar to that in the proof of Theorem 2.3.

Proof of Theorem 2.4. Given $\gamma > 0$, let

$$d = \frac{\gamma}{20}\,, \quad \varepsilon' = \frac{d^3}{30}\,, \quad \beta = \frac{1}{3} + \gamma\,, \quad t_0 = \frac{1}{\varepsilon'}\,. \tag{2.11}$$

Let ε_0, $C_{\mathrm{L2.9}}$ be the outputs if Lemma 2.9 is applied with ε' and d. We take $\varepsilon = \min\{\varepsilon_0, \varepsilon'\}$ and let t_1 be the output if Lemma 2.7 is applied with β, ε and t_0. We require as well that $t_1 \geqslant 10$. We choose $c = 2C_{\mathrm{L2.9}}t_1$ (which is needed for the application of Lemma 2.9). Finally we choose

$$M = 2t_1\,, \quad r = 2^{t_1} + 1\,, \quad C' = 10^4 \cdot 2^{10t_1}\varepsilon^{-3}\,, \quad C = \max(rC'^2, c^2)\,. \tag{2.12}$$

As in the proof of Theorem 2.3, if $p \leqslant n^{-7/4}$ a.a.s. $G(n,p)$ is bipartite and the statement is trivially true, while for any graph G a maximum r-partition of G contains at least $\frac{r-1}{r}e(G)$ edges, so that when $p \geqslant n^{-7/4}$ a.a.s. we can make any subgraph of $G(n,p)$ r-partite by deleting at most $\big(\frac{1}{2r}+\gamma\big)pn^2$ edges. Again, this leaves the hard case when $p \geqslant cn^{-1/2}$.

Now sample $\Gamma = G(n,p)$. Since $p > cn^{-1/2} = \omega(\frac{\ln n}{n})$ we can assume that Γ satisfies the properties of Lemmas 2.7, 2.14, 2.15, and 2.16 with the parameters chosen above.

Let H be a triangle-free spanning subgraph of Γ with $\delta(H) \geqslant \big(\frac{1}{3}+\gamma\big)n$. By Lemma 2.7 there is an (ε, d, p)-regular partition $V_0 \cup V_1 \cup \cdots \cup V_t$ of H with $t \leqslant t_1$ such that the reduced graph R has $\delta(R) \geqslant \big(\frac{1}{3}+\gamma-d-3\varepsilon\big)v(R) \geqslant \big(\frac{1}{3}+\frac{\gamma}{2}\big)v(R)$, and such that for each i and each $v \in V_i$, the vertex v has at most $(d+\varepsilon)pn$ neighbours in $\bigcup_{j:ij\notin R} V_j$.

Let W consist of all vertices which either have more than $(1+\varepsilon)p|V_i|$, Γ-neighbours in V_i for some i, or more than $2\varepsilon pn$, Γ-neighbours in V_0. By Lemma 2.14(d) we have

$|W| \leqslant 10M(t+1)\varepsilon^{-2}p^{-1}$, and by Lemma 2.14($b$) the number of edges in W is therefore at most $\max\big(100M^2(t+1)^2\varepsilon^{-4}p^{-1}, 9n\big) \leqslant 10p^{-1}n$, where the inequality holds for all sufficiently large n. Now for each $I \subseteq [t]$, let N_I be the set of vertices of H with many H-neighbours exactly in the clusters V_i with $i \in I$, that is,

$$N_I = \{v \in V(H) \colon |N_H(v) \cap V_i| > 10dp|V_i| \text{ if and only if } i \in I\}\,.$$

Claim 2.22. *$\{N_I \colon |I| > \frac{t}{3}\}$ partitions $V(H) \setminus W$.*

Proof. The sets $\{N_I : I \subseteq [t]\}$ are disjoint and partition $V(H)\setminus W$ by definition. If $|I| \leqslant \frac{t}{3}$ then any vertex $v \in N_I$ has at most $\sum_{i\in I}(1+\varepsilon)p|V_i| + \sum_{i\notin I} 10dp|V_i| + 2\varepsilon pn < \big(\frac{1}{3}+\gamma\big)pn$ neighbours since $v \notin W$ and by definition of N_I, which is a contradiction, so $N_I = \emptyset$ if $|I| \leqslant \frac{t}{3}$. □

Our goal is thus to show that $e_H(N_I) \leqslant C'^2p^{-1}n$ for any I with $|I| > \frac{t}{3}$, since this implies that H can be made r-partite with $r = 2^{t_1}+1$ by removing at most $rC'^2p^{-1}n \leqslant Cp^{-1}n$ edges. This is established by the following two claims.

Claim 2.23. *If $R[I]$ contains an edge, then $e_H(N_I) \leqslant C'^2p^{-1}n$.*

Proof of Claim 2.23. Suppose that $ij \in R[I]$. If $v \in N_I$ is such that $\big(N_\Gamma(v,V_i), N_\Gamma(v,V_j)\big)$ is (ε', d, p)-lower-regular in H, since $v \notin W$, the pair $\big(N_H(v,V_i), N_H(v,V_j)\big)$ is $\big(\varepsilon'\frac{1+\varepsilon}{10d}, d, p\big)$-lower-regular in H. Since $d > \varepsilon'\frac{1+\varepsilon}{10d}$, there is an edge of H in this latter pair and hence H contains a triangle, a contradiction.

We conclude that there are no such vertices in N_I, so by Lemma 2.9 we have $|N_I| \leqslant C'\max\big(p^{-2}, p^{-1}\log n\big)$. By Lemma 2.14($b$) the number of edges in N_I is therefore at most $\max\big(C'^2p^{-3}, C'^2p^{-1}\log^2 n, 9n\big) \leqslant C'^2p^{-1}n$ by choice of p and C'. □

Claim 2.24. *If $R[I]$ is independent, then $e_H(N_I) \leqslant C'p^{-1}n$.*

Proof of Claim 2.24. Since $\delta(R) \geqslant \big(\frac{1}{3}+\frac{\gamma}{2}\big)t$, if $R[I]$ is independent then $|I| < \frac{2t}{3}$. Let $S_I := \bigcup_{i\in I} V_i$. We first show that S_I and N_I are disjoint. Indeed, if $v \in N_i$ were in some V_i with $i \in I$, then by definition of N_I the vertex v has at least $\sum_{j\in I} 10dp|V_j| \geqslant 5dpn/3$ neighbours in $\bigcup_{j\in I} V_j$, where the inequality follows since $|I| > t/3$. Since ij is not an edge of R for any $j \in I$, this is in contradiction to the guarantee that v has at most $(d+\varepsilon)pn$ neighbours in $\bigcup_{j:ij\notin R} V_j$.

We now delete some 'atypical' edges from $H[N_I]$. Remove from $H[N_I]$ each edge uv with $\deg_\Gamma(u,v,S_I) < (1-\varepsilon)|S_I|p^2$. to obtain the graph H'. By Lemma 2.15 this accounts for at most $10^3 \cdot 4\varepsilon^{-2}p^{-1}n \leqslant \frac{\varepsilon}{10}C'p^{-1}n$ edges.

Let Z be the set of vertices $v \in N_I$ such that $\deg_H(v) - \deg_{H'}(v) \geqslant \varepsilon pn$. By double counting we have $|Z| \leqslant \frac{\varepsilon C' p^{-1} n}{5\varepsilon pn} = \frac{1}{5}C'p^{-2}$.

We now proceed similarly as in the proof of Theorem 2.3. We orient the edges uv in $H'[N_I]$ towards u if $|N_\Gamma(u,v,S_I) \setminus N_{H'}(u,S_I)| \geqslant \frac{1}{2}\deg_\Gamma(u,v,S_I)$ and towards v otherwise. Again, for $s = 10^3 q^{-1}\varepsilon^{-2}p^{-1}$ and $q = (1-2\varepsilon)\frac{1}{2}$ any s-in-star with centre x not in Z is (q,ε)-bad with respect to S_I. Indeed, otherwise, analogously to the proof of (2.10), we have $|N_\Gamma(x,S_I) \setminus N_{H'}(x,S_I)| > qp|S_I|$, which implies

$$\deg_{H'}(x,S_I) < (1+\varepsilon)p|S_I| - qp|S_I| = \frac{1}{2}p|S_I| \leqslant \frac{1}{2}p\frac{2}{3}n = \frac{1}{3}pn$$

Since $x \notin Z$, we have $\deg_H(x) \leqslant \deg_{H'}(x) + \varepsilon pn < \left(\frac{1}{3}+\gamma\right)pn$, a contradiction.

We now pick greedily vertex disjoint s-in-stars whose centres are not in Z until no more remain. By Lemma 2.16, since S_I and N_I are disjoint, this process terminates having found less than $\frac{1}{2}p^{-1}$ such stars. Let Y be the set of vertices contained in all these stars; then $|Y| \leqslant \frac{1}{2}p^{-1}s \leqslant 10^3 q^{-1}\varepsilon^{-2}p^{-2}$. Now $e_{H'}\big(N_I \setminus (Y \cup Z)\big) \leqslant s|N_I|$ since $N_I \setminus (Y \cup Z)$ contains no s-in-star, so we conclude

$$e_H(N_I) \leqslant (1+\varepsilon)pn|Y \cup Z| + s|N_I| + \tfrac{1}{10}C'p^{-1}n \leqslant C'p^{-1}n\,,$$

as desired. □

Finally, these claims show that deleting all edges internal to any of the sets W and N_I for $I \subseteq [t]$ yields a $2^t + 1 = r$-partite graph, and that the number of edges deleted is at most $Cp^{-1}n$, as desired. □

3

Partite saturation problems

3.1 Introduction

The Turán problem of asking for the maximum number of edges a graph on a fixed number of vertices can have without containing some fixed subgraph H is one of the oldest and most famous questions in extremal graph theory, see [70],[84],[39].

Since the corresponding minimisation problem - asking how few edges an H-free graph can have - trivially gives the answer zero, if we want an interesting complementary question to the Turán problem we can require that our H-free graph G also has the property that it *nearly* contains a copy of H. By this we mean that the addition of any new edge to G creates an copy of H as a subgraph. Such a graph G is called *H-saturated* and over H-saturated graphs on n vertices the minimum number of edges is called the saturation number, $\mathrm{sat}(H, n)$. The study of saturation numbers was initiated by Erdős, Hajnal and Moon [36] when they proved that $\mathrm{sat}(K_r, n) = (r-2)(n-\frac{1}{2}(r-1))$. It was later shown by Kászonyi and Tuza in [57] that cliques have the largest saturation number of any graph on r vertices which in particular implies that for any H the saturation number $\mathrm{sat}(H, n)$ grows linearly in n.

These saturation questions can be generalised to require our H-free graph G to be a subgraph of another fixed graph F. Here we insist that adding any new edge of F to G would create a copy of H in G. The minimum number of edges in such a G we denote by $\mathrm{sat}(H, F)$. One natural class of host graphs are complete r-partite graphs. In the bipartite case Bollobás [15, 16] and Wessel [85, 86] independently determined the saturation number $\mathrm{sat}(K_{a,b}, K_{c,d})$. Working in the r-partite setting with $r \geqslant 3$, Ferrara, Jacobson, Pfender, and Wenger determined in [41] the value of $\mathrm{sat}(K_3, K_r[n])$ for sufficiently large n and showed that $\mathrm{sat}(K_3, K_3[n]) = 6n-6$ for all n, where $K_r[n]$ denotes the complete balanced r-partite graph on parts of size n.

In this chapter we consider the saturation problem when the host graph is a blow-up of the forbidden subgraph H. For any graph H and any $n \in \mathbb{N}$ let $H[n]$ denote the graph obtained from H by replacing each vertex with an independent set of size n and each edge with a complete bipartite graph between the corresponding independent sets. A copy of H in $H[n]$ is called *partite* if it has exactly one vertex in each part of $H[n]$. For a subgraph G of $H[n]$ we say G is *H-partite-free* if there is no partite copy of H in G. We say G is *$(H, H[n])$-partite-saturated* if G is H-partite-free but for any $uv \in E(H[n] \setminus G)$ the graph $G \cup uv$ is not H-partite-free. We consider the problem of determining the value

$$\mathrm{sat_p}(H, H[n]) := \min\left\{e(G) : G \subseteq H[n] \text{ is } (H, H[n])\text{-partite-saturated}\right\}$$

for graphs H.

Note that for a graph H with no homomorphism onto any proper subgraph of itself we have by definition $\mathrm{sat_p}(H, H[n]) = \mathrm{sat}(H, H[n])$. In this way we know that $\mathrm{sat_p}(K_3, K_3[n]) = 6n-6$ from [41] and can drop the partite requirement when considering cliques.

Our main result is the following looking at $(K_4, K_4[n])$-saturation.

Theorem 3.1. *For all large enough $n \in \mathbb{N}$ we have*

$$\mathrm{sat}(K_4, K_4[n]) = 18n - 21\,.$$

Furthermore we determine the unique graph achieving equality.

In addition we calculate the partite-saturation numbers of stars and paths proving the following two results.

Theorem 3.2. *For any $r \geqslant 2$ and $n \in \mathbb{N}$ all $(K_{1,r}, K_{1,r}[n])$-partite-saturated graphs have exactly $(r-1)n^2$ edges.*

Theorem 3.3. *For any $r \geqslant 4$ and $n \geqslant 2r$ we have the following.*

$$\mathrm{sat_p}(P_r, P_r[n]) = \begin{cases} (\frac{r}{2}-1)n^2 + (r-2)n + 3 - r, & \text{for } r \text{ even} \\ (\frac{r}{2}-\frac{1}{2})n^2 + (r-4)n + 5 - r, & \text{for } r \text{ odd} \end{cases}$$

In the original paper by Erdős, Hajnal and Moon they did not in fact require the graph G to be H-free but only required that the addition of any edge would create an extra copy of H. Interestingly for the problem they studied this did not have an effect as the extremal graphs were K_r-free even without requiring this restriction. We consider a similar notion in the partite setting. For $G \subseteq H[n]$ and $n \in \mathbb{N}$ we say G is $(H, H[n])$*-partite-over-saturated* if for any $uv \in E(H[n] \setminus G)$ the graph $G \cup uv$ has more partite copies of H than G. We also ask, given a graph H and $n \in \mathbb{N}$, the value of

$$\operatorname{exsat}_{\mathrm{p}}(H, H[n]) := \min\left\{e(G) : G \subseteq H[n] \text{ is } (H, H[n])\text{-partite-over-saturated}\right\}.$$

We observe some interesting differences in behaviour between these partite-saturation numbers and the saturation numbers studied by Erdős, Hajnal and Moon. Whilst for graphs on r vertices cliques gave the largest values of $\operatorname{sat}(H, n)$ we find that cliques are not the graphs which maximise $\operatorname{sat}_{\mathrm{p}}(H, H[n])$. In fact we prove the following theorem which shows that $\operatorname{sat}_{\mathrm{p}}(H, H[n])$ grows quadratically for graphs H which are not 2-connected whilst it grows linearly for those which are.

Theorem 3.4. *For any graph H with $e(H) \geqslant 2$ and no isolated vertices, if H is 2-connected then $\operatorname{sat}_{\mathrm{p}}(H, H[n]) = \Theta(n)$ and if H is not 2-connected then $\operatorname{sat}_{\mathrm{p}}(H, H[n]) = \Theta(n^2)$.*

On the other-hand we show in Theorem 3.5 that cliques do maximise the partite-over-saturation numbers and that all partite-over-saturation numbers are linear.

Theorem 3.5. *For any integer $r \geqslant 4$ and all large enough $n \in \mathbb{N}$ we have*

$$\operatorname{exsat}_{\mathrm{p}}(K_r, K_r[n]) = (2n-1)\binom{r}{2}.$$

Finally we determine the partite-over-saturation numbers of trees.

Theorem 3.6. *For any tree T on at least 3 vertices and any natural number $n \geqslant 4$ we have $\operatorname{exsat}_{\mathrm{p}}(T, T[n]) = (|T| - 1)n$.*

Organisation Section 3.2 is dedicated to determining the partite-saturation number of K_4. In Section 3.3 we then determine the partite-saturation numbers of paths and stars. We look at the link between 2-connectivity and the order of magnitude of partite-saturation numbers in Section 3.4 before focusing on partite-over-saturation numbers in Section 3.5. Finally in Section 3.6 we give some further remarks and open problems.

3.2 The partite-saturation number of K_4

Theorem 3.1 For all large enough $n \in \mathbb{N}$ we have

$$\operatorname{sat}(K_4, K_4[n]) = 18n - 21.$$

Furthermore we determine the unique graph for which equality holds.

We first give a construction of a graph $G \subseteq K_4[n]$ that is $(K_4, K_4[n])$-saturated and has $18n - 21$ edges.

Let X_1, X_2, X_3, X_4 be the parts of $K_4[n]$. Choose vertices x_i and x_i' in each X_i. Let Z denote the set of these 8 vertices. Include in G the following 15 edges x_1x_2, x_1x_2', x_1x_3', x_1x_4', $x_1'x_2'$, $x_1'x_3$, $x_1'x_4$, x_2x_3, x_2x_4, x_2x_4', $x_2'x_3'$, $x_2'x_4$, x_3x_4', $x_3'x_4$, $x_3'x_4'$. These are the edges drawn in the figure below. We now only add edges between Z and $V(G)\setminus Z$. Include all edges between $X_1 \setminus Z$ and each of x_2, x_3, x_3' and x_4. Attach all vertices in $X_2 \setminus Z$ to x_1', x_3, x_3', x_4 and x_4'. Join all of $X_3 \setminus Z$ to each of x_1, x_1', x_2 and x_4 and finally add all edges from $X_4 \setminus Z$ to x_1, x_1', x_2, x_2' and x_3.

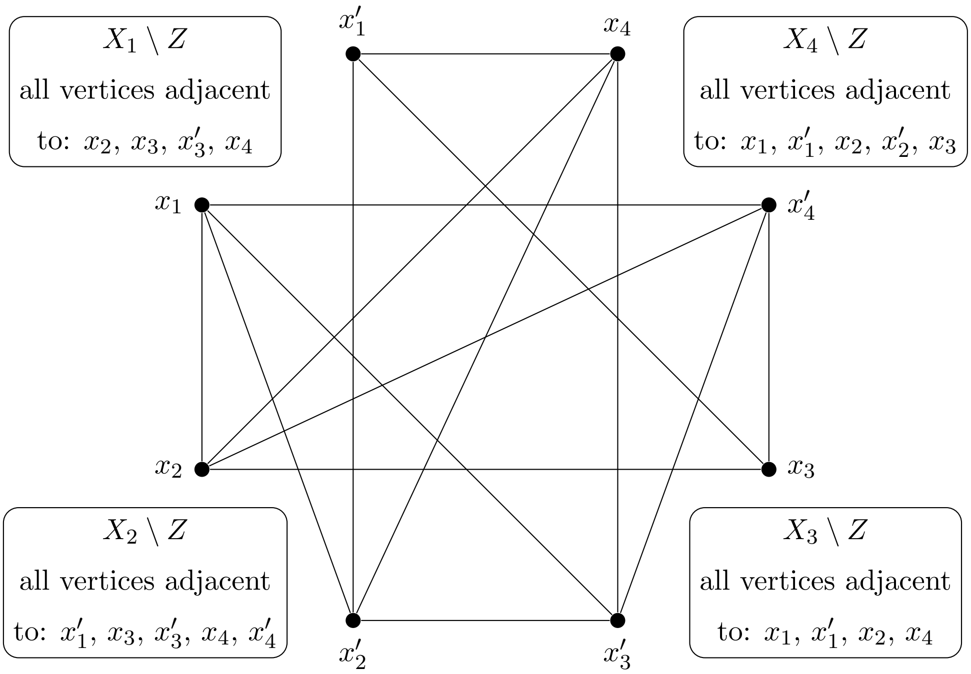

Figure 3.1: K_4-Partite-Saturation Construction

Proposition 3.7. *G is a $(K_4, K_4[n])$-saturated graph with $18n - 21$ edges.*

Proof. To see that this graph is K_4-free note that the graph induced on $V(G) \setminus Z$ has no edges so any K_4 would have to come from a triangle in Z extended to a vertex outside of Z. There are just six triangles induced on Z and none of them extend to a K_4.

To see that G is $(K_4, K_4[n])$-saturated we first observe that for any pair i, j there is an edge in Z such that $(X_i \cup X_j) \setminus Z$ is contained in the common neighbourhood of the ends of that edge. Therefore we could only add an edge with at least one end in Z.

For a vertex $v \in X_1 \setminus Z$ the only incident edges we could add are vx_2' or vx_4'. These additional edges would create a K_4 on $vx_2'x_3'x_4$ or $vx_2x_3x_4'$ respectively. For a vertex $v \in X_2 \setminus Z$ the only incident edge we could add is vx_1 but this would create a K_4 on $x_1vx_3'x_4'$. Similar arguments show we cannot add edges incident to $X_3 \setminus Z$ and $X_4 \setminus Z$. Adding any edge to Z that has either x_1 or x_3' as an endpoint will create a K_4 in Z. Adding any other edge of Z will create a triangle on Z that extends to a K_4 with a vertex outside of Z. That G has $18n - 21$ edges is easy to check. □

Before proving a matching lower bound we need the following lemmas.

Lemma 3.8. *Any $(K_4, K_4[n])$-saturated graph G with $n \geqslant 2$ has minimum degree at least 4.*

Proof. Let G be a $(K_4, K_4[n])$-saturated graph on $X_1 \cup \cdots \cup X_4$. Suppose for contradiction that there exists $a_1 \in X_1$ with at most 3 neighbours. If a_1 has no neighbours in one part, say X_2, then by saturation it must be adjacent to all vertices in the other parts, which for $n \geqslant 2$ contradicts the fact that $\deg(a_1) \leqslant 3$. So a_1 must have exactly three neighbours with one in each of the parts. Call these $x_i \in X_i$ for $i = 2, 3, 4$. Then for any $i = 2, 3, 4$ adding the edge a_1y_i for some $y_i \in X_i \setminus x_i$ must create a K_4. This implies that x_2x_3, x_2x_4 and x_3x_4 are all edges of G but along with a_1 this gives a K_4. □

We can also say more about the neighbourhoods of vertices with degree exactly 4.

Lemma 3.9. *Let G be a $(K_4, K_4[n])$-saturated graph on $X_1 \cup \cdots \cup X_4$ with $n \geqslant 3$ and let v be a vertex of degree exactly 4. Then v has one neighbour in each of two parts and two neighbours in one part. The neighbourhood of v induces a path beginning and ending with the vertices in the same part. All neighbours of v have degree at least $n - 2$.*

Proof. Suppose $v \in X_1$. If v had no neighbour in some X_i ($i \neq 1$) it would be adjacent to all vertices in other parts meaning it would have degree greater than 4. Suppose without loss of generality that the neighbours of v are x_2, x_3, x_3' and x_4 with the subscripts denoting the parts containing each vertex. By considering the effect of adding the edge vy_3 for some $y_3 \in X_3 \setminus \{x_3, x_3'\}$ we see that the edge x_2x_4 is present. We also see that all vertices in $X_3 \setminus \{x_3, x_3'\}$ are adjacent to x_2 and x_4. Similarly by considering a vertex in $X_2 \setminus \{x_2\}$ we see that there must be an edge between x_4 and one of x_3 or x_3'. Without loss of generality assume x_4x_3' is present. Finally by considering a vertex in $X_4 \setminus \{x_4\}$ we see that x_2 is adjacent to either x_3 or x_3'. In order not to create a K_4 it must be that x_2x_3 is present. We now cannot have the edges x_4x_3 or x_2x_3'. We then see that all vertices in $X_4 \setminus \{x_4\}$ are adjacent to x_3 and all vertices in $X_2 \setminus \{x_2\}$ are adjacent to x_3'. Hence the neighbours of v all have degree at least $n-2$. □

It follows that when $n > 6$ vertices of degree exactly 4 cannot be adjacent.

The following lemma gives us minimum degree conditions that more reflect those of the upper bound construction.

Lemma 3.10. *Let G be a $(K_4, K_4[n])$-saturated graph with $n \geqslant 22$ on $X_1 \cup \cdots \cup X_4$. There cannot be two degree 4 vertices, $a_i \in X_i$ and $a_j \in X_j$ with $i \neq j$ such that a_i has just one neighbour in X_j. Furthermore there are at most two parts with minimum degree 4.*

Proof. Suppose for contradiction that $a_1 \in X_1$ and $a_2 \in X_2$ are degree 4 vertices such that a_1 has just one neighbour in X_2 and let x_2, x_3, x_3', x_4 denote the neighbours of a_1. Then by Lemma 3.9 (up to switching between x_3' and x_3) the edges $x_2x_3, x_2x_4, x_3'x_4$ are all present. We also know that x_2 is adjacent to all of $(X_3 \cup X_4) \setminus x_3'$, that x_3 is adjacent to all of $X_4 \setminus x_4$, that x_3' is adjacent to all of $X_2 \setminus x_2$, and x_4 is adjacent to all of $(X_2 \cup X_3) \setminus x_3$. In particular this implies we have the edges a_2x_3' and a_2x_4. The vertex a_2 also has some neighbour $x_1 \in X_1 \setminus a_1$. As a_2 has degree 4 it must have one more neighbour. We split into cases depending on where this final neighbour is and show that each case leads to a contradiction. The possible cases are:

(i) a_2 has another neighbour $v \in (X_1 \cup X_3) \setminus \{a_1, x_1, x_3, x_3'\}$.

(ii) a_2 is adjacent to x_3.

(iii) a_2 has another neighbour $x_4' \in X_4 \setminus x_4$.

Case i) Since x_3 is not adjacent to a_2 it must be adjacent to x_4 as $X_4 \cap N(y_2) = \{x_4\}$ and hence $x_1x_2x_3x_4$ forms a K_4.

Case ii) By considering vertices in $X_3 \setminus N(a_2)$ we must have the edge x_1x_4 and we see that x_1 is adjacent to all of $X_3 \setminus \{x_3, x_3'\}$. We also see that all vertices in $X_1 \setminus N(a_2)$ are adjacent to x_3' and x_4. This means that in fact all vertices in $(X_1 \cup X_2) \setminus \{x_1, x_2\}$ are adjacent to x_3' and x_4 and hence all edges in $X_1 \cup X_2$ have one end in $\{x_1, x_2\}$. In fact all edges in $X_1 \cup X_2$ have exactly one end in $\{x_1, x_2\}$ as if the edge x_1x_2 were present this would create a K_4 with x_4 and any vertex in $X_3 \setminus \{x_3, x_3'\}$.

If all vertices in $X_3 \setminus \{x_3, x_3'\}$ were adjacent to all of $X_4 \setminus x_4$ this would give at least $(n-2)(n-1)$ edges which is greater than $18n$ for $n \geqslant 22$. Therefore consider some vertex $v_3 \in X_3 \setminus \{x_3, x_3'\}$ which is non-adjacent to some $v_4 \in X_4 \setminus x_4$. As v_4 is non-adjacent to v_3 it must be adjacent to both ends of an edge in $N(v_3) \cap (X_1 \cup X_2)$. We know that this edge has exactly one end in $\{x_1, x_2\}$ but this creates a K_4 with v_3 and x_4.

Case iii) As x_4' is not adjacent to a_1 it is adjacent to x_2 and x_3. By considering vertices in $X_4 \setminus N(a_2)$ we see that x_1x_3' is an edge of G and all vertices in $X_4 \setminus N(a_2)$ are adjacent to x_1 and x_3'. By considering vertices in $X_3 \setminus N(a_2)$ we see that x_1x_4' is an edge of G (as x_1x_4 would create a K_4) and all vertices in $X_3 \setminus N(a_2)$ are adjacent to x_1 and x_4'. Finally by considering vertices in $X_1 \setminus N(a_2)$ we observe that all vertices in $X_1 \setminus x_1$ are adjacent to x_3' and x_4 (as x_4' cannot be adjacent to x_3'). Now we know that all vertices in $(X_1 \cup X_2) \setminus \{x_1, x_2\}$ are adjacent to both ends of the edge $x_3'x_4$ and so there are no edges in $(X_1 \cup X_2) \setminus \{x_1, x_2\}$. Furthermore $x_1x_2 \notin E(G)$ as this would create a K_4 with x_4' and any vertex in $X_3 \setminus \{x_3, x_3'\}$. If all vertices in $X_3 \setminus \{x_3, x_3'\}$ were adjacent to all of $X_4 \setminus \{x_4, x_4'\}$ there would be at least $(n-2)^2$ edges in G which is more than $18n$ edges for $n \geqslant 22$. Therefore we can assume there is a vertex $v_3 \in X_3 \setminus \{x_3, x_3'\}$ and a vertex $v_4 \in X_4 \setminus \{x_4, x_4'\}$ which is not adjacent to v_3. Then v_4 must be adjacent to both ends of an edge e in $N(v_3) \cap (X_1 \cup X_2)$. This edge has exactly one end in $\{x_1, x_2\}$. If the edge e is incident to x_2 but not x_1 then it forms a K_4 with v_3 and x_4. If instead e is incident to x_1 but not x_2 it forms a K_4 with x_3' and v_4.

It follows from Lemma 3.9 and the above that there can be at most two parts with minimum degree exactly 4 otherwise we would have a degree 4 vertex with just one neighbour in the part containing another degree 4 vertex. □

Another distinctive feature of the upper bound construction is that low degree vertices are not adjacent to other low degree vertices. In proving the lower bound it is helpful to prove that at most a constant number of low degree vertices are adjacent to other low degree vertices. We do that in the following lemma.

Lemma 3.11. *For any $k \geqslant 5$ suppose G is a $(K_4, K_4[n])$-saturated graph on $X_1 \cup \cdots \cup X_4$. Then there are at most $24k^2(2k^2)^{2k^2}$ vertices v such that $5 \leqslant \deg(v) \leqslant k$ and v is adjacent to another vertex of degree between 5 and k.*

Proof. Call a vertex *bad* if it satisfies $5 \leqslant \deg(v) \leqslant k$ and is adjacent to another vertex with degree between 5 and k. Let $K = 24k^2(2k^2)^{2k^2}$ and suppose for contradiction that there are more than K bad vertices in G. Without loss of generality assume there are at least $\frac{K}{4}$ such vertices in X_1. Call the set of these vertices A_0 and let B_0 denote the set of bad vertices in $X_2 \cup X_3 \cup X_4$ which are adjacent to a bad vertex in A_0. By counting $e(A_0, B_0)$ from each side we see that $|A_0| \leqslant e(A_0, B_0) \leqslant k|B_0|$ and hence $|B_0| \geqslant \frac{K}{4k}$. By averaging we may assume without loss of generality that there are at least $\frac{K}{12k}$ bad vertices in X_2 adjacent to vertices in A_0. Let B_1 denote $B_0 \cap X_2$ and let A_1 be the vertices of A_0 which have a neighbour in B_1. Then every vertex in A_1 and B_1 has a neighbour in the other. By double counting we see that $|B_1| \leqslant e(A_1, B_1) \leqslant k|A_1|$ and so we know that both A_1 and B_1 contain at least $\frac{K}{12k^2}$ vertices.

For $i = 0, ..., k^2+1$ we construct a collection of sets $U_i \subseteq X_1$, $V_i \subseteq X_2$ such that $U_{i+1} \subseteq U_i$ and $V_{i+1} \subseteq V_i$. We also select vertices $u_i \in U_i$ and edges $e_i \in E(X_3, X_4)$ such that the following properties are satisfied for all $i = 0, ..., k^2+1$.

(i) All vertices in V_{i+1} are adjacent to both endpoints of e_{i+1}.

(ii) The vertex u_i is adjacent to both endpoints of e_{i+1}.

(iii) $|V_i| \geqslant \frac{K}{12k}(2k^2)^{-i} = 2k(2k^2)^{2k^2-i}$.

(iv) Each vertex in U_i has a neighbour in V_i.

(v) Each vertex in V_i has a neighbour in U_i.

(vi) $|U_i| \geqslant \frac{K}{12k^2}(2k^2)^{-i} = 2(2k^2)^{2k^2-i}$.

(vii) Each vertex in $U_i \cup V_i$ has degree at most k.

Before constructing these objects we show how they prove the lemma. Since $|V_i| \geqslant 2k(2k^2)^{2k^2-i}$ we see that the set V_{k^2+1} is non-empty. Any vertex in V_{k^2+1} is adjacent to both ends of all the edges $e_1, ..., e_{k^2}$. As vertices in V_{k^2+1} have at most k neighbours it must be that two of these edges are the same. If $e_s = e_t$ for some $s < t \leqslant k^2$ then we have that u_t is adjacent to some vertex v in V_s. As v is in V_s it is adjacent to both ends of e_s and so forms a K_4 along with u_t. This gives our contradiction.

We begin constructing these objects by letting $U_0 = A_1$ and $V_0 = B_1$. This ensures that property (vii) always holds. Given U_i and V_i satisfying the above properties we choose any $u_i \in U_i$ and will find U_{i+1}, V_{i+1}, and e_i satisfying the properties above. By saturation for any vertex v in $V_i \setminus N(u_i)$ there exists an edge $e \in E(X_3, X_4)$ such that both v and u_i are adjacent to both of the endpoints of e. Since u_i has at most k neighbours there are fewer than k^2 such candidates for e and hence at least $\frac{1}{k^2}|V_i \setminus N(u_i)|$ vertices of $V_i \setminus N(u_i)$ are adjacent to the endpoints of the same edge $e \in E(X_3, X_4)$. Let e_{i+1} be this edge and let V_{i+1} be the vertices of $V_i \setminus N(u_i)$ that are adjacent to both ends of e_{i+1}. From this we see that properties (i) and (ii) hold.

Using $|V_i| \geqslant \frac{K}{12k}(2k^2)^{-i} \geqslant 2k$ we then have

$$\begin{aligned} |V_{i+1}| \geqslant & \frac{1}{k^2}|V_i \setminus N(u_i)| \geqslant \frac{1}{k^2}(|V_i| - k) \\ \geqslant & \frac{1}{2k^2}|V_i| \geqslant \frac{K}{12k}(2k^2)^{-(i+1)}. \end{aligned}$$

This gives property (iii). We let $U_{i+1} = U_i \cap N(V_{i+1})$ which ensures (iv) and (v). Therefore by double counting $|V_{i+1}| \leqslant e(U_{i+1}, V_{i+1}) \leqslant k|U_{i+1}|$ and we see that $|U_{i+1}| \geqslant \frac{1}{k}|V_{i+1}| \geqslant \frac{K}{12k^2}(2k^2)^{-(i+1)}$ giving (vi). □

With these lemmas we are now ready to prove Theorem 3.1.

Proof of Theorem 3.1. Let G be a $(K_4, K_4[n])$-saturated graph.

We first make the following claim, the proof of which we postpone, about the minimum degree conditions of the parts of G.

Claim 3.12. *If G has at most $18n-21$ edges then G has precisely two parts of minimum degree exactly 4 and two parts of minimum degree exactly 5.*

From Lemma 3.10 we know that all degree 4 vertices in the two minimum degree 4 parts have two neighbours in the other minimum degree 4 part. We can now assume we have degree 4 vertices $a_1 \in X_1$ and $a_3 \in X_3$. Let the neighbours of a_1 be x_2, x_3, x_3' and x_4. We see that all vertices in $X_3 \setminus \{x_3, x_3'\}$ (including a_3) are adjacent to x_2 and x_4 and that x_2 and x_4 are adjacent. Let the other two neighbours of a_3 be x_1 and x_1'. Since any vertex v in $X_2 \setminus x_2$ is not adjacent to a_1, adding the edge a_1v must create a K_4 using v and a_1. Similarly, since any vertex v in $X_2 \setminus x_2$ is not adjacent to a_3, adding the edge a_3v must create a K_4 using v and a_3. This implies that v is adjacent to x_4 and that x_4 is adjacent to one of x_1 or x_1' and also one of x_3 or x_3'. Without loss of generality assume we have the edges $x_1'x_4$ and $x_3'x_4$. Similar arguments with a vertex in $X_4 \setminus x_4$ show that all vertices in X_4 are adjacent to x_2 and also that we have the edges x_1x_2 and x_2x_3.

We further see that by saturation every vertex of $(X_1 \cup X_3) \setminus \{x_1, x_1', x_3, x_3'\}$ is adjacent to x_2 and x_4. This means there are no edges with both ends lying in $(X_1 \cup X_3) \setminus \{x_1, x_1', x_3, x_3'\}$. All vertices in $X_2 \setminus x_2$ are adjacent to x_1', x_3' and x_4. All vertices of $X_4 \setminus x_4$ are adjacent to x_1, x_3 and x_2.

We now have that all vertices in $(X_1 \cup X_3) \setminus \{x_1, x_1', x_3, x_3'\}$ are adjacent to x_2 and x_4. All vertices in $X_2 \setminus x_2$ are adjacent to x_1', x_3' and x_4 whilst all vertices in $X_4 \setminus x_4$ are adjacent to all of x_1, x_2 and x_3.

The following claim, for which we again postpone the proof, gives us more conditions on the neighbourhoods of various vertices.

Claim 3.13. *All vertices in $X_1 \setminus \{x_1, x_1'\}$ are adjacent to x_3 and x_3'. All vertices in $X_3 \setminus \{x_3, x_3'\}$ are adjacent to x_1 and x_1'. All vertices in $(X_2 \cup X_4) \setminus \{x_2, x_4\}$ are adjacent to at least 3 of $\{x_1, x_1', x_3, x_3'\}$. Both x_1x_3' and $x_1'x_3$ are edges of G.*

Under the assumption of Claim 3.13 we now see that all vertices in $X_2 \setminus x_2$ are adjacent to x_1', x_3', x_4 and one of x_1 or x_3. Let A^1 denote the set of vertices in $X_2 \setminus x_2$ which are adjacent to x_1 but not x_3 and let A^3 denote the set of vertices in $X_2 \setminus x_2$ which are adjacent to x_3 but not x_1.

Similarly all vertices in $X_4 \setminus x_4$ are adjacent to x_1, x_3, x_2 and one of x_1' or x_3'. Let B^1 denote the set of vertices in $X_4 \setminus x_4$ which are adjacent to x_1' but not x_3' and let B^3 denote the set of vertices in $X_4 \setminus x_4$ which are adjacent to x_3' but not x_1'.

Adding any edge between A^1 and B^1 (likewise between A^3 and B^3) cannot create a K_4 so by saturation the induced graphs on (A^1, B^1) and (A^3, B^3) are complete. Any edge between A^1 and B^3 would create a K_4 with x_1x_3' whilst any edge between A^3 and B^1 would give a K_4 using $x_1'x_3$ therefore the bipartite graphs on (A^1, B^3) and (A^3, B^1) are empty.

Hence we see that there are at least

$$\begin{aligned} &5\big(2n - 2 - |A^1| - |A^3| - |B^1| - |B^3|\big)\\ &+4\big(|A^1| + |A^3| + |B^1| + |B^3|\big)\\ &+|A^1||B^1| + |A^3||B^3| + 4n - 4 + 1 \end{aligned}$$

edges with at least one end in $X_2 \cup X_4$. The $+1$ term comes from the edge x_2x_4 and the $+4n - 4$ term comes from the edges with one end in $\{x_2, x_4\}$ and the other end in $X_1 \cup X_3$. Along with the $4n - 6$ edges between X_1 and X_3 this gives a total of at least

$$18n - 21 + \big(|A^1| - 1\big)\big(|B^1| - 1\big) + \big(|A^3| - 1\big)\big(|B^3| - 1\big) \tag{3.1}$$

edges. We argue that either A^1 or B^1 being non-empty implies the other is non-empty.

Suppose there were a vertex in A^1. Then because it has degree at least 5 but is not adjacent to x_3 it has a neighbour v in $X_4 \setminus x_4$. This neighbour v cannot be adjacent to x_3'

or we would have a K_4. Therefore $v \in B^1$. Similarly for a vertex in B^1. Likewise either of A^3 or B^3 being non-empty implies the other is also non-empty.

This now means we have at least $18n - 21$ edges. Furthermore, since $A^1 \cup A^3 = X_2 \setminus x_2$ and $B^1 \cup B^3 = X_4 \setminus x_4$, equality in (3.1) is attained only if either $|A^1| = |B^3| = 1$ or $|A^3| = |B^1| = 1$. Letting x_2' and x_4' be the vertices in the sets of size 1 we have our extremal construction.

It remains to prove Claims 3.12 and 3.13.

Proof of Claim 3.12. We use Lemma 3.11 applied with $k = 180$. As in Lemma 3.11 we refer to vertices of degree between 5 and k which are adjacent to another such vertex as *bad*.

We now split our vertices into groups by their degrees and whether or not they are bad, and then count edges of G by counting edges between these groups.

We label our groups as follows

- $\mathrm{V_{bad}}$ is the set of bad vertices.
- $A := \{v : \deg(v) \geqslant k+1\}$.
- $B := \{v : 5 \leqslant \deg(v) \leqslant k\} \setminus \mathrm{V_{bad}}$.
- $C := \{v : \deg(v) = 4\}$.

We note that vertices in $B \cup C$ only have neighbours in A.

Now $e(G) \geqslant e(B, A) + e(C, A) \geqslant 5|B| + 4|C|$. We also have $e(G) \geqslant e(A, V(G)) \geqslant \frac{k+1}{2}|A|$. If $|A| \geqslant \frac{36n}{k+1}$ this gives at least $18n$ edges so we may assume $|A| < \frac{36n}{k+1}$.

Along with the fact that $|\mathrm{V_{bad}}| \leqslant K = 24k^2(2k^2)^{(2k^2)}$ we see that $|B| \geqslant 4n - |C| - K - \frac{36n}{k+1}$. Since $e(G) \geqslant 5|B| + 4|C|$ we have at least $20n - |C| - 5K - \frac{180n}{k+1}$ edges.

If we have at most one X_i with minimum degree 4 we know $|C| \leqslant n$. This implies that G has at least $19n - 5K - \frac{180n}{k+1}$ edges. For $k = 180$ and large enough n this is at least $18n$.

We can also rule out the possibility of there being a part with minimum degree greater than 5. With $\mathrm{V_{bad}}$, A, and C defined as above let $B^{(5)} := \{v \in B : \deg(v) = 5\}$ and let $B^{(6+)} := \{v \in B : \deg(v) \geqslant 6\}$. We still have that $|B| = |B^{(5)}| + |B^{(6+)}| \geqslant 4n - |C| - K - \frac{36n}{k+1}$ and $|C| \leqslant 2n$. If one part had minimum degree at least 6 that would imply that $|B^{(5)}| \leqslant n$ and so we would have

$$\begin{aligned} e(G) \geqslant & 6|B^{(6+)}| + 5|B^{(5)}| + 4|C| \\ = & 6|B| - |B^{(5)}| + 4|C| \\ \geqslant & 6\big(4n - |C| - K - \tfrac{36n}{k+1}\big) - |B^{(5)}| + 4|C| \\ = & 24n - 2|C| - |B^{(5)}| - 6K - \tfrac{216n}{k+1} \\ \geqslant & 19n - 6K - \tfrac{216n}{k+1}\,. \end{aligned}$$

For $k = 216$ and n large enough this is more than $18n$. □

Proof of Claim 3.13. We first consider a degree 5 vertex, a_2, in $X_2 \setminus x_2$. We consider separately the cases of whether a_2 is adjacent to neither, one, or both of x_1 and x_3.

Firstly we suppose the vertex a_2 is not adjacent to either of x_1 or x_3. Adding the edge a_2x_1 must create a K_4 using a_2, x_1 and a vertex in X_4. Since x_4 is not adjacent to x_1 it must be the case that a_2 has a neighbour $x'_4 \in X_4 \setminus x_4$. If a_2 had no neighbours in $(X_1 \cup X_3) \setminus \{x'_1, x'_3\}$ there would have to be an edge from x'_1 to x'_3 but this would create a K_4. Assume, without loss of generality, that a_2 has a neighbour $x''_1 \in X_1 \setminus \{x_1, x'_1\}$. By considering vertices in $X_4 \setminus N(a_2)$ we see that x'_3 is adjacent to x''_1. This means we now have a K_4 on the vertices x''_1, a_2, x'_3, x_4.

If instead a_2 had exactly one neighbour from $\{x_1, x_3\}$ then by symmetry we may assume it is adjacent to x_1 but not x_3. By saturation the addition of the edge a_2x_3 must create a K_4. Since x_3 is not adjacent to x_4 the vertex a_2 must have a neighbour x'_4 in $X_4 \setminus x_4$. Now a_2 is adjacent to x_1, x'_1, x'_3, x_4 and x'_4 and because a_2 has degree 5 these are all of its neighbours. As the only neighbour of a_2 in X_3 is x'_3 it must be the case that all vertices in $(X_1 \cup X_4) \setminus N(a_2)$ are adjacent to x'_3. We also see that if any vertex v in $X_3 \setminus \{x_3, x_3\}$ were not adjacent to x'_1 then, since adding the edge a_2v must create a K_4, we must have that v is adjacent to x_1 and x'_4 which would create a K_4 on $\{x_1, x_2, v, x'_4\}$. Therefore every vertex in $X_3 \setminus \{x_3, x'_3\}$ is adjacent to x'_1. By considering vertices on $X_4 \setminus N(a_2)$ it must also be the case that x'_3 is adjacent to x_1. From the fact that x_3 is not adjacent to a_2 we can see that x_3 must be adjacent to x'_1 and that x'_4 is also adjacent to x'_1. Now consider a degree 5 vertex, a_4 in $X_4 \setminus \{x_4, x'_4\}$. We know that a_4 is adjacent to x'_3 and we split into the case of when a_4 is adjacent to x'_1 or not.

If a_4 is not adjacent to x'_1 then a_4 has a neighbour $x'_2 \in X_2 \setminus x_2$. We know that x'_2 is adjacent to x'_1. In order to create a K_4 if $a_4x'_1$ were added it must be the case that x'_2 is adjacent to x_3. As x_1 is the only neighbour of a_4 in X_1 is must be the case that all vertices in $(X_2 \cup X_3) \setminus N(a_4)$ are adjacent to x_1. Now all vertices in $(X_3 \cup X_4) \setminus \{x_3, x_3', x_4\}$ are adjacent to both x_1 and x_2 which are themselves adjacent to each other. Therefore there are no edges between $X_3 \setminus \{x_3, x_3'\}$ and $X_4 \setminus x_4$. We also know that all vertices in $(X_2 \cup X_4) \setminus \{x_2, x_2', x_4, x_4'\}$ are adjacent to both ends of the edge x_1x_3'. Hence there are no edges between $X_2 \setminus \{x_2, x_2'\}$ and $X_4 \setminus \{x_4, x_4'\}$. Since all vertices in $(X_1 \cup X_2) \setminus \{x_1, x_1', x_2\}$ are adjacent to x_3' and x_4 there are no edges between $X_1 \setminus \{x_1, x_1'\}$ and $X_2 \setminus x_2$. In particular any vertex v in $X_1 \setminus \{x_1, x_1'\}$ is not adjacent to x_2' and by considering the K_4 created if a_4v were added we see that v is adjacent to x_3. Since v was arbitrary all vertices in $X_1 \setminus \{x_1, x_1'\}$ are adjacent to x_3. This proves the lemma for this case.

If instead a_4 is adjacent to x_1' then as a_4 is of degree 5 and is adjacent to x_1, x_1', x_2, x_3, and x_3' these are all of its neighbours. Any vertex in $X_1 \setminus \{x_1, x_1'\}$ is non-adjacent to a_4 and so must be adjacent to both ends of some edge in $N(a_4)$. This edge must be x_2x_3 and so all vertices in $X_1 \setminus \{x_1, x_1'\}$ are adjacent to x_3. Similarly vertices in $X_3 \setminus \{x_3, x_3'\}$ are non-adjacent to a_4 and so must be adjacent to x_1. All vertices in $X_2 \setminus x_2$ are non-adjacent to a_4 and hence must be adjacent to an edge in $N(a_4)$ implying each vertex in $X_2 \setminus x_2$ is adjacent to at least one of x_1 or x_3.

Finally we consider the case where a_2 is adjacent to both x_1 and x_3. We can assume all degree 5 vertices in X_4 are adjacent to both x_1' and x_3' or we would be in a situation symmetric to the last case we considered. Let a_4 be such a degree 5 vertex in X_4. Since all vertices in $X_1 \setminus \{x_1, x_1'\}$ and $X_3 \setminus \{x_3, x_3'\}$ are not adjacent to either a_2 or a_4 they must be adjacent to both ends of an edge in $N(a_2)$ and both ends of an edge in $N(a_4)$. This implies that vertices in $X_1 \setminus \{x_1, x_1'\}$ are adjacent to x_3 and x_3' and that vertices in $X_3 \setminus \{x_3, x_3'\}$ are adjacent to x_1 and x_1'. Similarly we see that vertices in $X_4 \setminus x_4$ are non-adjacent to a_2 and hence must be adjacent to an edge in $N(a_2)$. Therefore all vertices in $X_4 \setminus x_4$ are adjacent to one of x_1' or x_3'. Similarly all vertices in $X_2 \setminus x_2$ are adjacent to one of x_1 or x_3. This also shows that at least one of the edges x_1x_3' or $x_1'x_3$ exists. If one of them is not present, say $x_1x_3' \notin E(G)$ then by saturation there is some adjacent pair $b_2 \in X_2 \setminus x_2$, $b_4 \in X_4 \setminus x_4$ which are both adjacent to x_1 and x_3'. We also know, however, that b_2 and b_4 are both adjacent to x_1' and x_3 but this gives a K_4 on x_1', b_2, x_3, b_4. Therefore both x_1x_3' and $x_1'x_3$ exist. □

This completes the proof. □

3.3 Saturation numbers of paths and stars

We begin this section by determining the partite-saturation numbers of stars on at least three vertices.

Lemma 3.14. *For any $r \geqslant 2$, $n \in \mathbb{N}$ and any connected graph H which contains a vertex v such that $H \setminus v$ has r components we have* $\mathrm{sat_p}(H, H[n]) \geqslant (r-1)n^2$.

Theorem 3.2 For any $r \geqslant 2$ and $n \in \mathbb{N}$ all $(K_{1,r}, K_{1,r}[n])$-partite-saturated graphs have exactly $(r-1)n^2$ edges.

We show how Theorem 3.2 follows from Lemma 3.14 before proving Lemma 3.14 itself.

Proof of Theorem 3.2. The star $K_{1,r}$ has a vertex v such that $K_{1,r} \setminus v$ has r connected components and hence $\mathrm{sat_p}(K_{1,r}, K_{1,r}[n]) \geqslant (r-1)n^2$. For any $(K_{1,r}, K_{1,r}[n])$-partite-saturated graph G any vertex in the part corresponding to the centre of the star must have degree at most $(r-1)n$ or by the pigeonhole principle it would have a neighbour in each remaining part giving a partite copy of $K_{1,r}$. This maximum degree condition implies at most $(r-1)n^2$ edges. □

Proof of Lemma 3.14. Let v_1 be the cut-vertex of H and let $v_2, \ldots, v_{r+1}$ be neighbours of v_1 which are in distinct components of $H \setminus \{v_1\}$. Let X_i denote the part of $H[n]$ corresponding to v_i and let H_i denote the component of v_i in $H \setminus \{v_1\}$. Consider an $(H, H[n])$-partite-saturated graph G and an arbitrary vertex $x_1 \in X_1$. If x_1 has fewer than $(r-1)n$ neighbours then there are two parts, say X_2 and X_3, such that each has a vertex non-adjacent to x_1. Call these vertices x_2 and x_3. Since G is saturated adding the edge x_1x_2 must create a copy of H using x_1 and hence there must be a copy of $H \setminus H_2$ in G using x_1. Similarly adding the edge x_1x_3 must create a copy of H implying the existence of a copy of $H \setminus H_3$ at x_1. The union of these two subgraphs contains a partite copy of H which contradicts G being H-free. Hence each vertex in X_1 has at least $(r-1)n$ neighbours and so G has at least $(r-1)n^2$ edges. □

We now determine the partite-saturation numbers of paths on at least 4 vertices.

Theorem 3.3 For any $r \geqslant 4$ and $n \geqslant 2r$ we have the following.

$$\mathrm{sat_p}(P_r, P_r[n]) = \begin{cases} (\frac{r}{2} - 1)n^2 + (r-2)n + 3 - r, \text{ for } r \text{ even} \\ (\frac{r}{2} - \frac{1}{2})n^2 + (r-4)n + 5 - r, \text{ for } r \text{ odd} \end{cases} \tag{3.2}$$

Proof. Let $X_1, \dots, X_r$ be the parts of $P_r[n]$ with X_i adjacent to X_{i+1} for each i.

We first give an upper bound construction. Given subsets $A_i \subseteq X_i$ define the graph G on $\bigcup_i X_i$ to be the graph with precisely the edges that lie in (A_i, A_{i+1}) or $(X_i \setminus A_i, X_{i+1})$ for some $i \leqslant r-1$. For the upper bound if r is even consider the graph G created as above with $A_1 := X_1$, $A_r := \emptyset$, $|A_i| = 1$ for all even $i \leqslant r-2$ and $|A_i| = n-1$ for all odd $3 \leqslant i \leqslant r-1$. If r is odd consider the construction G given as above but with the A_i satisfying $A_1 := X_1$, $A_r = \emptyset$, $|A_{r-1}| = n-1$, $|A_i| = 1$ for all even $i \leqslant r-3$ and $|A_i| = n-1$ for all odd $3 \leqslant i \leqslant r-2$.

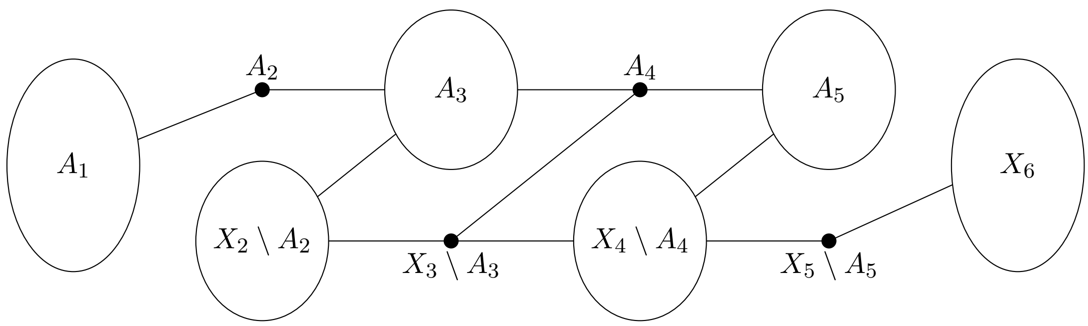

Figure 3.2: P_6-Partite-Saturation Construction

For the lower bound we assume that for some $r \geqslant 4$ and some $n \geqslant 2r$ equation (3.2) does not hold. Then consider the least such r and some $n \geqslant 2r$ for which (3.2) fails. In particular by this minimality and Theorem 3.2 (which gives the partite-saturation of $K_{1,2} = P_3$) we see that

$$\mathrm{sat_p}(P_{r-1}, P_{r-1}[n]) \geqslant \left(\tfrac{r-1}{2} - 1\right)n^2 . \tag{3.3}$$

Now consider a $(P_r, P_r[n])$-partite-saturated graph G on $X_1 \cup \cdots \cup X_r$. Let N_2 denote the set of vertices in X_2 which are adjacent to at least one vertex of X_1. For each $i \geqslant 3$ let N_i denote the set of vertices of X_i which are adjacent to at least one vertex of N_{i-1}. Since there can be no partite path on r vertices it must be the case that $N_r = \emptyset$. If $N_{r-1} = \emptyset$ then (X_{r-1}, X_r) must be complete in G as adding an edge to this pair cannot create a partite copy of P_r. If (X_{r-1}, X_r) is complete then $X_1 \cup \cdots \cup X_{r-1}$ is $(P_{r-1}, P_{r-1}[n])$-partite-saturated so by (3.3) there are at least $\frac{r-1}{2}n^2$ edges in G. This is at least as many as required. Therefore we may assume $N_i \neq \emptyset$ for all $2 \leqslant i \leqslant r-1$. If $N_i = X_i$ for some $i \geqslant 2$ then the pairs (X_j, X_{j+1}) are complete for all $1 \leqslant j \leqslant i-1$. Then $X_i \cup \cdots \cup X_r$ is $(P_{r-i+1}, P_{r-i+1}[n])$-partite-saturated so by (3.3) there are at least $(\frac{r-1}{2})n^2$ edges in G.

This is at least as many as required. We now assume $N_i \neq X_i$ for all $2 \leqslant i \leqslant r$ so for all $i = 2, \dots, r-1$ we have $1 \leqslant |N_i| \leqslant n-1$. For each $i \geqslant 2$ let $\overline{N_i}$ denote $X_i \setminus N_i$. We observe that (X_1, N_1) and $(\overline{N_{r-1}}, X_r)$ must be complete. As are (N_i, N_{i+1}) and $(\overline{N_i}, X_{i+1})$ for $2 \leqslant i \leqslant r-2$ because adding edges to either of these pairs cannot create a partite copy of P_r. Therefore we find that G has all possible edges except those in pairs $(X_1, \overline{N_2})$ or $(N_i, \overline{N_{i+1}})$ for $2 \leqslant i \leqslant r-1$ and so $e(G)$ is at least

$$(r-1)n^2 - n|\overline{N_2}| - \sum_{i=2}^{r-1} |N_i||\overline{N_{i+1}}| = (r-2)n^2 + n|N_2| - n\sum_{i=2}^{r-1}|N_i| + \sum_{i=2}^{r-2}|N_i||N_{i+1}| \,. \quad (3.4)$$

Suppose $N_2, ..., N_{r-1}$ have been chosen to minimise the above expression under the assumption that each $|N_i|$ is between 1 and $n-1$. The contribution to (3.4) from terms that include N_2 is exactly $|N_2||N_3|$ which (regardless of the value of $|N_3|$) is minimised by taking $|N_2| = 1$. For $3 \leqslant i \leqslant r-2$ the contribution to (3.4) from terms that include N_i is

$$|N_i|\big(|N_{i-1}| + |N_{i+1}| - n\big)\,.$$

When $|N_{i-1}| = 1$ the above expression is at most zero and so minimised by taking $|N_i| = n-1$. If $|N_{i-1}| = n-1$ it is at least zero and so minimised by taking $|N_i| = 1$. In this way using $|N_2| = 1$ we can see that for $2 \leqslant i \leqslant r-2$ we have the following.

$$|N_i| = \begin{cases} 1, \text{ for } i \text{ even} \\ n-1, \text{ for } i \text{ odd} \end{cases}$$

The contribution to (3.4) from the N_{r-1} terms is $|N_{r-1}|\big(|N_{r-2}| - n\big)$ which is always negative and so the expression is minimised when $|N_{r-1}| = n-1$. The graph given with the N_i taking these sizes is the same as our upper bound construction completing the proof. □

3.4 2-connectivity and the growth of saturation numbers

Recall that a graph is 2-*connected* if after the removal of any single vertex it is still connected. Observe that if H' can be obtained from H by adding or removing isolated vertices then $\mathrm{sat_p}(H, H[n]) = \mathrm{sat_p}(H', H'[n])$. It is also clear that $\mathrm{sat_p}(K_2, K_2[n]) = 0$.

Theorem 3.4 For any graph H with $e(H) \geqslant 2$ and no isolated vertices, if H is 2-connected then $\mathrm{sat_p}(H, H[n]) = \Theta(n)$ and if H is not 2-connected then $\mathrm{sat_p}(H, H[n]) = \Theta(n^2)$.

Proof. If H is connected but not 2-connected then there must be a cut vertex, v, of H such that $H \backslash v$ has at least two components. Then by Lemma 3.14 we have $\mathrm{sat_p}(H, H[n]) \geqslant n^2$.

We now consider the case when H is disconnected but has no isolated vertices. Let H_1 and H_2 be two connected components of H. If $G \subseteq H[n]$ is $(H, H[n])$-partite-saturated then by saturation the induced graph of G onto at least one of $H_1[n]$ or $H_2[n]$ must be complete. Since each H_i contains an edge this means G has at least n^2 edges.

Finally we consider the case when H is 2-connected. The fact that $\mathrm{sat_p}(H, H[n]) = \Omega(n)$ comes from the fact that in an $(H, H[n])$-saturated graph G every vertex, x, has degree at least one. If not adding an edge incident to x would not create a copy of H since H has minimum degree at least two by 2-connectivity.

We now give an upper bound construction. For each edge ij of H we define H_{ij} to be the graph obtained from H be removing all edges incident to i or j including the edge ij. We define $V_i(H_{ij})$ to be the vertices of $H_{ij} \setminus \{i, j\}$ which were incident to i in H. Similarly $V_j(H_{ij})$. For $n \geqslant e(H)$ we let $G_1 \subseteq H[n]$ be the disjoint union of a copy of H_{ij} for each edge ij of H. Create G_2 from G_1 by adjoining each vertex of $V_i(H_{ij})$ (in the copy of H_{ij} in G_1) to every vertex in $X_i \setminus V(G_1)$, and by adjoining each vertex of $V_j(H_{ij})$ to every vertex in $X_j \setminus V(G_1)$ for each edge ij of H. We then create G_3 from G_2 by arbitrarily adding edges until the graph is $(H, H[n])$-partite-saturated.

We claim that G_3 is $(H, H[n])$-partite-saturated and has at most $2e(H)^2 n - e(H)^3$ edges. To prove this it is sufficient to show that G_2 has no partite copy of H and that G_3 has at most $2e(H)^2 n - e(H)^3$ edges. We first note that there are no edges of G_2 or G_3 with both end points in $V(G_3) \setminus V(G_1)$ since any such edge $x_i x_j$ would form a copy of H with the H_{ij}. We can then bound the number of edges of G_3 by $E(H[n]) - E(H[n - e(H)]) = n^2 e(H) - (n - e(H))^2 e(H) = 2e(H)^2 n - e(H)^3$.

Suppose now for contradiction that G_2 has a partite copy of H. Denote the vertices of this copy of H by x_i for $i = 1, \ldots, |H|$. Since G_1 is H-free at least one of the x_i's lies in $V(G_2) \setminus V(G_1)$. Suppose without loss of generality that $x_1 \notin V(G_1)$. Let x_2 be a neighbour of x_1 in the partite copy of H. Since there are no edges with both end points in $V(G_2) \setminus V(G_1)$ it must be the case that $x_2 \in V(G_1)$. Since $x_1 x_2$ is an edge of G_2 it must be the case that $x_2 \in V_1(H_{1i})$ for some i adjacent to 1 in H. Suppose $x_2 \in V_1(H_{13})$. Then similarly $x_3 \in V_1(H_{1k})$ for some $k \neq 3$. Therefore x_2 and x_3 are in different H_{ij}'s and hence different connected components of G_1. Our copy of H is separated by following the set

$$\{x_i : x_i \notin V(G_1) \text{ and } x_i \text{ is adjacent to a vertex in } H_{13}\}.$$

Since H is 2-connected this set must contain at least two vertices, one of which is x_1. The only x_i's that vertices in H_{13} can be adjacent to outside of H_{13} are x_1 and x_3 but $x_3 \in V(G_1)$ which gives a contradiction. □

3.5 Over-saturation numbers

In this section we determine the partite-over-saturation numbers of cliques and trees, and show that of graphs on r vertices the cliques have the largest partite-over-saturation numbers.

Since it follows from the proof of $\mathrm{sat}(K_3, K_3[n]) = 6n-6$ in [41] that $\mathrm{exsat_p}(K_3, K_3[n]) = 6n-6$ we look only at cliques on at least 4 vertices. The proof of the following Theorem uses ideas from [41].

Theorem 3.5 For any integer $r \geqslant 4$ and all large enough $n \in \mathbb{N}$ we have

$$\mathrm{exsat_p}(K_r, K_r[n]) = (2n-1)\binom{r}{2}.$$

Proof. For the upper bound consider the graph G consisting of a copy of K_r with each vertex of this clique adjacent to all vertices in adjacent parts of $K_r[n]$. For the lower bound consider a $(K_r, K_r[n])$-partite-over-saturated graph G on $X_1 \cup \cdots \cup X_r$.

For all $i = 1, \ldots, r$ let $\delta_i := \min\{\deg(x) : x \in X_i\}$. Since for any i we have $e(G) \geqslant \delta_i n$ we must have $\delta_i < r^2$ or G would have more than $(2n-1)\binom{r}{2}$ edges. By the fact that any vertex which is not adjacent to some part must be incident to all vertices in the other parts we see that $\delta_i \geqslant r-1$ for all i.

Claim 3.15. *All vertices of degree $r-1$ are in a K_r.*

Proof. If $v \in X_1$ is a vertex of degree $r-1$ it must have a neighbour in each adjacent part. Denote these by $x_i \in X_i$ for $i = 2, ..., r$. For any $y_2 \in X_2 \setminus x_2$ adding the edge vy_2 must create a new K_r. This new clique must be on $\{v, y_2, x_3, x_4, ..., x_r\}$ so $x_3, ..., x_r$ must all be pairwise adjacent. Similarly for any $y_3 \in X_3 \setminus x_3$ adding the edge vy_3 must create a new K_r (which must be $\{v, x_2, y_3, x_4, x_5, ..., x_r\}$) so $x_2, x_4, x_5, ..., x_r$ must all be pairwise adjacent. This gives a K_r on $v, x_2, x_3, ..., x_r$. □

Let x_i be a vertex of degree δ_i for each i. For each i let $Y_i := \bigcup_{j\neq i} \big(N(x_j) \cap X_i\big)$ and let $Y := \bigcup_i Y_i = \bigcup_i N(x_i)$. Observe that $|Y| \leqslant r^3$.

Claim 3.16. *For all $i \neq j$, each vertex in $X_i \setminus Y_i$ has a neighbour in Y_j.*

Proof. Given some $i \neq j$ and a vertex $v \in X_i \setminus Y_i$ consider any $k \in \{1, ..., r\} \setminus \{i, j\}$. As v is not in Y_i it must be that v is not adjacent to x_k. Therefore, by saturation, adding vx_k creates a new K_r. This K_r must use a neighbour of x_k in X_j and hence this neighbour is both in Y_j and also adjacent to v. □

We can now lower bound the edges of G by

$$\begin{aligned}
e(G) &\geqslant e(Y, X\setminus Y) + e(X \setminus Y)\\
&\geqslant \sum_{v\in X\setminus Y} \Big(\deg(v,Y) + \tfrac{1}{2}\big(\deg(v, X\setminus Y)\big)\Big)\\
&\geqslant \sum_i |X_i \setminus Y|\big(r-1+\tfrac{1}{2}\big(\delta_i-(r-1)\big)\big)\\
&= \tfrac{1}{2}(r-1)|X\setminus Y| + \tfrac{1}{2}\sum_i |X_i\setminus Y|\delta_i\\
&\geqslant \tfrac{1}{2}n\Big(r(r-1)+\sum_i \delta_i\Big) - \tfrac{1}{2}r^3\Big(r-1+\sum_i \delta_i\Big)\\
&\geqslant \tfrac{1}{2}n\Big(r(r-1)+\sum_i \delta_i\Big) - r^6\\
&= (2n-1)\binom{r}{2} + \tfrac{1}{2}n\sum_i \big(\delta_i-(r-1)\big) + \binom{r}{2} - r^6\,.
\end{aligned} \tag{3.5}$$

The third inequality comes from the minimum degree condition. The first equality uses $|X \setminus Y| = \sum_i |X_i \setminus Y|$. The fourth inequality uses $|Y| \leqslant r^3$ whilst the fifth uses $\delta_i \leqslant r^2$.

By equation (3.5) for $n > 2r^6$ we have $\delta_i = r-1$ for all i. Each of the x_i's has one neighbour in each adjacent part and is in a copy of K_r. We see that by saturation for a vertex v of degree $r-1$ every vertex w in a different part from v which is not adjacent to v is incident to all neighbours of v outside of the part of w. Therefore vertices of degree $r-1$ are not adjacent. We also see that for any $i \neq j$ the vertices x_i and x_j have $r-2$ common neighbours and so with the sets Y_i and Y as before we find that $|Y_i| = 1$ for all i, so $|Y| = r$.

Using (3.5) we get

$$\begin{aligned} e(G) &= e(Y, X \setminus Y) + e(X \setminus Y) + e(Y) \\ &\geqslant \tfrac{1}{2}(r-1)|X \setminus Y| + \tfrac{1}{2}\sum_i |X_i \setminus Y|\delta_i + e(Y) \\ &\geqslant 2(n-1)\binom{r}{2} + e(Y)\,. \end{aligned}$$

Since there is a K_r on Y we have $e(Y) = \binom{r}{2}$ and the result follows. □

The upper bound construction can be generalised to any H by letting G consist of a copy of H with each vertex of this H adjacent to all vertices in adjacent parts of $H[n]$. This gives an upper bound of

$$\mathrm{exsat_p}(H, H[n]) \leqslant (2n-1)e(H)\,.$$

In particular this shows that over graphs H on r vertices the cliques give rise to the largest value of $\mathrm{exsat_p}(H, H[n])$ and also that all partite-over-saturation numbers of graphs with at least two edges are linear.

Next we determine the partite-over-saturation number of trees.

Theorem 3.6 For any tree T on at least 3 vertices and any natural number $n \geqslant 4$ we have $\mathrm{exsat_p}(T, T[n]) = (|T|-1)n$.

Proof. For an upper bound construction let G be the union of n disjoint partite copies of T.

Turning our attention to the lower bound we let L denote the set of leaves of T and call the vertices in $C = V(T) \setminus L$ *core* vertices.

Now suppose G is a $(T, T[n])$-partite-over-saturated graph with $n \geqslant 4$. Let x be a vertex of G lying in a part associated to a core vertex $v \in C$. In G the vertex x must either have a neighbour in each adjacent part of $T[n]$ or it must be that $\deg_G(x) \geqslant n(\deg_T(v) - 1) \geqslant 2\deg_T(v)$. This is because if x had no neighbour in some adjacent part it must be adjacent to all vertices in the other adjacent parts. Since $\deg_T(v) \geqslant 2$ and $n \geqslant 4$ this means x has at least $2\deg_T(v)$ neighbours. We let $L[n]$ and $C[n]$ denote the set of vertices in $T[n]$ that lie in parts corresponding to L and C respectively.

We have

$$\begin{aligned} e(G) &= \sum_{x\in C[n]} \left(\tfrac{1}{2}\deg_G(x, C[n]) + \deg_G(x, L[n])\right) \\ &= \frac{1}{2}\sum_{x\in C[n]} \big(\deg_G(x) + \deg_G(x, L[n])\big)\,. \end{aligned} \tag{3.6}$$

Let $x \in C[n]$ be a vertex associated in the part associated to a vertex $v \in C$. If x is adjacent to a vertex in each adjacent part then

$$\deg_G(x) + \deg_G(x, L[n]) \geqslant \deg_T(v) + \deg_T(v, L)$$

otherwise we also obtain

$$\deg_G(x) + \deg_G(x, L[n]) \geqslant \deg_G(x) \geqslant 2\deg_T(v) \geqslant \deg_T(v) + \deg_T(v, L)\,.$$

Using these and (3.6), we see that

$$\begin{aligned} e(G) &\geqslant \frac{n}{2}\sum_{v\in C} \big(\deg_T(v) + \deg_T(v, L)\big) \\ &= n\cdot e(T) = n\big(|T| - 1\big) \end{aligned}$$

completing the proof. □

3.6 Concluding remarks

It would be very nice to be able to determine the value of $\mathrm{sat}(K_r, K_r[n])$ for $r \geqslant 5$. Exact answers here would probably be very difficult though it may be possible to determine up to an error term of $o(n)$ or even $O(1)$. It would be helpful to be able to determine the following value in order to make progress on this problem.

For integers $r \geqslant s \geqslant 3$ let $m(r, s)$ denote the fewest vertices an r-partite graph G can have such that G is K_s-free but every set of $s-1$ parts contains a K_{s-1}.

We can use $m(r, r-1)$ and $m(r-1, r-1)$ to get upper and lower bounds respectively on $\mathrm{sat}(K_r, K_r[n])$.

For the upper bound let $F \subseteq K_r[n]$ be a K_{r-1}-free graph on $m(r, r-1)$ vertices such that any $r-2$ parts contain a K_{r-2}. Create a $(K_r, K_r[n])$-saturated graph $G \subseteq K_r[n]$ by attaching all vertices of F to all vertices outside of F which lie in a different part. Then if necessary add edges between vertices of F until the graph is $(K_r, K_r[n])$-saturated. This implies that $\mathrm{sat}(K_r, K_r[n])$ is less than $m(r, r-1)\cdot(r-1)n$. Using the fact that

$m(4, 3) = 6$ this shows that $\mathrm{sat}(K_4, K_4[n]) \leqslant 18n$ which we know from Theorem 3.1 to be close to the correct answer.

For the lower bound we prove a minimum degree condition in all $(K_r, K_r[n])$-saturated graphs. If G is a $(K_r, K_r[n])$-saturated graph note that any vertex in G is either adjacent to all vertices in one part of $K_r[n]$ or its neighbourhood induces an $(r-1)$-partite graph which is K_{r-1}-free but where there is a K_{r-2} on any $r-2$ parts. Therefore, for $n \geqslant m(r-1, r-1)$ we have $\delta(G) \geqslant m(r-1, r-1)$ and hence $\mathrm{sat}(K_r, K_r[n]) \geqslant m(r-1, r-1) \cdot rn/2$. When $r = 4$ this gives the minimum degree condition of $\delta(G) \geqslant m(3, 3) = 4$.

4

Size-Ramsey numbers of powers of paths

4.1 Introduction

Given graphs G and H and a positive integer q we say that G *is q-Ramsey for H*, denoted $G \to (H)_q$, if every q-colouring of the edges of G contains a monochromatic copy of H. When $q = 2$, we simply write $G \to H$. In its simplest form, the classical theorem of Ramsey [75] states that for any H there exists an integer N such that $K_N \to H$. The *Ramsey number* $R(H)$ of a graph H is defined to be the smallest such N. Ramsey problems have been well studied and many beautiful techniques have been developed to estimate Ramsey numbers. The survey by Conlon, Fox and Sudakov [24] provides a detailed summary of developments in the area.

A number of variants of the classical Ramsey problem have also been introduced and are under active study (the survey [24] also provides a good introduction to these related problems). In particular, Erdős, Faudree, Rousseau and Schelp [35] proposed the problem of determining the smallest number of edges in a graph G such that $G \to H$. More precisely, we define the *size-Ramsey number* $\hat{r}(H)$ of a graph H as $\hat{r}(H) = \min\{|E(G)| \colon G \to H\}$. Here, we are interested in problems involving estimating $\hat{r}(H)$.

For any graph H we have the obvious bound $\hat{r}(H) \leqslant \binom{R(H)}{2}$. A result due to Chvátal (see, e.g., [35]) implies that this is the right value for the size-Ramsey number of complete graphs, i.e., $\hat{r}(K_n) = \binom{R(K_n)}{2}$.

Considering the path P_n on n vertices, Erdős [34] asked the following question.

Question 4.1. *Is it true that*

$$\lim_{n\to\infty} \frac{\hat{r}(P_n)}{n} = \infty \quad \textit{and} \quad \lim_{n\to\infty} \frac{\hat{r}(P_n)}{n^2} = 0?$$

Using a probabilistic construction, Beck [11] proved that the size-Ramsey number of paths is linear, i.e., $\hat{r}(P_n) = O(n)$. Alon and Chung [7] provided an explicit construction of a graph G with $O(n)$ edges such that $G \to P_n$. Recently, Dudek and Prałat [31] gave a simple alternative proof for this result (see also [66]). More generally, Friedman and Pippenger [45] proved that the size-Ramsey number of bounded degree trees is linear (see also [28, 49, 58]) and it is shown in [50] that cycles also have linear size-Ramsey numbers.

Answering a question posed by Beck [12] (negatively), who asked whether $\hat{r}(G)$ is linear for all graphs G with bounded maximum degree, Rödl and Szemerédi showed that there exists a graph H with n vertices and maximum degree 3 such that $\hat{r}(H) = \Omega(n \log^{1/60} n)$. The current best upper bound for bounded degree graphs is proved in [63], where it is shown that, for every Δ, there is a constant c such that for any graph H with n vertices and maximum degree Δ we have

$$\hat{r}(H) \leqslant c n^{2-1/\Delta} \log^{1/\Delta} n.$$

For further results on size-Ramsey numbers the reader is referred to [13, 61, 76].

Given a graph H on n vertices and an integer $k \geqslant 2$, the kth power H^k of H is the graph with vertex set $V(H)$ and all edges $\{u, v\}$ such that the distance between u and v in H is at most k. Answering a question of Conlon [22] we prove that all powers of paths have linear size-Ramsey numbers. The following theorem is our main result.

Theorem 4.1. *For any integer $k \geqslant 2$,*

$$\hat{r}(P_n^k) = O(n). \tag{4.2}$$

Since $C_n^k \subseteq P_n^{2k}$, the following corollary follows directly from Theorem 4.1.

Corollary 1. For any integer $k \geqslant 2$,

$$\hat{r}(C_n^k) = O(n). \tag{4.3}$$

4.2 Proof of Theorem 4.1

To prove Theorem 4.1, we have to show that there is a graph G with $O(n)$ edges such that $G \to P_n^k$. The first result we need guarantees the existence of a bounded degree graph with two useful properties.

Lemma 4.2. *For every integer $k \geqslant 1$ and every $\varepsilon > 0$ there exists a_0 such that the following holds. For any $a \geqslant a_0$ there is a constant b such that, for any large enough n, there is a graph H with $v(H) = an$ and $\Delta(H) \leqslant b$ with the following two properties.*

1. *For every pair of disjoint sets S, $T \subseteq V(H)$ with $|S|, |T| \geqslant \varepsilon n$, we have $|E_H(S,T)| \geqslant 1$.*
2. *For any family of pairwise disjoint sets $A_1, \ldots, A_{k+1} \subseteq V(H)$ of size at least εan each, there is a path $P_n = (x_1, \ldots, x_n)$ in H with $x_i \in A_j$ for all i, where $j \equiv i \pmod{k+1}$.*

Proof. Fix $k \geqslant 1$ and $\varepsilon > 0$. Let

$$a_0 = 2 + \frac{4}{\epsilon(k+1)}, \tag{4.4}$$

and suppose $a \geqslant a_0$ is given. Let

$$c = \frac{4a}{\epsilon^2} \tag{4.5}$$

and

$$b = 4ac. \tag{4.6}$$

Let n be sufficiently large and $G = G(2an, p)$ be the binomial random graph with $p = c/n$. By Chernoff's inequality, with high probability we have $|E(G)| < (4a^2c)n$. Moreover, with high probability G satisfies Property 1 (with $H = G$) by the following reason: Let X be the number of pairs of disjoint subsets of $V(G)$ of size εn with no edges between them. Then, recalling (4.5), we have

$$\mathbb{E}[X] \leqslant \binom{2an}{\varepsilon n}^2 \left(1 - \frac{c}{n}\right)^{(\varepsilon n)^2} < 2^{4an} \cdot e^{-c\varepsilon^2 n} = o(1).$$

By Markov's inequality the probability that there exists such a pair of sets with no edges between also tends to zero. Thus, we can fix a graph G satisfying these properties.

Now let H be a subgraph of G obtained by iteratively removing a vertex of maximum degree until exactly an vertices remain. Then $\Delta(H) \leqslant b$, as otherwise we would have deleted more than $b \cdot an > |E(G)|$ edges from G during the iteration, which, in view of (4.6), is a contradiction. Moreover, as H is an induced subgraph of G, Property 1 is maintained.

It remains to prove that H also satisfies Property 2. To do so, we analyse a depth first search algorithm, adapting a proof idea in [13, Lemma 4.4]. The algorithm receives as input a graph H with $v(H) = an$ satisfying Property 1 and a family of pairwise disjoint sets $A_1, \ldots, A_{k+1} \subseteq V(H)$ with $|A_i| \geqslant \varepsilon an$ for all i. The output is a path $P_n = (x_1, \ldots, x_n)$ in H with $x_i \in A_j$ for all i, where $j \equiv i \pmod{k+1}$.

As it runs, the algorithm builds a path $P = (x_1, \ldots, x_r)$ with $x_i \in A_j$ for all i and j with $j \equiv i \pmod{k+1}$. Furthermore, it maintains sets U_j and $D_j \subseteq A_j$ for all j, with the property that U_j, D_j, and $V(P) \cap A_j$ form a partition of A_j for every j. The sets U_j decrease as the algorithm runs, while the D_j increase. At each step if the path is currently $P = (x_1, \ldots, x_r)$ we look to see if the path can be extended by adding a vertex $u \in U_{r+1}$ that is adjacent to x_r. If such a u exists we call it x_{r+1} extending the path. We also remove u (now called x_{r+1}) from U_{r+1}. We then repeat the procedure with $P = (x_1, \ldots, x_{r+1})$. If, on the other hand, no such $u \in U_{r+1}$ is adjacent to x_r we consider x_r to be a *dead end* and remove it from the path adding it to D_r. We then repeat the procedure for the path $P = (x_1, \ldots, x_{r-1})$. If at any stage the path is empty we simply choose a vertex from U_1 to be the new start of the path.

If at any point the path consists of n vertices, we stop having achieved our aim. We also artificially stop the algorithm if at any time there is an i such that $|D_i| \geqslant \varepsilon n$. We will show that in a graph with Property 1 this does not happen without the path having first had at least n vertices. Of course the algorithm will also stop if both the path and U_1 are empty. This however would have meant $|D_1| = |A_1| \geqslant \varepsilon n$ and so the algorithm would have already been stopped.

A useful observation is that there can never be any edges between D_i and U_{i+1} for any i otherwise the path would have been extended along such an edge and the endpoint in A_i would not have been added to D_i.

Now suppose that (for the first time in the process) there is some i such that $|D_i| \geqslant \varepsilon n$. Note that since there are no edges between D_i and U_{i+1} we must have $|U_{i+1}| < \varepsilon n$ by Property 1. We also know that $D_{i+1} < \varepsilon n$ as D_i was the first 'dead end set' to reach εn vertices. Therefore we see that

$$|V(P) \cap A_{i+1}| > |A_{i+1} - 2\varepsilon n| \geqslant (a-2)\varepsilon n \geqslant \frac{4}{k+1} n\,.$$

Since the path wraps around the sets $A_1, ..., A_{k+1}$ it intersects each set almost as often. In particular we see that $|V(P)| \geqslant (k+1)(\frac{4}{k+1}n - 1)$ which is greater than n when n is large compared to k. □

Remark 4.7. We remark that, in the proof above, we in fact proved that graphs that satisfy 1 also satisfy 2.

The following definition plays an important role in our proof.

Definition 4.3 (*Complete blow-up of H^k*)**.** Given a graph H and positive integers t and k, we denote by H^k_t the graph obtained by replacing each vertex v of the kth power H^k

of H by a complete graph with $R(K_t)$ vertices, the *cluster* $C(v)$, and by adding, for every $\{u, v\} \in E(H^k)$, all the edges between $C(u)$ and $C(v)$.

The simple fact stated below says that complete blow-ups of powers of bounded degree graphs have a linear number of edges.

Fact 4.4. *Let k, t, a and b be positive constants. If H is a graph with $|V(H)| = an$ and $\Delta(H) \leqslant b$, then $|E(H_t^k)| = O(n)$.*

Proof. Since $\Delta(H) \leqslant b$, we have $|E(H^k)| = O(n)$. Therefore, $|E(H_t^k)| \leqslant R(K_t)^2 \cdot |E(H^k)| + R(K_t)^2 an = O(n)$. □

We shall also make use of the following result in our proof.

Theorem 4.5 (Pokrovskiy [73, Theorem 1.7])**.** *Let $k \geqslant 1$. Suppose that the edges of K_n are coloured with red and blue. Then K_n can be covered by k vertex-disjoint blue paths and a vertex-disjoint red balanced complete $(k+1)$-partite graph.*

The proof of Theorem 4.5 is rather complex. We remark that we do not need the full strength of that result, in the sense that we do not need the complete $(k+1)$-partite graph to be balanced; it suffices for us to know that the vertex classes are of comparable cardinality. Such a result can be derived easily by iterating Lemma 1.5 in [73], for which Pokrovskiy gives a short and elegant proof (see also [72, Lemma 1.10]).

We shall also use the classical Kővári–Sós–Turán theorem [65], in the following simple form.

Theorem 4.6. *Let G be a balanced bipartite graph with t vertices in each vertex class. If G contains no $K_{s,s}$, then G has at most $4t^{2-1/s}$ edges.*

To prove Theorem 4.1, we fix a graph H as in Lemma 4.2 and consider its kth power H^k. We then consider, for a suitably large integer t, the complete blow-up H_t^k (see Definition 4.3). We then show that

$$H_t^k \to P_n^k. \tag{4.8}$$

Let us give a brief outline of the proof of (4.8). Suppose the edges of H_t^k have been coloured red and blue by a colouring χ. Recall that H_t^k is obtained by blowing up H^k; in particular, the vertices v of H^k become large complete graphs $C(v)$. By the choice of parameters, Ramsey's theorem tells us that each such $C(v)$ contains a monochromatic K_t. We suppose, without loss of generality, that at least half of the $C(v)$ contain a blue K_t

and let F be the subgraph of H induced by the corresponding vertices v. So, in particular, $|F| \geqslant \frac{1}{2}|H|$.

We shall define an auxiliary edge-colouring χ' of F^k and then we shall show that $F^k \to P_n$. If we find a blue P_n in F^k with the colouring χ', then we shall be able to find a blue P_n^k in H_t^k. On the other hand, if no such blue path P_n exists in F^k, then we shall be able to find a *red* P_n *in* $F \subseteq H$ (not in F^k), with certain additional properties. More precisely, such a red $P_n \subseteq F \subseteq H$ will be found as in 2 in Lemma 4.2, with the sets A_i being the vertex classes of a red $(k+1)$-partite subgraph of F^k as given by Theorem 4.5, applied to a suitable red/blue coloured complete graph (we complete F^k with its auxiliary colouring χ' to a red/blue coloured complete graph by considering non-edges of F^k red). It will then be easy to find a red P_n^k in H_t^k.

Proof of Theorem 4.1. Fix $k \geqslant 1$ and let $\varepsilon = 1/3(k+1)$. Let a_0 be the constant given by an application of Lemma 4.2 with parameters k and ε. Set $a = \max\{6k, a_0\}$ and let b be given by Lemma 4.2 for this choice of a. Moreover, let H be a graph with $|V(H)| = an$ and $\Delta(H) \leqslant b$ be as in Lemma 4.2. Finally, put $t = (64k)^{2k}$ and $s = 2k$.

Let H_t^k be a complete blow-up of H^k, as in Definition 4.3, and let $\chi : E(H_t^k) \to \{\text{red}, \text{blue}\}$ be an edge-colouring of H_t^k. We shall show that H_t^k contains a monochromatic copy of P_n^k under χ. By the definition of H_t^k, any cluster $C(v)$ contains a monochromatic copy $B(v)$ of K_t. Without loss of generality, the set $W := \{v \in V(H) : B(v) \text{ is blue}\}$ has cardinality at least $v(H)/2$. Let $F := H[W]$ be the subgraph of H induced by W, and let F' be the subgraph of $F_t^k \subseteq H_t^k$ induced by $\bigcup_{w \in W} V(B(w))$.

Given the above colouring χ, we define a colouring χ' of F^k as follows. An edge $\{u, v\} \in E(F^k)$ is coloured *blue* if the bipartite subgraph $F'[V(B(u)), V(B(v))]$ of F' naturally induced by the sets $V(B(u))$ and $V(B(v))$ contains a blue $K_{s,s}$. Otherwise $\{u, v\}$ is coloured *red*.

Claim 4.7. *Any* 2*-colouring of* $E(F^k)$ *has either a blue* P_n *or a red* P_n^k*.*

Proof. We apply Theorem 4.5 to F^k, where if an edge is not present in F^k, then we consider it to be in the red colour class. If F^k contains a blue copy of P_n, then we are done. Hence we may assume F^k contains a balanced, complete $(k+1)$-partite graph K with parts $A_1, \dots, A_{k+1}$ on at least $v(F^k) - kn \geqslant an/2 - kn$ vertices, with no blue edges between any two parts. As $a \geqslant 6k$, each one of these parts has size at least

$$\frac{1}{k+1}\left(\frac{a}{2} - k\right)n \geqslant \varepsilon a n. \tag{4.9}$$

By Property 2 of Lemma 4.2 applied to the collection of sets of vertices $A_1, \ldots, A_{k+1}$ of $F \subseteq H$ (specifically F and not F^k), we see that $F[V(K)]$ contains a path with n vertices such that any consecutive $k+1$ vertices are in distinct parts of K. Therefore $F^k[V(K)]$ contains a copy of P_n^k in which every pair of adjacent vertices are in distinct parts of K. By definition of K, such a copy is red. $\square$

By Claim 4.7, F^k contains a blue copy of P_n or a red copy of P_n^k under the edge-colouring χ'. Thus, we can split our proof into these two cases.

(*Case 1*) First suppose F^k contains a blue copy $(x_1, \ldots, x_n)$ of P_n. Then, for every $1 \leqslant i \leqslant n-1$, the bipartite graph $F'[V(B(x_i)), V(B(x_{i+1}))]$ contains a blue copy of $K_{s,s}$, with, say, vertex classes $X_i \subseteq V(B(x_i))$ and $Y_{i+1} \subseteq V(B(x_{i+1}))$. As $|X_i| = |Y_i| = s = 2k$ for all $2 \leqslant i \leqslant n-1$, we can find sets $X_i' \subseteq X_i$ and $Y_i' \subseteq Y_i$ such that $|X_i'| = |Y_i'| = k$ and $X_i' \cap Y_i' = \emptyset$ for all $2 \leqslant j \leqslant n-1$. Let $X_1' = X_1$ and $Y_n' = Y_n$.

We now show that the set $U := \bigcup_{i=1}^{n-1} X_i' \cup \bigcup_{i=2}^{n} Y_i'$ provides us with a blue copy of P_{2kn}^k in $F' \subseteq H_t^k$. Note first that $|U| = 2k + 2k(n-2) + 2k = 2kn$. Let $u_1, \ldots, u_{2kn}$ be an ordering of U such that, for each i, every vertex in X_i' comes before any vertex in Y_{i+1}' and after every vertex in Y_i'. By the definition of the sets X_i' and Y_i' and the construction of $F' \subseteq F_t^k \subseteq H_t^k$, each vertex u_j is adjacent in blue to $\{u_{j'} \in U : 1 \leqslant |j - j'| \leqslant k\}$. Thus, U contains a blue copy of P_{2nk}^k, as claimed.

(*Case 2*) Now suppose F^k contains a red copy P of P_n^k. That is, F^k contains a set of vertices $\{x_1, \ldots, x_n\}$ such that x_i is adjacent in red to all x_j with $1 \leqslant |j - i| \leqslant k$. We shall show that, for each $1 \leqslant i \leqslant n$, we can pick a vertex $y_i \in V(B(x_i))$ so that $y_1, \ldots, y_n$ define a red copy of P_n^k in $F' \subseteq F_t^k \subseteq H_t^k$. This can be done greedily, by picking the y_i one by one in order. We can also do this by applying the local lemma [37, p. 616]. We show the latter argument.

We have to show that it is possible to pick the y_i $(1 \leqslant i \leqslant n)$ in such a way that $\{y_i, y_j\}$ is a red edge in F' for every i and j with $1 \leqslant |i - j| \leqslant k$. Let us choose $y_i \in V(B(x_i))$ $(1 \leqslant i \leqslant n)$ uniformly and independently at random. Let $e = \{x_i, x_j\}$ be an edge in $P \subseteq F^k$. We know that e is red. Let A_e be the event that $\{y_i, y_j\}$ is a *blue* edge in F'. Since the edge e is red, we know that the bipartite graph $F'[V(B(x_i)), V(B(x_j))]$ contains no blue $K_{s,s}$. Theorem 4.6 then tells us that $\mathbb{P}[A_e] \leqslant 4t^{-1/s}$.

The events A_e are not independent, but we can define a dependency graph D for the collection of events A_e $(e \in E(P))$ by adding an edge between A_e and A_f if and only if $e \cap f \neq \emptyset$. Then $\Delta(D) \leqslant 4k$. Given that

$$4\Delta\mathbb{P}[A_e] \leqslant 64kt^{-1/s} = 1 \tag{4.10}$$

for all e, the local lemma tells us that $\mathbb{P}\big[\bigcap_{e\in E(P)} \bar{A}_e\big] > 0$, and hence a simultaneous choice of the y_i ($1 \leqslant i \leqslant n$) as required is possible. This completes the proof of Theorem 4.1. $\square$

5

Independent sets, matchings and occupancy fractions

5.1 Independent sets

Let G be a d-regular graph on n vertices. The independence polynomial of G is

$$P_G(\lambda) = \sum_{I \in \mathcal{I}} \lambda^{|I|}$$

where $\mathcal{I}$ is the set of all independent sets of G. Note that by convention we consider the empty independent set to be a member of $\mathcal{I}$. The *hard-core model* with *fugacity* λ on G is a random independent set I drawn according to the distribution

$$\mathbb{P}_\lambda[I] = \frac{\lambda^{|I|}}{P_G(\lambda)}$$

$P_G(\lambda)$ is also called the partition function of the hard-core model on G.

In the hard-core model, the quantity $\alpha(G) = \frac{\lambda}{n}\frac{P'_G(\lambda)}{P_G(\lambda)}$ is the *occupancy fraction*: the expected fraction of vertices of G belonging to the random independent set I. In particular,

$$\alpha(G) = \frac{1}{n}\mathbb{E}\big[|I|\big] = \frac{1}{n}\sum_{v \in G}\mathbb{P}[v \in I] = \frac{1}{n}\frac{\sum_{I \in \mathcal{I}}|I|\lambda^{|I|}}{P_G(\lambda)} = \frac{1}{n}\frac{\lambda P'_G(\lambda)}{P_G(\lambda)}\,. \tag{5.1}$$

Note that $\alpha(G)$ does not denote the independence number of G.

We write $K_{d,d}$ for the complete bipartite graph with d vertices in each part. If $2d$ divides n, let $H_{d,n}$ denote the d-regular, n-vertex graph that is the disjoint union of $n/(2d)$ copies of $K_{d,d}$. Kahn [56] showed that $H_{d,n}$ maximises the total number of independent sets over all d-regular, n-vertex bipartite graphs. Galvin and Tetali [47] gave a broad generalisation of Kahn's result to counting homomorphisms from a d-regular, bipartite G to any graph

H. The case of H formed of two connected vertices, one with a self-loop, is that of counting independent sets. Via a modification of H and a limiting argument, they proved that in fact $P_G(\lambda)^{1/n}$ is maximised for any λ over d-regular bipartite G by $K_{d,d}$. Zhao [87] then removed the bipartite restriction in Galvin and Tetali's result for independent sets by reducing the general case to the bipartite case, in particular proving that $H_{d,n}$ has the largest number of independent sets of any d-regular graph on n vertices.

Here we prove a strengthening of the above results for independent sets.

Theorem 5.1. *For all d-regular graphs G and all $\lambda > 0$, we have*

$$\alpha(G) \leqslant \alpha(K_{d,d}) = \frac{\lambda(1+\lambda)^{d-1}}{2(1+\lambda)^d - 1}\,.$$

The maximum is achieved only by disjoint unions of $K_{d,d}$'s. That is, the quantity $\frac{1}{n}\frac{P_G'(\lambda)}{P_G(\lambda)}$ is uniquely maximised by $H_{d,n}$.

Corollary 2. For any d-regular graph G and any $\lambda > 0$

$$P_G(\lambda)^{1/v(G)} \leqslant P_{K_{d,d}}(\lambda)^{1/v(K_{d,d})}\,.$$

In particular Theorem 5.1 states that the derivative of $\log P_G(\lambda)/n$ is maximised over d-regular graphs for all λ by $K_{d,d}$. We next show how integrating this proves Corollary 2.

Proof of Corollary 2. Suppose G is a d-regular graph and $\lambda > 0$. Noting that $\alpha(G) = \frac{1}{v(G)}\frac{\lambda P_G'(\lambda)}{P_G(\lambda)} = \frac{\lambda}{v(G)}\frac{d}{d\lambda}\log(P_G(\lambda))$ we see that

$$\begin{aligned}\frac{1}{v(G)}\log(P_G(\lambda)) &= \int_0^\lambda \frac{\alpha(G)}{x}dx \\ &\leqslant \int_0^\lambda \frac{\alpha(K_{d,d})}{x}dx \\ &= \frac{1}{v(K_{d,d})}\log(P_{K_{d,d}}(\lambda))\,.\end{aligned}$$

The inequality comes from Theorem 5.1 and Corollary 2 then follows by exponentiating both sides. □

In particular by choosing $\lambda = 1$ this shows that when $2d$ divides into n the n vertex d-regular graph with the most independent sets is $H_{d,n}$. (Note that $P_{K_{d,d}}(\lambda)^{1/v(K_{d,d})} = P_{H_{d,n}}(\lambda)^{1/v(H_{d,n})}$.)

In Section 5.8 we use Theorem 5.1 to give new upper bounds on the number of independent sets of a given size in d-regular graphs.

5.2 Matchings

The matching polynomial of a graph G is

$$M_G(\lambda) = \sum_{H \in \mathcal{M}} \lambda^{|H|}$$

where $\mathcal{M}$ is the set of all matchings of G (including the empty matching) and $|H|$ is the number of edges in the matching H. Just as in the hard-core model above we can define a probability distribution over matchings:

$$\mathbb{P}_\lambda[H] = \frac{\lambda^{|H|}}{M_G(\lambda)}\,.$$

This defines the monomer-dimer model from statistical physics [51]: dimers are edges of the random matching H, and monomers are the unmatched vertices.

For a d-regular graph G, the *edge occupancy fraction*, or the *dimer density*, is the expected fraction of the edges of G in such a random matching:

$$\alpha_M(G) = \frac{1}{e(G)}\mathbb{E}\big[|M|\big] = \frac{2}{dn}\sum_{e \in G}\mathbb{P}[e \in H] = \frac{2\lambda M_G'(\lambda)}{dnM_G(\lambda)}\,.$$

Our next result is an upper bound on the edge occupancy fraction of any d-regular graph:

Theorem 5.2. *For all d-regular graphs G and all $\lambda > 0$, we have*

$$\alpha_M(G) \leqslant \alpha_M(K_{d,d})\,.$$

That is, the quantity $\frac{2}{dn}\frac{M_G'(\lambda)}{M_G(\lambda)}$ is maximised by $H_{d,n}$.

Corollary 3. For any d-regular graph G and any $\lambda > 0$

$$M_G(\lambda)^{1/e(G)} \leqslant M_{K_{d,d}}(\lambda)^{1/e(K_{d,d})}\,.$$

The proof of Corollary 3 follows from Theorem 5.2 via integration in exactly the same way that Corollary 2 follows from Theorem 5.1. Similarly with $\lambda = 1$, this shows that $H_{d,n}$ has the greatest total number of matchings of any d-regular graph on n vertices.

In Section 5.8 we use Theorem 5.2 to give new upper bounds on the number of matchings of a given size in d-regular graphs. Using this we prove the asymptotic upper matching conjecture of Friedland, Krop, Lundow, and Markström [43].

5.3 Off-diagonal Ramsey numbers

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 [81] later improved the constant to 1, asymptotically as $d \to \infty$, 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.

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 [59] show that $R(3, k) = \Theta(k^2/\log k)$. Shearer's result gives the current best upper bound, showing that $R(3, k) \leqslant (1 + o(1))k^2/\log k$. Independent work of Bohman and Keevash [14] and Fiz Pontiveros, Griffiths, and Morris [42] shows that $R(3, k) \geqslant (1/4 + o(1))k^2/\log k$. 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 $R(3, k)$.

Theorem 5.3. *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*

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

Moreover, the constant '1' is best possible.

This result is weaker than Shearer's [81] 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 analysing a randomised greedy algorithm).

The proof of Theorem 5.3 is almost identical to the proof of Theorem 5.1 restricted to triangle-free graphs. In particular both methods reduce the problem to the same optimisation problem over a family of random variables. For triangle-free graphs Theorem 5.3 follows from maximising this optimisation problem whilst Theorem 5.1 comes minimising. After sharing a draft of [27] with colleagues, we discovered that James Shearer also knew the proof of the lower bound in Theorem 5.3 and presented a sketch of it at the SIAM Conference on Discrete Mathematics in 1998, but never published it [80].

To see that Theorem 5.3 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 5.3 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.

5.4 Related work

The results of Kahn [56], Galvin and Tetali [47], and Zhao [87] culminating in the fact that $P_G(\lambda)^{1/n}$ is maximised over d-regular graphs by $K_{d,d}$ are based on the entropy method, a powerful tool for the type of problems we address here. Apart from the results mentioned above, see [74] and [46] for surveys of the method. A direct application of the method requires the graph G to be bipartite. Zhao [88] showed that in some, but not all applications of the method, this restriction can be avoided by using a 'bipartite swapping trick'. An entropy-free proof of Galvin and Tetali's general theorem on counting homomorphisms was recently given by Lubetzky and Zhao [68]. Our method also does not use entropy, but in contrast to the other proofs it works directly for all d-regular graphs, without a reduction to the bipartite case. The method deals directly with the hard-core model instead of counting homomorphisms and seems to require more problem-specific information than the entropy method; a question for future work is to extend the method to more general homomorphisms.

The technique of writing the expected size of an independent set in two ways (as we do here) was used by Alon [6] in proving lower bounds on the size of an independent set in a graph in which all vertex neighbourhoods are r-colourable. The idea of bounding the occupancy fraction instead of the partition function comes in part from work of Will Perkins [71] in improving, at low densities, the bounds on matchings of a given size in Ilinca and Kahn [52] and independent sets of a given size in Carroll, Galvin, and Tetali [20]. The use of linear programming for counting graph homomorphisms appears in Kopparty and Rossman [64], where they use a combination of entropy and linear programming to compute a related quantity, the homomorphism domination exponent, in chordal and series-parallel graphs.

For matchings, Carroll, Galvin, and Tetali [20] used the entropy method to give an upper bound of $(1+d\lambda)^{1/2}$ on $M_G(\lambda)^{1/n}$. It was previously conjectured (eg. [44, 46]) that $K_{d,d}$ maximises $M_G(\lambda)^{1/n}$ over all d-regular graphs G. This is an implication of our

Theorem 5.2.

5.5 Independent sets in triangle-free graphs

In this section we introduce our method by giving a unified proof of the Theorems of Kahn and Shearer. Specifically we shall prove Theorem 5.1 under the assumption that G is triangle-free as well as Theorem 5.3.

In what follows I will denote the random independent set drawn according to the hard-core model with fugacity λ on an n-vertex graph G with maximum degree d. For Theorem 5.1 we require that G in fact be d-regular however most of the argument is the same and we will just note the distinction in the one step where it appears.

We say a vertex v is *occupied* if $v \in I$ and *uncovered* if none of its neighbours are in I: $N(v) \cap I = \emptyset$. In particular any vertex that is occupied is necessarily uncovered. Let p_v be the probability v is occupied and q_v be the probability v is uncovered. The idea of considering the q_v's appears in Kahn's paper [56].

We will show that for every $\lambda > 0$, over the set of triangle-free d-regular graphs G, $\alpha(G)$ is maximised by $K_{d,d}$. For any graph H the occupancy of H is the same as the occupancy fraction of multiple disjoint copies of H. In this way the occupancy fraction $\alpha(G)$ is maximised not just by $K_{d,d}$ but also by any number of copies of $K_{d,d,}$. We will also show that for graphs with maximum degree d and suitable choice of λ the value $\alpha(G)$ is at least $(1 + o_d(1))\frac{\log d}{d}n$.

Letting $\alpha = \alpha(G)$, we write

$$\begin{aligned}
\alpha &= \frac{1}{n}\sum_{v\in G} p_v \\
&= \frac{1}{n}\sum_{v\in G} \frac{\lambda}{1+\lambda} q_v && (5.2) \\
&= \frac{\lambda}{1+\lambda}\cdot\frac{1}{n}\sum_{v\in G}\sum_{j=0}^{d} \mathbb{P}[j \text{ neighbours of } v \text{ are uncovered}]\cdot(1+\lambda)^{-j} && (5.3) \\
&= \frac{\lambda}{1+\lambda}\cdot\mathbb{E}[(1+\lambda)^{-Y}]
\end{aligned}$$

where Y is the random variable that counts the number of uncovered neighbours of a uniformly chosen vertex from G, with respect to the random independent set I. Y is an integer valued random variable bounded between 0 and d. Equation (5.2) follows since v must be uncovered if it is to be occupied; conditioning on being uncovered, v is occupied with probability $\frac{\lambda}{1+\lambda}$. Equation (5.3) is similar: conditioned on $u_1, \ldots, u_j$ all uncovered, where the u_i's are neighbours of v, the probability that none are occupied is $(1+\lambda)^{-j}$. This is where we use triangle-freeness: we know there are no edges between the u_i's. So we are asking how likely a vertex v is to be in the independent set conditioning on the number of uncovered neighbours of v and averaging over all vertices in the graph.

The next step is to come up with another expression relating Y and $\alpha(G)$. This second way asks how many neighbours of v we expect to be in the independent set conditioned on the number of uncovered neighbours of v.

$$\mathbb{E}Y = \frac{1}{n}\sum_{v\in G}\sum_{u\sim v} q_u \leqslant d\cdot\frac{1+\lambda}{\lambda}\alpha$$

since each u appears in the double sum at most d times as G has maximum degree d. This gives the inequality

$$\mathbb{E}Y \leqslant d\cdot\mathbb{E}[(1+\lambda)^{-Y}]\,. \tag{5.4}$$

If G is d-regular then the two inequalities above hold with equality. From here we no longer consider the problem to be one of optimising over graphs but instead we are optimising over distributions Y that take values in $[0,d]$. For any such distribution Y we can now calculate α. Of course α has only been defined for graphs however we extend the definition to distributions Y by setting $\alpha = \frac{\lambda}{1+\lambda}\cdot\mathbb{E}[(1+\lambda)^{-Y}]$. We additionally have the constraint that we only consider distributions Y that satisfy equation (5.4).

We now let

$$\alpha^{max} = \frac{\lambda}{1+\lambda}\cdot\sup_{0\leqslant Y\leqslant d}\left\{\mathbb{E}[(1+\lambda)^{-Y}] : \mathbb{E}[(1+\lambda)^{-Y}] = \tfrac{1}{d}\mathbb{E}Y\right\}$$

and

$$\alpha^{min} = \frac{\lambda}{1+\lambda}\cdot\inf_{0\leqslant Y\leqslant d}\left\{\mathbb{E}[(1+\lambda)^{-Y}] : \mathbb{E}[(1+\lambda)^{-Y}] \geqslant \tfrac{1}{d}\mathbb{E}Y\right\}$$

where in both cases the sup is over all distributions of integer-valued random variables Y bounded between 0 and d.

For any λ and d there is a unique distribution Y supported only on 0 and d that satisfies the constraint $\mathbb{E}Y = d\cdot\mathbb{E}[(1+\lambda)^{-Y}]$. We claim that the supremum is achieved by this distribution. The claim follows from convexity, but we defer details to the proof of a more general statement in Section 5.6. Since the distribution Y associated to $H_{d,n}$ satisfies the constraint and is supported on 0 and d, it must maximise α. Since disjoint unions of $K_{d,d}$'s are the only graphs whose associated distribution is supported on 0 and d, they uniquely achieve the maximum.

Turning to the lower bound we use Jensen's inequality to see that

$$\begin{aligned}\alpha(G) &\geqslant \frac{\lambda}{1+\lambda}\inf_{0\leqslant Y\leqslant d}\max\left\{\mathbb{E}[(1+\lambda)^{-Y}], \tfrac{1}{d}\mathbb{E}Y\right\}\\ &\geqslant \frac{\lambda}{1+\lambda}\inf_{0\leqslant Y\leqslant d}\max\left\{(1+\lambda)^{-\mathbb{E}Y}, \tfrac{1}{d}\mathbb{E}Y\right\}\\ &\geqslant \frac{\lambda}{1+\lambda}\min_{x\in\mathbb{R}^+}\max\left\{(1+\lambda)^{-x}, \tfrac{x}{d}\right\}.\end{aligned}$$

Since x/d is increasing in x whilst $(1+\lambda)^{-x}$ is decreasing in x the min-max is attained when $\frac{x}{d} = (1+\lambda)^{-x}$. We can rearrange this to

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

from which we can see that

$$x = \frac{W(d\cdot\log(1+\lambda))}{\log(1+\lambda)}\,,$$

where W is the Lambert-W function; the inverse function of ze^z. We choose λ tending to zero with d such that $d\cdot\log(1+\lambda)$ tends to infinity (eg $\lambda = 1/\log(d)$) and use that fact that for $y \geqslant e$, $W(y) \geqslant \log(y) - \log(\log(y))$ which is equal to $(1+o(1))\log(y)$ if y tends to infinity. For λ tending to zero we also have $1/\log(1+\lambda) = (1+o(1))/\lambda$. We then see that

$$x \geqslant (1+o(1))\log(d\cdot\log(1+\lambda))/\lambda = (1+o(1))\log(d)/\lambda\,.$$

Therefore

$$\alpha(G) \geqslant \frac{\lambda}{1+\lambda}\frac{x}{d} \geqslant (1+o(1))\frac{\log d}{d}\,.$$

Since the occupancy fraction at this λ is at least $(1+o(1))\log d/d$ there must be an independent set of G of size at least

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

If G is a triangle-free graph with no independent set of size k it must have maximum degree less than k. Therefore it has an independent set of size $(1+o(1))\frac{\log k}{k}|G|$ and so this must also be less than k. Therefore $|G| \leqslant (1+o(1))\frac{k^2}{\log k}$ matching the asymptotic of Shearer's bound on off-diagonal Ramsey numbers.

In Section 5.6 we give the full proof of Theorem 5.1. We turn to matchings and Theorem 5.2 in Section 5.7 before giving new bounds on the number of independent sets and matchings of a given size in Section 5.8.

5.6 Independent sets in d-regular graphs

Here we show that Theorem 5.1 holds for all d-regular graphs. For a vertex $v \in G$ and an independent set I, we define the *free neighbourhood* of v to be the subgraph of G induced by the neighbours of v which are not adjacent to any vertex in $I \setminus N(v)$. We use the convention $v \notin N(v)$. The vertices in the free neighbourhood may be uncovered or covered, but if they are covered it must be from another vertex in the free neighbourhood. In a triangle-free graph the free neighbourhood is always a set (possibly empty) of isolated vertices. Note that if $v \in I$, then the free neighbourhood of v is necessarily empty.

Let C be the random free neighbourhood of v when we draw I according to the hard-core model and choose vertex v uniformly at random from G. For any graph F, let p_F be the probability that C is isomorphic to F. Also let $P_C = P_C(\lambda)$ be the independence polynomial of C at fugacity λ. Then we can write α in two ways:

$$\alpha = \frac{\lambda}{1+\lambda}\mathbb{E}\left[\frac{1}{P_C(\lambda)}\right] \tag{5.5}$$

and

$$\alpha = \frac{\lambda}{d}\mathbb{E}\left[\frac{P_C'(\lambda)}{P_C(\lambda)}\right] \tag{5.6}$$

where in both equations the expectations are over the random free neighbourhood C. Equation (5.5) follows since v itself is uncovered if and only if all vertices in its free neighbourhood are unoccupied. Given that the free neighbourhood is isomorphic to C, all vertices in the free neighbourhood are unoccupied with probability $\frac{1}{P_C(\lambda)}$. Equation (5.6) follows by counting the expected number of occupied neighbours of v and dividing by d: only vertices in the free neighbourhood can be occupied, and, given C, the expected number of occupied vertices in the free neighbourhood is $\frac{\lambda P_C'(\lambda)}{P_C(\lambda)}$.

Now let

$$\alpha^* = \frac{\lambda}{1+\lambda} \cdot \sup\left\{\mathbb{E}\left[\frac{1}{P_C(\lambda)}\right] : \frac{d}{1+\lambda} \cdot \mathbb{E}\left[\frac{1}{P_C(\lambda)}\right] = \mathbb{E}\left[\frac{P_C'(\lambda)}{P_C(\lambda)}\right]\right\} \tag{5.7}$$

where the sup is over all distributions of random free neighbourhoods C supported on graphs of at most d vertices. From (5.5) and (5.6), the distribution obtained from G satisfies the constraint above.

We claim that for any $\lambda > 0$, α^* is achieved uniquely by a distribution supported only on the empty graph and the graph consisting of d isolated vertices, $\overline{K_d}$. The theorem follows since a disjoint union of $K_{d,d}$'s is the only graph for which the free neighbourhood can only be the empty set or $\overline{K_d}$. To prove this claim we use the language of linear programming, see e.g. [17].

5.6.1 The linear program

Let p_C be the probability of a given free neighbourhood C, and let $\mathcal{C}_d$ be the set of all graphs on at most d vertices. Equation (5.7) defines a linear program with the decision variables $\{p_C\}_{C\in\mathcal{C}_d}$. We write the linear program in standard form as

$$
\begin{aligned}
\alpha^* = \max \frac{\lambda}{2(1+\lambda)} &\sum_{C\in\mathcal{C}_d} p_C(a_C+b_C) \text{ s.t.}\\
&\sum_{C\in\mathcal{C}_d} p_C = 1\\
&\sum_{C\in\mathcal{C}_d} p_C(a_C-b_C) = 0\\
&p_C \geqslant 0 \ \ \forall C\in\mathcal{C}_d
\end{aligned}
$$

where $a_C = \frac{1}{P_C(\lambda)}$ and $b_C = \frac{(1+\lambda)P_C'(\lambda)}{dP_C(\lambda)}$. We can calculate $a_\emptyset = 1$, $b_\emptyset = 0$, $a_{\overline{K}_d} = (1+\lambda)^{-d}$, $b_{\overline{K}_d} = 1$. The solution $p_\emptyset = \frac{1-(1+\lambda)^{-d}}{2-(1+\lambda)^{-d}}$ and $p_{\overline{K}_d} = \frac{1}{2-(1+\lambda)^{-d}}$ is the unique feasible solution supported only on $\emptyset$ and $\overline{K}_d$, and gives the objective value $\frac{\lambda(1+\lambda)^{d-1}}{2(1+\lambda)^d-1}$. Our claim is that this is the unique maximum.

The dual linear program is

$$
\begin{aligned}
\alpha^* = &\min \frac{\lambda}{2(1+\lambda)}\Lambda_1 \text{ s.t.}\\
&\Lambda_1 + \Lambda_2(a_C-b_C) \geqslant a_C+b_C \ \ \forall C\in\mathcal{C}_d
\end{aligned}
$$

where Λ_1, Λ_2 are the decision variables.

Guided by the candidate solution above we set $\Lambda_1 = \frac{2}{2-(1+\lambda)^{-d}}$, and $\Lambda_2 = 1-\Lambda_1$. With these values, the dual constraints corresponding to $C = \emptyset, \overline{K}_d$ hold with equality, and the objective value is $\frac{\lambda}{2(1+\lambda)}\Lambda_1 = \frac{\lambda(1+\lambda)^{d-1}}{2(1+\lambda)^d-1}$. To finish the proof we claim that Λ_1, Λ_2 are feasible for the dual program, which means showing that

$$\Lambda_1 + \Lambda_2(a_C-b_C) > a_C+b_C$$

for all $C\in\mathcal{C}_d\setminus\{\emptyset,\overline{K}_d\}$. Substituting our values of Λ_1, Λ_2, this inequality reduces to

$$\frac{\lambda P_C'(\lambda)}{P_C(\lambda)-1} < \frac{\lambda d(1+\lambda)^{d-1}}{(1+\lambda)^d-1}\,. \tag{5.8}$$

The LHS of (5.8) is the expected size of the random independent set from the hard-core model on C conditioned on it being non-empty. The RHS is the same quantity for $\overline{K}_d$.

Inequality (5.8) follows directly from the observation that, over all $C\in\mathcal{C}_d$, the graph $\overline{K}_d$ maximises the ratio of subsequent terms in the polynomial P_C. Let $a_i = \binom{d}{i}$ be the coefficient of λ^i in $P_{\overline{K}_d}$ and write $P_C = 1 + \sum_{i=1}^d b_i\lambda^i$. We have $(i+1)a_{i+1} = (d-i)a_i$ and $(i+1)b_{i+1} \leqslant (d-i)b_i$ by counting independent sets of size $i+1$.

To verify (5.8) we show that for each $1 \leqslant k \leqslant d$ the coefficient c_k of λ^k in the polynomial $(\lambda P'_{\overline{K}_d})(P_C - 1) - (\lambda P'_C)(P_{\overline{K}_d} - 1)$ is non-negative. We have

$$\begin{aligned} c_k &= \sum_{i=1}^{k-1} i a_i b_{k-i} + \sum_{i=1}^{k-1} i a_{k-i} b_i \\ &= \sum_{i=1}^{\lfloor k/2 \rfloor} (k-2i)(a_{k-i}b_i - a_i b_{k-i})\,. \end{aligned}$$

Observe that term-by-term the above sum giving c_k is non-negative by comparing the ratio of successive coefficients in $P_{\overline{K}_d}$ and P_C. Furthermore, if $P_C \neq P_{\overline{K}_d}$ then at least one c_k must be positive, which completes the claim.

To see the optimiser is unique note that there is a unique distribution supported on $\emptyset$ and $\overline{K}_d$ satisfying the primal constraints, and fixing $\Lambda_1 = \alpha(K_{d,d})$ in the dual gives a unique feasible value for Λ_2, since its coefficient $a_C - b_C$ takes different signs on $C = \emptyset, \overline{K}_d$. Therefore this is the unique optimal solution in the dual, and since all other dual constraints hold with strict inequality, any primal optimal solution must be supported on $\emptyset$ and $\overline{K}_d$. Disjoint unions of $K_{d,d}$'s are the only graphs whose distributions have this support.

5.7 Matchings in d-regular graphs

Recall that we use the notation $M_G(\lambda)$ for the matching polynomial of a graph G, and let H be a matching drawn from the monomer-dimer model at fugacity λ.

We refer to an edge as covered if an incident edge is in the random matching H. Let e be an edge of G chosen uniformly at random, with an arbitrary left/right orientation also chosen at random. In applying the method to matchings we introduce a subtle change of presentation. We now define the *free neighbourhood* C to be the subgraph of G containing all the incident edges to e that are not covered by edges outside of both e and its incident edges. When considering independent sets, the free neighbourhood was empty if the random vertex v was in the independent set. Here the presence or absence in the matching of e or an edge adjacent to e does not affect C. Given e and C, we use the term *externally uncovered neighbour* to refer to an edge of C incident to e.

The possible free neighbourhoods C are completely defined by three parameters: $L, R, K \in \{0, 1, \ldots, d-1\}$, counting the number of left and right neighbouring edges in C with an endpoint of degree 1, and the number of triangles formed by e and C. An example is pictured below.

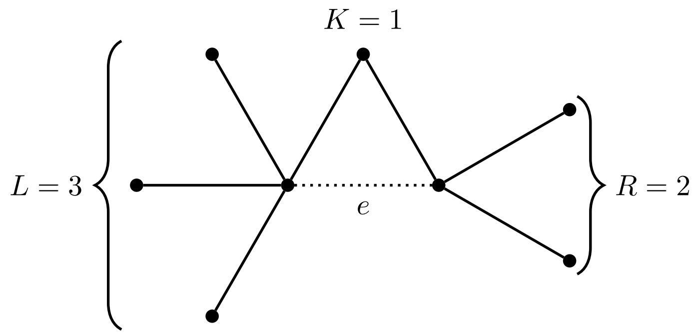

Let $q(i,j,k) = \mathbb{P}[L = i, R = j, K = k]$, and denote the matching polynomial for such a free neighbourhood by $M_{i,j,k}$, where we can compute

$$M_{i,j,k}(\lambda) = 1 + (i + j + 2k)\lambda + \left[k^2 + k(i + j - 1) + ij\right]\lambda^2 .$$

Conditioned on the event that the free neighbourhood of e is C, the random matching H restricted to e and its incident edges is distributed according to the monomer-dimer model on the graph C with the edge e added; the partition function of this model is $\lambda + M_C(\lambda)$, with the term λ corresponding to the event that $e \in H$.

We write α_M as the expected fraction of edges incident to e that are in the matching. In order for there to be any such edges e must be unoccupied. Conditioned on the free neighbourhood of e being C, the probability that e is unoccupied is

$$\frac{M_C(\lambda)}{\lambda + M_C(\lambda)} ,$$

and conditioned on C and on e being unoccupied the expected number of occupied neighbours of e is

$$\frac{\lambda M_C'(\lambda)}{M_C(\lambda)} .$$

As each edge in a d-regular graph is incident to exactly $2(d-1)$ other edges:

$$\begin{aligned}
\alpha_M &= \frac{2}{dn} \sum_{e} \sum_{f \sim e} \frac{1}{2(d-1)} \mathbb{P}[f \in H] \\
&= \mathbb{E}\left[\frac{\lambda M_C'(\lambda)}{2(d-1)(\lambda + M_C(\lambda))}\right] \\
&= \sum_{i,j,k} q(i,j,k) \frac{\lambda M_{i,j,k}'(\lambda)}{2(d-1)(\lambda + M_{i,j,k}(\lambda))} ,
\end{aligned}$$

where the expectation in the second line is over the random free neighbourhood C resulting from the two-part experiment described above. If we write the expected fraction of occupied neighbours of e in a configuration as $\overline{\alpha}_M(i,j,k) = \frac{1}{2(d-1)}\frac{\lambda M'_{i,j,k}}{\lambda+M_{i,j,k}}$, the above expression can be written $\alpha_M = \sum_{i,j,k} q(i,j,k)\overline{\alpha}_M(i,j,k)$.

5.7.1 The linear program for matchings

We now introduce additional constraints before optimising α_M over distributions of free neighbourhoods. We could write multiple expressions for α_M, equate them, and solve the maximization problem as we did for independent sets. Using three expressions for α_M we were able to prove Theorem 5.2 for the case $d = 3$, in which the optimal distribution is supported on only three values: $q(0,0,0)$, $q(1,1,0)$, $q(2,2,0)$. But in general we need at least $d-1$ constraints (in addition to the constraint that the $q(i,j,k)'s$ sum to one) as the distribution induced by $K_{d,d}$ is supported on d values.

Instead, we write, for all t, two expressions for the marginal probability that the number of uncovered neighbours on a randomly chosen side of a random edge is equal to t. We find the two expressions by choosing uniformly: a random edge e, a random side left or right, and f, a random neighbouring edge of e from the given side. We first calculate the probability that e has t uncovered neighbours on the side containing f, then we calculate the probability that f has t uncovered neighbours on the side containing e.

Given a free neighbourhood C with $L = i$, $R = j$, and $K = k$, e can have $0, 1, i+k-1$, or $i+k$ uncovered left neighbours; an edge f to the left of e can have $0, 1, i+k-2, i+k-1, i+k$, or $i+k+1$ uncovered right neighbours (depending on whether f itself is in the free neighbourhood C).

Let $\gamma^e_{i,j,k}(t) = \mathbb{P}[e \text{ has } t \text{ uncovered left neighbours } | L = i, R = j, K = k]$ and $\gamma^f_{i,j,k}(t) = \mathbb{P}[f \text{ has } t \text{ uncovered right neighbours } | L = i, R = j, K = k]$, where f is a uniformly chosen left neighbour of e.

Claim 5.4. *Let $\beta_t = 1 + t\lambda$. Then we have*

$$\begin{aligned}\gamma^e_{i,j,k}(t) = \frac{1}{\lambda+M_{i,j,k}}\Big(&\mathbf{1}_{t=0}\cdot\lambda + \mathbf{1}_{t=1}\cdot[i\lambda\beta_{j+k} + k\lambda\beta_{j+k-1}] \\ &+ \mathbf{1}_{t=i+k}\cdot\beta_j + \mathbf{1}_{t=i+k-1}\cdot k\lambda\Big)\end{aligned} \tag{5.9}$$

$$\begin{aligned}\gamma^f_{i,j,k}(t) = \frac{1}{(d-1)(\lambda+M_{i,j,k})}\Big(&\mathbf{1}_{t=0}\cdot[i\lambda\beta_{j+k} + k\lambda\beta_{j+k-1}] \\ &+ \mathbf{1}_{t=1}\cdot[(d-1)\lambda + (d-2)(i\lambda\beta_{j+k} + k\lambda\beta_{j+k-1})] \\ &+ \mathbf{1}_{t=i+k-2}\cdot[(i+k-1)k\lambda] + \mathbf{1}_{t=i+k-1}\cdot[(d-i-k)k\lambda + (i+k)j\lambda] \\ &+ \mathbf{1}_{t=i+k}\cdot[(d-1-i-k)j\lambda + (i+k)] + \mathbf{1}_{t=i+k+1}\cdot[d-1-i-k]\Big).\end{aligned} \tag{5.10}$$

Proof. To compute the functions $\gamma^e_{i,j,k}(t)$ we consider the following disjoint events: 1) no left edge and no right edge from a triangle is in the matching 2) e is in the matching 3) a left edge is in the matching 4) no left edge is in the matching, but a right edge from a triangle is in the matching. These events happen with probability $\frac{\beta_j}{\lambda+M_{i,j,k}}$, $\frac{\lambda}{\lambda+M_{i,j,k}}$, $\frac{i\lambda\beta_{j+k}+k\lambda\beta_{j+k-1}}{\lambda+M_{i,j,k}}$, and $\frac{k\lambda}{\lambda+M_{i,j,k}}$ respectively. Under these events the number of uncovered neighbours of e is $i+k, 0, 1$, and $i+k-1$ respectively. This gives (5.9).

To compute the functions $\gamma^f_{i,j,k}(t)$ we refine the above events to include the possible choices of f: f can be an edge outside the free neighbourhood with probability $(d-1-i-k)/(d-1)$; an edge in the free neighbourhood but not in a triangle with probability $i/(d-1)$; in the free neighbourhood and in a triangle with probability $k/(d-1)$. If a left edge is in the matching we choose it as f with probability $1/(d-1)$, and if a right edge in a triangle is in the matching we choose f adjacent to it with probability $1/(d-1)$. Computing the number of uncovered neighbours of f in each case gives (5.10). □

We now define a linear program with constraints imposing that the two different ways of writing the marginal probabilities are equal. The marginal probability constraint for $t = d-1$ is redundant and we omit it. To account for the equal chance that f is chosen from the left side of e and the right side of e, we average $\gamma^f_{i,j,k}(t)$ and $\gamma^f_{j,i,k}(t)$, and $\gamma^e_{i,j,k}(t)$ and $\gamma^e_{j,i,k}(t)$.

$$
\begin{aligned}
\alpha^*_M = \max \sum_{i,j,k} q(i,j,k)\overline{\alpha}_M(i,j,k) \quad &\text{subject to} \\
q(i,j,k) \geqslant 0 \;\forall\; i,j,k& \\
\sum_{i,j,k} q(i,j,k) = 1& \\
\sum_{i,j,k} q(i,j,k)\frac{1}{2}\left[\gamma^f_{i,j,k}(t)+\gamma^f_{j,i,k}(t)-\gamma^e_{i,j,k}(t)-\gamma^e_{j,i,k}(t)\right] = 0 \;\;\forall\, t&=0,\dots,d-2\,.
\end{aligned}
$$

Disjoint unions of copies of $K_{d,d}$ are the only graphs that induce a distribution $q(i,j,k)$ supported on triples with $i = j$ and $k = 0$. This gives us a candidate solution to the linear program.

The dual program is

$$
\begin{aligned}
\alpha^*_M = \;& \min \Lambda_p \;\; \text{subject to} \\
& \Lambda_p - \overline{\alpha}_M(i,j,k) + \sum_{t=0}^{d-2}\Lambda_t\frac{1}{2}\left[\gamma^f_{i,j,k}(t)+\gamma^f_{j,i,k}(t)-\gamma^e_{i,j,k}(t)-\gamma^e_{j,i,k}(t)\right] \geqslant 0 \;\forall\; i,j,k\,.
\end{aligned}
$$

To show that $K_{d,d}$ is optimal, we find values for the dual variables $\Lambda_0,\dots,\Lambda_{d-2}$ so that the dual constraints hold with $\Lambda_p = \alpha^M_{K_{d,d}}(\lambda)$. To find such values, we solve the system of equations generated by setting equality in the constraints corresponding to $i = j$ and $k = 0$ and solve for the variables Λ_t, $t = 0, \dots, d-2$.

With this choice of values for the dual variables, we start by simplifying the form of the dual constraints with a substitution coming from equality in the $(i, j, k) = (0, 0, 0)$ constraint. The $(0, 0, 0)$ dual constraint has the simple form

$$\Lambda_0 - \Lambda_1 = \alpha^M_{K_{d,d}} .$$

Moreover, observe that from the $\mathbf{1}_{t=0}$ and $\mathbf{1}_{t=1}$ terms in $\gamma^e_{i,j,k}(t)$ and $\gamma^f_{i,j,k}(t)$, every dual constraint contains the term

$$\left[\overline{\alpha}_M(i,j,k) - \frac{\lambda}{(\lambda + M_{i,j,k})}\right](\Lambda_0 - \Lambda_1) = \left[\overline{\alpha}_M(i,j,k) - \frac{\lambda}{(\lambda + M_{i,j,k})}\right]\alpha^M_{K_{d,d}} .$$

With this simplification, we multiply through by $2(d-1)(\lambda + M_{i,j,k})$ and expand $\overline{\alpha}_M(i, j, k)$ terms to obtain the following form of the dual constraints:

$$\begin{aligned}
\alpha^M_{K_{d,d}}&\big[\lambda M'_{i,j,k} + 2(d-1)M_{i,j,k}\big] - \lambda M'_{i,j,k} \\
&+ \Lambda_{i+k-2} \cdot (i+k-1)k\lambda \\
&+ \Lambda_{i+k-1} \cdot [(d-i-k)k\lambda + (i+k)j\lambda - (d-1)k\lambda] \\
&+ \Lambda_{i+k} \cdot [(d-1-i-k)j\lambda + i + k - (d-1)\beta_j] \\
&+ \Lambda_{i+k+1} \cdot (d-1-i-k) \\
&+ \Lambda_{j+k-2} \cdot (j+k-1)k\lambda \\
&+ \Lambda_{j+k-1} \cdot [(d-j-k)k\lambda + (j+k)i\lambda - (d-1)k\lambda] \\
&+ \Lambda_{j+k} \cdot [(d-1-j-k)i\lambda + j + k - (d-1)\beta_i] \\
&+ \Lambda_{j+k+1} \cdot (d-1-j-k) \geqslant 0 .
\end{aligned} \tag{5.11}$$

Recalling that we use β_t to denote $1 + t\lambda$ the $(i, i, 0)$ equality constraints now read

$$\alpha^M_{K_{d,d}}\beta_i\big(\beta_i + \tfrac{i\lambda}{d-1}\big) - \tfrac{i\lambda\beta_i}{d-1} + \Lambda_{i-1}\tfrac{i^2\lambda}{d-1} - \Lambda_i\tfrac{d-1-i+i^2\lambda}{d-1} + \Lambda_{i+1}\tfrac{d-1-i}{d-1} = 0 . \tag{5.12}$$

With this we can write Λ_{i+k+1} in terms of Λ_{i+k} and Λ_{i+k-1}, and similarly for Λ_{j+k+1}. Substituting this into (5.11) and dividing by λ we derive the simplified form of the dual constraints:

$$\begin{aligned}
\lambda&\big[(i-j)^2 + 2k\big](1 - d\alpha^M_{K_{d,d}}) \\
&+ \Lambda_{i+k-2}(i+k-1)k + \Lambda_{i+k-1}[k + (i+k)(j-i-2k)] \\
&+ \Lambda_{i+k}(i+k)(i+k-j) \\
&+ \Lambda_{j+k-2}(j+k-1)k + \Lambda_{j+k-1}[k + (j+k)(i-j-2k)] \\
&+ \Lambda_{j+k}(j+k)(j+k-i) \geqslant 0 .
\end{aligned} \tag{5.13}$$

Write $L(i,j,k)$ for the LHS of this inequality.

The marginal constraint for $t=d-1$ was omitted, but we nonetheless introduce $\Lambda_{d-1}:=0$ in order to simplify the presentation of the argument. The $(d-1,d-1,0)$ equality constraint gives Λ_{d-2} directly:

$$\Lambda_{d-2}=\frac{1}{(d-1)\lambda}\left[\lambda+(d-1)\lambda^2-\alpha^M_{K_{d,d}}\beta_{d-1}\beta_d\right].$$

With Λ_{d-1}, Λ_{d-2}, and the recurrence relation (5.12) the dual variables are fully determined. We do not give a closed-form expression for Λ_t as the values are used in an induction below. Using Λ_{d-1}, Λ_{d-2}, and (5.12) suffices for the proof.

We now reduce the problem of showing that the dual constraints (5.13) corresponding to triples (i,j,k) with $k>0$ or $i\neq j$ hold with strict inequality to showing that a particular function is increasing. We go on to prove this fact in Claims 5.5 and 5.6.

Putting $k=0$ into (5.13) gives:

$$\begin{aligned}\frac{L(i,j,0)}{(j-i)}&=\lambda(j-i)(1-d\alpha^M_{K_{d,d}})+i\Lambda_{i-1}-i\Lambda_i-j\Lambda_{j-1}+j\Lambda_j\\&=F_d(j)-F_d(i)\end{aligned}$$

where

$$F_d(t):=t\left[\lambda(1-d\alpha^M_{K_{d,d}})+\Lambda_t-\Lambda_{t-1}\right]. \tag{5.14}$$

From (5.13) we obtain

$$L(i-1,j-1,k+1)-L(i,j,k)=F_d(i+k)-F_d(i+k-1)+F_d(j+k)-F_d(j+k-1).$$

Therefore if $F_d(t)$ is strictly increasing, we have $L(i,j,0)>0$ for $i\neq j$, and $L(i-1,j-1,k+1)>L(i,j,k)>\cdots>L(i+k,j+k,0)\geqslant 0$.

We first find an explicit expression for $F_d(t)$. Recall that we write $M_{K_{t,t}}$ for the matching polynomial of the graph $K_{t,t}$.

Claim 5.5. *For all $d\geqslant 2$ and $1\leqslant t\leqslant d-1$,*

$$F_d(t)=\frac{t(d-1)}{M_{K_{d,d}}}\sum_{\ell=t-1}^{d-2}\frac{(d-1-t)!}{(\ell+1-t)!}\lambda^{d-\ell}M_{K_{\ell,\ell}}. \tag{5.15}$$

Proof. We will use the following two facts:

$$M_{K_{d,d}}-\beta_{2d-1}M_{K_{d-1,d-1}}+(d-1)^2\lambda^2M_{K_{d-2,d-2}}=0 \tag{5.16}$$

$$\alpha^M_{K_{d,d}}=\frac{\lambda M_{K_{d-1,d-1}}}{M_{K_{d,d}}}. \tag{5.17}$$

The first is a Laguerre polynomial identity, verifiable by hand; the second is a short calculation. The equality dual constraint (5.12) implies:

$$(d-1-t)F_d(t+1) = (t+1)[t\lambda F_d(t) + (d-1)\lambda - (d-1)\alpha^M_{K_{d,d}}\beta_{d+t}]\,. \tag{5.18}$$

We first show that the right hand side of (5.15) satisfies the above recurrence relation. Using (5.17) this amounts to showing that the following expression is equal to zero for all $d \geqslant 2$ and $1 \leqslant t \leqslant d-1$:

$$\Phi_d(t) := (d-1-t)!\left(\sum_{\ell=t}^{d-2}\frac{\lambda^{d-\ell}M_{K_{\ell,\ell}}}{(\ell-t)!} - t^2\sum_{\ell=t-1}^{d-2}\frac{\lambda^{d+1-\ell}M_{K_{\ell,\ell}}}{(\ell+1-t)!}\right) - \lambda(M_{K_{d,d}} - \beta_{d+t}M_{K_{d-1,d-1}})\,.$$

We proceed by induction on d. Note that when $d=2$, $\Phi_2(1)$ is easily verified to be zero. Note that

$$\Phi_{d+1}(t) = \lambda\Big((d-t)\Phi_d(t) - M_{K_{d+1,d+1}} + \beta_{2d+1}M_{K_{d,d}} - d^2\lambda^2 M_{K_{d-1,d-1}}\Big)\,.$$

By the induction hypothesis and (5.16) the result follows. To complete the proof of the claim it suffices to show that (5.15) holds for $t = d-1$. Recalling that

$$\begin{aligned}\Lambda_{d-1} &= 0\\ \Lambda_{d-2} &= \frac{1}{d-1} + \lambda - \frac{\alpha^M_{K_{d,d}}}{(d-1)\lambda}\beta_d\beta_{d-1}\,,\end{aligned}$$

substituting into (5.14), and using (5.16) and (5.17) we have

$$\begin{aligned}F_d(d-1) &= (d-1)\left[\lambda(1-d\alpha^M_{K_{d,d}}) - \frac{1}{d-1} - \lambda + \frac{\alpha^M_{K_{d,d}}}{(d-1)\lambda}\beta_d\beta_{d-1}\right]\\ &= \frac{\alpha^M_{K_{d,d}}}{\lambda}\beta_{2d-1} - 1\\ &= \frac{1}{M_{K_{d,d}}}\left[\beta_{2d-1}M_{K_{d-1,d-1}} - M_{K_{d,d}}\right]\\ &= \frac{(d-1)^2\lambda^2 M_{K_{d-2,d-2}}}{M_{K_{d,d}}}\,,\end{aligned}$$

verifying (5.15) for $t = d-1$. □

Using Claim 5.5 we prove the following.

Claim 5.6. *$F_d(t)$ is strictly increasing as a function of t.*

Proof. To prove that $F_d(t)$ is increasing, we show that

$$\begin{aligned}R_d(t) &:= \frac{M_{K_{d,d}}}{(d-1)}\cdot\frac{F_d(t+1)-F_d(t)}{(d-2-t)!}\\ &= (t+1)\sum_{\ell=t}^{d-2}\frac{\lambda^{d-\ell}}{(\ell-t)!}M_{K_{\ell,\ell}} - t(d-1-t)\sum_{\ell=t-1}^{d-2}\frac{\lambda^{d-\ell}}{(\ell+1-t)!}M_{K_{\ell,\ell}}\end{aligned}$$

is positive for each t with $1 \leqslant t \leqslant d-2$. We do this by fixing t and performing induction on d from $t+2$ upwards. A useful inequality will be $M_{K_{t,t}} > t\lambda M_{K_{t-1,t-1}}$ which comes from only counting matchings of $K_{t,t}$ that use a specific vertex. Iterating this inequality we obtain

$$M_{K_{t,t}} > \frac{t!}{\ell!}\lambda^{t-\ell} M_{K_{\ell,\ell}} \quad \text{for } 0 \leqslant \ell \leqslant t-1\,. \tag{5.19}$$

For the base case of our induction, $d = t+2$, we have

$$R_d(d-2) = \lambda^2\left[M_{K_{d-2,d-2}} - (d-2)\lambda M_{K_{d-3,d-3}}\right],$$

which by (5.19) is positive.

For the inductive step we have

$$R_{d+1}(t) = \lambda\left[R_d(t) + \frac{\lambda}{(d-1-t)!}M_{K_{d-1,d-1}} - \sum_{\ell=t-1}^{d-2}\frac{t\lambda^{d-\ell}}{(\ell-t+1)!}M_{K_{\ell,\ell}}\right],$$

and so it is sufficient to show

$$\sum_{\ell=t-1}^{d-2}\frac{t\lambda^{d-\ell}}{(\ell+1-t)!}M_{K_{\ell,\ell}} < \frac{\lambda}{(d-1-t)!}M_{K_{d-1,d-1}}\,. \tag{5.20}$$

We use the inequality (5.19) in each term of the sum to see that the LHS of (5.20) is less than

$$\sum_{\ell=t-1}^{d-2}\frac{t\ell!\lambda}{(\ell+1-t)!(d-1)!}M_{K_{d-1,d-1}},$$

and so

$$\begin{aligned}
\sum_{\ell=t-1}^{d-2}\frac{t\lambda^{d-\ell}}{(\ell+1-t)!}M_{K_{\ell,\ell}} &< \sum_{\ell=t-1}^{d-2}\frac{t\ell!\lambda}{(\ell+1-t)!(d-1)!}M_{K_{d-1,d-1}} \\
&= \frac{\lambda M_{K_{d-1,d-1}}}{(d-1-t)!}\cdot\sum_{\ell=t-1}^{d-2}\frac{t\ell!(d-1-t)!}{(\ell+1-t)!(d-1)!} \\
&= \frac{\lambda M_{K_{d-1,d-1}}}{(d-1-t)!}\cdot\binom{d-1}{t}^{-1}\cdot\sum_{\ell=t-1}^{d-2}\binom{\ell}{t-1} \\
&= \frac{\lambda M_{K_{d-1,d-1}}}{(d-1-t)!}\,.
\end{aligned}$$

Therefore (5.20) holds as required. □

This completes the proof of dual feasibility and shows our candidate solution to the primal program is optimal. The uniqueness of the solution follows from two facts. First, strict inequality in the dual constraints outside of the $(i,i,0)$ constraints implies, by complementary slackness, that the support of any optimal solution in the primal is contained in the set of $(i,i,0)$ configurations. Second, the distribution induced by $K_{d,d}$ is the unique distribution satisfying the constraints with such a support. This follows from the fact that Λ_i is uniquely determined by (5.12) where we have set the $(i,i,0)$ dual constraints to hold with equality, which in turn shows that the relevant $d \times d$ submatrix of the constraint matrix is full rank. This proves Theorem 5.2.

5.8 Independent sets and matchings of a given size

Let $i_k(G)$ be the number of independent sets of size k in a graph G, and $m_k(G)$ the number of matchings of size k. Kahn [56] conjectured that $i_k(G)$ is maximised over d-regular, n-vertex graphs by $H_{d,n}$ for all k (when $2d$ divides n), and Friedland, Krop, and Markström [44] conjectured the same for $m_k(G)$. Previous bounds towards these conjectures were given in [20, 52, 71]. Here we adapt the method of Carroll, Galvin, and Tetali (and use the above result on the matching polynomial) to give bounds for both problems that are tight up to a multiplicative factor of $2\sqrt{n}$, for all d and all k. Previous bounds had been off by an exponential factor in n.

Theorem 5.7. *For all d-regular graphs G on n vertices (where $2d$ divides n),*

$$i_k(G) \leqslant 2\sqrt{n} \cdot i_k(H_{d,n})$$

and

$$m_k(G) \leqslant 2\sqrt{n} \cdot m_k(H_{d,n})\,.$$

The general idea to get from the result that $H_{d,n}$ maximises $P_G(\lambda)$ and $M_G(\lambda)$ among d-regular graphs on n vertices is to use the fact that for any $k \in \{0,\ldots,n/2\}$ there is a value of λ such that independent sets of size k are the most common size when running the hard-core model on $H_{d,n}$. There is also a (potentially different) value of λ such that the most common size of matching in the monomer-dimer model on $H_{d,n}$ is k. By comparing the partition functions of an arbitrary graph G with $H_{d,n}$ we see that the number of independent sets of size k in G is at most $\frac{n}{2}+1$ times that of $H_{d,n}$. Similarly for matchings. The $\frac{n}{2}+1$ factor comes from the fact that the partition functions have this many non-zero coefficients, as there are this many possible sizes or independent set or matching. For Theorem 5.7 we prove something slightly stronger; that for every k there is a λ such that in the hard-core model on $H_{d,n}$ independent sets of size k are selected at least a $\frac{1}{2\sqrt{n}}$ fraction of the time. The same holds for matchings.

Lemma 5.8. *For all $1 \leqslant k \leqslant n/2$, there exists a λ so that*

$$\frac{i_k(H_{d,n})\lambda^k}{P_{H_{d,n}}(\lambda)} = \mathbb{P}_{H_{d,n}}[|I| = k] > \frac{1}{2\sqrt{n}}$$

and a λ so that

$$\frac{m_k(H_{d,n})\lambda^k}{M_{H_{d,n}}(\lambda)} = \mathbb{P}_{H_{d,n}}[|H| = k] > \frac{1}{2\sqrt{n}}\,.$$

Proof. The distribution of the size of a random independent set I drawn from the hard-core model on $H_{d,n}$ is log-concave; that is,

$$\mathbb{P}_{H_{d,n}}[|I| = j]^2 > \mathbb{P}_{H_{d,n}}[|I| = j+1] \cdot \mathbb{P}_{H_{d,n}}[|I| = j-1]$$

for all $1 < j < n/2$. This follows from two facts: the size distribution of the hard-core model on $K_{d,d}$ is log-concave, and the convolution of two log-concave distributions is again log-concave. The first fact is simply the calculation

$$\binom{d}{j}^2 > \binom{d}{j-1}\binom{d}{j+1}\,.$$

Now choose λ so that $\mathbb{P}_{H_{d,n}}[|I| = k] = \mathbb{P}_{H_{d,n}}[|I| = k+1]$. Log-concavity then implies that $\mathbb{P}_{H_{d,n}}[|I| = k]$ is maximal. Some explicit computations for the variance for a single $K_{d,d}$ give that the variance of $|I|$ is at most $n/8$; then via Chebyshev's inequality, with probability at least $2/3$ the size of I is one of at most $\frac{4}{3}\sqrt{n}$ values, and thus the largest probability of a single size is greater than $\frac{1}{2\sqrt{n}}$.

The proof for $m_k(H_{d,n})$ is the same: the variance of the size of a random matching is also at most $n/8$ (see, e.g. [55]), and log-concavity of the size distribution on $K_{d,d}$ is verified via the inequality

$$\binom{d}{j}^4 j!^2 > \binom{d}{j-1}^2 (j-1)!\binom{d}{j+1}^2 (j+1)!$$

□

Proof of Theorem 5.7. Assume for sake of contradiction that $m_k(G) > 2\sqrt{n} \cdot m_k(H_{d,n})$. Choose λ according to Lemma 5.8. We have:

$$M_G(\lambda) \geqslant m_k(G)\lambda^k > 2\sqrt{n} \cdot m_k(H_{d,n})\lambda^k > M_{H_{d,n}}(\lambda)\,,$$

but this contradicts Theorem 5.2. The case of independent sets is identical. □

The above proof is essentially the same as the proofs in Carroll, Galvin, and Tetali [20] with the small observation that λ can be chosen so that k is the most likely size of a matching (or independent set) drawn from $H_{d,n}$. The factor $2\sqrt{n}$ in both cases can surely be improved by using some regularity of the independent set and matchings sequence of a general d-regular graph.

As a consequence, we prove the asymptotic upper matching conjecture of Friedland, Krop, Lundow, and Markström [43]. Fix d and consider an infinite sequence of d-regular graphs $\mathcal{G}_d = G_1, G_2, \ldots$ where each G_n has n vertices if such a graph exists. For any $\varrho \in [0, 1/2]$, the ϱ-monomer entropy is

$$h_{\mathcal{G}_d}(\varrho) = \sup_{\{k_n\}} \limsup_{n\to\infty} \frac{\log m_{k_n}(G_n)}{n},$$

where the supremum is taken over all integer sequences $\{k_n\}$ with $k_n/n \to \varrho$. Let $h_d(\varrho) = \lim_{n\to\infty} \frac{\log m_{\lfloor \varrho n\rfloor}(H_{d,n})}{n}$, where the limit is take over the sequences of integers divisible by $2d$. Then the conjecture states that for all $\mathcal{G}_d$ and all $\varrho \in [0, 1/2]$, $h_{\mathcal{G}_d}(\varrho) \leqslant h_d(\varrho)$.

To prove this, first assume $\varrho > 0$ since for $\varrho = 0$ the result is trivially true. Assume for the sake of contradiction that $\limsup \frac{\log m_{k_n}(G_n)}{n} > h_d(\varrho) + \varepsilon$ for some $\varepsilon > 0$. Take N_0 large enough that for all $n_1 \geqslant N_0$, divisible by $2d$, $\frac{\log m_{\lfloor \varrho n_1\rfloor}(H_{d,n_1})}{n_1} < h_d(\varrho) + \varepsilon/2$. Now take some $n \geqslant N_0$ with $\frac{\log m_{k_n}(G_n)}{n} > h_d(\varrho) + \varepsilon$, and let $n_1 = 2d \cdot \lceil n/(2d)\rceil$. Choose λ so that $m_{\lfloor \varrho n_1\rfloor}(H_{d,n_1}) > \frac{1}{2\sqrt{n_1}} M_{H_{d,n_1}}(\lambda)$. Note that since $\varrho > 0$, such λ is bounded away from 0 as $n_1 \to \infty$. Then we have

$$\begin{aligned}\frac{\log M_{G_n}(\lambda)}{n} \geqslant \frac{\log m_{k_n}(G_n)\lambda^{k_n}}{n} &> \frac{k_n}{n}\log\lambda + h_d(\varrho) + \varepsilon \\ &= \varrho\log\lambda + h_d(\varrho) + \varepsilon + o(1) \text{ as } n\to\infty\end{aligned}$$

and

$$\begin{aligned}\frac{\log M_{K_{d,d}}(\lambda)}{2d} = \frac{\log M_{H_{d,n_1}}(\lambda)}{n_1} &< \frac{\log\left(2\sqrt{n_1}\cdot m_{\lfloor \varrho n_1\rfloor}(H_{d,n_1})\lambda^{\lfloor \varrho n_1\rfloor}\right)}{n_1} \\ &< \frac{\log(2\sqrt{n_1})}{n_1} + \frac{\lfloor \varrho n_1\rfloor}{n_1}\log\lambda + h_d(\varrho) + \varepsilon/2 \\ &= \varrho\log\lambda + h_d(\varrho) + \varepsilon/2 + o(1),\end{aligned}$$

but this contradicts Theorem 5.2. With the same proof, the analogous statement for independent set entropy holds.

Bibliography

[1] M. Ajtai, J. Komlós, and E. Szemerédi. A note on Ramsey numbers. *Journal of Combinatorial Theory, Series A*, 29(3):354–360, 1980.

[2] P. Allen, J. Böttcher, J. Ehrenmüller, and A. Taraz. The Bandwidth Theorem in sparse graphs. submitted (arXiv:1612.00661), 2016.

[3] P. Allen, J. Böttcher, S. Griffiths, Y. Kohayakawa, and R. Morris. The chromatic thresholds of graphs. *Adv. Math.*, 235:261–295, 2013.

[4] P. Allen, J. Böttcher, H. Hàn, Y. Kohayakawa, and Y. Person. Blow-up lemmas for sparse graphs. submitted (arXiv:1612.00622), 2016.

[5] P. Allen, J. Böttcher, B. Roberts, and Y. Kohayakawa. Triangle-free subgraphs of random graphs. *Combinatorics, Probability and Computing. To appear*, 2017.

[6] N. Alon. Independence numbers of locally sparse graphs and a ramsey type problem. *Random Structures and Algorithms*, 9(3):271–278, 1996.

[7] N. Alon and F. R. K. Chung. Explicit construction of linear sized tolerant networks. *Discrete Math.*, 72(1-3):15–19, 1988.

[8] B. Andrásfai, P. Erdős, and V. T. Sós. On the connection between chromatic number, maximal clique and minimal degree of a graph. *Discrete Math.*, 8:205–218, 1974.

[9] L. Babai, M. Simonovits, and J. Spencer. Extremal subgraphs of random graphs. *Journal of Graph Theory*, 14(5):599–622, 1990.

[10] J. Balogh, R. Morris, and W. Samotij. Independent sets in hypergraphs. *Journal of the American Mathematical Society*, 28(3):669–709, 2015.

[11] J. Beck. On size Ramsey number of paths, trees, and circuits. I. *J. Graph Theory*, 7(1):115–129, 1983.

[12] J. Beck. On size Ramsey number of paths, trees and circuits. II. In *Mathematics of Ramsey theory*, volume 5 of *Algorithms Combin.*, pages 34–45. Springer, Berlin, 1990.

[13] I. Ben-Eliezer, M. Krivelevich, and B. Sudakov. The size Ramsey number of a directed path. *J. Combin. Theory Ser. B*, 102(3):743–755, 2012.

[14] T. Bohman and P. Keevash. Dynamic concentration of the triangle-free process. *arXiv preprint arXiv:1302.5963*, 2013.

[15] B. Bollobás. Determination of extremal graphs by using weights. *Wiss. Z. Techn. Hochsch. Ilmenau*, 13:419–421, 1967.

[16] B. Bollobás. On a conjecture of Erdős, Hajnal and Moon. *American Mathematical Monthly*, pages 178–179, 1967.

[17] S. Boyd and L. Vandenberghe. *Convex optimization.* Cambridge university press, 2004.

[18] S. Brandt and S. Thomassé. Dense triangle-free graphs are four-colorable: A solution to the Erdős-Simonovits problem. *preprint*, 2005.

[19] G. Brightwell, K. Panagiotou, and A. Steger. Extremal subgraphs of random graphs. *Random Structures Algorithms*, 41(2):147–178, 2012.

[20] T. Carroll, D. Galvin, and P. Tetali. Matchings and independent sets of a fixed size in regular graphs. *Journal of Combinatorial Theory, Series A*, 116(7):1219–1227, 2009.

[21] D. Clemens, M. Jenssen, Y. Kohayakawa, N. Morrison, G.O. Mota, D Reding, and B. Roberts. The size-Ramsey number of powers of paths. submitted (arXiv:1707.04297), 2017.

[22] D. Conlon. Question suggested for the ATI–HIMR Focused Research Workshop: *Large-scale structures in random graphs*, Alan Turing Institute, December 2016.

[23] D. Conlon. Combinatorial theorems relative to a random set. *Proceedings of the 2014 ICM*, 4:303–328.

[24] D. Conlon, J. Fox, and B. Sudakov. Recent developments in graph Ramsey theory. In *Surveys in combinatorics 2015*, volume 424 of *London Math. Soc. Lecture Note Ser.*, pages 49–118. Cambridge Univ. Press, Cambridge, 2015.

[25] D. Conlon, W. T. Gowers, W. Samotij, and M. Schacht. On the KŁR conjecture in random graphs. *Israel Journal of Mathematics*, 203(1):535–580, 2014.

[26] E. Davies, M. Jenssen, W. Perkins, and B. Roberts. Independent sets, matchings, and occupancy fractions. *Journal of the London Mathematical Society (to appear in print)*, 2017.

[27] E. Davies, M. Jenssen, W. Perkins, and B. Roberts. On the average size of independent sets in triangle-free graphs. *Proceedings of the American Mathematical Society (to appear in print)*, 2017.

[28] D. Dellamonica, Jr. The size-Ramsey number of trees. *Random Structures Algorithms*, 40(1):49–73, 2012.

[29] B. DeMarco and J. Kahn. Mantel's theorem for random graphs. *Random Structures & Algorithms*, 47(1):59–72, 2015.

[30] B. Descartes. A three-colour problem. *Eureka*, 9:21, April 1947. Solution March 1948.

[31] A. Dudek and P. Prałat. An alternative proof of the linearity of the size-Ramsey number of paths. *Combin. Probab. Comput.*, 24(3):551–555, 2015.

[32] P. Erdős. Some remarks on the theory of graphs. *Bulletin of the American Mathematical Society*, 53(4):292–294, 1947.

[33] P. Erdős. Graph theory and probability. *Canad. J. Math*, 11(11):34–38, 1959.

[34] P. Erdős. On the combinatorial problems which I would most like to see solved. *Combinatorica*, 1(1):25–42, 1981.

[35] P. Erdős, R. J. Faudree, C. C. Rousseau, and R. H. Schelp. The size Ramsey number. *Period. Math. Hungar.*, 9(1-2):145–161, 1978.

[36] P. Erdős, A. Hajnal, and J. Moon. A problem in graph theory. *American Mathematical Monthly*, pages 1107–1110, 1964.

[37] P. Erdős and L. Lovász. Problems and results on 3-chromatic hypergraphs and some related questions. In *Infinite and finite sets (Colloq., Keszthely, 1973; dedicated to P. Erdős on his 60th birthday), Vol. II*, pages 609–627. Colloq. Math. Soc. János Bolyai, Vol. 10. North-Holland, Amsterdam, 1975.

[38] P. Erdős and M. Simonovits. On a valence problem in extremal graph theory. *Discrete Mathematics*, 5(4):323–334, 1973.

[39] P. Erdős and A. H. Stone. On the structure of linear graphs. *Bull. Amer. Math. Soc*, 52(1087-1091):3, 1946.

[40] P. Erdős and G. Szekeres. A combinatorial problem in geometry. *Compositio Mathematica*, 2:463–470, 1935.

[41] M. Ferrara, M. Jacobson, F. Pfender, and P. Wenger. Graph saturation in multipartite graphs. *J. Comb.*, 19:1–19, 2016.

[42] G. Fiz Pontiveros, S. Griffiths, and R. Morris. The triangle-free process and $R(3, k)$. *Memoirs of the American Mathematical Society (to appear in print)*, 2017.

[43] S. Friedland, E. Krop, P. H. Lundow, and K. Markström. On the validations of the asymptotic matching conjectures. *Journal of Statistical Physics*, 133(3):513–533, 2008.

[44] S. Friedland, E. Krop, and K. Markström. On the number of matchings in regular graphs. *the electronic journal of combinatorics*, 15(1):R110, 2008.

[45] J. Friedman and N. Pippenger. Expanding graphs contain all small trees. *Combinatorica*, 7(1):71–76, 1987.

[46] D. Galvin. Three tutorial lectures on entropy and counting. *arXiv preprint arXiv:1406.7872*, 2014.

[47] D. Galvin and P. Tetali. On weighted graph homomorphisms. *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, 63:97–104, 2004.

[48] R. Häggkvist. Odd cycles of specified length in nonbipartite graphs. In *Graph theory (Cambridge, 1981)*, volume 62 of *North-Holland Math. Stud.*, pages 89–99. North-Holland, Amsterdam-New York, 1982.

[49] P. E. Haxell and Y. Kohayakawa. The size-Ramsey number of trees. *Israel J. Math.*, 89(1-3):261–274, 1995.

[50] P. E. Haxell, Y. Kohayakawa, and T. Łuczak. The induced size-Ramsey number of cycles. *Combin. Probab. Comput.*, 4(3):217–239, 1995.

[51] O. J. Heilmann and E. H. Lieb. Theory of monomer-dimer systems. *Communications in Mathematical Physics*, 25(3):190–232, 1972.

[52] L. Ilinca and J. Kahn. Asymptotics of the upper matching conjecture. *Journal of Combinatorial Theory, Series A*, 120(5):976–983, 2013.

[53] S. Janson, T. Łuczak, and A. Ruciński. Random graphs. 2000. *Wiley–Intersci. Ser. Discrete Math. Optim*, 2000.

[54] G. Jin. Triangle-free four-chromatic graphs. *Discrete Mathematics*, 145(1-3):151–170, 1995.

[55] J. Kahn. A normal law for matchings. *Combinatorica*, 20(3):339–391, 2000.

[56] J. Kahn. An entropy approach to the hard-core model on bipartite graphs. *Combinatorics, Probability and Computing*, 10(03):219–237, 2001.

[57] L. Kászonyi and Zs. Tuza. Saturated graphs with minimal number of edges. *Journal of graph theory*, 10(2):203–210, 1986.

[58] X. Ke. The size Ramsey number of trees with bounded degree. *Random Structures Algorithms*, 4(1):85–97, 1993.

[59] J. H. Kim. The Ramsey number $R(3,t)$ has order of magnitude $t^2/\log t$. *Random Structures & Algorithms*, 7(3):173–207, 1995.

[60] Y. Kohayakawa, T. Łuczak, and V. Rödl. On K^4-free subgraphs of random graphs. *Combinatorica*, 17(2):173–213, 1997.

[61] Y. Kohayakawa, T. Retter, and V. Rödl. The size-Ramsey number of short subdivisions of bounded degree graphs. Submitted, 2016.

[62] Y. Kohayakawa, V. Rödl, M. Schacht, and E. Szemerédi. Sparse partition universal graphs for graphs of bounded degree. *Adv. Math.*, 226(6):5041–5065, 2011.

[63] Yoshiharu Kohayakawa. Szemerédi's regularity lemma for sparse graphs. *Foundations of computational mathematics*, pages 216–230, 1997.

[64] S. Kopparty and B. Rossman. The homomorphism domination exponent. *European Journal of Combinatorics*, 32(7):1097–1114, 2011.

[65] T. Kővári, V. T. Sós, and P. Turán. On a problem of K. Zarankiewicz. *Colloquium Math.*, 3:50–57, 1954.

[66] S. Letzter. Path Ramsey number for random graphs. *Combin. Probab. Comput.*, 25(4):612–622, 2016.

[67] J. Liouville. Sur des classes très étendues de quantités dont la valeur nést ni algébrique, nimême réductible des irrationnelles algébriques. *J. Math. Pures et Appl.*, (18):883–885 and 910–911, 1844.

[68] E Lubetzky and Y. Zhao. On replica symmetry of large deviations in random graphs. *Random Structures & Algorithms*, 47(1):109–146, 2015.

[69] J. Lyle. On the chromatic number of H-free graphs of large minimum degree. *Graphs and Combinatorics*, pages 1–14, 2010.

[70] W. Mantel. Problem 28. *Wiskundige Opgaven*, 10:60–61, 1907. (solution by H. Gouwentak, W. Mantel, J. Texeira de Mattes, F. Schuh and W.A. Wythoff).

[71] W. Perkins. Birthday inequalities, repulsion, and hard spheres. *Proceedings of the American Mathematical Society*, 144(6):2635–2649, 2016.

[72] A. Pokrovskiy. Partitioning edge-coloured complete graphs into monochromatic cycles and paths. *J. Combin. Theory Ser. B*, 106:70–97, 2014.

[73] A. Pokrovskiy. Calculating Ramsey numbers by partitioning colored graphs. *Journal of Graph Theory*, 84(4):477–500, 2017.

[74] J. Radhakrishnan. Entropy and counting. *Computational Mathematics, Modelling and Algorithms*, page 146, 2003.

[75] F. P. Ramsey. On a problem of formal logic. *Proc. London Math. Soc.*, S2-30(1):264, 1930.

[76] D. Reimer. The Ramsey size number of dipaths. *Discrete Math.*, 257(1):173–175, 2002.

[77] B. Roberts. Partite saturation problems. *Journal of Graph Theory*, 85(2):429–445, 2017.

[78] J. K. Roberts. Some properties of adsorbed films of oxygen on tungsten. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 152(876):464–477, 1935.

[79] D. Saxton and A. Thomason. Hypergraph containers. *Invent. Math.*, 201(3):925–992, 2015.

[80] J. B. Shearer. Personal communication.

[81] J. B. Shearer. A note on the independence number of triangle-free graphs. *Discrete Mathematics*, 46(1):83–87, 1983.

[82] E. Szemerédi. Regular partitions of graphs. In *Problèmes combinatoires et théorie des graphes (Colloq. Internat. CNRS, Univ. Orsay, Orsay, 1976)*, pages 399–401. CNRS, Paris, 1978.

[83] C. Thomassen. On the chromatic number of triangle-free graphs of large minimum degree. *Combinatorica*, 22(4):591–596, 2002.

[84] P. Túran. On an extremal problem in graph theory. *Matematikai s Fizikai Lapok (in Hungarian)*, 1941.

[85] W. Wessel. Uber eine klasse paarer graphen, i: Beweis einer vermutung von Erdös, Hajnal, und Moon. *Wiss. Z. Tech. Hochsch. Ilmenau*, 12:253–256, 1966.

[86] W. Wessel. Uber eine klasse paarer graphen. ii. bestimmung der minimalgraphen. *Wiss. Z. Techn. Hochsch. Ilmenau*, 13:423–426, 1967.

[87] Y. Zhao. The number of independent sets in a regular graph. *Combinatorics, Probability and Computing*, 19(02):315–320, 2010.

[88] Y. Zhao. The bipartite swapping trick on graph homomorphisms. *SIAM Journal on Discrete Mathematics*, 25(2):660–680, 2011.