Partition regularity and other combinatorial problems

Ben Barber

Trinity College

A dissertation submitted to the University of Cambridge
for the degree of Doctor of Philosophy
July 2013

Declaration

This dissertation is the result of my own work and includes nothing which is the outcome of work done in collaboration except where specifically indicated in the text. No part of this dissertation has been submitted for any other qualification.

Abstract

A system of linear equations is called *partition regular* if, whenever the natural numbers are finitely coloured, there is a monochromatic solution. The finite partition regular systems were characterised by Rado in 1933 in terms of a simple property of their matrix of coefficients. The infinite case is much harder, and to date only a few example of infinite partition regular systems are known. The main contribution of the first part of this thesis is a new family of infinite partition regular systems. We go on to use these systems to settle a number of long-standing open problems, including showing that there are systems that are partition regular over the rationals but not the natural numbers.

The Erdős–Ko–Rado theorem is a basic result in extremal combinatorics. Borg considered a variant of this result where we count the sizes of intersecting families in a different way. In the second part of this thesis we answer a question of Borg concerning an extension of his theorem.

A graph is *quasirandom* if it resembles a random graph in a particular sense. In the third part of this thesis we examine random walks on quasirandom graphs. We show that the subgraph of edges traversed by a random walk on a quasirandom graph is very likely to itself be quasirandom. This answers a question of Böttcher, Hladký, Piguet and Taraz.

The maximum sized independent sets in the discrete hypercube are precisely the set of all odd subsets of $[n]$ and all even subsets of $[n]$. In the fourth part of this thesis we answer a question of Ramras that asked for the maximum size of a *balanced* independent set, containing equal numbers of sets of odd and even size.

Acknowledgements

I would like to take this opportunity to thank several people who have had an effect on this thesis or my work over the last three years.

Neil Hindman and Imre Leader provided the seed from which my interest in partition regularity grew. It was a pleasure to explore quasirandomness with Eoin Long. The workshop on probabilistic methods in graph theory organised by Daniela Kühn, Richard Mycroft and Deryk Osthus at the University of Birmingham in March 2012 was doubly useful: as well as featuring Jan Hladký's problem session presentation, it was the backdrop against which the first partition regularity proofs were worked out. Josh Erde contributed three years of excellent miniseminars and wide-ranging mathematical discussion.

Above all I am hugely grateful for the enthusiastic supervision of Imre Leader, whose ready supply of problems, advice, criticism and encouragement made everything else possible.

Contents

Notation

We collect here some common notation that will be used throughout this thesis.

$$\begin{aligned}
\mathbb{N} &= \{1, 2, \ldots\} \\
[m, n] &= \{m, \ldots, n\} \\
[n] &= \{1, \ldots, n\} \\
X^{(r)} &= \{A \subseteq X : |A| = r\}
\end{aligned}$$

A graph G has vertex set $V(G)$, edge set $E(G)$ and $e(G)$ edges. The degree of a vertex v is $d(v)$, and its neighbourhood is $N(v)$. The subgraph induced by a set X of vertices is $G[X]$.

Chapter 1

Introduction

In this introductory chapter we give an overview of the problems considered in this thesis and the results we shall obtain. Each of the remaining chapters gives a self-contained presentation of one of the four bodies of work described below.

In Chapter 2 we consider some problems in partition regularity. This is a part of Ramsey theory concerned with finding sets with arithmetic structure. Many classic results in Ramsey theory have such a form.

Schur's theorem (1916)**.** *Whenever the natural numbers are finitely coloured there is a monochromatic solution to the equation* $x + y = z$*.*

Van der Waerden's theorem (1921)**.** *Whenever the natural numbers are finitely coloured there is a colour class that contains arbitrarily long arithmetic progressions.*

Van der Waerden's theorem can easily be strengthened to give that we can always find a monochromatic arithmetic progression whose common difference is also the same colour. In this form the conclusion of the theorem is equivalent to the assertion that, for any n, the system of equations

$$\begin{aligned} x_1 &= a + d \\ x_2 &= a + 2d \\ &\vdots \\ x_n &= a + nd \end{aligned}$$

has a monochromatic solution.

If a system of linear equations has monochromatic solutions for every finite colouring of the natural numbers we call it *partition regular*. Rado characterised the finite partition regular systems in terms of a simple property of their matrix of coefficients.

Given an $m \times n$ matrix A, let $c^{(1)}, \ldots, c^{(n)}$ be its columns. We say that A has the *columns property* if there is a partition $[n] = I_1 \cup I_2 \cup \cdots \cup I_s$ of the columns of A such that $\sum_{i \in I_1} c^{(i)} = 0$, and, for each t,

$$\sum_{i \in I_t} c^{(i)} \in \langle c^{(i)} : i \in I_1 \cup \cdots \cup I_{t-1} \rangle,$$

where $\langle \cdot \rangle$ denotes (rational) linear span.

Rado's theorem (1933). *A finite system of equations $Ax = 0$, with integer coefficients, is partition regular if and only if A has the columns property.*

Neither direction is obvious. Rado's theorem is very powerful, as it replaces an infinite condition—is there a monochromatic solution for *every* finite colouring of the natural numbers?—by one which is easy to check in finite time.

In the infinite case, even examples of partition regular systems are hard to come by. For instance, it is hopeless to ask for a monochromatic infinite arithmetic progression: because there are only countably many of these, we can choose two elements from each and ensure that they get different colours. It was not until much later that Hindman found the first non-trivial example of an infinite partition regular system.

Hindman's theorem (1974). *Whenever the natural numbers are finitely coloured, there exist* $x_1, x_2, \ldots$ *such that all finite sums* $\sum_{i \in I} x_i$, *where* $I \neq \emptyset$, *are the same colour.*

Hindman's theorem has been generalised in two main directions. The first is due independently to Milliken and Taylor.

Milliken–Taylor theorem (1975). *Whenever the natural numbers are finitely coloured, there exist* $x_1 < x_2 < \cdots$ *such that all finite sums* $\sum_{i \in I} x_i + \sum_{j \in J} 2x_j$, *where* $I, J \neq \emptyset$ *and* $\max I < \min J$, *are the same colour.*

This is the $(1, 2)$ version of Milliken and Taylor's result; there are corresponding versions for any finite string of natural numbers.

Deuber–Hindman theorem (1987). *For any sequence* $E_1, E_2, \ldots$ *of finite partition regular systems of equations, whenever the natural numbers are finitely coloured there is a sequence of corresponding solutions sets* $S_1, S_2, \ldots$ *such that all finite sums of the form* $\sum_{i \in I} x_i$, *where* $I \neq \emptyset$ *and* $x_i \in S_i$ *for all* $i \in I$, *are the same colour.*

Here a *solution set* S_i is the set of values taken by the variables in some solution to E_i.

The Milliken–Taylor theorem can be proved by mimicking the proof of Ramsey's theorem, replacing appeals to the pigeonhole principle with appeals to Hindman's theorem. The Deuber–Hindman theorem can be proved by looking inside the proof of Hindman's theorem and finding more structure: it can be viewed as Hindman's theorem crossed with Rado's theorem.

The preceding three theorems are almost everything that is known about partition regularity of infinite systems. In particular, in every known example every variable appears with only a finite set of coefficients.

Unbounded coefficients. *Is there a partition regular system in which some variable appears with an unbounded set of coefficients?*

Any example could not come from a 'local' modification of Hindman's theorem, so would necessarily be of a new and different kind.

Partition regularity can also be defined over the rationals.[1] In the finite case, the two notions of partition regularity coincide, but in the infinite case the corresponding question remained open.

Rationals versus naturals. *Is there a system of equations which is partition regular over the rationals but not over the natural numbers?*

For every non-partition regular example that had been considered, there turned out to be some way to extend the 'bad' colouring of the natural numbers to the rationals.

These questions were asked by Hindman, Leader and Strauss in their "Survey of open problems in partition regularity", *Combinatorics, Probability and Computing* 12(5–6) 571–583 (2003).

The main contribution of Chapter 2 is a new family of infinite partition regular systems which provide a positive answer to these two questions. These are the first examples of partition regular systems that do not arise in the same way as Hindman's theorem.

Theorem. *For any sequence (a_n) of rational coefficients, the system of equations*

$$\begin{gathered} x_{11} + a_1 y = z_{11} \\ x_{21} + x_{22} + a_2 y = z_{21} + z_{22} \\ \vdots \\ x_{n1} + \cdots + x_{nn} + a_n y = z_{n1} + \cdots + z_{nn} \\ \vdots \end{gathered}$$

is partition regular over the rationals. Moreover, if the a_n are integers, then the system of equations is partition regular over the natural numbers.

Taking $a_n = 2^n$ provides an example for the first question; taking $a_n = 1/n$ provides an example for the second.

[1] We give formal definitions of this and other concepts in Chapter 2.

Rather than seeking monochromatic solutions to systems of linear equations (*kernels* of a matrix A), we could instead ask for monochromatic *images*. A variant of the above theorem provides a ready source of examples.

Theorem. *For any sequence (a_n) of integer coefficients, the system of equations*

$$\begin{aligned} x_{11} + a_1 y &= z_1 \\ x_{21} + x_{22} + a_2 y &= z_2 \\ &\vdots \\ x_{n1} + \cdots + x_{nn} + a_n y &= z_n \\ &\vdots \end{aligned}$$

is partition regular. It follows that the system of expressions

$$\begin{gathered} x_{11} + a_1 y \\ x_{21} + x_{22} + a_2 y \\ \vdots \\ x_{n1} + \cdots + x_{nn} + a_n y \\ \vdots \end{gathered}$$

is partition regular.

We use this system to examine the connection between image and kernel partition regularity. Given an image partition regular matrix A, there is a particular process that produces a kernel partition regular matrix B. (The matrix B encodes the linear dependences between the rows of A.)

Image versus kernel. *If we apply this process to any matrix A, and the resulting matrix B is kernel partition regular, must A be image partition regular?*

One question that arises when considering this problem is the following.

Congruence conditions. *Let A be a matrix that is image partition regular and (d_i) be a sequence of integers. Is it true that, whenever the natural* *numbers are finitely coloured, there is a monochromatic image Ax such that the variables x_i satisfy $x_i \equiv 0 \pmod{d_i}$?*

These two questions were also asked by Hindman, Leader and Strauss in their survey. Using our new image partition regular system we give a negative answer to these questions.

De and Hindman considered partition regularity over different subsets of $\mathbb{R}$. They sought a matrix that was image partition regular over $\mathbb{N}$ but not image partition regular over $\mathbb{R}$ near zero. Our new partition regular system provides an example. By adapting our proof we are also able to show that the matrix suggested by De and Hindman is image partition regular over $\mathbb{N}$, providing a second example.

The results of Sections 3.1–3.4 are joint work with with Neil Hindman and Imre Leader, and have been published as "Partition regularity in the rationals", *Journal of Combinatorial Theory, Series A* 120 1590–1599 (2013). The results of Section 4 are joint work with Imre Leader and are due to appear as "Partition regularity with congruence conditions", *Journal of Combinatorics* (2013). The results in Section 3.5 are my own work.

In Chapter 3 we consider a problem from extremal combinatorics. One of the fundamental results in this area is the Erdős–Ko–Rado theorem. A family of sets $\mathcal{A} \subseteq [n]^{(r)}$ is *intersecting* if $A \cap B \neq \emptyset$ for all $A, B \in \mathcal{A}$.

Erdős–Ko–Rado theorem. *If $\mathcal{A} \subseteq [n]^{(r)}$ is intersecting and $n \geq 2r$, then $|\mathcal{A}| \leq \binom{n-1}{r-1}$.*

This bound is attained by the *star* consisting of all r-sets containing a fixed element of $[n]$, and for $n > 2r$ this extremal family is unique.

In his 2011 paper "Maximum hitting of a set by compressed intersecting families", Borg considered how the Erdős–Ko–Rado theorem is affected if we only count sets containing one of a fixed set X of elements when calculating the size of our family. That is, let

$$\mathcal{A}(X) = \{A \in \mathcal{A} : A \cap X \neq \emptyset\}.$$

Given $X \subseteq [n]$, which families $\mathcal{A}$ maximise $|\mathcal{A}(X)|$? For any $x \in X$, $|\mathcal{A}(X)|$

is maximised by the star at x, so Borg restricted attention to families that are *left-compressed*, ruling out stars centred anywhere but 1. He then asked for which X it was true that the star remained the optimal family in this new setting. He obtained a complete answer to this question for $|X| \geq r$, and a partial answer when $|X| < r$. In Chapter 3 we answer Borg's question for $|X| < r$, provided n is sufficiently large.

The results in this chapter follow from a structural result on left-compressed intersecting families. The statement and proof of this result were suggested by examining small cases using a computer. Annotated program listings and output are included at the end of the chapter.

The results in Chapter 3 are due to appear as "Maximum hitting for n sufficiently large", *Graphs and Combinatorics* (2013).

In Chapter 4 we study random walks on quasirandom graphs: graphs that look random, provided we don't look too closely.

There are many natural definitions of what it might mean for a graph G on n vertices to look like a random graph of density p.

- Every not too small induced subgraph of G has density about p.
- The density of the bipartite subgraph of G induced by any two sets of vertices has density about p.
- Let H be any graph on k vertices. Then the number of labelled induced subgraphs of G isomorphic to H is about $p^{e(H)}(1-p)^{\binom{k}{2}-e(H)}$.

Surprisingly, all of these properties, and several more, turn out to be equivalent in a particular sense. We call graphs satisfying these properties *quasirandom.*

Given a quasirandom graph G we can obtain a new quasirandom graph G_p by retaining each edge of G independently with probability p. (The random graph $G_{n,p}$ can be viewed as the result of applying this process to the complete graph K_n.) What happens if we choose a random subgraph in some other way?

One way to choose a random subgraph is to take a random walk on G. That is, let $W = (W_0, W_1, \ldots, W_{\alpha n^2})$ be a random walk on G of length

αn^2, and let $G_{walk}(\alpha)$ be the subgraph of G consisting of the edges W_iW_{i+1} traversed by W. Is $G_{walk}(\alpha)$ also likely to be quasirandom?

It is easy to see that it is if G is complete. In this case each W_i is (very nearly) a vertex selected uniformly at random from $V(G)$. Then $W_0W_1, W_2W_3, W_4W_5, \dots$ and $W_1W_2, W_3W_4, W_5W_6, \dots$ are (very nearly) two independent sequences of uniform random edges from G, so their union is also a random set of edges.

Morally, we expect the same argument to work for any quasirandom graph G: the edges of G are evenly distributed, so we expect the random walk to again consist of a sequence of edges selected almost uniformly at random. If we try to make this precise we encounter a problem: quasirandom graphs only resemble random graphs provided we do not look too closely, so can have small configurations of vertices that can trap the random walk for long periods of time. These difficulties can, however, be overcome, and in this chapter we give a proof that $G_{walk}(\alpha)$ is very likely to be quasirandom roughly following the outline above.

This problem was suggested by Böttcher, Hladký, Piguet and Taraz whose real interest was in packing lots of small trees into a complete graph. A random walk can be viewed as a random homomorphism of a path; in the last part of the chapter we extend our result on random walks to random homomorphisms of trees.

The results in Chapter 4 are joint work with Eoin Long.

In Chapter 5 we answer a question in extremal graph theory. The n-dimensional hypercube is the graph with vertex set the power set of $[n]$ and edges between sets which differ at exactly one element. The maximum size of an independent set in this graph is 2^{n-1}, and this is attained only by the set of all odd-sized subsets of $[n]$ and the set of all even-sized subsets of $[n]$.

Ramras asked how large an independent set we can find that contains equal number of sets of odd and even size. He conjectured that the best we can do is to take even-sized sets starting from the bottom of the cube and odd-sized sets starting from the top of the cube with an appropriate gap left in the middle layers. We give a short proof of this result using a (known) vertex isoperimetric inequality on the cube.

The results in Chapter 5 were published as “A note on balanced independent sets in the cube”, *Australasian Journal of Combinatorics* 52 (2012), 205–207.

Chapter 2

Partition regularity in the integers and rationals

1 Introduction

An *r-colouring* of a set X is a partition of X into r parts $A_1, \ldots, A_r$. Each A_i is a *colour class*, and a subset S of X is *monochromatic* if $S \subseteq A_i$ for some i. Equivalently, an r-colouring is a function c from X to a set of r *colours*, and a subset S of X is monochromatic if c is constant on S.

The basic question of Ramsey theory is "What sort of monochromatic sets are we guaranteed to find?"

Theorem 1 (Ramsey [Ram30])**.** *Whenever $\mathbb{N}^{(m)}$ is finitely coloured, there is an infinite subset X of $\mathbb{N}$ such that $X^{(m)}$ is monochromatic.*

That is, whenever we finitely colour the edges of a complete infinite (hyper-) graph, there is an infinite monochromatic clique. We say that the collection of infinite cliques is *partition regular*.

Ramsey's theorem applies equally well if we replace $\mathbb{N}$ by any infinite set, but there are also Ramsey-type theorems that use the additive structure of $\mathbb{N}$.

Theorem 2 (Van der Waerden [VdW21])**.** *Let k be a natural number. Whenever $\mathbb{N}$ is finitely coloured, there is monochromatic arithmetic progression of length k.*

We say that "the set of arithmetic progressions is partition regular."

Theorem 3 (Schur [Sch16]). *Whenever $\mathbb{N}$ is finitely coloured we can find x and y in $\mathbb{N}$ such that x, y and $x+y$ are the same colour. Equivalently, whenever $\mathbb{N}$ is finitely coloured there is a monochromatic solution to the equation $x+y=z$.*

We say that a system of linear equations $Ax = 0$ with integer coefficients is *partition regular* (*over* $\mathbb{N}$) if, whenever $\mathbb{N}$ is finitely coloured, the equations have a monochromatic solution; that is, there is a vector x with entries in $\mathbb{N}$ such that $Ax = 0$ and each entry of x is the same colour. We say that a system of linear expressions Ax with integer coefficients is *partition regular* (*over* $\mathbb{N}$) if, whenever $\mathbb{N}$ is finitely coloured, there is a vector x with entries in $\mathbb{N}$ such that each entry of Ax is in $\mathbb{N}$ and has the same colour. Partition regularity over $\mathbb{Z}$ or $\mathbb{Q}$ is defined similarly, with $\mathbb{N}$ replaced by $\mathbb{Z} \setminus \{0\}$ or $\mathbb{Q} \setminus \{0\}$ throughout.

Partition regularity is traditionally ascribed to the matrix A rather than the system of equations $Ax = 0$ or the system of expressions Ax. In the first case we call A *kernel partition regular*; in the second case we call A *image partition regular*. Then Schur's theorem can be interpreted as saying either that the matrix

$$\begin{pmatrix} 1 & 1 & -1 \end{pmatrix}$$

is kernel partition regular, or that the matrix

$$\begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 1 & 1 \end{pmatrix}$$

is image partition regular, and the length 3 version of Van der Waerden's theorem is the statement that the matrix

$$\begin{pmatrix} 1 & 0 \\ 1 & 1 \\ 1 & 2 \end{pmatrix}$$

is image partition regular.

The matrix notation can be useful when discussing general systems, but for concrete examples it is usually easier to describe the equations or expressions directly. We switch between the two perspectives freely.

Later we shall chiefly be concerned with infinite systems. We remark here that, although the corresponding matrices will have infinitely many rows and columns, there will only be finitely many non-zero entries in each row, so there will be no problem with the implicit infinite sums in the matrix multiplications.

1.1 Finite partition regular systems

Rado [Rad33] gave a simple characterisation of the finite kernel partition regular systems. Let A be an $m \times n$ matrix and let $c^{(1)}, \ldots, c^{(n)}$ be the columns of A. We say that A has the *columns property* if there is a partition $[n] = I_1 \cup I_2 \cup \cdots \cup I_s$ of the columns of A such that $\sum_{i \in I_1} c^{(i)} = 0$, and, for each t,

$$\sum_{i \in I_t} c^{(i)} \in \langle c^{(j)} : j \in I_1 \cup \cdots \cup I_{t-1} \rangle,$$

where $\langle \cdot \rangle$ denotes (rational) linear span.

Theorem 4 ([Rad33])**.** *A finite matrix A with integer coefficients is kernel partition regular over $\mathbb{N}$ if and only if it has the columns property.*

Rado's theorem has the following immediate consequence. If finite matrices A and B are kernel partition regular then they each have the columns property, so the diagonal sum

$$\begin{pmatrix} A & 0 \\ 0 & B \end{pmatrix}$$

also has the columns property, hence by Rado's theorem is kernel partition regular. So whenever $\mathbb{N}$ is finitely coloured the systems of equations $Ax = 0$ and $By = 0$ can be solved inside the same colour class; we say that A and B are *consistent*.

What about kernel partition regularity over different spaces? If A is kernel partition regular over $\mathbb{N}$ then it is certainly kernel partition regular over $\mathbb{Z}$: since a colouring of $\mathbb{Z} \setminus \{0\}$ induces a colouring of $\mathbb{N}$, we can find a monochromatic solution to $Ax = 0$ inside the positive part of $\mathbb{Z}$. Similarly, if A is kernel partition regular over $\mathbb{Z}$ then it must also be kernel partition regular over $\mathbb{Q}$. What about the converses?

It is easy to see that if A is kernel partition regular over $\mathbb{Z}$ then it must also be kernel partition regular over $\mathbb{N}$. Indeed, given a colouring of $\mathbb{N}$ we can reflect the colouring to $\mathbb{Z}_{<0}$ using a new set of colours; that is, if $c : \mathbb{N} \to [r]$ is an r-colouring of $\mathbb{N}$ then extend it to a $2r$-colouring of $\mathbb{Z}_{<0}$ by setting $c(-x) = -c(x)$. Since A is kernel partition regular over $\mathbb{Z}$, $Ax = 0$ has a monochromatic solution under c. By construction, either all of the variables in this solution are positive, or all of the variables are negative. If all of the variables are negative then we can multiply them by -1 to get a monochromatic solution in which all of the variables are positive. Hence there is always a monochromatic solution in $\mathbb{N}$.

It is less obvious, but still true, that if A is kernel partition regular over $\mathbb{Q}$ then it is kernel partition regular over $\mathbb{Z}$. Let A be kernel partition regular over $\mathbb{Q}$ and let c be an r-colouring of $\mathbb{Z}$. We will show that A is kernel partition regular over $\mathbb{Z}$ using a compactness argument. We think of $\mathbb{Q}$ as a nested union of copies of $\mathbb{Z}$: we have $\mathbb{Q} = \bigcup_{n=1}^{\infty} \frac{1}{n!} \cdot \mathbb{Z}$ where the dot denotes pointwise multiplication and $\frac{1}{n!} \cdot \mathbb{Z}$ is the nth *level* of $\mathbb{Q}$.

If q is in the nth level of $\mathbb{Q}$, write $c_n(q) = c(p)$ where $q = \frac{p}{n!}$. (If q is not in the nth level of $\mathbb{Q}$ we leave $c_n(q)$ undefined.) We obtain a colouring c' of $\mathbb{Q} \setminus \{0\}$ as follows. Enumerate $\mathbb{Q} \setminus \{0\}$ as $q_1, q_2, \ldots$. There is some colour k_1 such that $c_n(q_1) = k_1$ for infinitely many n. Let $c'(q_1) = k_1$, and let $S_1 = \{n : c_n(q_1) = k_1\}$. Now there is some colour k_2 such that $c_n(q_2) = k_2$ for infinitely many $n \in S_1$. Let $c'(q_2) = k_2$, and let $S_2 = \{n \in S_1 : c_n(q_2) = k_2\}$. Then there is some colour k_3 such that $c_n(q_3) = k_3$ for infinitely many $n \in S_2$. Let $c'(q_3) = k_3$, and let $S_3 = \{n \in S_2 : c_n(q_3) = k_3\}$. Continuing in this way we get an r-colouring c' of $\mathbb{Q} \setminus \{0\}$.

Since A is kernel partition regular over $\mathbb{Q}$ there is a monochromatic solution to $Ax = 0$ in $\mathbb{Q} \setminus \{0\}$ coloured by c'. Let $(q_{i_1}, \ldots, q_{i_m})$ be such a solution and let $t = \max_{1\le j\le m} i_j$. Let n be any element of S_t and write $q_{i_j} = \frac{p_{i_j}}{n!}$. Then the p_{i_j} are all given the same colour by c, and $(p_{i_1}, \dots, p_{i_m})$ is also a solution to $Ax = 0$. So we have found a monochromatic solution in $\mathbb{Z} \setminus \{0\}$ coloured by c, and A is kernel partition regular over $\mathbb{Z}$.

The finite image partition regular systems were characterised by Hindman and Leader [HL93]. We shall return to the question of how the two notions of partition regularity are related later.

1.2 Infinite partition regular systems

In the finite case partition regularity is very well understood. In the infinite case only a few examples of partition regular systems are known.

Some trivial examples of partition regular systems can be obtained from Ramsey's theorem.

Corollary 5. *Let $a_1, a_2, \dots, a_m$ be natural numbers. Whenever $\mathbb{N}$ is finitely coloured, there exist $x_1 < x_2 < \cdots$ such that all sums $a_1x_{i_1} + \cdots + a_mx_{i_m}$, where $i_1 < \cdots < i_m$, are the same colour.*

Proof. Given a colouring c of $\mathbb{N}$, obtain a colouring c' of $\mathbb{N}^{(m)}$ by setting $c'(x_1 < \cdots < x_m) = c(a_1x_1 + \cdots + a_mx_m)$ and apply Ramsey's theorem. □

The first non-trivial example of an infinite partition regular system is due to Hindman.

Theorem 6 ([Hin74])**.** *Whenever $\mathbb{N}$ is finitely coloured, there exist $x_1, x_2, \dots$ such that all finite sums $\sum_{i\in I} x_i$, where $I \neq \emptyset$, are the same colour.*

Hindman's theorem has been generalised in two directions. The first is due independently to Milliken and Taylor.

Theorem 7 ([Mil75] and [Tay76])**.** *Whenever $\mathbb{N}$ is finitely coloured, there exist $x_1 < x_2 < \cdots$ such that all finite sums $\sum_{i\in I} x_i + \sum_{j\in J} 2x_j$, where $I, J \neq \emptyset$ and $\max I < \min J$, are the same colour.*

This is the $(1, 2)$ version of Milliken and Taylor's result; there are corresponding versions for any finite string of natural numbers.

The Milliken–Taylor theorem is proved via the stronger result that, whenever $\mathbb{N}^{(2)}$ is finitely coloured there exist $x_1 < x_2 < \cdots$ in $\mathbb{N}$ such that all pairs of finite sums $\{\sum_{i\in I} x_i, \sum_{j\in J} x_j\}$, where $I, J \neq \emptyset$ and $\max I < \min J$, are the same colour. Thus the Milliken–Taylor theorem can be viewed as Hindman's theorem crossed with Ramsey's theorem.

The second generalisation is due to Deuber and Hindman.

Theorem 8 ([DH87])**.** *For any sequence $E_1, E_2, \ldots$ of finite partition regular systems of equations, whenever $\mathbb{N}$ is finitely coloured there is a sequence of corresponding solutions sets $S_1, S_2, \ldots$ such that all finite sums of the form $\sum_{i\in I} x_i$, where $I \neq \emptyset$ and $x_i \in S_i$ for all $i \in I$, are the same colour.*

Here a *solution set* S_i is the set of values taken by the variables in some solution to E_i.

Hindman's theorem has a proof using ultrafilters due to Galvin and Glazer (see [Com77]). It involves proving the existence of an ultrafilter on $\mathbb{N}$ with strong combinatorial properties. Deuber and Hindman's generalisation goes by finding an ultrafilter with even stronger properties, and can be viewed as Hindman's theorem crossed with Rado's theorem.

There is no analogue of Rado's theorem for infinite systems. There is not even a consistency result: indeed, Deuber, Hindman, Leader and Lefmann ([DHLL95]) showed that, apart from trivial cases, every pair of Milliken–Taylor systems is inconsistent. They also gave examples of related infinite kernel partition regular systems for which consistency fails.

In order to recover some of the properties of finite partition regular systems, Hindman, Leader and Strauss [HLS03a] studied a stronger notion of partition regularity concerning central sets, a type of set with very strong combinatorial properties which is related to ultrafilters on $\mathbb{N}$. This allowed them to construct some new partition regular systems. For example, they showed that 'bounded parts' of Milliken–Taylor systems are consistent.

Theorem 9 ([HLS03a])**.** *Whenever $\mathbb{N}$ is finitely coloured, there exist $x_1 < x_2 < \cdots$ such that all finite sums $\sum_{i\in I} x_i + \sum_{j\in J} 2x_j$ or $\sum_{i\in I} 2x_i + \sum_{j\in J} x_j$, where $I, J \neq \emptyset$, $\max I < \min J$ and $|I| + |J| \leq 100$, are the same colour.*

Together with some other systems from [HLS03a], these systems are essentially everything that is known.

2 New results

In [HLS03b] Hindman, Leader and Strauss surveyed the outstanding open problems in partition regularity. The main aim of this chapter is to settle a number of open problems from [HLS03b]. In the rest of this section we describe and motivate the open problems and state the new results.

2.1 New partition regular systems

A complete classification of the infinite partition regular systems seems out of reach, but all the known systems are strongly related to the Hindman's theorem. Must every example fall into this framework, or can partition regular systems arise in different ways?

Hindman, Leader and Strauss [HLS03b] observed that an essential feature of the known partition regular systems is that each variable only ever appears with a bounded set of coefficients. So one way to search for new partition regular systems unrelated to Hindman's theorem is to look at systems in which a single variable appears with an unbounded set of coefficients.

If an infinite system of equations is to have any chance of being partition regular, then by Rado's theorem every finite subset of the equations must have the columns property. The easiest way to ensure this is to have, for each equation, some variables with matching coefficients on each side of the equation, with these variables not appearing in any other equation. So the most basic candidate for a new infinite kernel partition regular system is

$$\begin{aligned} x_1 + a_1 y &= z_1 \\ x_2 + a_2 y &= z_2 \\ &\vdots \\ x_n + a_n y &= z_n \\ &\vdots \end{aligned} \tag{1}$$

where (a_n) is some strictly increasing sequence of integer coefficients.

Unfortunately, System 1 is not partition regular. There are several ways to see this; we describe one based on a generalised 'mod 3' colouring.

For $\lambda \in \mathbb{R}$, let

$$c_\lambda(x) = \begin{cases} 0 & \text{if } 0 \le \lfloor \lambda x \rfloor < \frac{1}{3}; \\ 1 & \text{if } \frac{1}{3} \le \lfloor \lambda x \rfloor < \frac{2}{3}; \\ 2 & \text{if } \frac{2}{3} \le \lfloor \lambda x \rfloor < 1. \end{cases}$$

We will show that there is a choice of λ for which System 1 has no monochromatic solution under c_λ. By passing to a subsequence if necessary we may assume that (a_n) grows as fast as we please: say $a_n \ge 4a_{n-1}$.

Suppose we have a monochromatic solution to System 1. Then for every n, $\lfloor \lambda x_n \rfloor$ and $\lfloor \lambda z_n \rfloor$ are in the same interval of length 1/3, so

$$\lfloor \lambda a_n y \rfloor = \lfloor \lambda z_n - \lambda x_n \rfloor \notin [\tfrac{1}{3}, \tfrac{2}{3}).$$

So it suffices to choose λ such that $\lfloor \lambda a_n n \rfloor \in [\frac{1}{3}, \frac{2}{3})$ for every n, as then the nth equation has no monochromatic solution with $y = n$.

For $n = 1$, any value of λ in $[\frac{1}{3a_1}, \frac{2}{3a_1})$ will do. For $n = 2$, any value of λ in an interval of the form $[\frac{1+3k}{3a_2 \cdot 2}, \frac{2+3k}{3a_2 \cdot 2})$ will do. But since $a_2 \ge 4a_1$, at least one such interval is contained entirely within $[\frac{1}{3a_1}, \frac{2}{3a_1})$. Similarly, this interval contains one of length $\frac{1}{3a_3 \cdot 3}$ which is good for $n = 3$. Continuing in this fashion we obtain at the nth stage an interval of length $\frac{1}{3a_n n}$, any element of which will rule out $y \le n$. Taking λ to be the unique point in the intersection of these intervals gives a colouring c_λ under which System 1 has no monochromatic solution.

To increase the chance of obtaining a partition regular system we can add more matched variables to each side of our equations. For example, we might try the following system.

$$\begin{gathered} x_{11} + x_{12} + a_1 y = z_{11} + z_{12} \\ x_{21} + x_{22} + a_2 y = z_{21} + z_{22} \\ \vdots \\ x_{n1} + x_{n2} + a_n y = z_{n1} + z_{n2} \\ \vdots \end{gathered}$$

However, this system is not partition regular either. Indeed, the above argument applies essentially unchanged if we break into 5 pieces instead of 3. Similarly, we can find a bad colouring if we sum any fixed number of x's and z's on each side.

What if we allow a growing number of sums? The following systems of equations were suggested by Neil Hindman and Imre Leader.

Theorem 10. *For any sequence (a_n) of integer coefficients, the system of equations*

$$\begin{gathered} x_{11} + a_1 y = z_{11} \\ x_{21} + x_{22} + a_2 y = z_{21} + z_{22} \\ \vdots \\ x_{n1} + \cdots + x_{nn} + a_n y = z_{n1} + \cdots + z_{nn} \\ \vdots \end{gathered} \tag{2}$$

is partition regular.

We give the proof in Section 3.1. Theorem 10 is the central result of this chapter. We will go on to generalise it in two directions and give a number of applications of this result.

The proof of Theorem 10 uses the symmetry of System 2 (that there are the same number of x's and z's on each side of the equation). For applications it is convenient to have a version for which there is only a single z on the right hand side of each equation.

Theorem 11. *For any sequence (a_n) of integer coefficients, the system of equations*

$$\begin{aligned} x_{11} + a_1 y &= z_1 \\ x_{21} + x_{22} + a_2 y &= z_2 \\ &\vdots \\ x_{n1} + \cdots + x_{nn} + a_n y &= z_n \\ &\vdots \end{aligned} \tag{3}$$

is partition regular.

The proof of Theorem 11 follows in Section 3.3. Some consequences for image partition regularity are described in Section 3.4.

2.2 Partition regularity over different spaces

We saw in Section 1.1 that in the finite case the notions of kernel partition regularity over $\mathbb{N}$, $\mathbb{Z}$ and $\mathbb{Q}$ coincide. In the infinite case the two trivial implications still hold, and the reflection argument again shows that kernel partition regularity over $\mathbb{N}$ and kernel partition regularity over $\mathbb{Z}$ are the same. But the proof that kernel partition regularity over $\mathbb{Q}$ implied kernel partition regular over $\mathbb{Z}$ really did use the finiteness of the system when we said that any solution lay in some level $\frac{1}{n!} \cdot \mathbb{Q}$. In summary, we know that the following relationships hold.

$$\text{KPR}/\mathbb{N} \iff \text{KPR}/\mathbb{Z} \implies \text{KPR}/\mathbb{Q}$$

Question 12 ([HLS03b, Q6])**.** *If a system of equations is partition regular over $\mathbb{Q}$, must it be partition regular over $\mathbb{Z}$?*

The easiest way to provide a negative answer to Question 12 would be to find a system which is partition regular over $\mathbb{Q}$ but does not even have any solutions over $\mathbb{Z}$. For example, we could look at System 1 with $a_n = 1/n$.[1]

$$\begin{aligned} x_1 + y &= z_1 \\ x_2 + \tfrac{1}{2}y &= z_2 \\ &\vdots \\ x_n + \tfrac{1}{n}y &= z_n \\ &\vdots \end{aligned} \tag{4}$$

However, as for the integer case, System 4 is not kernel partition regular over $\mathbb{Q}$. Indeed, Hindman, Leader and Strauss [HLS06] showed that System 1 is not kernel partition regular over $\mathbb{Q}$ for any choice of rational coefficients a_n.

Applying the same trick to System 2 *does* provide an example.

Theorem 13. *For any sequence (a_n) of rational coefficients, the system of equations*

$$\begin{aligned} x_{11} + a_1 y &= z_{11} \\ x_{21} + x_{22} + a_2 y &= z_{21} + z_{22} \\ &\vdots \\ x_{n1} + \cdots + x_{nn} + a_n y &= z_{n1} + \cdots + z_{nn} \\ &\vdots \end{aligned} \tag{5}$$

is partition regular over $\mathbb{Q}$.

What happens for image partition regularity? We again have the trivial implications

$$\mathrm{IPR}/\mathbb{N} \implies \mathrm{IPR}/\mathbb{Z} \implies \mathrm{IPR}/\mathbb{Q},$$

but can we say anything about the converses?

This time $\mathbb{N}$ and $\mathbb{Z}$ are not equivalent. We can see this by considering the

[1] Formally, we should multiply the nth equation by n so that our system has integer coefficients.

following system.

$$\begin{gathered} x_1 - x_2 \\ x_2 - x_3 \\ x_3 - x_4 \\ \vdots \end{gathered}$$

This system is image partition regular over $\mathbb{Z}$: we can take $x_n = n$ and then every expression takes the value -1, so they are certainly all the same colour. But it is not image partition regular over $\mathbb{N}$: we cannot even ensure that all the variables and images are in $\mathbb{N}$, as the x_n would have to form a strictly decreasing sequence, which is impossible.

Question 14 ([HLS03b, Q9]). *If a system of expressions is partition regular over $\mathbb{Q}$, must it be partition regular over $\mathbb{Z}$?*

We describe a counterexample for this question in Section 4.

2.3 Relationship between image and kernel partition regularity

Schur's theorem has equivalent statements in terms of image and kernel partition regularity. Hindman's theorem is naturally a statement about image partition regularity, but it can also be viewed as the statement that the system of equations

$$\left\{ y_I = \sum_{i \in I} x_i \right\},$$

where the y_I are new variables indexed by the finite non-empty subsets of $\mathbb{N}$, is partition regular. Is this rephrasing possible in general?

Given a matrix A (with only finitely many non-zero entries in each row) we define a matrix B, expressing the linear dependences between the rows of A, as follows. Let $\{r_i : i \in I\}$ be a maximal linearly independent set of rows of A, and write each of the remaining rows $\{s_j : j \in J\}$ as a linear combination of the r_i. Let B be the matrix corresponding to these linear equations. That is, for each $j \in J$ write $s_j = \sum_{i \in I} c_{ji} r_i$, and let B be the $J \times (I \cup J)$ matrix with left-hand side (c_{ji}) and right-hand side -1 times the $J \times J$ identity matrix.[2]

If A is image partition regular, then certainly B is kernel partition regular, because any monochromatic image of A is in the kernel of B. Conversely:

Question 15 ([HLS03b, Q7])**.** *If B is kernel partition regular, must A be image partition regular?*

Over $\mathbb{Q}$ this question has a positive answer. The construction of B ensures that, given a monochromatic kernel vector y of B, the entries of y form a consistent set of values for the entries of Ax. The equation $Ax = y$ can then be solved over $\mathbb{Q}$ by Gaussian elimination.

This inverse procedure fails over $\mathbb{Z}$ because we may need to perform divisions that produce non-integer results. This motivated the following question.

Question 16 ([HLS03b, Q8])**.** *Let A be a matrix that is image partition regular over $\mathbb{Z}$ and (d_i) be a sequence of integers. Is it true that, whenever the natural numbers are finitely coloured, there is a monochromatic image Ax such that the variables x_i satisfy $x_i \equiv 0 \pmod{d_i}$?*

Note that we cannot ask for $x_i \equiv a_i \pmod{d_i}$ for arbitrary a_i. For example, in Schur's theorem we cannot ask for x to be odd; for then colouring the odd numbers red and the even numbers blue forces all of x, y and $x + y$ to be odd.

We describe a counterexample for Question 16 in Section 4. The same system turns out to also provide a counterexample for Question 15.

[2] If the rows of A are linearly independent, then B is an empty matrix, which is automatically kernel partition regular.

3 Proofs of main theorems

3.1 Proof of Theorem 10

System 2 has a solution in colour class A if and only if there is some $y \in A$ such that

$$\begin{aligned} a_1 y &\in A - A \\ a_2 y &\in 2A - 2A \\ &\vdots \\ a_n y &\in nA - nA \\ &\vdots \end{aligned}$$

where we use the standard notation for sumsets and difference sets

$$\begin{aligned} A + B &= \{a + b : a \in A, b \in B\} \\ A - B &= \{a - b : a \in A, b \in B\} \\ kA &= \underbrace{A + \cdots + A}_{k \text{ times}}. \end{aligned}$$

We expect iterated sumsets and difference sets to have some additive structure, so this rewriting of System 2 suggests that a sensible first question might be 'what kind of structure can we find inside $nA - nA$ when n is large?'

We shall need a notion of density. For a subset S of $\mathbb{N}$, its (*upper*) *density* is

$$d(S) = \limsup_{n \to \infty} \frac{|S \cap [n]|}{n}.$$

For a subset S of $\mathbb{Z}$ we write $d(S) = d(S \cap \mathbb{N})$ (that is, we measure its density in one direction). We call S *dense* if $d(S) > 0$. Density has the following properties:

- If $A \subseteq B$, then $d(A) \leq d(B)$.
- For any A and B, $d(A \cup B) \leq d(A) + d(B)$. In particular, whenever $\mathbb{N}$ is finitely coloured, at least one colour class is dense.

- If $A + x_1, \dots, A + x_k$ are disjoint translates of a set A, then

$$d\left(\bigcup_{i=1}^{k}(A + x_i)\right) = kd(A).$$

  This is because A and $A + x_i$ are dense in roughly the same intervals. (It is not true in general that if A and B are disjoint then $d(A \cup B) = d(A) + d(B)$: indeed, it is easy to construct infinitely many pairwise disjoint sets that each have density 1 by having the sets take turns to have density close to 1 on initial segments.)

For dense sets things work as well as we could hope for. Write $m \cdot S = \{ms : s \in S\}$ for the set obtained from S by pointwise multiplication by m. The following is a slight generalisation of a result of Stewart and Tijdeman [ST83] (whose argument covers the case where k is a power of 2).

Lemma 17. *Let S be a dense, symmetric subset of $\mathbb{Z}$ containing 0. Then there is an $m \in \mathbb{N}$ such that, for $k \geq 2/d(S)$, $kS = m \cdot \mathbb{Z}$.*

In particular, if A is a dense subset of $\mathbb{N}$, then there is an $m \in \mathbb{N}$ such that, for $k \geq 2/d(A)$, $kA - kA = m \cdot \mathbb{Z}$.

Proof. We will show that there is some $j \leq 1/d(S)$ such that $(2j + 1)S = (2j)S$, and so $(k + 1)S = kS$ for $k \geq 2/d(S)$. Once we have shown this, let $k = \lceil 2/d(S) \rceil$ and let $X = kS$. We have that $X + X = kS + kS = (2k)S = X$. Since S is symmetric, we also have that $X = -X$, so X is closed under addition and the taking of inverses, hence is a subgroup of $\mathbb{Z}$ as required.

Suppose instead that, for each $j \leq 1/d(S)$, $(2j)S \subset (2j + 1)S$. We claim that, for each such j, $(2j + 1)S$ contains $j + 1$ disjoint translates of S. This is a contradiction for $j = \lfloor 1/d(S) \rfloor$.

The claim is true for $j = 0$. For $j > 0$, choose $x \in (2j + 1)S \setminus (2j)S$. Then $x = s_1 + \cdots + s_{2j+1}$ with $s_i \in S$ for each i. We have

$$S + x - s_1 = S + s_2 + \cdots + s_{2j+1} \subseteq (2j + 1)S$$
$$\text{and } (2j - 1)S - s_1 \subseteq (2j)S \subseteq (2j + 1)S.$$

Hence it suffices to show that $S + x - s_1$ and $(2j-1)S - s_1$ are disjoint. But if they intersect then $t_0 + x - s_1 = t_1 + \cdots + t_{2j-1} - s_1$ for some $t_i \in S$, from which it follows that $x = t_1 + \cdots + t_{2j-1} - t_0 \in (2j)S$, contradicting the choice of x.

For the 'in particular' statement, note that $d(A - A) \geq d(A)$. □

If there is a $y \in A$ with y divisible by m then Lemma 17 tells us that we can solve System 2 'eventually' inside A. That leaves only finitely many equations unsolved, which we hope to solve using Rado's theorem. The problem is that the colour class obtained from Rado's theorem might not be dense. To avoid this situation we work inside a long homogeneous arithmetic progression that is disjoint from the non-dense colour classes.

Lemma 18. *Let $\mathbb{N} = A \cup B$ where $d(B) = 0$. Then, for any $l \in \mathbb{N}$, there is a $d \in \mathbb{N}$ such that $d \cdot [l] \subseteq A$.*

Proof. Try every $d = 1, 2, 3, \ldots$ in turn. If $d \cdot [l]$ is not contained in A then it contains an obstruction $b \in B$. Each b can obstruct at most l arithmetic progressions, so if every $d \leq D$ fails then we can find at least D/l elements of B inside $[Dl]$. Suppose that there is no good choice of d. Then D can be taken arbitrarily large, so the density of B is at least

$$\frac{D/l}{Dl} = \frac{1}{l^2},$$

which contradicts the assumption that B is not dense. □

We can now show that System 2 is partition regular.

Theorem 19. *Let $\mathbb{N} = A_1 \cup \cdots \cup A_r$ be an r-colouring of $\mathbb{N}$. Then there is a colour class A_i containing a solution to System 2.*

Proof. For each dense colour class A_i, apply Lemma 17 to obtain m_i and K_i such that, for $k \geq K_i$, $kA_i - kA_i = m_i \cdot \mathbb{Z}$. Let m be the least common multiple of the m_i, and let K be the maximum of the K_i. For every dense colour class A_i, and $k \geq K$, we have $kA_i - kA_i \supseteq m \cdot \mathbb{Z}$.

Write P for the system consisting of the first $K-1$ equations of System 2. It is easy to check that P has the columns property, so by Rado's theorem it is partition regular. It follows that there exists an l such that whenever $[l]$, or any progression $c \cdot [l]$, is r-coloured it contains a monochromatic solution to P.

Apply Lemma 18 to get d with $d \cdot [ml]$ disjoint from the non-dense colour classes. Then $md \cdot [l] \subseteq d \cdot [ml]$ is also disjoint from the non-dense colour classes, and by the choice of l there is a dense colour class A_i such that $A_i \cap (md \cdot [l])$ contains a solution

$$\{y, x_{1,1}, \ldots, z_{K-1,K-1}\}$$

to P, where every element of the solution set is divisible by m. Then, for $k \geq K$, $a_k y \in kA_i - kA_i$, so this solution to P can be extended to a solution for System 2 inside A_i. □

Many results in Ramsey theory have corresponding density versions. For example, Szemerédi's theorem [Sze75] is the density version of Van der Waerden's theorem: any dense subset of $\mathbb{N}$ contains arbitrarily long arithmetic progressions. Although our argument used density in an essential way, there is no density version of Theorem 10; indeed, if $a_1 = 1$ then we cannot even solve the first equation inside the odd numbers.

3.2 Proof of Theorem 13

Given a finite colouring of $\mathbb{Q}$ we seek a colour class A and a $y \in A$ such that

$$\begin{aligned} a_1 y &\in A - A \\ a_2 y &\in 2A - 2A \\ &\vdots \\ a_n y &\in nA - nA \\ &\vdots \end{aligned}$$

The idea is to view $\mathbb{Q}$ as infinitely many nested copies of $\mathbb{Z}$ and apply the methods of the previous section.

Recall that $\mathbb{Q} = \bigcup_{j=1}^{\infty}(\frac{1}{j!} \cdot \mathbb{Z})$, where $(\frac{1}{j!} \cdot \mathbb{Z})$ is the jth *level* of $\mathbb{Q}$. For a set $S \subseteq \mathbb{Q}$, let

$$d_j(S) = \limsup_{n\to\infty} \frac{|S \cap (\frac{1}{j!} \cdot [n])|}{n}$$

be the density of S in the jth level, and let

$$d^*(S) = \limsup_{j\to\infty} d_j(S).$$

If $\mathbb{Q} = A_1 \cup \cdots \cup A_r$, then, for each j,

$$1 \leq d_j(A_1) + \cdots + d_j(A_r),$$

so

$$1 \leq d^*(A_1) + \cdots + d^*(A_r),$$

and $d^*(A_i) > 0$ for at least one i.

The proof follows the same pattern as before.

Lemma 20. *Let $A \subseteq \mathbb{Q}$ with $d^*(A) > 0$. Then there is an $m \in \mathbb{N}$ such that, for $k \geq 4/d^*(A)$, $kA - kA \supseteq \mathbb{Q}$.*

Proof. There are infinitely many levels j such that $d_j(A) > d^*(A)/2$. Let j be one such level and let $A_j = A \cap \frac{1}{j!} \cdot \mathbb{Z}$. Then by Lemma 17 inside level j, there is an $m_j \in \mathbb{N}$ such that, for $k \geq 4/d^*(A) > 2/d_j(A) = 2/d_j(A_j)$,

$$kA - kA \supseteq kA_j - kA_j = \frac{m_j}{j!} \cdot \mathbb{Z}\,.$$

So far we have said nothing about m_j. But m_j cannot be too large, else the density of the subgroup produced by Lemma 17 will be less than the density of the set A_j to which we applied it. So there are only finitely many possible values of m_j, and some m_j occurs infinitely often: call it m. Choose an infinite set J of levels such that, for $k \geq 4/d^*(A)$, $kA - kA \supseteq \frac{m}{j!} \cdot \mathbb{Z}$. Then, for $k \geq 4/d^*(A)$,

$$\begin{aligned} kA - kA &\supseteq \bigcup_{j \in J} \left(\frac{m}{j!} \cdot \mathbb{Z} \right) \\ &= \bigcup_{j=0}^{\infty} \left(\frac{m}{j!} \cdot \mathbb{Z} \right) \\ &= m \cdot \mathbb{Q} \\ &= \mathbb{Q}, \end{aligned}$$

since the levels of $\mathbb{Q}$ are nested and J is infinite. □

Lemma 21. *Let $\mathbb{Q} = A \cup B$ where $d^*(B) = 0$. Then, for any $l \in \mathbb{N}$, there exist j and d such that $\frac{d}{j!} \cdot [l] \subseteq A$.*

Proof. Since $d^*(B) = 0$ there is a j such that $d_j(B) < \dfrac{1}{4l^2}$. We then apply the argument of Lemma 18 inside the jth level of $\mathbb{Q}$. □

Theorem 22. *Let $\mathbb{Q} = A_1 \cup \cdots \cup A_r$ be an r-colouring of $\mathbb{Q}$. Then there is a colour class A_i containing a solution to System 5.*

Proof. For each dense colour class A_i, apply Lemma 20 to obtain K_i such that, for $k \geq K_i$, $kA_i - kA_i \supseteq \mathbb{Q}$. Let K be the maximum of the K_i. For every dense colour class A_i, and $k \geq K$, we have $kA_i - kA_i \supseteq \mathbb{Q}$.

Write P for the system consisting of the first $K-1$ equations of System 5. It is easy to check that P has the columns property, so by Rado's theorem it is partition regular. It follows that there exists an l such that, whenever a progression $\frac{c}{j!} \cdot [l]$ is r-coloured, it contains a monochromatic solution to P.

Apply Lemma 21 to get j and d with $\frac{d}{j!} \cdot [l]$ disjoint from the colour classes with $d^*(A_i) = 0$. Then by the choice of l there is a dense colour class A_i such that $A_i \cap \left(\frac{d}{j!} \cdot [l] \right)$ contains a solution

$$\{y, x_{1,1}, \ldots, z_{K-1,K-1}\}$$

to P. Then, for $k \geq K$, $a_k y \in kA_i - kA_i$, so this solution to P can be extended to a solution for System 5 inside A_i. □

3.3 Proof of Theorem 11

We now turn our attention to the less symmetric System 3. Given a finite colouring of $\mathbb{N}$ we seek a colour class A and a $y \in A$ for which

$$\begin{aligned} a_1 y &\in A - A \\ a_2 y &\in A - 2A \\ &\vdots \\ a_n y &\in A - nA \\ &\vdots \end{aligned}$$

The next two lemmas generalise Lemma 17 to the asymmetric case.

Lemma 23. *Let S be a dense subset of $\mathbb{Z}$ with $0 \in S$. Then there is an $X \subseteq \mathbb{Z}$ such that, for any $k \geq 2/d(S)$, we have $S - kS = X$.*

Proof. As in the proof of Lemma 17, we suppose to the contrary that $S - (2j)S \subset S - (2j+1)S$ for all $j \leq 1/d(S)$ and show that $S - (2j+1)S$ contains $j+1$ disjoint translates of S for each such j, which is impossible for $j = \lfloor \frac{1}{d(S)} \rfloor$.

The claim is true for $j = 0$. For $j > 0$, choose $x \in \big(S - (2j+1)S\big) \setminus \big(S - (2j)S\big)$. Then $x = s_0 - s_1 - \cdots - s_{2j+1}$ with $s_i \in S$ for each i. We have

$$\begin{aligned} S + x - s_0 &\subseteq S - (2j+1)S \text{ and} \\ S - (2j-1)S - s_0 \subseteq S - (2j)S &\subseteq S - (2j+1)S, \end{aligned}$$

hence it suffices to show that $S + x - s_0$ and $S - (2j-1)S - s_0$ are disjoint. But if they intersect then $t_0 + x - s_0 = t_1 - t_2 - \cdots - t_{2j} - s_0$ for some $t_i \in S$, whence $x = t_1 - t_2 - \cdots - t_{2j} - t_0 \in S - (2j)S$, contradicting the choice of x. □

What can we say about X?

Lemma 24. *Let A be a dense subset of $\mathbb{N}$. Then there is an m such that, for $k \geq 2/d(A)$, $A - kA$ is a union of cosets of $m \cdot \mathbb{Z}$.*

Proof. Let $k \geq 2/d(A)$, and let $X = A - kA$. For any $a \in A$, we have by Lemma 23 that

$$(A - a) - k(A - a) = (A - a) - (k + 1)(A - a),$$

and so

$$X = X - A + a.$$

Since $a \in A$ was arbitrary it follows that $X = X + A - A$, whence $X = X + l(A - A)$ for all l. Taking l sufficiently large, by Lemma 17 there is an $m \in \mathbb{Z}$ such that $X = X + m \cdot \mathbb{Z}$. Thus X is a union of cosets of $m \cdot \mathbb{Z}$. □

The example of the odd numbers shows that it is not necessarily the case that $A - kA \supseteq m \cdot \mathbb{Z}$ for large k. However, the obstruction is clear: in that case A does not contain even a single element of $m \cdot \mathbb{Z}$. But then we can pass down to a copy of $\mathbb{N}$ coloured with one fewer colour and apply induction. Combining this idea with the argument for System 2 concludes the proof that System 3 is partition regular.

Theorem 25. *Let* $\mathbb{N} = A_1 \cup \cdots \cup A_r$ *be an* r*-colouring of* $\mathbb{N}$*. Then there is a colour class* A_i *containing a solution to System 3.*

Proof. Suppose first that there is an m and an i such that A_i is disjoint from $m \cdot \mathbb{Z}$. Then $m \cdot \mathbb{N}$ is $(r - 1)$-coloured by the remaining colour classes, so by induction on r we can find a monochromatic solution to System 3 inside $m \cdot \mathbb{N}$.

Otherwise we may assume that every colour class meets every subgroup of $\mathbb{Z}$. Apply Lemma 24 to each dense colour class A_i to obtain m_i and K_i such that, for $k \geq K_i$, $A_i - kA_i$ is a union of cosets of $m_i \cdot \mathbb{Z}$. Let m be the least common multiple of the m_i, and let K be the maximum of the K_i. Then, for every dense colour class A_i, and $k \geq K$, $A_i - kA_i$ is a union of cosets of $m \cdot \mathbb{Z}$. Since every colour class contains a multiple of m, one of those cosets is $m \cdot \mathbb{Z}$ itself.

Write P for the system consisting of the first $K - 1$ equations of System 3. Again, it is easy to check that P has the columns property, so by Rado's theorem it is partition regular. It follows that there exists an l such that whenever $[l]$, or a progression $c\cdot[l]$, is r-coloured it contains a monochromatic solution to P.

Apply Lemma 18 to get d with $d\cdot[ml]$ disjoint from the non-dense colour classes. Then $md\cdot[l]\subseteq d\cdot[ml]$ is also disjoint from the non-dense colour classes, and by the choice of l there is a dense colour class A_i such that $A_i\cap(md\cdot[l])$ contains a solution

$$\{y, x_{1,1}, \ldots, z_{K-1}\}$$

to P, where every element of the solution set is divisible by m. Then, for $k\geq K$, $a_k y\in A_i - kA_i$, so this solution to P can be extended to a solution for System 3 inside A_i. □

3.4 Consequences for image partition regularity

Theorem 11 has the following immediate consequence.

Theorem 26. *For any sequence (a_n) of integer coefficients, the system of expressions*

$$\begin{array}{rcc} x_{11}+a_1y & x_{11} & y, \\ x_{21}+x_{22}+a_2y & x_{21} & \\ x_{31}+x_{32}+x_{33}+a_3y & x_{22} & \\ \vdots & \vdots & \end{array} \tag{6}$$

consisting of the left-hand sides of System 3 and the variables they contain, is partition regular over $\mathbb{N}$.

Proof. The expressions in the first column are precisely the values of the z_n in any solution to System 3. □

De and Hindman [DH09] studied image partition regularity over different subsets of $\mathbb{R}$. We call a matrix A *image partition regular over* $\mathbb{R}$ *near zero* if, for all $\delta>0$, whenever $(-\delta,\delta)\setminus\{0\}$ is finitely coloured there is a vector x with real entries such that Ax is monochromatic and contained in $(-\delta,\delta)\setminus\{0\}$.

Seeking to establish whether a diagram of implications in [DH09] involving nineteen properties had any missing implications, De and Hindman asked the following question.

Question 27 ([DH09, Q3.10])**.** *If a system of expressions is partition regular over $\mathbb{N}$, must it be partition regular over $\mathbb{R}$ near zero?*

We now show that System 6 provides a counterexample.

Theorem 28. *For any strictly increasing sequence (a_n) of integer coefficients, System 6 is not partition regular over $\mathbb{R}$ near zero.*

Proof. Choose $\delta > 0$ and partition $(-\delta, \delta)\setminus\{0\}$ as $(-\delta, 0)\cup(0, \delta)$. If System 6 is to be monochromatic then y and each of the x_{ij} must lie in the same part: without loss of generality they are all positive. But then

$$x_{n1} + \cdots + x_{nn} + a_n y > a_n y > \delta$$

for n sufficiently large. □

System 6 is not the original system proposed by De and Hindman as a possible counterexample for Question 27. In Section 3.5 we show by a more complicated argument that De and Hindman's system is also partition regular over $\mathbb{N}$.

3.5 A system of De and Hindman

De and Hindman considered the following system.

$$
\begin{array}{lll}
x_{21} + x_{22} & x_{21} + 2y & y, \\
 & x_{22} + 2y & \\
 & & \\
x_{41} + x_{42} + x_{43} + x_{44} & x_{41} + 4y & \\
 & x_{42} + 4y & \\
 & x_{43} + 4y & \\
 & x_{44} + 4y & \\
 & \vdots & \\
x_{2^n 1} + \cdots + x_{2^n 2^n} & x_{2^n 1} + 2^n y & \\
 & \vdots & \\
 & x_{2^n 2^n} + 2^n y & \\
 & \vdots &
\end{array}
\tag{7}
$$

Using the argument of Theorem 28 it is easy to show that System 7 is not partition regular over $\mathbb{R}$ near zero. In this section we show that System 7 is partition regular over $\mathbb{N}$.

Theorem 29. *For any sequence (a_n) of integer coefficients, the system of expressions*

$$\begin{array}{lll} x_{11} & x_{11}+a_1y & y, \\ \\ x_{21}+x_{22} & x_{21}+a_2y & \\ & x_{22}+a_2y & \\ \\ x_{31}+x_{32}+x_{33} & x_{31}+a_3y & \\ & x_{32}+a_3y & \\ & x_{33}+a_3y & \\ & \vdots & \end{array} \tag{8}$$

is partition regular over $\mathbb{N}$*.*

Taking $a_n = n$ implies that System 7 is partition regular over $\mathbb{N}$.

The partition regularity of System 8 *almost* follows directly from Theorem 11. Indeed, by Theorem 11 we can always find a monochromatic solution to the following system of equations.

$$\begin{aligned} \tilde{x}_{11} - a_1y &= z_1 \\ \tilde{x}_{21} + \tilde{x}_{22} - 2a_2y &= z_2 \\ &\vdots \\ \tilde{x}_{n1} + \cdots + \tilde{x}_{nn} - na_ny &= z_n \\ &\vdots \end{aligned}$$

Set $x_{ni} = \tilde{x}_{ni} - a_ny$. Then, for each n and i,

$$x_{n1} + \cdots + x_{nn} = \tilde{x}_{n1} + \cdots + \tilde{x}_{nn} - na_ny = z_n,$$

and

$$x_{ni} + a_ny = \tilde{x}_{ni},$$

so we have found a monochromatic image for System 8. The problem is that we cannot be sure that the variables $x_{ni} = \tilde{x}_{ni} - a_n y$ are positive. To fix this we look inside the proof of Theorem 11 to show that we can take the x_{ni} to be as large as we please.

Much of the following lemma has been used implicitly already. Write $A_{>t} = \{a \in A : a > t\}$.

Lemma 30. *Let A be a dense subset of $\mathbb{N}$ that meets every subgroup of $\mathbb{Z}$, and let m be the least common multiple of $1, 2, \ldots, \lfloor 1/d(A) \rfloor$. Then, for $n \geq 2/d(A)$ and any t,*

$$A_{>t} - nA_{>t} \supseteq m \cdot \mathbb{Z}.$$

Proof. First observe that, for any t, $d(A_{>t}) = d(A)$. Let $n \geq 2/d(A)$, and let $X = A_{>t} - nA_{>t}$. For any $a \in A_{>t}$, we have by Lemma 23 that

$$(A_{>t} - a) - n(A_{>t} - a) = (A_{>t} - a) - (n+1)(A_{>t} - a),$$

and so

$$X = X - A_{>t} + a.$$

Since $a \in A_{>t}$ was arbitrary it follows that $X = X + A_{>t} - A_{>t}$, whence $X = X + l(A_{>t} - A_{>t})$ for all l. By Lemma 17 there is an $m_t \in \mathbb{Z}$ such that, for $l \geq 2/d(A)$, $l(A_{>t} - A_{>t}) = m_t \cdot \mathbb{Z}$. Hence $X = X + m_t \cdot \mathbb{Z}$, and X is a union of cosets of $m_t \cdot \mathbb{Z}$. Since A contains arbitrarily large multiples of m_t, one of these cosets is $m_t \cdot \mathbb{Z}$ itself.

Since $lA_{>t} - lA_{>t}$ contains a translate of $A_{>t}$,

$$1/m_t = d(m_t \cdot \mathbb{Z}) \geq d(A),$$

and $m_t \leq 1/d(A)$. So m_t divides m and

$$A_{>t} - nA_{>t} \supseteq m \cdot \mathbb{Z}. \qquad \square$$

We can now show that System 8 is partition regular.

Proof of Theorem 29. Let $\mathbb{N}$ be r-coloured. Suppose first that some colour class does not meet every subgroup of $\mathbb{Z}$; say some class contains no multiple of m. Then $m \cdot \mathbb{N}$ is $(r-1)$-coloured by the remaining colour classes, so by induction on r we can find a monochromatic image. So we may assume that every colour class meets every subgroup of $\mathbb{Z}$.

Let d be the least density among the dense colour classes, and let m be the least common multiple of $1, 2, \ldots, \lfloor 1/d \rfloor$. Then by Lemma 30, for any dense colour class A, any t and $n \geq 2/d$,

$$A_{>t} - nA_{>t} \supseteq m \cdot \mathbb{Z}.$$

Now let $N = \lceil 2/d \rceil - 1$. We will find a monochromatic image for the the expressions containing only y and x_{ni} for $n \leq N$ using Rado's theorem. Indeed, consider the following system of linear equations.

$$\begin{aligned}
&u_1 = x_{11} && v_{11} = x_{11} + a_1 y && y, \\
\\
&u_2 = x_{21} + x_{22} && v_{21} = x_{21} + a_2 y \\
& && v_{22} = x_{22} + a_2 y \\
& && \vdots \\
&u_N = x_{N1} + \cdots + x_{NN} && v_{N1} = x_{N1} + a_N y \\
& && \vdots \\
& && v_{NN} = x_{NN} + a_N y
\end{aligned} \tag{9}$$

It is easy to check that the matrix corresponding to this system has the columns property, so by Rado's theorem there is an l such that, whenever a progression $c \cdot [l]$ is r-coloured, it contains a monochromatic solution to System 9.

Apply Lemma 18 to get c with $c \cdot [ml]$ disjoint from the non-dense colour classes. Then $mc \cdot [l] \subseteq c \cdot [ml]$ is also disjoint from the non-dense colour classes, and by the choice of l there is a dense colour class A such that $A \cap (md \cdot [l])$ contains a solution to System 9. Since the u_n, v_{ni} and y are all in A, y and the corresponding x_{ni} make the first part of System 8 monochromatic.

Now y is divisible by m, so for $n > N$ we have that

$$-na_n y \in A_{>a_n y} - nA_{>a_n y},$$

so there are $\tilde{x}_{ni}$ and z_n in $A_{>a_n y}$ such that

$$-na_n y = z_n - \tilde{x}_{n1} - \cdots - \tilde{x}_{nn}.$$

Set $x_{ni} = \tilde{x}_{ni} - a_n y$. Then

$$x_{n1} + \cdots + x_{nn} = \tilde{x}_{n1} + \cdots + \tilde{x}_{nn} - na_n y = z_n,$$

and

$$x_{ni} + a_n y = \tilde{x}_{ni},$$

for each $n > N$ and $1 \leq i \leq n$. Since $\tilde{x}_{ni}$ and z_n are in A it follows that the whole of System 8 is monochromatic.

It remains only to check that all of the variables are positive. But for y and x_{ni} with $n \leq N$ this is guaranteed by Rado's theorem; for $n > N$ it holds because $\tilde{x}_{ni} > a_n y$. $\square$

4 Counterexamples

In this section we will first describe a system which provides a counterexample for Question 16. We will then go on to show how this system can be adapted to provide counterexamples for Questions 14 and 15.

We seek a matrix A, and a sequence of natural numbers (d_i), such that A is image partition regular but we cannot always find a monochromatic image Ax satisfying $x_i \equiv 0 \pmod{d_i}$ for each i. Our system will come from Theorem 26. Roughly speaking, our idea is to simulate the condition that a variable be odd by giving some conditions on the other variables. This will be achieved by a particular choice of the constants a_n and the divisibility constraints d_{ij} on the x_{ij}.

For every $n \in \mathbb{N}$, we choose $a_n \in \mathbb{N}$ such that

$$a_n n \equiv 2^{n-1} \pmod{2^n}.$$

To see that this is possible, write $n = 2^k p$, where p is odd and $k < n$. Then we seek a_n such that

$$a_n p \equiv 2^{n-k-1} \pmod{2^{n-k}}.$$

But p is odd, hence invertible mod 2^{n-k}, so these congruences have solutions. For example, we can take $a_1 = 1$, $a_2 = 1$, $a_3 = 4$, $a_4 = 2$ and $a_5 = 16$.

We now show that the System 6 with this choice of coefficients provides a counterexample for Question 16. Note that, since every variable is in the image of System 6, we may assume without loss of generality that all variables are positive by using disjoint sets of colours for the positive and negative integers. So it suffices to prove the following result.

Proposition 31. *There is a 2-colouring of* $\mathbb{N}$ *such that there are no natural numbers* y *and* x_{ij} *such that* $x_{ij} \equiv 0 \pmod{2^i}$ *and System 6 is monochromatic.*

Proof. We shall define the colouring in stages, so that at the nth stage we force n and the nth expression in the first column to have different colours. This rules out the possible values of y one by one.

For $n = 1$, we colour 1 red and all the other odd numbers blue. Since x_{11} must be even and positive, and a_1 is odd, the first expression must be blue.

For $n = 2$, we colour 2 red and all the other numbers that are 2 mod 4 blue. Since x_{21} and x_{22} are 0 mod 4 and positive, and $2a_2 \equiv 2 \pmod{4}$, the second expression must be blue.

For $n = 3$, the number 3 has already been coloured blue. So we colour every number that is 4 mod 8 red. Since $x_{31} \equiv x_{32} \equiv x_{33} \equiv 0 \pmod{8}$, and $3a_3 \equiv 4 \pmod{8}$, the third expression must be red.

Continue. At the nth stage, n has already received some colour. Give all numbers which are 2^{n-1} mod 2^n the opposite colour to n. Since $x_{n1} \equiv \cdots \equiv$

$x_{nn} \equiv 0 \pmod{2^n}$, and $a_n n \equiv 2^{n-1} \pmod{2^n}$, the nth expression must have the opposite colour to n. □

We now turn to Question 15. We will use a reformulation of Proposition 31 to obtain a counterexample. Consider the system below, obtained by reparameterising System 6 by setting $x_{ij} = 2^i z_{ij}$.

$$\begin{array}{rcc} 2z_{11} + a_1 y & 2z_{11} & y \\ 4z_{21} + 4z_{22} + a_2 y & 4z_{21} & \\ 8z_{31} + 8z_{32} + 8z_{33} + a_3 y & 4z_{22} & \\ \vdots & \vdots & \\ 2^n z_{n1} + \ldots + 2^n z_{nn} + a_n y & 2^i z_{ij} & \\ \vdots & \vdots & \end{array} \tag{10}$$

System 10 with the previous choice of a_n is not partition regular over $\mathbb{Z}$, as if it were then taking $x_{ij} = 2^i z_{ij}$ would contradict Proposition 31. But if System 6 is represented by the matrix A_1 and System 10 is represented by the matrix A_2, then $B_1 = B_2$, so image partition regularity of A cannot be determined by examining the matrix B. (In fact, B_1 *is* kernel partition regular by Theorem 11.)

Finally, observe that System 6 and System 10 have the same images over $\mathbb{Q}$. So System 10 is not partition regular over $\mathbb{Z}$ but is partition regular over $\mathbb{Q}$, providing a counterexample for Question 14.

The results of Sections 3.1–3.4 of this chapter are joint work with Neil Hindman and Imre Leader, and have been published as [BHL13]. The results of Section 4 of this chapter are joint work with Imre Leader, and have been published as [BL13]. The results in Section 3.5 are my own work.

Chapter 3

Maximum hitting for n sufficiently large

1 Introduction

A *family* of sets is a subset of $[n]^{(r)}$ for some n and r. We think of a set A as an increasing sequence of elements $a_1 a_2 \dots a_r$. The *compression order* on $[n]^{(r)}$ has $A \leq B$ if and only if $a_i \leq b_i$ for $1 \leq i \leq r$. A family $\mathcal{A}$ is *left-compressed* if $A \in \mathcal{A}$ whenever $A \leq B$ for some $B \in \mathcal{A}$. The corresponding notion of left-compression is described in Section 2.

A family $\mathcal{A}$ is *intersecting* if $A \cap B \neq \emptyset$ for all $A, B \in \mathcal{A}$. (If $n < 2r$ then every family is intersecting.) The most basic result about intersecting families is the Erdős-Ko-Rado Theorem. For any n and r, write $\mathcal{S} = \{A \in [n]^{(r)} : 1 \in A\}$ for the *star* at 1.

Theorem 1 (Erdős-Ko-Rado [EKR61])**.** *If $n \geq 2r$ and $\mathcal{A} \subseteq [n]^{(r)}$ is intersecting, then $|\mathcal{A}| \leq |\mathcal{S}|$.*

Borg considered a variant problem where we only count members that meet some fixed set X. For a family $\mathcal{A}$ and a non-empty set X, write

$$\mathcal{A}(X) = \{A \in \mathcal{A} : A \cap X \neq \emptyset\}.$$

Theorem 1 tells us that we can maximise $|\mathcal{A}(X)|$ by taking $\mathcal{A}$ to consist of all r-sets containing some fixed element of X. To avoid this trivial case we insist that $\mathcal{A}$ be left-compressed, which rules out stars centred anywhere but 1. The star at 1 remains the optimal family if $1 \in X$, so we assume further that $X \subseteq [2, n]$.

Question 2. *For which X do we have $|\mathcal{A}(X)| \leq |\mathcal{S}(X)|$ for all left-compressed intersecting families $\mathcal{A}$?*

Borg asked this question in [Bor11], giving a complete answer for the case $|X| \geq r$ and a partial answer for the case $|X| < r$. Call X *good* (for n and r) if for every left-compressed intersecting family $\mathcal{A} \subseteq [n]^{(r)}$ we have $|\mathcal{A}(X)| \leq |\mathcal{S}(X)|$.

Theorem 3 (Borg [Bor11])**.** *Let $r \geq 2$, $n \geq 2r$ and $X \subseteq [2, n]$.*

(a) *If $|X| > r$, then X is good.*

(b) *If X is good and $X \leq X'$, then X' is good.*

(c) *For any $k \leq r$, $\{2k, 2k+2, \ldots, 2r\}$ is good.*

(d) *If $n = 2r$ and $|X| = r$, then X is good if and only if $\{2, 4, \ldots, 2r\} \leq X$.*

(e) *If $n > 2r$, $|X| = r$ and either*

  (i) *$r \geq 4$ and $X \neq [2, r+1]$,*

  (ii) *$r = 3$ and $\{2, 3\} \not\subseteq X$, or*

  (iii) *$r = 2$ and $\{2, 3\} \neq X$,*

  then X is good. Otherwise, X is not good.

It is not true that all X are good. For example, consider the *Hilton–Milner* family $\mathcal{T} = \mathcal{S}([2, r+1]) \cup \{[2, r+1]\}$. The family $\mathcal{T}$ is left-compressed, and for any $X \subseteq [2, r+1]$, $|\mathcal{T}(X)| = |\mathcal{S}(X)| + 1$, so X is not good.

Our main result is that, surprisingly, for large n and $|X| \geq 4$ this turns out to be the only obstruction.

Theorem 4. *Let $r \geq 3$, $n \geq 2r$ and $X \subseteq [2, n]$ with $|X| \leq r$. If $X \not\subseteq [2, r+1]$ and either*

(i) $|X| \geq 4$,

(ii) $|X| = 3$ *and* $\{2, 3\} \not\subseteq X$,

(iii) $|X| = 2$ *and* $2, 3 \notin X$, *or*

(iv) $|X| = 1$,

then, for n sufficiently large, X is good. Otherwise, X is not good.

For $r = 2$, condition (iii) is unnecessary. In this case the result can be checked easily by hand or read out of Theorem 3 in conjunction with the Hilton–Milner example, so we assume $r \geq 3$ for simplicity.

Our proof uses Ahlswede and Khachatrian's notion of generating sets to express the sizes of maximal left-compressed intersecting families, and their restrictions under X, as polynomials in n. It turns out to be sufficient to consider only leading terms, reducing a question about intersecting families of r-sets to a question about intersecting families of 2-sets, which have a very simple structure.

Section 2 sets out the basic properties of compressions and generating sets that we shall use. Section 3 describes a way of thinking about maximal left-compressed intersecting families and proves the lemma that allows us to compare coefficients of polynomials instead of set sizes. Section 4 completes the proof of Theorem 4. Section 5 discusses possible improvements and generalisations.

The statements of the results in this chapter were suggested by examining small cases on a computer. The programs used to do this are described in Section 6. Our proof does not rely on a computer in any way.

2 Compressions and generating sets

In this section we describe the notion of left-compression corresponding to $\leq$ on $[n]^{(r)}$ and the use of generating sets.

2.1 Compressions

For a set A, and $i < j$, the *ij-compression* of A is

$$C_{ij}(A) = \begin{cases} A - j + i & \text{if } j \in A, i \notin A, \\ A & \text{otherwise;} \end{cases}$$

that is, replace j by i if possible. Observe that $A \leq B$ if and only if A can be obtained from B by a sequence of ij-compressions.

For a set family $\mathcal{A}$, define

$$C_{ij}(\mathcal{A}) = \{C_{ij}(A) : A \in \mathcal{A} \text{ and } C_{ij}(A) \notin \mathcal{A}\} \cup \{A : A \in \mathcal{A} \text{ and } C_{ij}(A) \in \mathcal{A}\};$$

that is, compress A if possible. Observe that $\mathcal{A}$ is left-compressed if and only if $C_{ij}(\mathcal{A}) = \mathcal{A}$ for all $i < j$. The following basic result is well known.

Lemma 5. *If $\mathcal{A}$ is intersecting then $C_{ij}(\mathcal{A})$ is intersecting.*

Proof. Let $A, B \in \mathcal{A}$ and let A', B' be their images in $C_{ij}(\mathcal{A})$. We will show that $A' \cap B' \neq \emptyset$.

If neither A nor B moved, or if both did, then $A' \cap B' \neq \emptyset$, so we may assume that $A' = A - j + i$ and $B' = B$. Then A' and B' intersect unless $A \cap B = \{j\}$. Now if $i \in B$, then $A' \cap B' = \{i\} \neq \emptyset$. But if $i \notin B$, then $B - j + i \in \mathcal{A}$, as B didn't move. So

$$A' \cap B' = (A - j + i) \cap B = A \cap (B - j + i) \neq \emptyset. \qquad \square$$

Lemma 5 means that we can always compress an intersecting family to a left-compressed intersecting family of the same size by repeatedly applying ij-compressions. We eventually reach a left-compressed family as $\sum_{A \in \mathcal{A}} \sum_{i=1}^{r} a_i$ is positive and strictly decreases with each successful compression.

2.2 Generating sets

For any r and n, and a collection $\mathcal{G}$ of sets, the family *generated* by $\mathcal{G}$ is

$$\mathcal{F}(r,n,\mathcal{G}) = \{A \in [n]^{(r)} : A \supseteq G \text{ for some } G \in \mathcal{G}\}.$$

Generating sets were introduced by Ahlswede and Khachatrian [AK97], and are useful for the study of intersecting families because they give a restricted number of sets on which all the intersecting actually happens.

Lemma 6 ([AK97]). *Let $\mathcal{G}$ be a collection of subsets of $[n]$ of size at most r. For $n \geq 2r$, $\mathcal{F}(r,n,\mathcal{G})$ is intersecting if and only if $\mathcal{G}$ is.*

Proof. If $\mathcal{G}$ is intersecting then certainly $\mathcal{F}(r,n,\mathcal{G})$ is. Conversely, if $\mathcal{G}$ contains two disjoint sets then (since $n \geq 2r$) they can be completed to disjoint r-sets in $\mathcal{F}(r,n,\mathcal{G})$. □

If $\mathcal{G}$ generates a left-compressed intersecting family, then

$$\mathcal{G}' = \{G' : G' \leq G \text{ for some } G \in \mathcal{G}\}$$

generates the same family, so we may assume that $\mathcal{G}$ is 'left-compressed' (overlooking non-uniformity) and can therefore be described by listing its maximal elements. It is convenient to take

$$\mathcal{F}(r,n,\mathcal{G}) = \{A \in [n]^{(r)} : A \prec G \text{ for some } G \in \mathcal{G}\},$$

where $A \prec G$ ('A is generated by G') if and only if $|G| \leq |A|$ and $a_i \leq g_i$ for $1 \leq i \leq |G|$. We can think of $\prec$ as an extension of $\leq$ to the non-uniform case, where 'missing' elements are assumed to take the value ∞. Thus

$$123 \prec 12 \;(=12\infty);$$
$$(12\infty =)\; 12 \not\prec 123.$$

The following weaker form of Lemma 6 is better suited to our new definition and is sufficient for our purposes.

Corollary 7. *Let $n \geq 2r$ and $\mathcal{G}$ be a collection of subsets of $[2s]$ of size at most s. If $\mathcal{F}(s, 2s, \mathcal{G})$ is intersecting, then so is $\mathcal{F}(r, n, \mathcal{G})$.* □

3 Maximal left-compressed intersecting families

We say an intersecting family $\mathcal{A} \subseteq [n]^{(r)}$ is *maximal* if no other set can be added to $\mathcal{A}$ while preserving the intersecting property. The maximal objects in the set of left-compressed intersecting families are maximal intersecting families (otherwise an extension could be compressed to a left-compressed extension), so the ordering of 'maximal' and 'left-compressed' is unimportant.

The maximal left-compressed intersecting subfamilies of $[n]^{(2)}$ are $\{12, 13, \ldots, 1n\}$ and $\{12, 13, 23\}$, and we can already distinguish between these families when $n = 4$. In fact, the same phenomenon occurs for all r.

Lemma 8. *Let $\mathcal{A} \subseteq [2r]^{(r)}$ be a maximal left-compressed intersecting family and $n \geq 2r$. Then $\mathcal{A}$ extends uniquely to a maximal left-compressed intersecting subfamily of $[n]^{(r)}$. Moreover, every maximal left-compressed intersecting subfamily of $[n]^{(r)}$ arises in this way.*

Proof. Since $\mathcal{A}$ is left-compressed, it can be completely described by listing its $\leq$-maximal elements $A_1, \ldots, A_k$. Some of these sets might contain final segments of [2r]. The idea is that the elements of these final segments would take larger values if they were allowed to, so we obtain a generating set by 'replacing them by ∞'.

For $A = A_i$, take s greatest with $a_s < r + s$ (s exists since $[r+1, 2r]$ is not a member of any left-compressed intersecting family), and let $A' = a_1 \ldots a_s$. Then $\mathcal{G} = \{A'_1, \ldots, A'_k\}$ generates $\mathcal{A}$, as the sets generated by A'_i are precisely those lying below A_i. Since $\mathcal{G}$ is a collection of subsets of $[2r]$ of size at most r and $\mathcal{A} = \mathcal{F}(r, 2r, \mathcal{G})$ is intersecting, Corollary 7 tells us that $\mathcal{F}(r, n, \mathcal{G})$ is a left-compressed intersecting family for every n.

Now let $\mathcal{B}$ be any extension of $\mathcal{A}$ to a left-compressed intersecting subfamily of $[n]^{(r)}$. We will show that $\mathcal{B} \subseteq \mathcal{F}(r, n, \mathcal{G})$. Indeed, if $\mathcal{B} \not\subseteq \mathcal{F}(r, n, \mathcal{G})$

then there is a $B \in \mathcal{B} \setminus \mathcal{F}(r, n, \mathcal{G})$. We claim that there is a $B' \in [2r]^{(r)}$ with $B' \leq B$ and $B' \notin \mathcal{F}(r, 2r, \mathcal{G})$, contradicting the maximality of $\mathcal{A}$.

We obtain B' from B by compressing as little as possible to get $B' \subseteq [2r]$; that is, we take $B' = (B \cap [2r]) \cup [q, 2r]$ with q chosen such that $|B'| = r$. Explicitly, $b'_i = \min(b_i, r + i)$. Now take $G \in \mathcal{G}$. Since $B \notin \mathcal{F}(r, n, \mathcal{G})$, there is an i with $b_i > g_i$. By construction, $r + i > g_i$. So $b'_i = \min(b_i, r + i) > g_i$, and G does not generate B'. Hence $\mathcal{A}$ extends uniquely to a maximal left-compressed intersecting subfamily of $[n]^{(r)}$.

It remains to show that every maximal left-compressed intersecting subfamily of $[n]^{(r)}$ arises in this way. So suppose $\mathcal{C} \subseteq [n]^{(r)}$ is a maximal left-compressed intersecting family with $\mathcal{C} \cap [2r]^{(r)}$ not maximal. Let $\mathcal{D}_0$ be an extension of $\mathcal{C} \cap [2r]^{(r)}$ to a maximal left-compressed intersecting subfamily of $[2r]^{(r)}$, and let $\mathcal{D}$ be the unique maximal extension of $\mathcal{D}_0$ to $[n]^{(r)}$. Since $\mathcal{C}$ is maximal and $\mathcal{D} \setminus \mathcal{C} \neq \emptyset$, there is a $C \in \mathcal{C} \setminus \mathcal{D}$. As above, we obtain $C' \in [2r]^{(r)}$ with $C' \leq C$ and $C' \notin \mathcal{D}_0$. But then $C' \notin \mathcal{C}$, contradicting the assumption that $\mathcal{C}$ is left-compressed. □

Lemma 8 allows a compact description of maximal left-compressed intersecting families. For example, $\{1\}$ generates the star and $\{1(r+1), [2, r+1]\}$ generates the Hilton–Milner family. Enumerating the generating sets using a computer is feasible for small r; Table 1 lists some possibilities. We describe the programs used to generate Table 1 in Section 6.

r	2	3	4	5
	1	1	1	1
	23	23	23	23
		345	345	345
		$14, 234$	4567	4567
		$13, 235, 145$	$15, 2345$	56789
		$12, 245$		$16, 23456$
			72 families in total	37145 families in total

Table 1: Generators for intersecting families of r-sets

In view of Lemma 8, our key tool is the following.

Lemma 9. *Let $n \geq 2$, $X \subseteq [2, 2r]$. Then*

$$|\mathcal{F}(r,n,\mathcal{G})(X)| = \sum_{i=1}^{r} |\mathcal{F}(i,2r,\mathcal{G})(X)| \binom{n-2r}{r-i}.$$

Proof. How do we construct a member of $\mathcal{F}(r,n,\mathcal{G})(X)$? We first choose an initial segment for our set that is contained in $[2r]$ and witnesses the membership of $\mathcal{F}(r,n,\mathcal{G})(X)$ (i.e. meets X and is $\prec$ some $G \in \mathcal{G}$). We then complete our set by taking as many elements as we need from outside $[2r]$. This gives rise to the size claimed. □

4 Proof of Theorem 4

We first show that X is not good if the given conditions do not hold. We have already seen that for $X \subseteq [2, r+1]$ the Hilton–Milner family shows that X is not good for any n. In each of the remaining cases we claim that the family generated by $\{23\}$ shows that X is not good for any n.

So take $X = 23k$ with $k \geq r+2$. We have

$$|\mathcal{F}(r,n,\{1\})(23k)| = \binom{n-2}{r-2} + \binom{n-3}{r-2} + \binom{n-4}{r-2},$$

where the first term counts the sets containing 1 and 2, the second term the sets containing 1 and 3 but not 2, and the third term the sets containing 1 and k but neither 2 nor 3. Similarly,

$$|\mathcal{F}(r,n,\{23\})(23k)| = \binom{n-2}{r-2} + \binom{n-3}{r-2} + \binom{n-3}{r-2},$$

where the terms count the sets containing 1 and 2, the sets containing 1 and 3 but not 2, and the sets containing 2 and 3 but not 1 respectively. Since $r \geq 3$, $|\mathcal{F}(r,n,\{23\})(23k)| > |\mathcal{F}(r,n,\{1\})(23k)|$ and $23k$ is not good.

Next take $X = 3j$ with $j \geq r+2$. We have

$$|\mathcal{F}(r, n, \{1\})(3j)| = \binom{n-2}{r-2} + \binom{n-3}{r-2},$$

where the terms count the sets containing 1 and 3, and the sets containing 1 and j but not 3 respectively. Similarly,

$$|\mathcal{F}(r, n, \{23\})(3j)| = \binom{n-2}{r-2} + \binom{n-3}{r-2} + \binom{n-4}{r-3},$$

where the terms count the sets containing 1 and 3, the sets containing 2 and 3 but not 1, and the sets containing 1, 2 and j but not 3 respectively. Again, since $r \geq 3$, $|\mathcal{F}(r, n, \{23\})(3j)| > |\mathcal{F}(r, n, \{1\})(3j)|$ and $3j$ is not good. It follows from Theorem 3(b) that $2j$ is not good either.

Now we take X satisfying the conditions of the theorem and show that X is good for n sufficiently large. We will show that, for any $\mathcal{G} \neq \{1\}$, $|\mathcal{F}(2, 2r, \mathcal{G})(X)| < |\mathcal{F}(2, 2r, \{1\})(X)| = |X|$. Note that, for any $\mathcal{G}$, $|\mathcal{F}(1, 2r, \mathcal{G})(X)| = 0$ as the only possible singleton generator is 1, which does not meet X. So by Lemma 9, $\mathcal{F}(2, n, \mathcal{G})(X)$ has size polynomial in n with leading coefficient $|\mathcal{F}(2, 2r, \mathcal{G})(X)|$, from which the result will follow.

There are two maximal left-compressed intersecting families of 2-sets, and $\mathcal{F}(2, 2r, \mathcal{G})(X)$ must be contained in one of them. We handle each case separately.

Suppose first that $\mathcal{F}(2, 2r, \mathcal{G})(X) \subseteq \{12, 13, 23\}$. Then it is enough to show that

$$|\{12, 13, 23\}(X)| < |X|.$$

This is clearly true for $|X| \geq 4$. If $|X| = 3$, then it is true because one of 2 or 3 is missing from X so that $|\{12, 13, 23\}(X)| \leq 2$. If $|X| = 2$, then it is true because both 2 and 3 are missing from X, so that $|\{12, 13, 23\}(X)| = 0$. Finally, if $|X| = 1$, then it is true because $X = \{i\}$ with $i \geq r+2 \geq 4$.

Next suppose that $\mathcal{F}(2, 2r, \mathcal{G})(X) \subseteq \{12, 13, \ldots, 1(2r)\}$. Since $\mathcal{F}(r, 2r, \mathcal{G})$ is left-compressed and has a member not containing the element 1, it has $[2, r+1]$ as a member. Hence by the intersecting property of the generators,

$\mathcal{F}(2, 2r, \mathcal{G})(X)$ cannot contain $1j$ for any $j \geq r+2$. But $X \not\subseteq [2, r+1]$, so there is such a $j \in X \setminus [2, r+1]$ and $|\mathcal{F}(2, 2r, \mathcal{G})(X)| < |X|$. □

5 Improvements and generalisations

What happens for small n? Theorem 3(d) tells us that our characterisation cannot be correct for all $n \geq 2r$.

Question 10. *How large is 'sufficiently large' for n in Theorem 4?*

Answering a question in the published version of this chapter, Bond [Bon12] showed that $n \geq (\phi^2 - o(1))r$ (where $\phi^2 = (3+\sqrt{5})/2$ is the square of the golden ratio) is necessary for $X = 24(r+2)$ and conjectured that $n \geq \phi^2 r$ suffices for all X that are good 'eventually'.

r	2	3	4	5
	4	6	8	T
	24	46	68	7T
		246	468	58T
			2468	468T
				2468T

Table 2: Minimal good sets for $n = 2r$ (T = 10)

Table 2 shows the $\leq$-minimal good sets in the case $n = 2r$ for $2 \leq r \leq 5$. The case $r = 5$ rules out the natural conjecture that $\{2k, 2k+2, \ldots, 2r\}$ is the unique minimal good set of its size when $n = 2r$. In fact, it is not even obvious that there *is* a unique minimal good set of each size, although this is true for $r \leq 5$.

Question 11. *Is there a 'nice' characterisation of the good sets for $n = 2r$ when r is sufficiently large?*

It seems unlikely that a good explicit description exists for intermediate values of r and n. The following may be easier.

Question 12. *Is there a short list of families, one of which maximises $|\mathcal{A}(X)|$ for any X?*

Versions of Lemma 8 hold for any property that is preserved under left-compression and can be detected on generating sets. The most obvious candidate is that of being t-intersecting (a family $\mathcal{A}$ is *t-intersecting* if $|A \cap B| \geq t$ for all $A, B \in \mathcal{A}$). Indeed, an identical argument gives the corresponding result that, for large n, a set $X \subseteq [t+1, n]$ with $|X| \geq t+3$ is good if and only if $X \not\subseteq [t+1, r+1]$. (For smaller X the form of good X is again decided by the need to prevent problems caused when $\mathcal{F}(t+1, 2r-t+1, \mathcal{G})(X) \subseteq [t+2]^{(t+1)}$.)

In the context of t-intersecting families it may be more natural to consider

$$\mathcal{A}(s, X) = \{A \in \mathcal{A} : |A \cap X| \geq s\}.$$

For $s = 1$ the argument relies on the fact that maximal left-compressed t-intersecting families of $(t+1)$-sets have one of two very simple forms. For $s = 2$, even the $t = 1$ case is complicated by the larger number of structures of intersecting families of 3-sets (more generally, $(t+s)$-sets); this problem seems likely to get worse for larger s and t.

6 Program listings and results

Many of the results in this chapter were suggested by examining small cases using a computer. The following programs are written in Haskell, and were compiled and run using GHC version 7.0.4.

In this section sets are represented by lower case letters for consistency with the Haskell code.

6.1 Initial exploration

Program listing

We use some list functions.

```
import Data.List
```

The compression order on $[n]^{(r)}$:

```
x `leftOf` y = all id $ zipWith (<=) x y
```

The r-sets from $[n]$, ordered lexicographically:

```
choose  _ 0      = [[]] -- one way to choose empty set
choose [] _      = []   -- no other subsets of empty set
choose (x:xs) r = first ++ second where
    first  = map (x:) $ choose xs (r-1)
    second = choose xs r
```

We build up left-compressed intersecting families by accepting or rejecting each r-set as it is presented to us. Observe that the lexicographic order extends the compression order, so that if we accept a set then we have already accepted everything to its left. We work with lists of pairs (partial family, remaining choices) and sort families lexicographically so that we can use that property later.

```
compressedIntersectingFamilies r n = map sort $
        cif [([], choose [1..n] r)]

cif [] = []
cif partialFamilies = map fst done ++ cif take ++ cif leave where

    -- stop processing families with no further options
    (done, rest) = partition (([]==) . snd) partialFamilies

    -- if we use c then every other set we use must meet c
    take = map take' rest
    take' (f, (c:cs)) = (c:f, [d | d <- cs, c `meets` d])

    -- if we do not use c then we can use no set lying above c
    leave = map leave' rest
    leave' (f, (c:cs)) = (f, [d | d <- cs, not (c `leftOf` d)])
```

We only want to keep maximal intersecting families. We detect non-maximal families by testing all pairs for inclusion.

```
maximalCompressedIntersectingFamilies r n =
    maximalsBy (strict subset) $
    compressedIntersectingFamilies r n

-- maximal elements in generic partial order
maximalsBy (<|) xs = [x | x <- xs, not $ any (x <|) xs]

-- the partial order <| should be strict
strict order x y = order x y && x /= y
```

A listing of the elements of each family is hard to understand, so instead we filter out all but the maximal elements in the compression order.

```
niceList r n = map (maximalsBy (strict leftOf)) $
    maximalCompressedIntersectingFamilies r n
```

Utility functions

Do two ordered lists intersect?

```
[] `meets` _ = False
_ `meets` [] = False
(x:xs) `meets` (y:ys) = case compare x y of
                             LT ->    xs  `meets` (y:ys)
                             EQ ->         True
                             GT -> (x:xs) `meets`    ys
```

Are two ordered lists nested?

```
    [] `subset`  _ = True  -- empty list contained in everything
     _ `subset` [] = False -- nothing else contained in empty list
(x:xs) `subset` (y:ys) = case compare x y of
              LT ->         False      -- x is missing from Y
              EQ ->    xs  `subset` ys -- in both
              GT -> (x:xs) `subset` ys -- y not in X but don't mind
```

Sample output

The function `niceList` lists the $\leq$-maximal elements of each maximal left-compressed intersecting subfamily of $[n]^{(r)}$.

```
*Main> mapM_ print $ niceList 3 6
[[1,5,6]]
[[1,4,6],[2,3,4]]
[[1,3,6],[1,4,5],[2,3,5]]
[[2,3,6]]
[[1,2,6],[2,4,5]]
[[3,4,5]]

*Main> mapM_ print $ niceList 3 7
[[1,6,7]]
[[1,4,7],[2,3,4]]
[[1,3,7],[1,4,5],[2,3,5]]
[[2,3,7]]
[[1,2,7],[2,4,5]]
[[3,4,5]]

*Main> mapM_ print $ niceList 3 8
[[1,7,8]]
[[1,4,8],[2,3,4]]
[[1,3,8],[1,4,5],[2,3,5]]
[[2,3,8]]
[[1,2,8],[2,4,5]]
[[3,4,5]]
```

These results suggest that there are exactly 6 maximal left-compressed intersecting subfamily of $[n]^{(3)}$ for any n. We can also observe how the form of the maximal elements varies with n, allowing us to guess the proof as well as the statement of Lemma 8.

6.2 An optimisation

By Lemma 8, if we want to find the possible generators of maximal left-compressed intersecting subfamilies of $[n]^{(r)}$, it suffices to consider the cases when $n = 2r$. This permits a number of optimisations to the code in the previous section.

The maximal intersecting subfamilies of $[2r]^{(r)}$ are precisely those that contain exactly one of each complementary pair of r-sets. So provided we make sure that we select exactly one set from each complementary pair of r-sets we need not check separately that our families are either intersecting or maximal.

Program listing

We work with pairs (partial family, remaining choices) where choices are now pairs (x, y) of complementary r-sets.

```
complement r = ([1..2*r] \\)

pairs r = [(x, complement r x) | x <- choose [1..2*r-1] r]
```

For some pairs (x, y) our choice is forced as x is to the left of y. We want all pairs to represent genuine choices, so we process these cases first.

```
startpoint r = [(f, cs)] where
    choices = pairs r
    forced = [(x,y) | (x,y) <- choices, x `leftOf` y]
    f = map fst forced
    cs = choices \\ forced
```

The following observation is useful. Suppose that $a \leq b$. Then a has more elements than b in any initial segment of $[n]$. Taking complements, a^c has fewer elements than b^c in any initial segment of $[n]$. Hence $b^c \leq a^c$.

From each complementary pair (x, y) we must choose one set to accept, and one set to reject. If we accept x then we will already have accepted everything to the left of x (and so already rejected everything to the right of y), so there is nothing further to do.

```
takeFirst (f, ((x,y):cs)) = (x:f, cs)
```

If instead we accept y then we must also accept everything to the left of y. By the observation, this will also ensure that we reject everything to the right of x.

```
takeSecond (f, ((x,y):cs)) = (added ++ f, cs') where
    firsts  = [(a,b) | (a,b) <- cs, a `leftOf` y]
    seconds = [(a,b) | (a,b) <- cs, b `leftOf` y]
    added = y : map fst firsts ++ map snd seconds
    cs' = cs \\ (firsts ++ seconds)

families r = map sort $ cif $ startpoint r

cif [] = []
cif partialFamilies = map fst done ++
                      (cif $ map takeFirst rest) ++
                      (cif $ map takeSecond rest) where
    (done, rest) = partition (([]==) . snd) partialFamilies
```

Finally, we want to extract the generators of each family.

```
generatingSets r = map (generators r) $ families r

generators r = map (reverse .
                    strip [2*r, 2*r-1..] .
                    reverse) . maximalsBy (strict leftOf)

strip [] ys = ys
strip _  [] = []
strip (x:xs) (y:ys) = if x == y then strip xs ys else y:ys
```

This program was used to produce the list of generators in Table 1.

6.3 Good sets

Program listing

We use generators to construct the maximal left-compressed intersecting families.

```
_  `genBy` [] = True  -- infinity greater than everything
[] `genBy`  _ = False -- nothing else greater than infinity
(x:xs) `genBy` (y:ys) = x <= y && xs `genBy` ys

fromGenerators r n gs =
    [x | x <- choose [1..n] r, any (x `genBy`) gs]

maximalCompressedIntersectingFamilies r n =
    map (fromGenerators r n) $ generatingSets r
```

After restricting to sets that meet x, is the star still biggest?

```
star r n = map (1:) $ choose [2..n] (r-1)

good r n x = all p $ maximalCompressedIntersectingFamilies r n where
    p f    = (length $ filter (meets x) f) <= target
    target =  length $ filter (meets x) $ star r n
```

Which sets are good?

```
minGoodSets r n = map (minGoodSets' r n) [1..r]

-- which s-sets are good?
minGoodSets' r n s = f $ choose [2..n] s where
    f [] = []
    f (x:xs) | good r n x = x : f [y | y <- xs, not (x `leftOf` y)]
             | otherwise = f xs
```

Sample output

In light of Theorem 3(b), we list only $\leq$-minimal good sets. Curiously, for $r \leq 5$ there is only one of each size.

```
*Main> mapM_ print $ minGoodSets 2 4
[[4]]
[[2,4]]

*Main> mapM_ print $ minGoodSets 3 6
[[6]]
[[4,6]]
[[2,4,6]]

*Main> mapM_ print $ minGoodSets 4 8
[[8]]
[[6,8]]
[[4,6,8]]
[[2,4,6,8]]

*Main> mapM_ print $ minGoodSets 5 10
[[10]]
[[7,10]]
[[5,8,10]]
[[4,6,8,10]]
[[2,4,6,8,10]]
```

The results of this chapter have previously been published as [Bar13].

Chapter 4

Random walks on quasirandom graphs

1 Introduction

We call a graph *quasirandom* if it resembles a random graph in some way. There are many things this could mean. We are interested in dense graphs, so let G be a graph on n vertices with $\rho\binom{n}{2}$ edges. We call ρ the *density* of G.

A random-looking graph should have its edges uniformly spread out. We cannot of course ask that G contains ρk edges from any set of k potential edges, but we can ask for this we restrict ourselves to a small number of natural sets of edges specified in advance: for example, the edges of each large clique in G.

$P_1(\epsilon)$: For every $A \subseteq V(G)$ with $|A| \geq \epsilon n$,

$$\left| e(A) - \binom{|A|}{2} \right| < \epsilon \binom{n}{2},$$

where $e(A) = e(G[A])$ is the number of edges of G spanned by the vertices in A.

Another possibility is to look at the number of edges between large sets of vertices.

$P_2(\epsilon)$: For all $A, B \subseteq V(G)$ with $|A|, |B| \geq \epsilon n$,

$$|e(A, B) - \rho|A||B|| < \epsilon|A||B|,$$

where $e(A, B) = |\{(a, b) \in A \times B : ab \in E(G)\}|$.

Properties P_1 and P_2 tell us about the large scale structure of the graph. Alternatively, we might be interested in the number of copies of small subgraphs present in G.

$P_3^{(s)}(\epsilon)$: For every graph H on s vertices,

$$\left| N_G^*(H) - n^s \rho^{e(H)} (1-\rho)^{\binom{s}{2} - e(H)} \right| < \epsilon n^s,$$

where $N_G^*(H)$ is the number of labelled induced subgraphs of G isomorphic to H; that is, the number of injections $f : V(H) \to V(G)$ such that $uv \in E(H) \iff f(u)f(v) \in E(G)$.

$P_4(\epsilon)$: The number of labelled 4-cycles in G (sequences of distinct vertices $xyzw$ with $xy, yz, zw, wx \in E(G)$) satisfies

$$|C_4(G) - \rho^4 n^4| < \epsilon n^4.$$

Property P_4 looks much weaker than any of P_1, P_2 or $P_3^{(s)}$ (for fixed $s \geq 4$). However, the surprising fact is that they all equivalent, in the following sense: for each pair of properties P and Q and every $\epsilon > 0$, there is a $\delta > 0$ such that, if G satisfies $P(\delta)$, then G satisfies $Q(\epsilon)$.

To see why P_1 and P_2 might be equivalent in this sense, note that $e(A) = \frac{1}{2}e(A, A)$ for every A and, in the other direction,

$$e(A, B) = e(A \cup B) + e(A \cap B) - e(A \setminus B) - e(B \setminus A)$$

for every A and B.

There are many other notions of quasirandomness that are equivalent in this sense. We mention just one.

$P_5(\epsilon)$: Let A be the adjacency matrix of G, and let $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n$ be its eigenvalues. Then $|\lambda_1 - \rho n| < \epsilon n$ and $\max(|\lambda_2|, |\lambda_n|) < \epsilon n$.

Quasirandom graphs were first studied by Thomason [Tho87] (under the name "jumbled graphs"), and the large number of equivalences between various quasirandomness properties was observed by Chung, Graham and Wilson [CGW89].

We take P_2 as our definition of quasirandomness. Thus a graph G on n vertices with $\rho\binom{n}{2}$ edges will be *ϵ-quasirandom* if

$$|e(A,B) - \rho|A||B|| < \epsilon|A||B|,$$

for all sets $A, B \subseteq V(G)$ with $|A|, |B| \geq \epsilon n$.

The basic philosophy of quasirandom graphs is that they resemble random graphs, provided we don't look too closely. For example, we expect most vertices to have degrees close to ρn, but it is too much to hope for that every vertex should have this property. Indeed, G could easily have a small number of isolated vertices, as quasirandomness is not sensitive to what happens on small parts of the graph.

Proposition 1. *Let $X \subseteq V(G)$ with $|X| \geq \epsilon n$ and let $Y = \{v \in V(G) : |e(v,X) - \rho|X|| \geq \epsilon|X|\}$. Then $|Y| < 2\epsilon n$.*

Proof. We have $Y = Y^+ \cup Y^-$ where

$$Y^+ = \{v \in V(G) : e(v,X) \geq \rho|X| + \epsilon|X|\},$$
$$Y^- = \{v \in V(G) : e(v,X) \leq \rho|X| - \epsilon|X|\},$$

hence

$$\left|e(X,Y^+) - \rho\,|X|\,\left|Y^+\right|\right| \geq \epsilon\,|X|\,\left|Y^+\right|,$$
$$\left|e(X,Y^-) - \rho\,|X|\,\left|Y^-\right|\right| \geq \epsilon\,|X|\,\left|Y^-\right|.$$

But then, since G is ϵ-quasirandom and $|X| \geq \epsilon n$, we must have $|Y^+|, |Y^-| < \epsilon n$. $\square$

In particular, taking $X = V(G)$ there are at least $(1-2\epsilon)n$ vertices v of G with $|d(v) - \rho n| \le \epsilon n$. We will call such vertices *balanced.*

Explicit constructions of quasirandom graphs are hard to come by. They are typically based on some algebraic structure; for example, for a prime q congruent to 1 modulo 4, the *Paley graph* G_q has vertex set the finite field $\mathbb{F}_q$ and edges between i and j whenever $i - j$ is a quadratic residue modulo q. Paley graphs are quasirandom.

It is more typical for quasirandomness to appear as a result of other processes. For example, Szemerédi's regularity lemma [Sze78] produces partitions of a graph such that most pairs of parts span a graph that satisfies a bipartite quasirandomness condition.

1.1 New graphs from old

By design, the random graph $G_{n,p}$ (in which edges appear independently with probability p) is quasirandom with high probability (that is, with probability tending to 1 as $n \to \infty$). More generally, given a quasirandom graph G we can, with high probability, obtain a new quasirandom graph $G_{edge}(p)$ by retaining edges of G with some fixed probability p. (The random graph $G_{n,p}$ can be thought of as the result of applying this process to the complete graph K_n.)

Another way to choose a random set of edges from G is to take a random walk on G. A *random walk* W on G is a sequence of vertices $W_0, W_1, \ldots, W_l$ where W_0 is chosen from some initial distribution and W_{i+1} is selected uniformly from the neighbours of W_i, with all choices made independently. We want to obtain a dense set of edges, so we take the length l of W to be αn^2 for some constant $\alpha > 0$. Let $G_{walk}(\alpha)$ denote the random subgraph of G consisting of those edges traversed by W.

Question 2. *Is $G_{walk}(\alpha)$ quasirandom with high probability?*

This question was asked by Böttcher, Hladký, Piguet and Taraz [Hla12] in relation to their work on tree packing, but it is also a natural question in its own right.

Observe that the answer is positive if G is the complete graph K_n. Indeed, in this case $G_{walk}(\alpha)$ is very close to $G_{n,p}$ for some p. This is because the sequence $W_0, W_1, \ldots$ is very nearly a sequence of independent random vertices of G ('very nearly' because consecutive terms of the sequence are forced to be distinct). Then $W_0W_1, W_2W_3, \ldots$ and $W_1W_2, W_3W_4, \ldots$ are very nearly two sequences of independent random edges of G, so $G_{walk}(\alpha)$ is very close to a random subgraph of G.

Why might we expect $G_{walk}(\alpha)$ to be quasirandom in general?

1. The graph G is approximately regular, of degree ρn, so the equilibrium distribution of W is approximately uniform. If the random walk W mixes rapidly, then most sequence terms will also be uniformly distributed and W will visit each vertex around αn times.

2. At each visit to a vertex v, W picks up a random edge leaving v, so W collects a random set of αn edges, chosen with replacement, from the ρn edges at v. Taking the union of these sets should then give a random subgraph of G.

The main obstacle to turning this outline into a proof is related to what exactly we mean by saying W 'mixes rapidly'. Since quasirandomness does not say anything about small parts of the graph, G might have small configurations of low degree vertices that can trap the random walk for long periods of time. We handle this possibility in two different ways.

The first way is to sidestep it by only considering graphs with a linear bound on their minimum degree. This allows us to get a strong result which applies in some common situations, such as when G is regular.

Theorem 3. *Let $\alpha, \epsilon, \rho, \eta > 0$ with $\eta > \epsilon$ and let $\gamma = C\epsilon^{1/4}$ for some absolute constant $C > 0$. Let G be an ϵ-quasirandom graph on n vertices with $\rho\binom{n}{2}$ edges and minimum degree at least γn, and let W be a random walk on G of length αn^2. Then, with probability $1 - o(1)$, the graph $G_{walk}(\alpha)$ is η-quasirandom with $(1 - e^{-2\alpha/\rho} + o(1))\rho\binom{n}{2}$ edges.*

This result cannot hold if we remove the condition on the minimum degree. Indeed, let G be the disjoint union of a small clique on $\epsilon^2 n/2$ vertices and a large clique on $(1-\epsilon^2/2)n$ vertices, and start W from a vertex of G selected uniformly at random. Then we do not even have concentration of the density of $G_{walk}(\alpha)$, as with positive probability (depending on ϵ but not on n) W will start in the small clique and so remain trapped there for all time. (We describe a similar, connected example in more detail at the start of Section 4.) So we must necessarily allow our failure probability to depend on ϵ as well as n. Write $o_\epsilon(1)$ for a quantity which is less than $f(\epsilon)$ for n sufficiently large, for some $f(\epsilon)$ tending to zero with ϵ.

Theorem 4. *Given $\alpha, \rho, \eta > 0$ there exists $\epsilon > 0$ such that the following holds. Let G be an ϵ-quasirandom graph on n vertices with $\rho\binom{n}{2}$ edges, and let W be a random walk on G of length αn^2 starting at any vertex W_0 of G with $(\rho-\epsilon)n \leq d(W_0) \leq (\rho+\epsilon)n$. Then, with probability $1-o_\epsilon(1)$, the graph $G_{walk}(\alpha)$ is η-quasirandom with $(1-e^{-2\alpha/\rho}+o_\epsilon(1))\rho\binom{n}{2}$ edges.*

In Section 2 we define an explicit model for our random walks and show that Step 2 works straightforwardly in this setting. In Section 3 we carry out Step 1 for the case where we have a bound on the minimum degree of G. This proves Theorem 3, and illustrates why we get the weaker conclusion in Theorem 4. In Section 4 we use a more elaborate argument to perform Step 1 in the general case, proving Theorem 4.

The problems considered in this chapter were suggested by Böttcher, Hladký, Piguet and Taraz [Hla12] after they encountered similar problems in connection with their work on tree packing. Suppose that we are trying to pack many trees into a copy of K_n. One approach is to embed some of the trees randomly. If we succeed in packing a small number of trees, then it would be good to know that the subgraph consisting of unused edges has nice enough properties that we can iterate the argument and therefore pack a much larger number of trees. If H is a subgraph of G, and both graphs are quasirandom, then $G-H$ is also quasirandom. So it would be useful to have a result like Theorem 4, but for random images of trees rather than paths. We consider such a generalisation in Section 5.

Since we will only prove asymptotic results we make a number of simplifying assumptions. We assume that ϵ is sufficiently small compared to the other parameters, and are only interested in statements for n sufficiently large. We omit notation indicating the taking of integer parts, and ignore questions of divisibility when breaking walks into pieces of a given size.

2 The list model

We now define a third model of a random subgraph to act as a staging post between $G_{walk}(\alpha)$ and $G_{edge}(p)$. The subgraph $G_{list}(\nu)$ of G is obtained by selecting $\nu d(v)$ edges at each vertex v of G to be retained, with all choices made independently and with replacement. We give a rather elaborate formal definition in order to introduce some ideas which will be useful later.

For each $v \in V(G)$, let L_v be an infinite list of uniform selections from the neighbourhood of v, with all choices made independently. The entry u on the list L_v corresponds naturally to the edge uv of G, and we define

$$G_{list}(\nu) = \bigcup_{v \in V(G)} \{uv : u \text{ appears in the first } \nu d(v) \text{ entries of } L_v\}.$$

In this section we will show that $G_{list}(\nu)$ is very close to $G_{edge}(p)$ for some p, in the sense that large subgraphs have similar densities in each model. It will then follow that $G_{list}(\nu)$ is quasirandom with high probability.

We first calculate the expected density of $G_{list}(\nu)$ in G. Throughout this section G will be an ϵ-quasirandom graph on n vertices with $\rho\binom{n}{2}$ edges and minimum degree at least γn. The corresponding arguments for the general case are similar and appear in Section 4.

Lemma 5. *For all $A, B \subseteq V(G)$,*

$$\mathbb{E}\left(e_{G_{list}(\nu)}(A,B)\right) = (1 - e^{-2\nu} + o(1))e_G(A,B).$$

Proof. The edge uv of G appears in $G_{list}(\nu)$ if and only if u appears in the first $\nu d(v)$ entries of L_v, or v appears in the first $\nu d(u)$ elements of L_u. The probability of this occurring is

$$1 - (1 - 1/d(v))^{\nu d(v)}(1 - 1/d(u))^{\nu d(u)} = 1 - e^{-2\nu} + o(1),$$

since $d(v), d(u) \geq \gamma n$. Hence

$$\mathbb{E}\left(e_{G_{list}(\nu)}(A, B)\right) = (1 - e^{-2\nu} + o(1))e_G(A, B). \qquad \square$$

To show that the number of edges retained in any subgraph is close to its expectation we use Talagrand's concentration inequality [Tal95]. In its usual form Talagrand's inequality is asymmetric and bounds a random variable in terms of its median. We use the following symmetric version (see [MR02, Chapter 20]) that gives concentration of the random variable about its mean.

Theorem 6 (Talagrand's inequality)**.** *Let $\Omega = \prod_{i=1}^{N} \Omega_i$ be a product of probability spaces with the product measure. Let X be a random variable on Ω such that*

(i) there is a constant $c > 0$ such that $|X(\omega) - X(\omega')| \leq c$ whenever ω and ω' differ on only a single coordinate;

(ii) whenever $X(\omega) \geq r$ there is a set $I \subseteq \{1, \ldots, N\}$ with $|I| = r$ such that $X(\omega') \geq r$ for all $\omega' \in \Omega$ with $\omega'_i = \omega_i$ for all $i \in I$.

Then, for $0 \leq s \leq \mathbb{E}(X)$,

$$\mathbb{P}\left(|X - \mathbb{E}(X)| \geq s + 60c\sqrt{\mathbb{E}(X)}\right) \leq 4e^{-s^2/8c^2\mathbb{E}(X)}.$$

Lemma 7. *Let $\nu > 0$. Then, with probability $1 - o(1)$,*

$$e_{G_{list}(\nu)}(A, B) = (1 - e^{-2\nu} + o(1))e_G(A, B),$$

for all $A, B \subseteq V(G)$ with $|A|, |B| \geq \epsilon n$.

Proof. We apply Theorem 6 to the space $\Omega = \prod_{v \in V(G)} \prod_{i=1}^{\nu d(v)} N(v)$, where each neighbourhood has the uniform probability measure; we can view Ω as the space of choices for the first $\nu d(v)$ entries of each list L_v. For $A, B \subseteq V(G)$ with $|A|, |B| \geq \epsilon n$, let $X_{A,B} = e_{G_{list}(\nu)}(A, B)$. It is easy to see that $X_{A,B}$ satisfies the conditions of Talagrand's inequality. Indeed, (i) holds since changing a list entry can change $X_{A,B}$ by at most $c = 2$. Furthermore, (ii) holds since, if $X_{A,B} \geq s$, then there are s list entries witnessing this fact. Therefore, by Theorem 6, for $120\sqrt{\mathbb{E}\left(X_{A,B}\right)} \leq t \leq \mathbb{E}\left(X_{A,B}\right)$ we have

$$\mathbb{P}\left(\left|X_{A,B} - \mathbb{E}\left(X_{A,B}\right)\right| \geq 2t\right) \leq 4e^{-t^2/32\mathbb{E}\left(X_{A,B}\right)}.$$

By Lemma 5 we have $\mathbb{E}\left(X_{A,B}\right) = (1-e^{-2\nu}+o(1))e_G(A,B)$. Since $e_G(A,B) \geq (\rho-\epsilon)\epsilon^2 n^2$, taking $t = C'\sqrt{n\mathbb{E}\left(X_{A,B}\right)}$ $(= o(\mathbb{E}\left(X_{A,B}\right)))$ for large enough $C' > 0$ gives that

$$\mathbb{P}\left(\left|X_{A,B} - (1-e^{-2\nu}+o(1))e_G(A,B)\right| \geq 2t\right) \leq 8^{-n}.$$

But there are at most 2^n choices for A and 2^n choices for B. Therefore, with probability at least $1-2^{-n}$, we have that $X_{A,B} = (1-e^{-2\nu}+o(1))e_G(A,B)$ for all $A, B \subseteq V(G)$ with $|A|, |B| \geq \epsilon n$. □

This is enough to ensure that $G_{list}(\nu)$ is quasirandom with high probability.

Theorem 8. *Let $\nu > 0$. Then, with probability $1-o(1)$, $G_{list}(\nu)$ is ϵ-quasirandom.*

Proof. By Lemma 7, with probability $1-o(1)$,

$$|e_{G_{list}(\nu)}(A,B) - (1-e^{-2\nu})e_G(A,B)| = o(1)e_G(A,B),$$

for all A, B with $|A|, |B| \geq \epsilon n$. By the definition of quasirandomness, we also have that

$$|e_G(A,B) - \rho|A||B|| < \epsilon|A||B|.$$

So by the triangle inequality,

$$\begin{aligned}|e_{G_{list}(\nu)}(A,B) - (1-e^{-2\nu})\rho|A||B|| &< o(1)e_G(A,B) + (1-e^{-2\nu})\epsilon|A||B| \\ &\leq (o(1) + (1-e^{-2\nu})\epsilon)|A||B| \\ &< \epsilon|A||B|,\end{aligned}$$

for n sufficiently large. □

Having shown that $G_{list}(\nu)$ is quasirandom with high probability, it suffices to show that $G_{walk}(\alpha)$ is close to $G_{list}(\nu)$ for some ν. The construction of the random walk W requires, at each visit to a vertex v, a choice of a random neighbour of v. We obtain a coupling of $G_{walk}(\alpha)$ and $G_{list}(\nu)$ by, at the jth visit to v, taking this choice to be the jth entry of the list L_v. Then $G_{walk}(\alpha)$ and $G_{list}(\nu)$ both consist of the edges corresponding to some initial segments of the lists L_v, and it is enough to show that we can choose ν such that the lengths of these initial segments are similar: that is, that the number of times the random walk W visits each vertex of G is roughly proportional to its degree.

We give two arguments. The first, appearing in Section 3, applies when we have a good lower bound on the minimum degree of G. The second, appearing in Section 4, applies to a general quasirandom graph G, but necessarily gives a weaker result. We include the argument for the special case where G has large minimum degree for two reasons: it is the more natural argument and, when it applies, it shows that there is essentially no loss of quasirandomness when we pass from G to $G_{walk}(\alpha)$, which could be useful for some applications.

3 Bounded minimum degree

To begin this section we recall some useful facts. A random walk W on a graph G is a Markov chain with transition matrix P given by

$$P_{uv} = \begin{cases} 1/d(u) & \text{if } uv \in E(G); \\ 0 & \text{if } uv \notin E(G). \end{cases}$$

Thus P is a normalised version of the adjacency matrix A, where each row has been scaled by the degree of the corresponding vertex. The eigenvalues of P are all real; let these be $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n$ and write $\lambda = \max(|\lambda_2|, |\lambda_n|)$. The first eigenvalue λ_1 of P is always equal to 1 and has a corresponding eigenvector $\pi = (\pi_v)$ given by $\pi_v = \frac{d(v)}{2e(G)}$. This vector π is called the *stationary distribution* of the walk W. If G is connected and non-bipartite then, for any initial distribution of W_0, the distribution of W_i converges to π as

$i \to \infty$ (i.e. $\mathbb{P}(W_i = v) \to \pi_v$ as $i \to \infty$ for each v). The following standard result, which can read out of Jerrum and Sinclair [JS89], gives control on the *rate* of this convergence.

Lemma 9. *For any graph G on n vertices with minimum degree at least γn, and any initial distribution on W_0, we have*

$$\max_{v \in V(G)} |\mathbb{P}(W_i = v) - \pi_v| \leq c_\gamma \lambda^i,$$

for some c_γ depending on γ.

If G is ϵ-quasirandom then property P_5 tells us that all but the first eigenvalue of the adjacency matrix A are small. If G is regular then, since the transition matrix P is a scalar multiple of A, λ will also be small. For a general ϵ-quasirandom graph this need not be true: for example, if G contains a small connected component, then $\lambda = 1$ (the 1-eigenspace is spanned by the stationary distributions of each connected component of G). Similarly, λ can be very close to 1 if there is a small set of vertices that is only weakly connected to the rest of the graph. However, a lower bound on the minimum degree of G is enough to recover an upper bound on λ.

Lemma 10. *Let G be an ϵ-quasirandom graph on n vertices with $\rho\binom{n}{2}$ edges and minimum degree at least γn, where $\gamma \geq C\epsilon^{1/4}$ for some absolute constant $C > 0$. Then, for n sufficiently large, $\lambda \leq 1/2$.*

Proof of Lemma 10. The proof follows the standard argument for showing that property P_4 implies property P_5. We first estimate the number of labelled copies of C_4 in G, and then evaluate the trace of P^4 in two different ways. Note that the implicit constants in our use of $O(\cdot)$ notation here are absolute.

The number of labelled copies of C_4 in G is

$$\begin{aligned}
C_4(G) &= 2 \sum_{u\in V(G)} \sum_{v\in V(G)} \binom{|N(u)\cap N(v)|}{2} \\
&= 2\cdot(1+O(\epsilon))n\cdot(1+O(\epsilon))n\cdot\binom{(\rho+O(\epsilon))^2 n}{2} + O(\epsilon)n^2\binom{n}{2} \\
&= (\rho+O(\epsilon))^4 n^4 + O(\epsilon)n^4 \\
&= \big(1+O\big(\epsilon/\rho^4\big)\big)\,\rho^4 n^4.
\end{aligned}$$

where the main term here accounts for balanced vertices u and v with close to $\rho^2 n$ common neighbours, and the error term bounds the contribution to the sum from each other pair by $\binom{n}{2}$. (The number of these pairs is small by Proposition 1).

Now the trace of P^4 is a weighted sum of the closed walks of length 4 in G, where the weight of the closed walk $uvwx$ is $1/(d(u)d(v)d(w)d(x))$. Thus

$$\begin{aligned}
\sum_{v\in V(G)} (P^4)_{vv} &= \frac{(1+O(\epsilon/\rho^4))\rho^4 n^4}{((\rho+O(\epsilon))n)^4} + \frac{O(\epsilon)n^4}{(\gamma n)^4} + \frac{O(n^3)}{(\gamma n)^4} \\
&= 1 + O\big(\epsilon/\rho^4\big) + O\big(\epsilon/\gamma^4\big) + O\big(1/(\gamma^4 n)\big),
\end{aligned}$$

where the main term counts the contribution from 4-cycles containing only balanced vertices and the error terms account for the contributions from 4-cycles with at least one unbalanced vertex and from closed walks of length 4 which are not 4-cycles respectively. (The lower bound on the minimum degree of G gives an upper bound of $1/(\gamma n)^4$ for the weight of any one walk.) But we also have

$$\sum_{v\in V(G)} (P^4)_{vv} = \sum_{i=1}^{n} \lambda_i^4 = 1 + \sum_{i=2}^{n} \lambda_i^4,$$

from which it follows that

$$\lambda^4 \le \sum_{i=2}^{n} \lambda_i^4 = O\big(\epsilon/\rho^4\big) + O\big(\epsilon/\gamma^4\big) + O\big(1/(\gamma^4 n)\big) \le 1/16,$$

for $\rho \geq \gamma \geq C\epsilon^{1/4}$ and n sufficiently large. □

For a finite probability space Ω, the *total variation distance* between two probability measures μ_1 and μ_2 is defined by

$$d_{TV}(\mu_1, \mu_2) = \frac{1}{2} \sum_{\omega \in \Omega} |\mu_1(\omega) - \mu_2(\omega)|.$$

This is the amount of probability mass that would have to be moved to turn one distribution into the other. Lemmas 9 with 10 tell us that the total variation distance between W_t and a vertex sampled from the stationary distribution is small when t is a little larger than $\log n$. We can obtain much more by breaking W into pieces with large gaps.

Let $L = (\log n)^2$, and let $K = \alpha n^2/L$. Given $i < L$, let $W^{(i)}$ denote the subsequence of W obtained by starting from W_i and taking L steps at a time: that is, $W^{(i)} = (W_1^{(i)}, \ldots, W_K^{(i)})$ where $W_j^{(i)} = W_{i+jL}$ for all $j < K$. For each $v \in V(G)$, let $X_v^{(i)}$ be the random variable which counts the number of times $W^{(i)}$ visits v. Our next lemma shows that, with high probability, $X_v^{(i)}$ takes a value close to what we expect.

Lemma 11. *Let G be a graph satisfying the conditions of Lemma 10 and let $v \in V(G)$. Then we have*

$$\mathbb{P}\left(\left|X_v^{(i)} - K\pi_v\right| \geq \sqrt{\frac{8 \log n}{K\pi_v}} K\pi_v \right) = O\big(n^{-3}\big).$$

Proof. Let $\mu = \pi^K$ be the K-fold product measure of π on $V(G)^K$; that is, $\mu(w) = \prod_{i=1}^K \pi_{w_i}$ for $w \in V(G)^K$. By Lemma 9 and Lemma 10, we have

$$\begin{aligned}
\mathbb{P}(W^{(i)} = w) &= \mathbb{P}(W_1^{(i)} = w_1)\mathbb{P}(W_2^{(i)} = w_2 | W_1^{(i)} = w_1) \cdots \mathbb{P}(W_K^{(i)} = w_K | W_{K-1}^{(i)} = w_{K-1}) \\
&= \left(\pi_{w_1} + O\Big(2^{-(\log n)^2}\Big)\right) \left(\pi_{w_2} + O\Big(2^{-(\log n)^2}\Big)\right) \cdots \left(\pi_{w_K} + O\Big(2^{-(\log n)^2}\Big)\right) \\
&= \left(\pi_{w_1} + O\big(n^{-6}\big)\right) \left(\pi_{w_2} + O\big(n^{-6}\big)\right) \cdots \left(\pi_{w_K} + O\big(n^{-6}\big)\right) \\
&= \left(1 + O\big(n^{-3}\big)\right) \mu(w),
\end{aligned}$$

since $\frac{\gamma}{\rho n} \leq \pi_v \leq \frac{1}{\rho n}$ for all v and $K = O(n^2)$. Summing over all w gives that

$$d_{TV}(\mathbb{P}, \mu) = O\big(n^{-3}\big),$$

where $\mathbb{P}$ is the measure on $V(G)^K$ induced by $W^{(i)}$. Now let

$$A = \left\{ w \in V(G)^K : |X_v^{(i)}(w) - K\pi_v| \geq \sqrt{2 \log nK\pi_v} \right\}.$$

By Chernoff's inequality (see [AS08, A.1.11 and A.1.13]),

$$\mu(A) \leq 2e^{-(4+o(1)) \log n} = O\big(n^{-3}\big).$$

Since $\mathbb{P}(A) \leq \mu(A) + d_{TV}(\mathbb{P}, \mu)$, the result follows. □

Now let $X_v = \sum_{i=0}^{L-1} X_v^{(i)}$ be the number of visits W makes to vertex v. Observing that $LK\pi_v = \alpha n^2 \pi_v = (1 + \frac{1}{n-1})\frac{\alpha}{\rho} d(v)$, we obtain the following corollary by summing over i and v.

Corollary 12. *Let $\alpha, \epsilon, \rho, \gamma > 0$ with $\rho \geq \gamma \geq C\epsilon^{1/4}$ for some absolute constant $C > 0$. Let G be an ϵ-quasirandom graph on n vertices with $\rho\binom{n}{2}$ edges and minimum degree at least γn, and let W be a random walk on G of length αn^2. Then*

$$\mathbb{P}\left(\left| X_v - \frac{\alpha}{\rho} d(v) \right| \geq \sqrt{\frac{8 \log n}{K\pi_v}} \frac{\alpha}{\rho} d(v) \text{ for some } v \right) = O\big(n^{-1}\big). \quad \square$$

Hence, with high probability, the number of visits W makes to each $v \in V(G)$ is $\left(\frac{\alpha}{\rho} + o(1)\right) d(v)$. We can now complete the proof of Theorem 3.

Proof of Theorem 3. By Corollary 12, we have that, with probability $1-o(1)$,

$$G_{list}(\alpha/\rho - o(1)) \subseteq G_{walk}(\alpha) \subseteq G_{list}(\alpha/\rho + o(1)).$$

From the proof of Theorem 8, we have that, with probability $1 - o(1)$,

$$|e_{G_{walk}(\alpha)}(A, B) - (1 - e^{-2\frac{\alpha}{\rho}})\rho|A||B|| < (1 - e^{-2\frac{\alpha}{\rho}} + o(1))\epsilon|A||B|,$$

for all $A, B \subseteq V(G)$ with $|A|, |B| \geq \epsilon n$. Since $1 - e^{-2\frac{\alpha}{\rho}} < 1$, $G_{walk}(\alpha)$ is ϵ-quasirandom with probability $1 - o(1)$. □

4 General case

We now move to the case of a general ϵ-quasirandom graph G with edge density ρ. Such G must always contain a connected component of order at least $(1-\epsilon)n$ (as otherwise we can find two sets of at least ϵn vertices with no edges between them), so by restricting our walk to this component (and increasing ϵ slightly) we may assume that G is connected.

The extra difficulty in the general case is that there might be small sets of vertices that are only weakly connected to the rest of the graph in which the random walk can get stuck. For example, let G consist of a small clique on $\epsilon^2 n/2$ vertices joined to a large clique on $(1-\epsilon^2/2)n$ vertices by a single edge. Then G is ϵ-quasirandom but it is not even true that the number of edges in $G_{walk}(\alpha)$ is concentrated near some value. Indeed, the worst case has $|A| = |B| = \epsilon n$ with all $\epsilon^2 n/2$ vertices of the small clique in A. Then G has density $\rho \approx 1$ and

$$\begin{aligned} e(A, B) &\approx (\epsilon - \epsilon^2/2)n \times \epsilon n \\ &= (1 - \epsilon/2)\epsilon^2 n^2 \\ &= (1 - \epsilon/2)|A||B|, \end{aligned}$$

so

$$|e(A, B) - |A||B|| < \epsilon |A||B|,$$

and G is ϵ-quasirandom.

Now suppose that we start our random walk in the large clique. If we remain in the large clique for all time then we will visit the end of the bridge to the small clique around αn times. The probability that we do not cross on any of these occasions is approximately

$$\left(1 - \frac{1}{n}\right)^{\alpha n} \approx e^{-\alpha}.$$

Similarly, the probability that we do not cross to the small clique in the first ϵn^2 steps is approximately

$$\left(1 - \frac{1}{n}\right)^{\epsilon n} \approx e^{-\epsilon},$$

so with probability $1 - e^{-\epsilon} \approx \epsilon$ we will cross to the small clique within the first ϵn^2 steps. The probability that we then remain there for all time is around

$$\left(1 - \frac{1}{\epsilon^2 n/2}\right)^{\frac{\alpha n^2}{\epsilon^2 n/2}} \approx e^{-\frac{4\alpha}{\epsilon^4}}.$$

So with positive probability we remain in the large clique for all time, but also with positive probability (depending on ϵ but not on n) we will quickly cross to the small clique and then remain there.

So for general quasirandom graphs we cannot hope for as strong a result as Theorem 3, and our assertions about high probability will necessarily depend on ϵ as well as n. In this section we use 'with high probability' to mean 'with probability $1 - o_\epsilon(1)$', with $o_\epsilon(1)$ small (depending on ϵ) for large n as defined in Section 1.

Our main task in this section is to find a weaker replacement for Corollary 12 in Section 3. Instead of saying that the random walk visits every vertex v around $\frac{\alpha}{\rho}d(v)$ times, we ask instead that the random walk visits *most* vertices of G around $\frac{\alpha}{\rho}d(v)$ times. Recall that we call a vertex v *balanced* if $|d(v) - \rho n| \leq \epsilon n$. We will show that, if W is a random walk of length αn^2 on G with W_0 balanced, then, with high probability, W hits most vertices of G about the right number of times. We can then use analogues of the results in Section 2 to prove Theorem 4 in the same way that Theorem 3 was deduced from Corollary 12.

Our first lemma gives a lower bound on the probability that a given step of a random walk W is in a set $S \subseteq V(G)$. Write $\mathbf{1}_X$ for the indicator function of a set X and $\mathbf{1}_v$ for the indicator function of the set $\{v\}$. Note that if the initial distribution for W_0 is π then $\mathbb{P}(W_i \in S) = \sum_{v \in S} \pi_v = \pi \cdot \mathbf{1}_S$ for any set $S \subseteq V(G)$ when $i \geq 0$. The next result shows that this is still almost true if W starts from a balanced vertex, S is large and $i \geq 2$.

Lemma 13. *Let G be a connected ϵ-quasirandom graph on n vertices with $\rho\binom{n}{2}$ edges, let v be a balanced vertex, and let $S \subseteq V(G)$. Then, for a random walk W starting at v, we have*

$$\mathbb{P}\left(W_i \in S\right) \geq \pi \cdot \mathbf{1}_S - 8\sqrt{\epsilon}/\rho \geq |S|/n - 9\sqrt{\epsilon}/\rho,$$

for $i \geq 2$ and n sufficiently large.

Proof. We first show that the random walk is quite well mixed after only two steps. Let A be the set of neighbours of v with degree at most $(\rho+\epsilon)n$ and B be the set of vertices with at least $(\rho-\epsilon)|A|$ neighbours in A. Thus A and B are the 'well-behaved' first and second neighbourhoods of v. By ϵ-quasirandomness, $|A| \geq d(v) - \epsilon n \geq (\rho - 2\epsilon)n$ and $|B| \geq (1-\epsilon)n$. We have

$$\begin{aligned}\mathbf{1}_v P &= \frac{1}{d(v)}\mathbf{1}_{N(v)} \\ &\geq \frac{1}{(\rho+\epsilon)n}\mathbf{1}_A,\end{aligned}$$

where the inequality holds in each coordinate. For $x \in B$,

$$\begin{aligned}\left(\mathbf{1}_A P\right)_x &= \sum_{\substack{y\in A \\ xy\in E(G)}} \frac{1}{d(y)} \\ &\geq \frac{(\rho-\epsilon)(\rho-2\epsilon)n}{(\rho+\epsilon)n} \\ &\geq \rho(1-4\epsilon/\rho),\end{aligned}$$

where the first inequality holds because each $y \in A$ has degree at most $(\rho+\epsilon)n$, x has at least $(\rho-\epsilon)|A|$ neighbours in A and $|A| \geq (\rho-2\epsilon)n$. Hence

$$\mathbf{1}_A P \geq \rho(1-4\epsilon/\rho)\mathbf{1}_B.$$

Since the entries of P are non-negative we can compose these inequalities to obtain

$$\mathbf{1}_v P^2 \geq \frac{(1-5\epsilon/\rho)}{n}\mathbf{1}_B.$$

Let $\mathbf{b} = \frac{(1-5\epsilon/\rho)}{n}\mathbf{1}_B$. Since $\pi_x = \frac{d(x)}{2\rho\binom{n}{2}}$, if x is a balanced vertex, then $\frac{(1-\epsilon/\rho)}{n-1} \leq \pi_x \leq \frac{(1+\epsilon/\rho)}{n-1}$; if x is not balanced, then we only have the trivial bound $\pi_x \leq \frac{1}{\rho n}$. Since at most $2\epsilon n$ vertices are unbalanced and at most ϵn vertices are not in B,

$$\begin{aligned}\|\mathbf{b} - \pi\|_2 &\leq \left(n\left(\frac{7\epsilon}{\rho n}\right)^2 + 3\epsilon n\left(\frac{2}{\rho n}\right)^2\right)^{1/2} \\ &\leq \left(\frac{64\epsilon}{\rho^2 n}\right)^{1/2},\end{aligned}$$

where we increased the constants slightly to account for the change from $n-1$ to n in the denominator. Then, for $i \geq 2$,

$$\begin{aligned}\mathbb{P}\left(W_i \in S\right) &= \mathbf{1}_v P^i \mathbf{1}_S \\ &= \mathbf{1}_v P^2 \cdot P^{i-2}\mathbf{1}_S \\ &\geq \mathbf{b}P^{i-2}\mathbf{1}_S \\ &= \pi P^{i-2}\mathbf{1}_S + (\mathbf{b} - \pi)P^{i-2}\mathbf{1}_S.\end{aligned}$$

By Cauchy-Schwarz, and the fact that the eigenvalues of P are at most 1,

$$\begin{aligned}\|(\mathbf{b} - \pi)P^{i-2}\mathbf{1}_S\|_2 &\leq \|\mathbf{b} - \pi\|_2 \|\mathbf{1}_S\|_2 \\ &\leq \left(\frac{64\epsilon|S|}{\rho^2 n}\right)^{1/2} \\ &\leq 8\sqrt{\epsilon}/\rho,\end{aligned}$$

and so

$$\mathbb{P}\left(W_i \in S\right) \geq \pi \cdot \mathbf{1}_S - 8\sqrt{\epsilon}/\rho,$$

proving the first inequality. Since at least $|S|-2\epsilon n$ elements of S are balanced,

$$\begin{aligned}\pi \cdot \mathbf{1}_S &= \sum_{x\in S} \frac{d(x)}{2\rho\binom{n}{2}} \\ &\geq \frac{(|S|-2\epsilon n)(\rho-\epsilon)}{\rho n} \\ &\geq |S|/n - 2\epsilon - \epsilon/\rho \\ &\geq |S|/n - \sqrt{\epsilon}/\rho,\end{aligned}$$

which proves the second inequality. □

We now consider the following variant of the list model for constructing a random walk. Fix some small length L and let $K = \alpha n^2/L$. By a *block rooted at v* we mean a random walk of length L starting at v. For each vertex v, let Λ_v be an infinite list of blocks rooted at v. We can construct a random walk of length αn^2 as follows. Choose W_0 from the given initial distribution, and, at each stage $s = 1, \ldots, K$, let $W_{(s-1)L} \cdots W_{sL}$ be the first unused block rooted at $W_{(s-1)L}$. At the end of the construction we have examined K blocks in total from the top of the n lists. Let M be the set of blocks examined (equivalently, the *multiset* of roots of blocks used).

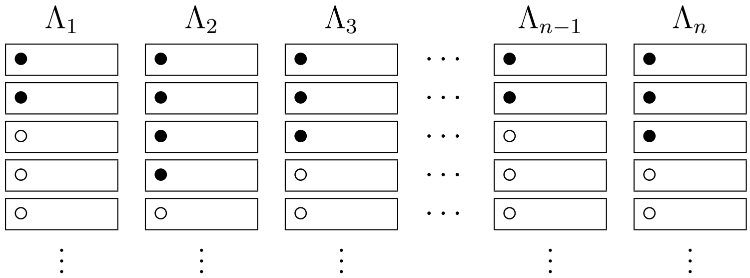

Figure 1: The construction examines K blocks from the top of the lists Λ_v, but we cannot tell in advance which blocks these will be.

This construction generalises the simple list model (which corresponds to the case $L = 1$), and we again hope to exploit the independence of blocks by applying standard concentration inequalities. There are two main obstacles. One is that we do not know anything about the distribution of a block rooted at a vertex v which is not balanced. We therefore first show that most of the root vertices are balanced. The second obstacle is that we do not know in advance which set of blocks we will examine. We handle this by approaching the problem from the other direction: for a given multiset M, what is the probability that the corresponding blocks do not contain an even distribution of the vertices? This turns out to be small enough that summing over all possible M gives the bound we require.

Lemma 14. *Let G be a connected ϵ-quasirandom graph on n vertices with $\rho\binom{n}{2}$ edges, and let W be a random walk of length αn^2 starting at a balanced vertex of G. Let $\delta = 3\sqrt[4]{\epsilon}/\sqrt{\rho}$ and suppose that n is sufficiently large. Then with probability at least $1-3\delta$ there exists a set $A \subseteq V(G)$, with $|A| \geq (1-\delta)n$, such that each vertex in A is hit at least $(1-4\delta)\alpha n$ times by W.*

Proof. Take $L = \omega_n$ for any $\omega_n \ll n/\log n$ which tends to infinity as $n \to \infty$, and let $K = \alpha n^2/L$. Construct a random walk W as described above and let $x_1, \ldots, x_K$ be the roots of the K blocks used. We first show that, with high probability, most of the vertices $x_1, \ldots, x_K$ are balanced.

Let U be the number of x_i that are unbalanced. By Lemma 13, for $i \geq 2$,

$$\begin{aligned}\mathbb{P}\left(x_i \text{ is unbalanced}\right) &\leq 1 - \big((1-2\epsilon) - \delta^2\big)\\ &\leq 2\delta^2,\end{aligned}$$

since there are at least $(1-2\epsilon)n$ balanced vertices and $\delta^2 > 2\epsilon$. By Markov's inequality,

$$\begin{aligned}\mathbb{P}\left(U \geq \delta K\right) &\leq \frac{\mathbb{E}\left(U\right)}{\delta K}\\ &\leq \frac{2\delta^2 K}{\delta K}\\ &= 2\delta.\end{aligned}$$

Now let M be a multiset of $(1-\delta)K$ balanced vertices and let $W^{(1)}$, $W^{(2)}, \ldots, W^{((1-\delta)K)}$ be the corresponding blocks. We will show that the probability that these blocks contain most balanced vertices about the right number of times is large.

Let $S \subseteq V(G)$ with $|S| = \delta n$. By Lemma 13, for every $1 \leq i \leq (1-\delta)K$ and every $j \geq 2$ we have $\mathbb{P}\left(W_j^{(i)} \in S\right) \geq \delta - \delta^2$. Let X_{ij} be the indicator of the event $W_j^{(i)} \in S$, let $X_j = \sum_{i=1}^{K} X_{ij}$ and let $X_{M,S} = \sum_{j=1}^{L} X_j$. For fixed j the X_{ij} are independent, so by Chernoff's inequality (see [AS08, Appendix A]),

$$\begin{aligned}\mathbb{P}\big(X_j < (\delta - 2\delta^2)|M|\big) &\leq e^{-2(\delta^2|M|)^2/|M|} \\ &= e^{-2\delta^4|M|}.\end{aligned}$$

Hence

$$\begin{aligned}\mathbb{P}\big(X_{M,S} < (\delta - 4\delta^2)\alpha n^2\big) &\leq \mathbb{P}\big(X_{M,S} < (\delta - 3\delta^2)(1-\delta)KL\big) \\ &= \mathbb{P}\big(X_{M,S} < (\delta - 3\delta^2)|M|L\big) \\ &\leq \mathbb{P}\big(X_j < (\delta - 2\delta^2)|M| \text{ for some } 2 \leq j \leq L\big) \\ &\leq Le^{-2\delta^4|M|},\end{aligned}$$

where the second inequality holds for large n because the contribution from X_1 is negligible as $L \to \infty$.

Let $A = \{x \in V(G) : x \text{ is visited at least } (1-4\delta)n \text{ times by } W\}$. If $|A| < (1-\delta)n$, then either δK of the x_i are unbalanced, or there is set S of δn vertices and a set of blocks M such that $X_{M,S} < (\delta - 4\delta^2)\alpha n^2$. Hence

$$\begin{aligned}\mathbb{P}\left(|A| < (1-\delta)n\right) &\leq \mathbb{P}\left(U \geq \delta K\right) + \sum_M \sum_S Le^{-2\delta^4|M|} \\ &\leq 2\delta + \binom{K+n-1}{n-1}\binom{n}{\delta n} Le^{-2\delta^4(1-\delta)K} \\ &\leq 2\delta + O(K)^n \cdot 2^n \cdot L \cdot e^{-2\delta^4(1-\delta)K} \\ &\leq 2\delta + \exp\big(O(n\log n) + O(n) + O(\log n) - 2\delta^4(1-\delta)K\big) \\ &\leq 3\delta,\end{aligned}$$

for n sufficiently large, since $K \gg n \log n$. □

We can now show that $G_{walk}(\alpha)$ is close to $G_{list}(\alpha/\rho)$.

Lemma 15. *Let G be a connected ϵ-quasirandom graph on n vertices with $\rho\binom{n}{2}$ edges, and let $\delta = 3\sqrt[4]{\epsilon}/\sqrt{\rho}$. Then, with probability $1 - o_\epsilon(1)$,*

$$|E(G_{walk}(\alpha)) \triangle E(G_{list}(\alpha/\rho))| = o_\epsilon(1)n^2.$$

Proof. For n sufficiently large, by Lemma 14 there is probability at least $1 - 3\delta$ that there is a set A of $(1-\delta)n$ vertices such that each vertex in A is visited at least $(1-4\delta)\alpha n$ times by the random walk. Since $\delta > 2\epsilon$, there is a subset $B \subseteq A$ of balanced vertices with $|B| \geq (1-2\delta)n$. This accounts for

$$(1-2\delta)n \cdot (1-4\delta)\alpha n \geq (1-6\delta)\alpha n^2$$

of the list entries examined, so $G_{walk}(\alpha)$ and $G_{list}(\alpha/\rho)$ can differ on at most $12\delta\alpha n^2$ edges. □

Hence to prove Theorem 4 it suffices to extend the results of Section 2 to the case were do not have any lower bound on the minimum degree of G. That turns out to be too much to ask, but the weaker statements that hold are still enough to prove Theorem 4.

The shape of the argument is the same as that of Section 2.

Lemma 16. *Let G be an ϵ-quasirandom graph on n vertices with $\rho\binom{n}{2}$ edges, and let $\nu > 0$. Let $A, B \subseteq V(G)$ with $|A|, |B| \geq \epsilon^{0.99}n$. Then*

$$\mathbb{E}\left(e_{G_{list}(\nu)}(A,B)\right) = (1 - e^{-2\nu} + o_\epsilon(1))e_G(A,B).$$

The exact lower bound on the sizes of A and B is unimportant; any value asymptotically larger than ϵ would work equally well.

Proof. Write $S = \{v \in V(G) : d(v) \geq (\rho - \epsilon)n\}$. Then

$$|e_G(V(G), S^c) - \rho n|S^c|| > \epsilon n|S^c|,$$

so, by ϵ-quasirandomness, $|S^c| < \epsilon n$.

The edge uv of G appears in $G_{list}(\nu)$ if and only if u appears in the first $\nu d(v)$ entries of L_v, or v appears in the first $\nu d(u)$ elements of L_u. If $u, v \in S$, then the probability of this occurring is

$$1-(1-1/d(v))^{\nu d(v)}(1-1/d(u))^{\nu d(u)} = 1-e^{-2\nu}+o(1),$$

since $d(v), d(u) \geq (\rho-\epsilon)n$. So

$$\begin{aligned}\mathbb{E}\left(e_{G_{list}(\nu)}(A,B)\right) &= (1-e^{-2\nu}+o(1))e_G(A,B)+O(\epsilon n(|A|+|B|)) \\ &= (1-e^{-2\nu}+o_\epsilon(1))e_G(A,B),\end{aligned}$$

since $e_G(A,B) \geq (\rho-\epsilon)|A||B|$ and $|A|,|B| \geq \epsilon^{0.99}n$. □

Lemma 17. *Let G be an ϵ-quasirandom graph on n vertices with $\rho\binom{n}{2}$ edges and let $\nu > 0$. Then with probability $1-o(1)$,*

$$e_{G_{list}(\nu)}(A,B) = (1-e^{-2\nu}+o_\epsilon(1))e_G(A,B),$$

for all $A, B \subseteq V(G)$ with $|A|,|B| \geq \epsilon^{0.99}n$.

Proof. We again apply Talagrand's inequality (Theorem 6) to the space $\Omega = \prod_{v\in V(G)}\prod_{i=1}^{\nu d(v)} N(v)$, where each neighbourhood has the uniform probability measure. For $A, B \subseteq V(G)$ with $|A|,|B| \geq \epsilon^{0.99}n$, let $X_{A,B} = e_{G_{list}(\nu)}(A,B)$. As before, $X_{A,B}$ satisfies the conditions of Talagrand's inequality so, for $120\sqrt{\mathbb{E}(X_{A,B})} \leq t \leq \mathbb{E}(X_{A,B})$, we have

$$\mathbb{P}(|X_{A,B}-\mathbb{E}(X_{A,B})| \geq 2t) \leq 4e^{-t^2/32\mathbb{E}(X_{A,B})}.$$

By Lemma 16 we have $\mathbb{E}(X_{A,B}) = (1-e^{-2\nu}+o_\epsilon(1))e_G(A,B)$. Since $e_G(A,B) \geq (\rho-\epsilon)\epsilon^{1.98}n^2$, taking $t = C'\sqrt{n\mathbb{E}(X_{A,B})}$ ($= o(\mathbb{E}(X_{A,B}))$) for large enough $C' > 0$ gives that

$$\mathbb{P}\left(\left|X_{A,B}-(1-e^{-2\nu}+o_\epsilon(1))e_G(A,B)\right| \geq 2t\right) \leq 8^{-n}.$$

But there are at most 2^n choices for A and 2^n choices for B. Therefore, with probability at least $1-2^{-n}$, we have that $X_{A,B} = (1-e^{-2\nu}+o_\epsilon(1))e_G(A,B)$, for all pairs (A,B) with $|A|,|B| \geq \epsilon^{0.99}n$. □

This is enough to ensure that $G_{list}(\nu)$ is quasirandom with high probability.

Theorem 18. *Let G be an ϵ-quasirandom graph on n vertices with $\rho\binom{n}{2}$ edges and let $\nu > 0$. Then, with probability $1 - o(1)$, $G_{list}(\nu)$ is $o_\epsilon(1)$-quasirandom.*

Proof. By Lemma 16, with probability $1 - o(1)$,

$$|e_{G_{list}(\nu)}(A, B) - (1 - e^{-2\nu})e_G(A, B)| = o_\epsilon(1)e_G(A, B),$$

for all A, B with $|A|, |B| \geq \epsilon^{0.99}n$. By the definition of quasirandomness, we also have that

$$|e_G(A, B) - \rho|A||B|| < \epsilon|A||B|.$$

So by the triangle inequality,

$$\begin{aligned} |e_{G_{list}(\nu)}(A, B) - (1 - e^{-2\nu})\rho|A||B|| &< o_\epsilon(1)e_G(A, B) + (1 - e^{-2\nu})\epsilon|A||B| \\ &\leq (o_\epsilon(1) + (1 - e^{-2\nu})\epsilon)|A||B| \\ &= o_\epsilon(1)|A||B|. \end{aligned}$$

Hence $G_{list}(\nu)$ is δ-quasirandom with $\delta = \max(\epsilon^{0.99}, o_\epsilon(1))$, where the $o_\epsilon(1)$ is taken from the last line. □

Taking ϵ sufficiently small completes the proof of Theorem 4.

5 Trees

A *homomorphism* from a graph H to a graph G is an edge-preserving map $\phi : V(H) \to V(G)$. A random walk can be viewed as a random homomorphism of a path; a natural generalisation is to consider a random homomorphism of some other tree T. (This is sometimes called a *tree-indexed random walk.*) Just as we traversed a path in one direction, our trees will be rooted and we think of them as directed 'downwards', away from the root. In this section we will explore to what extent the methods of Section 4 can be applied in this more general setting.

We generate a random homomorphism as follows. Enumerate the vertices of T as $v_0, v_1, \dots, v_k$ where, for each j, $T[v_0, \dots, v_j]$ is a connected subtree of T containing the root v_0. First choose $\phi(v_0)$ from a given initial distribution. Then, at each stage $j > 0$, let u be the parent of v_j in T and choose $\phi(v_j)$ uniformly at random from the neighbours of $\phi(u)$. All choices are made independently, and we can think of these choices as being taken from the lists L_v as before.

Suppose now that G is an ϵ-quasirandom graph on n vertices. Let ϕ be a random homomorphism from a tree T of size αn^2 to G, and let $G(T)$ be the subgraph of G consisting of the edges in the image of ϕ. Is $G(T)$ quasirandom with high probability? It is easy to see that in general the answer is no. For example, let $G = K_n$ and let T be an $n/2$-ary tree of depth 2 (here $\alpha = 1/4 + o(1)$). Then, with high probability, $\phi(T)$ contains a constant fraction of the edges of G. But all of these edges are incident on the neighbourhood of the root, which has only $(1 - e^{-1/2} + o(1))n$ vertices with high probability, so, with high probability, $G(T)$ is not quasirandom.

We seek conditions on T such that we can apply the approach taken in Section 4 with minimal changes. The condition we give here is an upper bound on the maximum degree of T.

We need an analogue of the block model for constructing a random walk. Instead of breaking our path into many short paths, we break our tree into many small edge-disjoint subtrees.

Lemma 19. *Let T be a rooted tree with N edges and let $L \leq N$. Then T can be written as an edge-disjoint union of rooted trees $R_1, \dots, R_K$, each of size between L and $3L$.*

Proof. Let v be a vertex of T furthest from the root such that v has at least L descendants. Then each branch of T lying below v has at most L edges, so some union of these branches has size between L and $2L$; let this be R_1. We obtain $R_2, \dots, R_K$ similarly until there are less than L edges of T remaining, which we add to R_K. □

Write $\mathcal{R} = \{R_1, \dots, R_K\}$ for the corresponding set of abstract rooted trees, up to isomorphism. In an abuse of notation we use R_i to refer to both the specific subtree of T and its isomorphism type.

It is convenient to number the R_i such that, for each j, $R_1 \cup \cdots \cup R_j$ is a subtree of T containing the root. We can then describe the block model for the construction of a random homomorphism as follows. For each $v \in V(G)$ and $R \in \mathcal{R}$, let $\Lambda_{v,R}$ be a list of independent random homomorphisms from R to G that map the root of R to v. Choose a vertex v_1 from the given distribution for the image of the root of T, and identify $\phi(R_1)$ with the first entry from Λ_{v_1,R_1}. (If R_1 has a non-trivial automorphism group then there is a choice of identification of R_1 with the reference copy in $\mathcal{R}$. The choice is unimportant provided the same choice is made every time.) Then at each stage j we have already determined the image v_j of the root of R_j, and we identify $\phi(R_j)$ with the first unused element of Λ_{v_j,R_j}.

Now let T be a rooted tree with αn^2 edges. As before we want to show that T 'visits' most vertices of G about the right number of times. We need to be careful here about what counts as a 'visit': what we want to count is the number of times an edge leaves a vertex, as that is the number of entries of the corresponding list that will be examined. So we say $\phi(T)$ *visits* $x \in V(G)$ whenever uv is an edge of T with u the parent of v and $\phi(u) = x$; the number of visits $\phi(T)$ makes to x is the number of edges uv for which this occurs.

There are three places where the argument in the proof of Lemma 14 needs modification or additional details need to be checked.

(i) In the path case the edges (or vertices) of the blocks had a natural order and the blocks were all the same size. In the tree case we are free to choose a labelling of the edges in each block, but the blocks might still have different sizes: when we look at the $2L$th edge from each block, are there enough blocks with $2L$ edges that Chernoff's inequality will give good concentration?

(ii) In the path case the set of list entries examined was parameterised by multisets of vertices of G. In the tree case the set of list entries examined is instead parameterised by multisets of pairs (v, R) with $v \in V(G)$ and $R \in \mathcal{R}$. So the factor $\binom{K+n-1}{n-1}$ in the final sum needs to be replaced by $\binom{K+n|\mathcal{R}|-1}{n|\mathcal{R}|-1}$, and we must restrict the size of $\mathcal{R}$ to prevent this becoming too large.

(iii) In the path case we had to ignore the first two vertices of each block as we needed to take two steps before we had good information about the distribution over vertices. This was safe because the ignored vertices were only a $o(1)$ fraction of the total number of vertices. In the tree case we must ignore the edges whose start point is the root of the block or is a child of the root. We need to ensure that the number of ignored edges is at most a small fraction of the total number of edges.

Problem (i) is avoided by throwing away the small number of edges that receive a label shared by few other edges. If we throw away all edges that receive a label which is used less that $\epsilon n^2/L^2$ times then the total number of edges thrown away is less than $3\epsilon n^2/L$ as there are at most $3L$ edges in each block.

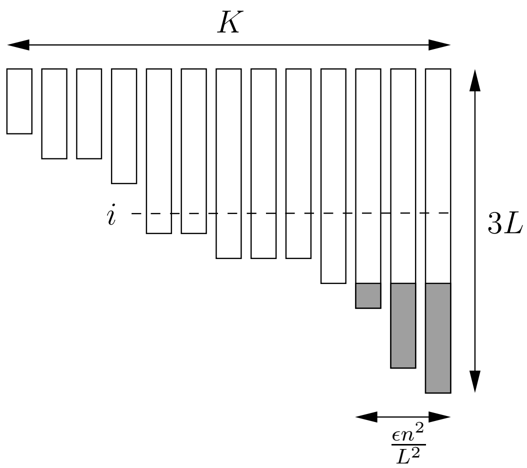

Figure 2: Deleting a $o(1)$ fraction of the edges ensures that the remaining labels i are each used in a large number of blocks.

Problem (ii) is avoided by taking L small: $L = \frac{\log n}{2\log 3}$ suffices. Indeed, since the number of rooted trees on L vertices is $O((2.9955\ldots)^L)$ (see [Ott48]) and $\frac{\alpha n^2}{3L} \leq K \leq \frac{\alpha n^2}{L}$, we have in this case that $n|\mathcal{R}| \ll n^{3/2} \ll K$, and

$$\binom{K+n|\mathcal{R}|-1}{n|\mathcal{R}|-1} \ll K^{n|\mathcal{R}|} \ll \exp\left(O(n^{3/2}\log n)\right),$$

which is small enough that it will not overpower the e^{-cK}-type decay.

Problem (iii) is avoided by having Δ^2, the square of the maximum degree of T small (depending on the desired level of quasirandomness) compared to L: so Δ can be as large as a small multiple of $\sqrt{\log n}$.

With these modifications to our earlier argument we obtain the following result.

Theorem 20. *Given $\alpha, \rho, \eta > 0$ there exists $\epsilon, c > 0$ such that the following holds. Let G be an ϵ-quasirandom graph on n vertices with $\rho\binom{n}{2}$ edges, let T be a rooted tree of size αn^2 with maximum degree $\Delta \leq c\sqrt{\log n}$, and let ϕ be a random homomorphism from T to G such that the image of the root is balanced. Then, with probability $1 - o_\epsilon(1)$, the subgraph $G(T)$ of G consisting of the edges of $\phi(T)$ is η-quasirandom with $(1 - e^{-2\alpha/\rho} + o_\epsilon(1))\rho\binom{n}{2}$ edges.*

It would be interesting to know how large $\Delta(T)$ can be taken in Theorem 20. By the example at the start of this section we must have $\Delta(T)$ small compared to n. Is this already enough?

The results of this chapter are joint work with Eoin Long.

Chapter 5

Balanced independent sets in the cube

The discrete hypercube $\mathcal{Q}_n$ is the graph with vertices the subsets of $[n]$ and edges between sets whose symmetric difference contains a single element. The cube $\mathcal{Q}_n$ is bipartite, with classes X_0 and X_1 consisting of the sets of even and odd size respectively. The maximum-sized independent sets in $\mathcal{Q}_n$ are precisely X_0 and X_1. Ramras [Ram10] asked: how large an independent set can we find with half its elements in X_0 and half in X_1? Call such an independent set *balanced*. The following result verifies the conjecture made by Ramras for the case where n is odd.

Theorem 1. *The largest balanced independent set in $\mathcal{Q}_n$ has size*

$$\begin{aligned} 2^{n-1} - 2\binom{n-2}{(n-2)/2} \qquad & \textit{if } n \textit{ is even,} \\ 2^{n-1} - \binom{n-1}{(n-1)/2} \qquad & \textit{if } n \textit{ is odd.} \end{aligned}$$

For a set A of vertices of $\mathcal{Q}_n$, write $N(A)$ for the set of vertices adjacent to an element of A. The maximal independent sets in $\mathcal{Q}_n$ all have the form $A \cup (X_1 \setminus N(A))$ for some $A \subseteq X_0$. So for a maximum-sized balanced independent set we seek the largest $A \subseteq X_0$ for which

$$|A| \leq |X_1 \setminus N(A)|.$$

We use the following isoperimetric theorem for even-sized sets, due independently to Bezrukov [Bez85] and Körner and Wei [KW84] (see also Tiersma [Tie85]). Recall that $x < y$ in the *simplicial order* on $\mathcal{Q}_n$ if either $|x| < |y|$, or $|x| = |y|$ and $x < y$ lexicographically.

Theorem 2 ([Bez85], [KW84]). *Let $A \subseteq X_0$, and let B be the initial segment of the simplicial order restricted to X_0 with $|B| = |A|$. Then $|N(B)| \leq |N(A)|$, and $X_1 \backslash N(B)$ is a terminal segment of the simplicial order restricted to X_1.*

Proof of Theorem 1. We will exhibit an initial segment A of the simplicial order restricted to X_0, and a terminal segment B of the simplicial order restricted to X_1, with $N(A) \cap B = \emptyset$ and $|A| = |B|$ as large as possible. It follows from Theorem 2 that $A \cup B$ will be a maximum-sized balanced independent set.

The form of A and B depends on the residue of n modulo 4. For $n = 4k$ we take

$$A = [n]^{(0)} \cup [n]^{(2)} \cup \cdots \cup [n]^{(2k-2)} \cup (12 + [3,n]^{(2k-2)})$$
$$B = (1 + [3,n]^{(2k)}) \cup [2,n]^{(2k+1)} \cup [n]^{(2k+3)} \cup \cdots \cup [n]^{(n-3)} \cup [n]^{(n-1)},$$

where, for instance,

$$12 + [3,n]^{(2k-2)} = \{\{1,2\} \cup x : x \subseteq \{3,4,\ldots,n\}, |x| = 2k-2\}\,.$$

For $n = 4k+1$ we take

$$A = [n]^{(0)} \cup [n]^{(2)} \cup \cdots \cup [n]^{(2k-2)} \cup (1 + [2,n]^{(2k-1)})$$
$$B = [2,n]^{(2k+1)} \cup [n]^{(2k+3)} \cup \cdots \cup [n]^{(n-2)} \cup [n]^{(n)}.$$

For $n = 4k+2$ we take

$$A = [n]^{(0)} \cup [n]^{(2)} \cup \cdots \cup [n]^{(2k-2)} \cup (1 + [2,n]^{(2k-1)}) \cup (2 + [3,n]^{(2k-1)})$$
$$B = [3,n]^{(2k+1)} \cup [n]^{(2k+3)} \cup \cdots \cup [n]^{(n-3)} \cup [n]^{(n-1)}.$$

Finally, for $n = 4k + 3$ we take

$$A = [n]^{(0)} \cup [n]^{(2)} \cup \cdots \cup [n]^{(2k)}$$
$$B = [n]^{(2k+3)} \cup \cdots \cup [n]^{(n-2)} \cup [n]^{(n)}.$$

Verifying that these sets have the claimed sizes, and that $|A| = |B|$ in each case, is a simple application of the identities $\binom{m}{r} = \binom{m-1}{r-1} + \binom{m-1}{r}$, $\binom{m}{r} = \binom{m}{m-r}$ and $\sum_{r=0}^{m} \binom{m}{r} = 2^m$. $\square$

The maximum-sized balanced independent sets constructed above are also maximal independent sets. For example, if $n = 4k + 3$, then any set not in the family is adjacent to a complete layer; the other cases are similar, with slight complications in the middle layers of the cube.

The results in this chapter have previously been published as [Bar12].

References

[AK97] Rudolf Ahlswede and Levon H. Khachatrian. The complete intersection theorem for systems of finite sets. *European J. Combin.*, 18(2):125–136, 1997.

[AS08] N. Alon and J. Spencer. *The probabilistic method*. John Wiley & Sons Inc., third edition, 2008.

[Bar12] Ben Barber. A note on balanced independent sets in the cube. *Australas. J. Combin.*, 52:205–207, 2012.

[Bar13] Ben Barber. Maximum hitting for n sufficiently large. *Graphs Combin.*, 2013.

[Bez85] S. L. Bezrukov. On the minimization of a neighborhood of subsets of Hamming space. In *Combinatorial-algebraic methods in applied mathematics (Russian)*, pages 45–58, 155. Gor′kov. Gos. Univ., Gorki, 1985.

[BHL13] Ben Barber, Neil Hindman, and Imre Leader. Partition regularity in the rationals. *J. Combin. Theory Ser. A*, 120:1590–1599, 2013.

[BL13] Ben Barber and Imre Leader. Partition regularity with congruence conditions. *J. Comb.*, 2013.

[Bon12] B. Bond. EKR sets for large n and r. 2012. arXiv:1210.7470, submitted.

[Bor11] Peter Borg. Maximum hitting of a set by compressed intersecting families. *Graphs Combin.*, 27(6):785–797, 2011.

[CGW89] F. R. K. Chung, R. L. Graham, and R. M. Wilson. Quasi-random graphs. *Combinatorica*, 9(4):345–362, 1989.

[Com77] W. W. Comfort. Ultrafilters: some old and some new results. *Bull. Amer. Math. Soc.*, 83(4):417–455, 1977.

[DH87] Walter Deuber and Neil Hindman. Partitions and sums of (m, p, c)-sets. *J. Combin. Theory Ser. A*, 45(2):300–302, 1987.

[DH09] Dibyendu De and Neil Hindman. Image partition regularity near zero. *Discrete Math.*, 309(10):3219–3232, 2009.

[DHLL95] Walter A. Deuber, Neil Hindman, Imre Leader, and Hanno Lefmann. Infinite partition regular matrices. *Combinatorica*, 15(3):333–355, 1995.

[EKR61] P. Erdős, Chao Ko, and R. Rado. Intersection theorems for systems of finite sets. *Quart. J. Math. Oxford Ser. (2)*, 12:313–320, 1961.

[Hin74] Neil Hindman. Finite sums from sequences within cells of a partition of N. *J. Combinatorial Theory Ser. A*, 17:1–11, 1974.

[HL93] Neil Hindman and Imre Leader. Image partition regularity of matrices. *Combin. Probab. Comput.*, 2(4):437–463, 1993.

[Hla12] J. Hladký. Problem session, Probabilistic methods in graph theory, University of Birmingham, 2012.

[HLS03a] Neil Hindman, Imre Leader, and Dona Strauss. Infinite partition regular matrices: solutions in central sets. *Trans. Amer. Math. Soc.*, 355(3):1213–1235, 2003.

[HLS03b] Neil Hindman, Imre Leader, and Dona Strauss. Open problems in partition regularity. *Combin. Probab. Comput.*, 12(5-6):571–583, 2003. Special issue on Ramsey theory.

[HLS06] Neil Hindman, Imre Leader, and Dona Strauss. Forbidden distances in the rationals and the reals. *J. London Math. Soc. (2)*, 73(2):273–286, 2006.

[JS89] M. Jerrum and A. Sinclair. Approximate counting, uniform generation and rapidly mixing Markov chains. *Inform. and Comput.*, 82(1):93–133, 1989.

[KW84] J. Körner and V. K. Wei. Odd and even Hamming spheres also have minimum boundary. *Discrete Math.*, 51(2):147–165, 1984.

[Mil75] Keith R. Milliken. Ramsey's theorem with sums or unions. *J. Combinatorial Theory Ser. A*, 18:276–290, 1975.

[MR02] M. Molloy and B. Reed. *Graph colouring and the probabilistic method*. Springer-Verlag, 2002.

[Ott48] Richard Otter. The number of trees. *Ann. of Math. (2)*, 49:583–599, 1948.

[Rad33] Richard Rado. Studien zur Kombinatorik. *Math. Z.*, 36(1):424–470, 1933.

[Ram30] F. P. Ramsey. On a problem of formal logic. *Proc. London Math. Soc.*, S2-30(1):264, 1930.

[Ram10] M. Ramras. Balanced independent sets in hypercubes. *Australas. J. Combin.*, 48:57–72, 2010.

[Sch16] Issai Schur. Über die Kongruenz $x^m + y^m \equiv z^m \pmod{p}$. *Jber. Deutschen Math. Verein.*, 25:114–117, 1916.

[ST83] C. L. Stewart and R. Tijdeman. On density-difference sets of sets of integers. In Paul Erdős, editor, *Studies in pure mathematics: To the memory of Paul Turán*, pages 701–710. Birkhäuser Verlag, Basel, 1983.

[Sze75] E. Szemerédi. On sets of integers containing no k elements in arithmetic progression. *Acta Arith.*, 27:199–245, 1975. Collection of articles in memory of Juriĭ Vladimirovič Linnik.

[Sze78] Endre Szemerédi. Regular partitions of graphs. In *Problèmes combinatoires et théorie des graphes (Colloq. Internat. CNRS, Univ. Orsay, Orsay, 1976)*, volume 260 of *Colloq. Internat. CNRS*, pages 399–401. CNRS, Paris, 1978.

[Tal95] M. Talagrand. Concentration of measure and isoperimetric inequalities in product spaces. *Inst. Hautes Études Sci. Publ. Math.*, (81):73–205, 1995.

[Tay76] Alan D. Taylor. A canonical partition relation for finite subsets of ω. *J. Combinatorial Theory Ser. A*, 21(2):137–146, 1976.

[Tho87] Andrew Thomason. Pseudorandom graphs. In *Random graphs '85 (Poznań, 1985)*, volume 144 of *North-Holland Math. Stud.*, pages 307–331. North-Holland, Amsterdam, 1987.

[Tie85] H. J. Tiersma. A note on Hamming spheres. *Discrete Math.*, 54(2):225–228, 1985.

[VdW21] B. L. Van der Waerden. Beweis einer baudetschen vermutung. *Nieuw Arch. Wisk.*, 15:212–216, 1921.