Extremal results for graphs and hypergraphs and other combinatorial problems

Barnabás Kristóf Janzer

University of Cambridge, Trinity College

May 2023

This thesis is submitted for the degree of
Doctor of Philosophy

Declaration

This thesis is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared here and specified in the text. It is not substantially the same as any work that has already been submitted before for any degree or other qualification.

Chapter 5 is based on joint work with Timothy Gowers.

Chapter 7 is based on joint work with Oliver Janzer.

Chapter 8 is based on joint work with Alberto Espuny Díaz, Gal Kronenberg and Joanna Lada.

Chapter 9 is based on joint work with J. Robert Johnson and Imre Leader.

Extremal results for graphs and hypergraphs and other combinatorial problems

Barnabás Kristóf Janzer

Abstract

In this dissertation we present several combinatorial results, primarily concerning extremal problems for graphs and hypergraphs, but also covering some additional topics.

In Chapter 2, we consider the following geometric problem of Croft. Let K be a convex body in $\mathbb{R}^d$ that contains a copy of another body S in every possible orientation. Is it always possible to continuously move any one copy of S into another, inside K? We prove that the answer is positive if S is a line segment, but, surprisingly, the answer is negative in dimensions at least four for general S.

In Chapter 3, we study the extremal number of tight cycles. Sós and Verstraëte raised the problem of finding the maximum possible size of an n-vertex r-uniform tight-cycle-free hypergraph. When $r = 2$ this is simply $n - 1$, and it was unknown whether the answer is $\Theta(n^{r-1})$ in general. We show that this is not the case for any $r \geq 3$ by constructing r-uniform hypergraphs with n vertices and $\Omega(n^{r-1} \log n/ \log \log n) = \omega(n^{r-1})$ edges which contain no tight cycles.

In Chapter 4, we study the following saturation question: how small can maximal k-wise intersecting set systems over $[n]$ be? Balogh, Chen, Hendrey, Lund, Luo, Tompkins and Tran resolved this problem for $k = 3$, and for general k showed that the answer is between $c_k \cdot 2^{n/(k-1)}$ and $d_k \cdot 2^{n/\lceil k/2\rceil}$. We prove that their lower bound gives the correct order of magnitude for all k.

In Chapter 5, we prove that for any r, s with $r < s$, there are n-vertex graphs containing $n^{r-o(1)}$ copies of K_s such that any K_r is contained in at most one K_s. This gives a natural generalisation of the Ruzsa–Szemerédi $(6, 3)$-problem. We also show that there are properly edge-coloured n-vertex graphs with $n^{r-1-o(1)}$ copies of K_r such that no K_r is rainbow, answering a question of Gerbner, Mészáros, Methuku and Palmer about generalised rainbow Turán numbers.

In Chapter 6, we continue the study of the generalised rainbow Turán problem: how many copies of H can a properly edge-coloured graph on n vertices contain if it has no rainbow copy of F? We determine the order of magnitude in essentially all cases when F is a cycle and H is a path or a cycle. In particular, we answer a question of Gerbner, Mészáros, Methuku and Palmer.

In Chapter 7, we consider the following problem. Let $g(n, H)$ be the smallest k such that we can assign a k-edge-colouring f_v of K_n to each vertex v in K_n with the property that for any copy H_0 of H in K_n, there is some $u \in V(H_0)$ such that H_0 is rainbow in f_u. This function was introduced by Alon and Ben-Eliezer, and we answer several of their questions. In particular, we determine all connected graphs H for which $g(n, H) = n^{o(1)}$, and show that for all $\varepsilon > 0$ there

exists r such that $g(n, K_r) = \Omega(n^{1-\varepsilon})$. We also prove a family of special cases of a conjecture of Conlon, Fox, Lee and Sudakov about the so-called hypergraph Erdős–Gyárfás function.

In Chapter 8, we study bootstrap percolation for hypergraphs. Consider the process in which, given a fixed r-uniform hypergraph H and starting with a given n-vertex r-uniform hypergraph G, at each step we add to G all edges that create a new copy of H. We are interested in maximising the number of steps that this process takes before it stabilises. For the case where $H = K_s^{(r)}$ with $s > r \geq 3$, we show that the number of steps of this process can be $\Theta(n^r)$. This answers a recent question of Noel and Ranganathan. We also demonstrate that different and interesting maximal running times can occur for other choices of H.

In Chapter 9, we study an extremal problem about permutations. How many random transpositions (meaning that we swap given pairs of elements with given probabilities) do we need to perform on a deck of cards to 'shuffle' it? We study several problems on this topic. Among other results, we show that at least $2n - 3$ such swaps are needed to uniformly shuffle the first two cards of the deck, proving a conjecture of Groenland, Johnston, Radcliffe and Scott.

In Chapter 10, we study the following extremal problem on set systems introduced by Holzman and Körner. We say that a pair $(\mathcal{A}, \mathcal{B})$ of families of subsets of an n-element set is cancellative if whenever $A, A' \in \mathcal{A}$ and $B \in \mathcal{B}$ satisfy $A \cup B = A' \cup B$, then $A = A'$, and whenever $A \in \mathcal{A}$ and $B, B' \in \mathcal{B}$ satisfy $A \cup B = A \cup B'$, then $B = B'$. Tolhuizen showed that there exist cancellative pairs with $|\mathcal{A}||\mathcal{B}|$ about 2.25^n, whereas Holzman and Körner proved an upper bound of 2.326^n. We improve the upper bound to about 2.268^n. This result also improved the then best known upper bound for a conjecture of Simonyi about 'recovering pairs' (the Boolean case of the 'sandglass conjecture'), although the upper bound for Simonyi's problem has since been further improved.

In Chapter 11 we study a continuous version of Sperner's theorem. Engel, Mitsis, Pelekis and Reiher showed that an antichain in the continuous cube $[0, 1]^n$ must have $(n - 1)$-dimensional Hausdorff measure at most n, and they conjectured that this bound can be attained. This was already known for $n = 2$, and we prove this conjecture for all n.

Chapter 12 has similar motivations to the preceding chapter. A subset A of $\mathbb{Z}^n$ is called a weak antichain if it does not contain two elements x and y satisfying $x_i < y_i$ for all i. Engel, Mitsis, Pelekis and Reiher showed that for any weak antichain A in $\mathbb{Z}^n$, the sum of the sizes of its $(n - 1)$-dimensional projections must be at least as large as its size $|A|$. They asked what the smallest possible value of the gap between these two quantities is in terms of $|A|$. We give an explicit weak antichain attaining the minimum for each possible value of $|A|$.

Finally, in Chapter 13, we study the following problem. Esperet, Gimbel and King introduced the orientation covering number of a graph G as the smallest k with the property that we can choose k orientations of G such that whenever x, y, z are vertices of G with $xy, xz \in E(G)$, then there is a chosen orientation in which both xy and xz are oriented away from x. We prove that the orientation covering number of G is the same as that of $K_{\chi(G)}$, answering a question of Esperet, Gimbel and King. We also determine the orientation covering numbers of complete graphs.

Acknowledgements

I am very grateful to Imre Leader, who has been a wonderful supervisor and helped me in many different ways over the last few years. His optimism and encouragement kept me going when I would have given up too early, and his feedback on earlier drafts of my papers greatly improved their quality. I have also really enjoyed recently doing research together.

I would also like to thank the other mathematicians I have had the pleasure to collaborate with: Alberto Espuny Díaz, Timothy Gowers, Oliver Janzer, J. Robert Johnson, Gal Kronenberg and Joanna Lada.

Finally, I feel very lucky to have such a loving family, and would like to thank them for all their support. I am particularly grateful to my brother Oliver for many useful mathematical discussions.

Contents

Chapter 1

Introduction

In this dissertation we study several problems in Combinatorics. While most of the chapters focus on extremal results for graphs and hypergraphs, we also consider some problems of slightly different flavours (for example, a geometric problem in Chapter 2, and an analytical question with combinatorial motivations in Chapter 11).

In Chapter 2, we study the following problem. Let K be a convex body in $\mathbb{R}^d$ that contains a copy of another body S in every possible orientation. Is it always possible to continuously move any one copy of S into another, inside K? As a stronger question, is it always possible to continuously select, for each orientation, one copy of S in that orientation? These questions were asked by Croft. We show that, in two dimensions, the stronger question always has an affirmative answer. We also show that in three dimensions the answer is negative, even for the case when S is a line segment – but that in any dimension the first question has a positive answer when S is a line segment. And we prove that, surprisingly, the answer to the first question is negative in dimensions at least four for general S. This chapter is based on [92].

From Chapter 3 onwards we deal with various extremal problems. One of the most fundamental results in Extremal Graph Theory is Turán's theorem, determining the largest possible size $\mathrm{ex}(n, K_t)$ of an n-vertex graph which does not have K_t as a subgraph. This result leads to many questions where we study extremal problems for graphs (or hypergraphs) avoiding certain substructures. The following several chapters deal with various related topics.

One of the simplest Turán-type results is the fact that an n-vertex cycle-free graph contains at most $n-1$ edges. It is very natural to ask for generalisations of this result for hypergraphs; this is the problem we study in Chapter 3, which is based on [91]. An r-uniform tight cycle of length $\ell > r$ is a hypergraph with vertices $v_1, \ldots, v_\ell$ and edges $\{v_i, v_{i+1}, \ldots, v_{i+r-1}\}$ (for all i), with the indices taken modulo ℓ. Sós (see [117, 137]) and (independently) Verstraëte [137] raised the problem of finding the maximum possible number of edges in an n-vertex r-uniform tight-cycle-free hypergraph. When $r = 2$ this is the usual graph case mentioned above, and the answer is $n-1$. It was originally unknown whether the answer is $\binom{n-1}{r-1}$ in general, which is attained by stars. Huang and Ma [86] showed that for $r \geq 3$ stars are not exactly extremal and we can have $(1+c_r)\binom{n-1}{r-1}$ edges for some $c_r > 0$, but (according to Sudakov and Tomon [135]) it was still widely believed that the extremal value should be $\Theta(n^{r-1})$.[1] However, we construct r-uniform hypergraphs with $\Omega(n^{r-1}\log n/\log\log n) = \omega(n^{r-1})$ edges which contain no tight cycles, showing that this is not the case. This lower bound is tight up to a factor of $(\log n)^{O(1)} = n^{o(1)}$ by recent results of Sudakov and Tomon [135] and Letzter [111].

In saturation problems we study the *smallest* possible size of a (hyper)graph which is maximal with respect to a certain property. Thus this topic is dual to the usual Turán-type problems, where we are concerned with finding the *largest* possible sizes. In Chapter 4, which is based on [93], we consider the following saturation problem: how small can maximal k-wise intersecting set systems over $[n]$ be? This problem was first mentioned (briefly) by Erdős and Kleitman [52], and recently studied in more detail by Hendrey, Lund, Tompkins and Tran [14], and later by Balogh, Chen and Luo [14].[2] For $k = 2$ the answer is simply 2^{n-1}, and for $k = 3$ Balogh, Chen, Hendrey, Lund, Luo, Tompkins and Tran [14] proved that the exact value of the minimum is $2^{n/2+1} - 3$ when n is sufficiently large. For general (fixed) k, the authors of [14] showed that the minimum $f_k(n)$ satisfies $\Omega(2^{n/(k-1)}) \leq f_k(n) \leq O(2^{n/\lceil k/2\rceil})$. We prove that for each k there are maximal k-wise intersecting families with size $O(2^{n/(k-1)})$, which is tight by the lower bound mentioned above and hence answers this natural saturation problem (up to a constant factor).

In Chapter 5, we study two related Turán-type problems. An equivalent formulation of the famous Ruzsa–Szemerédi problem (also known as the $(6,3)$-problem) is to determine how many triangles there can be in a graph on n vertices if no edge is contained in two different triangles. It was proved by Ruzsa and Szemerédi [126] that the answer is $o(n^2)$ but also at least $n^{2-o(1)}$. In Chapter 5 we prove the following natural generalization of this problem: for any r, s with $r < s$, there are graphs containing $n^{r-o(1)}$ copies of K_s such that any K_r is contained in at most one K_s. This is tight in the sense that the number of copies must be $o(n^r)$. We also use our construction to show that there are properly edge-coloured graphs with $n^{r-1-o(1)}$ copies of K_r such that no K_r is rainbow, answering a question of Gerbner, Mészáros, Methuku and Palmer [72] about generalised rainbow Turán numbers. (A subgraph is said to be rainbow if no two of its edges receive the same colour.) This chapter is based on [74] and is joint work with W. T. Gowers.

In Chapter 6, which is based on [94], we continue the study of the generalised rainbow Turán problem. Let $\mathrm{ex}(n, H, \text{rainbow-}F)$ denote the maximal number of copies of H that a properly edge-coloured graph on n vertices can contain if it has no rainbow copy of F. This function was introduced by Gerbner, Mészáros, Methuku and Palmer [72], who focused on the case $H = F$. In the case of cycles they showed that if $k \geq 2$, then $\mathrm{ex}(n, C_{2k+1}, \text{rainbow-}C_{2k+1}) = \Theta(n^{2k-1})$ and $\Omega(n^{k-1}) \leq \mathrm{ex}(n, C_{2k}, \text{rainbow-}C_{2k}) \leq O(n^k)$. They asked what the order of magnitude of

[1]In light of our results, this comment of Sudakov and Tomon only appears in the preprint version of [135].

[2]The results of these two groups of authors originally appeared as separate preprints, but were later combined into a single paper [14].

$\mathrm{ex}(n, C_{2k}, \text{rainbow-}C_{2k})$ is. We show that the answer is $\Theta(n^{k-1})$ if $k \geq 3$ and $\Theta(n^2)$ if $k = 2$. More generally, we determine the order of magnitude of $\mathrm{ex}(n, C_s, \text{ rainbow-}C_t)$ for all s, t with $s \neq 3$, as well as the order of magnitude of $\mathrm{ex}(n, P_\ell, \text{rainbow-}C_k)$ for all $\ell \geq 2$.

In Chapter 7, we study local rainbow colourings, a rather different problem about rainbow subgraphs in edge-coloured graphs. Let $g(n, H)$ be the smallest k such that we can assign a k-edge-colouring f_v of K_n to each vertex v in K_n with the property that for any copy H_0 of H in K_n, there is some $u \in V(H_0)$ such that H_0 is rainbow in f_u. Motivated by a problem in Theoretical Computer Science, this function was introduced by Alon and Ben-Eliezer [3]. We answer several of their questions: in particular, we determine all connected graphs H for which $g(n, H) = n^{o(1)}$, and show that for all $\varepsilon > 0$ there exists $r = r(\varepsilon)$ such that $g(n, K_r) = \Omega(n^{1-\varepsilon})$. We also show that local rainbow colourings are related to the so-called Erdős–Gyárfás function in Ramsey Theory, and prove a family of special cases of a conjecture of Conlon, Fox, Lee and Sudakov [40] about the hypergraph Erdős–Gyárfás function. This chapter is based on [95], which is joint work with O. Janzer.

Chapter 8 is based on [59], which is joint work with A. Espuny Díaz, G. Kronenberg and J. Lada. Consider the bootstrap percolation (or weak saturation) process in which, given a fixed r-uniform hypergraph H and starting with an r-uniform hypergraph G on n vertices, at each step we add to G all edges that create a new copy of H (together with the edges already present). Bollobás raised the problem of maximising the number of steps that this process takes before it stabilises. Several results are known in the graph case $r = 2$ when $H = K_s$ is a clique [15, 28, 114]; in particular, it is known that the answer is quadratic in n for $s \geq 6$, and subquadratic if $s \leq 4$ (and possibly for $s = 5$ as well). In the hypergraph case $H = K_s^{(r)}$ with $s > r \geq 3$, Noel and Ranganathan [119] recently showed that the process can take $\Theta(n^r)$ steps if $s \geq r + 2$, and asked about the case $s = r + 1$. We show that the number of steps of this process can be $\Theta(n^r)$ for $s = r + 1$ as well. To demonstrate that different running times can occur for hypergraphs too, we also prove that if H is $K_4^{(3)}$ minus an edge, then the maximum possible running time is $2n - \lfloor \log_2(n-2) \rfloor - 6$. However, if H is $K_5^{(3)}$ minus an edge, then we prove that the process can run for $\Theta(n^3)$ steps.

In Chapter 9, which is based on [96] and is joint work with J. R. Johnson and I. Leader, we consider the following extremal problems about permutations (related to a problem introduced by Fitzsimons [64] and Angel and Holroyd [12]). What is the smallest number of random transpositions (meaning that we swap given pairs of elements with given probabilities) that we can make on an n-point set to ensure that each element is uniformly distributed – in the sense that the probability that i is mapped to j is $1/n$ for all i and j? And what if we insist that each *pair* is uniformly distributed? We show that the minimum for the first problem is about $\frac{1}{2} n \log_2 n$, with this being exact when n is a power of 2. For the second problem, we show that, perhaps surprisingly, the answer is not quadratic: $O(n \log^2 n)$ random transpositions suffice. We also show that if we ask only that the pair $(1, 2)$ is uniformly distributed then the answer is $2n - 3$. This proves a conjecture of Groenland, Johnston, Radcliffe and Scott [76].

In Chapter 10, which is based on [87], we study the following extremal problem on set systems, introduced by Holzman and Körner [85]. We say that a pair $(\mathcal{A}, \mathcal{B})$ of families of subsets of an n-element set is cancellative if whenever $A, A' \in \mathcal{A}$ and $B \in \mathcal{B}$ satisfy $A \cup B = A' \cup B$, then $A = A'$, and whenever $A \in \mathcal{A}$ and $B, B' \in \mathcal{B}$ satisfy $A \cup B = A \cup B'$, then $B = B'$. Tolhuizen [136] showed that there exist cancellative pairs with $|\mathcal{A}||\mathcal{B}|$ about 2.25^n, whereas Holzman and Körner [85] proved an upper bound of 2.326^n. We improve the upper bound to about 2.268^n. This result also improved the then best known upper bound on a related conjecture of Simonyi [1] (which is the same as the 'sandglass conjecture' [1] for the Boolean lattice), although the upper bound for Simonyi's problem has since been further improved by Nair and Yazdanpanah [118].

In Chapter 11, we study a continuous version of Sperner's theorem. Engel, Mitsis, Pelekis and Reiher [51] showed that an antichain in the continuous cube $[0,1]^n$ must have $(n-1)$-dimensional Hausdorff measure at most n, and they conjectured that this bound can be attained. This was already known for $n = 2$, and we prove this conjecture for all n. This chapter is based on [88].

Chapter 12 is based on [89], and has similar motivations to the preceding chapter. A subset A of $\mathbb{Z}^n$ is called a weak antichain if it does not contain two elements x and y satisfying $x_i < y_i$ for all i. Engel, Mitsis, Pelekis and Reiher [51] showed that for any weak antichain A in $\mathbb{Z}^n$, the sum of the sizes of its $(n-1)$-dimensional projections must be at least as large as its size $|A|$. They asked what the smallest possible value of the gap between these two quantities is in terms of $|A|$, i.e., given the size of A how small its $(n-1)$-dimensional projections can be. We answer this question exactly by giving an explicit weak antichain attaining this minimum for each possible value of $|A|$.

Finally, in Chapter 13, which is based on [90], we study the following problem. Given a graph G, its orientation covering number $\sigma(G)$ is the smallest k with the property that we can choose k orientations of G such that whenever x, y, z are vertices of G with $xy, xz \in E(G)$, then there is a chosen orientation in which both xy and xz are oriented away from x. Orientation coverings were introduced by Esperet, Gimbel and King [58], who showed that there is a natural connection between orientation coverings and the minimal number of equivalence subgraphs (disjoint unions of cliques) needed to cover a line graph. The authors of [58] showed that $\sigma(G) \leq \sigma\left(K_{\chi(G)}\right)$ for any graph G (where χ denotes the chromatic number), and asked whether this upper bound is in fact tight for all G, i.e., whether the value of $\sigma(G)$ is determined just by the chromatic number $\chi(G)$ of G. We answer this question in the positive, and also determine for all n the orientation covering number of K_n exactly in terms of a sequence sometimes called 'Hoşten–Morris numbers'.

Chapter 2

Rotation inside convex Kakeya sets

2.1 Introduction

A subset K of $\mathbb{R}^d$ is called a *Kakeya set* (or *Besicovitch set*) if it contains a unit segment in all directions, i.e., whenever $v \in \mathbb{S}^{d-1}$ then there is some $w \in K$ such that $w+tv \in K$ for all $t \in [0,1]$. The main foundational results about Kakeya sets were proved by Besicovitch [19, 20], who showed that, surprisingly, there exist Kakeya sets of measure zero in the plane, and there are (Kakeya) sets in $\mathbb{R}^2$ of arbitrarily small measure in which a unit segment can be continuously moved and rotated around by 360°. Since then there has been a lot of interest in Kakeya sets and related problems, see, e.g., [33, 46, 106, 108]. The study of Kakeya sets is connected to surprisingly many different areas of mathematics, including harmonic analysis, arithmetic combinatorics and PDEs (see, e.g., [33, 63]). We mention the so-called Kakeya conjecture, which claims that if K is a compact Kakeya set in $\mathbb{R}^d$, then K has (Hausdorff) dimension d (see, e.g., [33]).

While there are many interesting problems about Kakeya sets in various areas of mathematics, in this chapter we will consider a geometric question that is more similar to the original problem and study when we can rotate a body around inside another body. Questions of this form have also attracted much interest. For example, van Alphen [9] showed that it is possible to construct sets of arbitrarily small area and bounded diameter in $\mathbb{R}^2$ in which a segment can be rotated around. Cunningham [43] showed that such a set can even be made simply connected. Csörnyei, Héra and Laczkovich [42] showed that if S is a closed and connected set in $\mathbb{R}^2$ such that any two copies of S can be moved into each other within a set of arbitrarily small measure, then S must be a segment, a circular arc, or a singleton. Järvenpää, Järvenpää, Keleti and Máthé [98] proved that for $n \geq 3$ it is possible to move a line around within a set of measure zero in $\mathbb{R}^n$ such that all directions are traversed; however, if $K \subseteq \mathbb{R}^n$ is such that we can choose a copy of a line in each direction simultaneously in a continuous way (parametrized by $\mathbb{S}^{n-1}$), then the complement of K must be bounded. There is a very large literature on Kakeya sets, and many other interesting problems have been studied, see, e.g., [49, 62, 83].

As hinted above, several results about Kakeya sets concern the stronger property of being able to continuously move and rotate around a segment (or some other set), as opposed to simply containing a segment in each direction (i.e., being Kakeya). It is then interesting to ask how strong the former property is compared to the latter: can we make some additional, natural assumption on our set such that the second property implies the first one? Without any such assumptions, being Kakeya does not imply the first property – for example, our set could consist of two, disconnected components that together cover all possible orientations of segments. It is also easy to see that being connected is not enough – but what happens if our set is *convex*? This question, in the following more general form, was asked by H. T. Croft (personal communication via Imre Leader, 2019).

Question 2.1.1 (Croft)**.** *If K is a convex and compact set in $\mathbb{R}^d$ that contains a copy of $S \subseteq \mathbb{R}^d$ in every possible orientation, is it necessarily possible to continuously transform any given copy of S into any other one within K?*

While it is very natural to study *convex* Kakeya sets, and they were already considered over a hundred years ago by Pál [120] (who proved that the minimal possible area of a convex Kakeya set in $\mathbb{R}^2$ is $1/\sqrt{3}$), it is important to point out that the question above is of a different flavour. Indeed, apart from focusing on convex sets, a significant difference between Question 2.1.1 and most of the known results about Kakeya sets is that here we are not interested in the measure of our Kakeya set (unlike in the papers mentioned earlier).

To formalise Question 2.1.1, we first need some definitions. By *body* we will mean a compact set (in Euclidean space). For any set $S \subseteq \mathbb{R}^d$, let us say that $K \subseteq \mathbb{R}^d$ is *S-Kakeya* if K contains a translate of every rotated copy of S, i.e., whenever $\rho \in \mathrm{SO}(d)$ then there is some $w \in \mathbb{R}^d$ such that $\rho(S) + w \subseteq K$. In particular, when S is a segment of length 1 then this is just the usual notion of being a Kakeya set. Let us also say that *any two S-copies can be rotated into each other within K* if whenever $\rho_0, \rho_1 \in \mathrm{SO}(d)$ and $w_0, w_1 \in \mathbb{R}^d$ are such that $\rho_i(S) + w_i \subseteq K$ $(i = 0, 1)$, then there are some $\gamma : [0, 1] \to \mathrm{SO}(d)$ and $\delta : [0, 1] \to \mathbb{R}^d$ continuous functions such that $\gamma(i) = \rho_i, \delta(i) = w_i$ for $i = 0, 1$ and $\gamma(t)(S) + \delta(t) \subseteq K$ for all t.

We mention that instead of having continuous γ, δ as above, we could define this notion in terms of a single continuous function ψ mapping each $t \in [0, 1]$ to a (rotated and translated) copy of S in a continuous way (with respect to the Hausdorff metric), our results below still hold in this alternative characterisation. Furthermore, in the case of usual Kakeya sets (i.e., when S is a unit segment), we can also parametrize the possible orientations of segments by the sphere $\mathbb{S}^{d-1}$ or by the projective space $\mathbb{PR}^d$ instead of $\mathrm{SO}(d)$, but these changes would make no difference.

Our first result shows that in the case of usual Kakeya sets, any two unit segments can be rotated into each other within K (if K is convex and compact).

Theorem 2.1.2. *Let $d \geq 2$ be a positive integer and let K be a convex Kakeya body in $\mathbb{R}^d$. Then any two unit segments can be rotated into each other within K.*

Given Theorem 2.1.2, one might expect that the corresponding statement is in fact true for any set S. Surprisingly, this is not the case.

Theorem 2.1.3. *There exist convex bodies S and K in $\mathbb{R}^4$ such that K is S-Kakeya but there are two S-copies which cannot be rotated into each other within K.*

While the result above is stated for $d = 4$, it is in fact easy to modify our construction to get a counterexample for any $d \geq 4$.

We mention that, in contrast with Theorem 2.1.3, if we replace the assumption 'K compact' by 'K open', then an easy connectedness argument shows that any two copies can be rotated into each other.

An alternative way to interpret Question 2.1.1 is to ask for a way to select a copy of S (in an S-Kakeya set) in each direction simultaneously in a continuous way. That is, we want the stronger property that there exists a continuous map $f : \mathrm{SO}(d) \to \mathbb{R}^d$ such that $\rho(S) + f(\rho) \subseteq K$ for all ρ. We show that this can be achieved for any shape in 2 dimensions.

Theorem 2.1.4. *Let K be a convex body in $\mathbb{R}^2$ and let $S \subseteq \mathbb{R}^2$, $S \neq \emptyset$. Assume that K is S-Kakeya. Then there is a continuous map $f : \mathrm{SO}(2) \to \mathbb{R}^2$ such that $\rho(S) + f(\rho) \subseteq K$ for all $\rho \in \mathrm{SO}(2)$.*

Again, in light of Theorem 2.1.4, one might expect that the corresponding statement is true in higher dimensions too, at least when S is a line segment. However, this strong property fails already when $d = 3$, even when S is a unit segment.

Theorem 2.1.5. *There exists a convex Kakeya body $K \subseteq \mathbb{R}^3$ such that there is no continuous function $\psi : \mathbb{S}^2 \to \mathbb{R}^3$ satisfying $\psi(v) + tv \in K$ for all $v \in \mathbb{S}^2, t \in [0, 1]$.*

As before, the fact that we chose to parametrize orientations of segments by the sphere $\mathbb{S}^{d-1}$ instead of $SO(d)$ or the projective space $\mathbb{PR}^d$ does not change anything, Theorem 2.1.5 would remain true for these parametrizations as well, and reason why the counterexample works is not topological.

The rest of the chapter is organised as follows. In Section 2.2, we prove Theorem 2.1.4 and Theorem 2.1.5 concerning the stronger property of being able to continuously select in all directions. In Section 2.3, we prove Theorem 2.1.2 about rotating in convex Kakeya sets in $\mathbb{R}^d$ for any d, and in Section 2.4 we give a counterexample for the corresponding statement for general bodies. We finish with some concluding remarks and open questions in Section 2.5.

The proofs in Section 2.2 are simpler than the ones in the later sections, but several elements of those proofs reappear or motivate our later approach. In particular, one of the main methods we will have for analysing different cases is to consider the dimensions of the sets $I_\rho = \{w \in \mathbb{R}^d : \rho(S) + w \subseteq K\}$. It is easy to deal with ρ (and its neighbourhood) if I_ρ has dimension d (i.e., has non-empty interior). One might initially expect that the larger the dimension of I_ρ is, the more room we have to move the copies around and hence the easier to deal with ρ. However, this is not entirely true, and the 0-dimensional case (when I_ρ is a single point) will be quite easy to deal with. For example, it is not difficult to prove that if $I_\rho = \{w_\rho\}$ is a single point for all ρ, then $\rho \mapsto w_\rho$ must be continuous. So the most difficult cases in Theorem 2.1.2 will come from the situation when some I_ρ has dimension between 1 and $d-1$, and these will also be the cases we use to obtain counterexamples in Theorems 2.1.5 and 2.1.3.

2.2 Continuous choice in each direction

In this section we prove Theorem 2.1.4 and Theorem 2.1.5 about selecting a copy in each direction in a continuous way. We begin with Theorem 2.1.4.

First we recall the definition of the Hausdorff metric. Given a point $p \in \mathbb{R}^d$ and a non-empty compact set $A \subseteq \mathbb{R}^d$, write

$$d(p, A) = \min_{a \in A} |p - a|.$$

Given two non-empty compact sets $X, Y \subseteq \mathbb{R}^d$, their distance in the Hausdorff metric d is defined as

$$d(X, Y) = \max\{\max_{x \in X} d(x, Y), \max_{y \in Y} d(y, X)\}.$$

It is well-known that this makes the set $\mathcal{C}_d$ of non-empty compact subsets of $\mathbb{R}^d$ a metric space. Let $\mathcal{K}_d$ denote the set of non-empty compact convex sets in $\mathbb{R}^d$ (so $\mathcal{K}_d \subseteq \mathcal{C}_d$).

We will prove the following result.

Lemma 2.2.1. *Let S be a non-empty compact subset of $\mathbb{R}^2$ and let K be convex, compact and S-Kakeya. For all $\rho \in \mathrm{SO}(2)$, let $I_\rho = \{v \in \mathbb{R}^2 : \rho(S) + v \subseteq K\}$. Then the map $\mathrm{SO}(2) \to \mathcal{K}_2$ given by $\rho \mapsto I_\rho$ is continuous.*

Given $I \in \mathcal{K}_d$, we say that I has Chebyshev centre c if $x = c$ minimises $\max_{p \in I} |x - p|$ among all points $x \in \mathbb{R}^d$. We will use the following properties of Chebyshev centres. (Much more general statements are known about Chebyshev centres in Banach spaces, but the next result is enough for our purposes.)

Lemma 2.2.2. *(See, e.g., [10, Theorem 5] and [11, subsection 7.1]) If $I \in \mathcal{K}_d$ then I has a unique Chebyshev centre c_I. Moreover, $c_I \in I$ for all I, and the map $\mathcal{K}_d \to \mathbb{R}^d$ given by $I \mapsto c_I$ is continuous.*

It is easy to see that Theorem 2.1.4 follows from Lemmas 2.2.1 and 2.2.2. So we now need to prove Lemma 2.2.1. In fact, we will prove the following stronger statement.

Lemma 2.2.3. *Let K be a compact convex set in $\mathbb{R}^2$. For any non-empty compact set S in $\mathbb{R}^2$, let $I_S = \{w \in \mathbb{R}^2 : S + w \subseteq K\}$. Let $\mathcal{A}_K$ be the set of all S with I_S non-empty. Then the map $\psi : \mathcal{A}_K \to \mathcal{K}_2$ given by $S \mapsto I_S$ is continuous (with respect to the Haudorff metric on both sides).*

Lemma 2.2.3 certainly implies Lemma 2.2.1, as $\rho \mapsto \rho(S)$ is easily seen to be continuous for any fixed S. Also, note that Lemma 2.2.1 and Lemma 2.2.3 are not true in dimensions greater than 2, by the construction in Theorem 2.1.5.

Let us start the proof of Lemma 2.2.3. The first lemma towards the proof essentially says that if I_S is a segment on the x axis (so I_S is one-dimensional), then the projections of K and S to the y axis have the same maximum values (and similarly minimum values). This is rather easy to see when S is a segment, and only slightly more complicated in general.

Lemma 2.2.4. *Suppose that $K \subseteq \mathbb{R}^2$ is compact and convex, $S \subseteq \mathbb{R}^2$ is non-empty and compact, and $\delta > 0$ is such that $\{v \in \mathbb{R}^2 : S + v \subseteq K\} \supseteq \{(a, 0) : |a| \leq \delta\}$. Let $p = (x_0, y_0)$ and $p' = (x_0', y_0')$ be points of S and K (respectively) with maximal second coordinates. Then either $y_0 = y_0'$, or there is some $\epsilon > 0$ such that $S + (0, \epsilon) \subseteq K$. Similarly, if $p'' = (x_0'', y_0'')$ and $p''' = (x_0''', y_0''')$ are points of S and K (respectively) with minimal second coordinates, then either $y_0'' = y_0'''$, or there is some $\epsilon > 0$ such that $S - (0, \epsilon) \subseteq K$*

Proof. We only prove the first claim, as the second one is similar. Certainly $y_0' \geq y_0$ as $S \subseteq K$. Let us assume that $y_0' > y_0$, we show that if $\epsilon > 0$ is sufficiently small then for any $q = (x_1, y_1) \in S$ we have $q + (0, \epsilon) \in K$. It is enough to consider the case $x_1 \geq x_0'$. Let $L > 0$ be such that $S \subseteq [-L, L]^2$. We know that $q' = (x_1 + \delta, y_1)$ is in K. By convexity, the line segment between q' and p' also lies in K and hence $\left(x_1, \frac{y_1 - y_0'}{x_1 + \delta - x_0'}(x_1 - x_0') + y_0'\right) \in K$. But we have

$$\left(\frac{y_1 - y_0'}{x_1 + \delta - x_0'}(x_1 - x_0') + y_0'\right) - y_1 = \frac{\delta}{x_1 + \delta - x_0'}(y_0' - y_1) \geq \frac{\delta}{2L + \delta}(y_0' - y_0).$$

It follows that $\epsilon = \frac{\delta}{2L+\delta}(y_0' - y_0)$ satisfies the conditions. □

The next lemma will be used to prove Hausdorff-continuity in the difficult case, i.e., when I_S is one-dimensional.

Lemma 2.2.5. *Suppose that $K \subseteq \mathbb{R}^2$ is compact and convex, and define the sets I_S and $\mathcal{A}_K$ as in Lemma 2.2.3. Assume that $u \in \mathbb{R}^2$, $\delta > 0$ and $S \in \mathcal{A}_K$ such that I_S has empty interior but $I_S \supseteq \{u + (a, 0) : |a| \leq \delta\}$. Then for all $\epsilon > 0$ there is some $\eta > 0$ such that whenever $S' \in \mathcal{A}_K$ satisfies $d(S, S') < \eta$ then there is some $w \in I_{S'}$ with $|w - u| < \epsilon$.*

Proof. We may assume $u = 0$ (by replacing K by $K - u$). Since I_S has empty interior, by Lemma 2.2.4 we have $y_0 = y_0'$ and $y_0'' = y_0'''$ (using the notation in the statement of that lemma). If $S' \in \mathcal{A}_K$ and $d(S, S') < \eta$, we know $S' + (0, z) \subseteq \mathbb{R} \times [y_0'', y_0]$ for some $z \in \mathbb{R}$ with $|z| < \eta$. (Indeed, we can pick z such that the largest second coordinate of a point in S' is $y_0 + z$.) We show that if η is small enough, then we must have $(0, z) \in I_{S'}$. (Then we are done, as we can choose $\eta < \epsilon$.) By replacing S' with $S' + (0, z)$ (and η by 2η), we may assume that $z = 0$.

So we need to show that for any point $q = (x_1, y_1) \in S'$, we have $q \in K$ (if η is small). We know there is some $q' = (x_2, y_2) \in S$ with $|x_1 - x_2| < \eta$, $|y_1 - y_2| < \eta$. We may assume that $y_1 \geq y_2$. We wish to show that for some $s \in [-\delta, \delta]$, $q - (s, 0)$ must lie on the line segment between $p = (x_0, y_0)$ and $q' = (x_2, y_2)$. (Then we are done, since $p, q' \in S$, $S + (s, 0) \subseteq K$, and K is convex.)

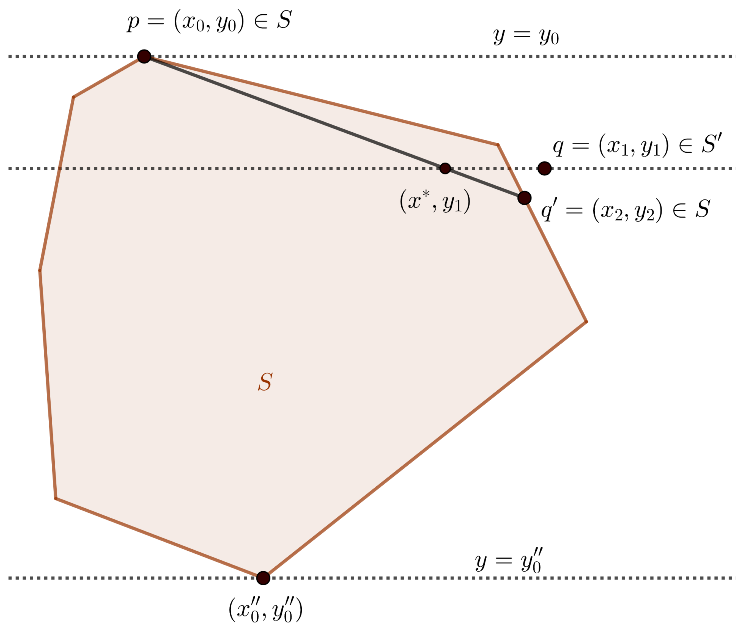

Figure 2.1: The points used in the proof of Lemma 2.2.5.

First assume that $y_2 = y_0$. So $y_0 = y_1 = y_2$. But then $q = q' + (s, 0)$ for some $s \in (-\eta, \eta)$, so our claim follows easily by picking $\eta < \delta$.

So let us now assume $y_2 \neq y_0$ (so $y_2 < y_0$). Observe that points (x, y) on the segment between p and q' are the ones satisfying the equation $x - x_0 = \frac{x_0 - x_2}{y_0 - y_2}(y - y_0)$ and have $y_2 \leq y \leq y_0$. It follows that (x^*, y_1) is on this segment, where $x^* = x_0 + \frac{x_0 - x_2}{y_0 - y_2}(y_1 - y_0)$. We have

$$\begin{aligned} |x^* - x_2| &= \left| x_0 + \frac{x_0 - x_2}{y_0 - y_2}(y_1 - y_0) - x_2 \right| \\ &= \left| \frac{x_0 - x_2}{y_0 - y_2}(y_1 - y_2) \right|. \end{aligned}$$

We will use the following claim to bound this quantity.

Claim. There is some $\mu > 0$ depending on S, δ only such that whenever $(x, y) \in S$ and $y > y_0 - \mu$, then there is some $\bar{x}_0$ such that $(\bar{x}_0, y_0) \in S$ and $|\bar{x}_0 - x| < \delta/2$.

Proof of Claim. If this is not true, then for all n we can find $(x(n), y(n)) \in S$ such that $y(n) > y_0 - 1/n$ and whenever $(\bar{x}_0, y_0) \in S$ then $|\bar{x}_0 - x(n)| \geq \delta/2$. By taking a subsequence, we may assume that $(x(n), y(n))$ converges to some $(\tilde{x}, \tilde{y}) \in S$. But then $\tilde{y} = y_0$ and $\tilde{x} - x(n) \to 0$, giving a contradiction and proving the claim. $\square$

By the claim above, we can modify x_0 if necessary so that either $y_0 - y_2 \geq \mu$ or $|x_0 - x_2| < \delta/2$.

In the first case we get $|x^* - x_2| \leq \frac{|x_0 - x_2|}{\mu}|y_1 - y_2|$. Let $L > 0$ be such that $S \subseteq [-L, L]^2$, then we get $|x^* - x_2| \leq \frac{2L}{\mu}\eta$ and hence $|x^* - x_1| \leq \eta + \frac{2L}{\mu}\eta$. This converges to 0 (independently of q, q') as $\eta \to 0^+$, as required.

On the other hand, if $|x_0 - x_2| < \delta/2$ then, using $y_0 - y_2 \geq y_0 - y_1$, we get $|x^* - x_2| \leq \delta/2$ and hence $|x^* - x_1| \leq \delta/2 + \eta$, which is less than δ for $\eta < \delta/2$. $\square$

Proof of Lemma 2.2.3. First note that all sets of the form I_S are convex and compact. Let $S \in \mathcal{A}_K$ be arbitrary, we show ψ is continuous at S, i.e., whenever $S_n \to S$ with $S_n \in \mathcal{A}_K$, then $d(I_{S_n}, I_S) \to 0$. First we show $\max_{x \in I_{S_n}} d(x, I_S) \to 0$. Indeed, if this is not true, then by taking an appropriate subsequence $(S_{k(n)})$ we get that there is a sequence (x_n) with $x_n \in I_{S_{k(n)}}$ such that $d(x_n, I_S) \not\to 0$ and $x_n \to x$ for some x. But we have $S_{k(n)} + x_n \subseteq K$ for all n. Hence $S + x \subseteq K$, i.e., $x \in I_S$. (Indeed, for any $s \in S$ we can take a sequence (s_n) with $s_n \in S_{k(n)}$ and $(s_n) \to s$. Then $s_n + x_n \in K$ for all n, so, by taking limits, $s + x \in K$.) But then $d(x_n, I_S) \to 0$, giving a contradiction. So $\max_{x \in I_{S_n}} d(x, I_S) \to 0$.

It remains to show that $\max_{x \in I_S} d(x, I_{S_n}) \to 0$. Observe that it suffices to show that $d(x, I_{S_n}) \to 0$ for any point $x \in I_S$. Indeed, the functions $x \mapsto d(x, I_{S_n})$ are 1-Lipschitz on the compact domain I_S, so pointwise convergence implies uniform convergence. We consider three cases: when I_S is a single point, when I_S is one-dimensional, i.e., $I_S = \{(1-t)a + tb : t \in [0,1]\}$ for some $a, b \in \mathbb{R}^2$ distinct, and when I_S is two-dimensional, i.e., has non-empty interior.

First assume that $I_S = \{p\}$ is a single point. Then trivially

$$d(p, I_{S_n}) = \min_{x \in I_{S_n}} d(p, x) \leq \max_{x \in I_{S_n}} d(p, x) = \max_{x \in I_{S_n}} d(x, I_S) \to 0,$$

giving the claim.

Next, assume that I_S is one-dimensional (i.e., a segment). By taking an appropriate rotation and translation, we may assume that $I_S = [-\delta, \delta] \times \{0\}$ for some $\delta > 0$. Let $x \in I_S$ and $\epsilon > 0$ be arbitrary, we show $d(x, I_{S_n}) < \epsilon$ for n large enough. We may assume that $\epsilon < \delta$. Let $x' \in [-\delta + \epsilon/2, \delta - \epsilon/2] \times \{0\}$ be such that $|x - x'| \leq \epsilon/2$. Since $I_S \supseteq \{x' + (a, 0) : |a| \leq \delta/2\}$, Lemma 2.2.5 shows that for all n large enough there is some $w \in I_{S_n}$ with $|w - x'| \leq \epsilon/4$. But then we also have $|w - x| < \epsilon$, as required.

Finally, assume that I_S is two-dimensional, i.e., has non-empty interior. Let $x \in I_S$ and $\epsilon > 0$ be arbitrary, we show $d(x, I_{S_n}) < \epsilon$ for n large enough. We can find $x' \in I_S$ with $|x' - x| < \epsilon$ such that x' is in the interior of I_S, i.e., I_S contains a ball of radius $r > 0$ around x'. Then whenever

$d(S', S) < r$, we have $x' \in I_{S'}$. (Indeed, $x' + S' \subseteq x' + B_r(0) + S \subseteq I_S + S \subseteq K$, where $B_r(0)$ denotes the ball of radius r centred at 0.) Hence $x' \in I_{S_n}$ for n large enough, giving the claim. $\square$

Proof of Theorem 2.1.4. Let S' be the the closure of S. Then $\{v \in \mathbb{R}^2 : \rho(S) + v \subseteq K\} = \{v \in \mathbb{R}^2 : \rho(S') + v \subseteq K\}$ for all $\rho \in \mathrm{SO}(2)$. By replacing S by S', we may assume that S is compact. Then the result follows easily from Lemma 2.2.3 and Lemma 2.2.2 by letting $f(\rho)$ be the Chebyshev centre of $I_{\rho(S)}$. $\square$

We finish this section by proving Theorem 2.1.5. Informally, the construction can be described as follows. Take a circle of diameter 1 in the xy plane, and start moving it in the x direction while simultaneously rotating it around the x axis. Stop when the rotated circle gets back to the xy plane, and take the convex hull of the points traversed. See Figure 2.2. The discontinuity will come at the direction $(0, 1, 0)$ by considering directions of the form $(0, y, \pm\sqrt{1 - y^2})$, $y \to 1^-$. The formal proof is given below.

(a) Some phases of the circle being rotated and translated.

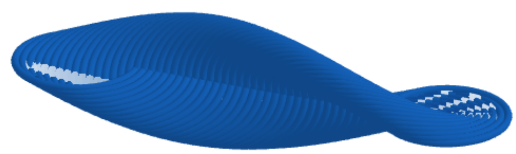

(b) The set of points traversed during the motion. The final construction is obtained by taking the convex hull of this set.

Figure 2.2: The counterexample in Theorem 2.1.5 is obtained by simultaneously translating and rotating a circle, and then taking convex hull of the points traversed.

Proof of Theorem 2.1.5. Define the function $f : [0, \pi] \times \mathbb{S}^1 \to \mathbb{R}^3$ by letting

$$f(t, x, y) = \frac{1}{2}(t + x, y \cos t, y \sin t).$$

Let K_0 be the image of f and let K be the convex hull of K_0. Observe that f is continuous and the domain of f is compact, hence K_0 is compact. It follows that K is convex and compact. Also, note that if $v \in \mathbb{S}^2$, then v can be written as $v = (r_1, r_2 \cos \varphi, r_2 \sin \varphi)$ for some $r_1, r_2 \in \mathbb{R}$ with $r_1^2 + r_2^2 = 1$ and $\varphi \in [0, \pi]$. Then $f(\varphi, r_1, r_2) - f(\varphi, -r_1, -r_2) = (r_1, r_2 \cos \varphi, r_2 \sin \varphi) = v$, so I_v is non-empty, where $I_v = \{u \in \mathbb{R}^2 : u, u + v \in K\}$. It remains to show that there is no continuous function $\psi : \mathbb{S}^2 \to K$ such that $\psi(v) \in I_v$ for all v.

Let $C = \{(a, b, c) \in \mathbb{R}^3 : b^2 + c^2 = 1/4\}$ and $C' = \{(a, b, c) \in \mathbb{R}^3 : b^2 + c^2 \leq 1/4\}$. Observe that $K_0 \subseteq C'$ and

$$K_0 \cap C = \left\{ \frac{1}{2}(t, s \cos t, s \sin t) : s = \pm 1, t \in [0, \pi] \right\}.$$

It is easy to deduce that $K \subseteq C'$ and

$$K \cap C = \left\{ \frac{1}{2}(t, s\cos t, s\sin t) : s = \pm 1, t \in [0,\pi] \right\} \cup \left\{ \frac{1}{2}(a, \pm 1, 0) : a \in [0,\pi] \right\}.$$

It is easy to deduce that if $v = (0, \cos\varphi, \sin\varphi)$ for some $\varphi \in (0,\pi)$, then I_v consists of the single point $\frac{1}{2}(\varphi, -\cos\varphi, -\sin\varphi)$, and if $v = (0, \cos\varphi, \sin\varphi)$ for some $\varphi \in (-\pi, 0)$, then I_v consists of the single point $\frac{1}{2}(\pi+\varphi, \cos(\pi+\varphi), \sin(\pi+\varphi)) = \frac{1}{2}(\pi+\varphi, -\cos\varphi, -\sin\varphi)$. It follows that if $\psi : \mathbb{S}^2 \to K$ such that $\psi(v) \in I_v$ for all v, then ψ cannot be continuous at $(0,1,0)$. $\square$

2.3 Segments in $\mathbb{R}^d$

2.3.1 Proof outline and some simple results

Our goal in this section is to prove Theorem 2.1.2 about Kakeya sets in $\mathbb{R}^d$. Throughout this section, we assume that $d \geq 3$ and K is a compact convex set in $\mathbb{R}^d$ such that for all $v \in \mathbb{S}^{d-1}$, the set $I_v = \{u \in \mathbb{R}^d : u, u+v \in K\}$ is non-empty. Note that I_v is a compact convex set for all v.

Given $v, v' \in \mathbb{S}^{d-1}$, $u \in I_v, u' \in I_{v'}$ and $\gamma : [0,1] \to \mathbb{S}^{d-1}$ continuous with $\gamma(0) = v, \gamma(1) = v'$, say that (v', u') is reachable from (v, u) along γ if there is a continuous $\delta : [0,1] \to K$ such that $\delta(t) \in I_{\gamma(t)}$ for all t, $\delta(0) = u$ and $\delta(1) = u'$. We say that v' is reachable from v along γ if there exist u, u' such that (v', u') is reachable from (v, u) along γ, and we say v' (or (v', u')) is reachable from v (respectively, (v, u)) if there exists a γ along which it is reachable. Given a subset $X \subseteq \mathbb{S}^{d-1}$, $\epsilon \geq 0$ and $\gamma : [0,1] \to \mathbb{S}^{d-1}$ we say that γ is ϵ-close to X if for all $t \in [0,1]$ there is some $p \in X$ such that $|p - \gamma(t)| \leq \epsilon$. Given $\epsilon \geq 0$ and $\gamma, \gamma' : [0,1] \to \mathbb{S}^{d-1}$ we say that γ is ϵ-close to γ' if it is ϵ-close to the image of γ'. (Note that this relation is not symmetric.)

So, using this terminology, our goal is to prove the following result.

Theorem 2.3.1. *Let $v, v' \in \mathbb{S}^{d-1}$ and $u \in I_v, u' \in I_{v'}$, and let $\gamma : [0,1] \to \mathbb{S}^{d-1}$ be continuous such that $\gamma(0) = v, \gamma(1) = v'$. Then for any $\epsilon > 0$, (v', u') is reachable from (v, u) along a path which is ϵ-close to γ.*

Note that the counterexample in Theorem 2.1.5 shows that it is not necessarily true that v' is reachable from v along γ (or along a path 0-close to γ).

We now briefly discuss our approach to proving Theorem 2.3.1. It is easy to see that if $p \in \mathbb{S}^{d-1}$ is such that I_p has non-empty interior, then every p' in some neighbourhood of p is reachable from p. Furthermore, it is not difficult to deal with points p such that I_p is a single point. This means that the complicated case is when I_p is not a single point, but has empty interior (i.e., its dimension is between 1 and $d-1$). We will prove (Lemma 2.3.6) that in the neighbourhood of such points p, there are 'many' points q with I_q having non-empty interior. Moreover, we will show that if for such a p we start moving on the sphere $\mathbb{S}^{d-1}$ from p in some direction, then for 'most' directions we initially only encounter points q such that I_q has non-empty interior, and that these q are reachable from p. We will deduce (Lemma 2.3.7) that Theorem 2.3.1 holds for γ if for all points v on γ such that I_v has empty interior and is not a single point, the tangent to γ at v is not in some special set of 'forbidden' directions. Finally, we will show that we can perturb γ slightly to make sure that we avoid such cases. We note that in some sense we can have 'many' points $p \in \mathbb{S}^{d-1}$ such that I_p is not a single point but has empty interior. For example, if $K = \{(x, y, z) \in \mathbb{R}^3 : x \in [-1, 1], y^2 + z^2 \leq 1/4\}$, then all p along a great circle have this property.

We believe the reader will not lose much by focusing on the case $d = 3$: some of the lemmas are easier to visualise and prove in that case, but the main ideas of the proof are the same.

Let us start with some simple observations.

Lemma 2.3.2. *Suppose that $v, v' \in \mathbb{S}^{d-1}$ and $\gamma : [0, 1] \to \mathbb{S}^{d-1}$ are such that v' is reachable from v along γ. Let $u \in I_v$, $u' \in I_{v'}$ be arbitrary. Then (v', u') is reachable from (v, u) along a path which has the same image as γ (and hence is* 0*-close to γ).*

Proof. Let $w \in I_v$, $w' \in I_{v'}$ and $\delta : [0, 1] \to K$ be such that δ is continuous, $\delta(t) \in I_{\gamma(t)}$ for all t, $\delta(0) = w$ and $\delta(1) = w'$. Define γ' and δ' by setting

$$\gamma'(t) = \begin{cases} v & \text{if } t \in [0, 1/3] \\ \gamma(3(t - 1/3)) & \text{if } t \in [1/3, 2/3] \\ v' & \text{if } t \in [2/3, 1] \end{cases}$$

and

$$\delta'(t) = \begin{cases} (1 - 3t)u + 3tw & \text{if } t \in [0, 1/3] \\ \delta(3(t - 1/3)) & \text{if } t \in [1/3, 2/3] \\ (1 - 3(t - 2/3))w' + 3(t - 2/3)u' & \text{if } t \in [2/3, 1]. \end{cases}$$

The statement of the lemma follows easily, using that $I_v, I_{v'}$ are convex. □

Note that Lemma 2.3.2 implies that if v' is reachable from v (along some path which is ϵ-close to X) and v'' is reachable from v' (along some path which is ϵ-close to Y) then v'' is reachable from v (along some path which is ϵ-close to $X \cup Y$).

Lemma 2.3.3. *Assume that $V \subseteq \mathbb{S}^{d-1}$ is such that for all $v \in V$, I_v has non-empty interior. Assume furthermore that $\gamma : [0, 1] \to V$ is continuous. Then $\gamma(1)$ is reachable from $\gamma(0)$ along a path which is* 0*-close to γ.*

Proof. For all $t \in [0, 1]$ we can find some $r_t > 0$, $p_t \in K$ such that $I_{\gamma(t)}$ contains an open ball of radius r_t around p_t, i.e., whenever $|z| < r_t$ then $p_t + z, p_t + z + \gamma(t) \in K$. It follows that whenever $|\gamma(s) - \gamma(t)| < r_t$ then $p_t, p_t + \gamma(s) \in K$, i.e., $p_t \in I_{\gamma(s)}$. Let $\eta_t > 0$ be such that $|\gamma(s) - \gamma(t)| < r_t$ whenever $|s - t| < \eta_t$. By compactness of $[0, 1]$, we can find some $r > 0$ such that whenever $s \in [0,1]$ then there is some $t_s \in [0,1]$ such that $|s - t_s| \leq \eta_{t_s} - r$. Pick some $N > 1/r$ integer, and let $x(i) = i/N$ $(i = 0, \ldots, N)$. Then $\gamma(x(i+1))$ is reachable from $\gamma(x(i))$ along a path which has the same image is $\gamma|_{[x(i),x(i+1)]}$ (the corresponding function δ is constant $p_{t_{x(i)}}$). Using Lemma 2.3.2 several times, and concatenating the appropriate paths, we get that $\gamma(1)$ is reachable from $\gamma(0)$ along a path with the same image as γ. □

In light of Lemma 2.3.3, finding points v such that I_v has non-empty interior is useful for proving reachability. The next lemma gives a convenient condition for checking that I_v has non-empty interior.

Lemma 2.3.4. *If $v \in \mathbb{S}^{d-1}$ and there is some $\lambda > 1$ and $u \in K$ such that $u + \lambda v \in K$, then I_v has non-empty interior.*

Proof. We may assume that $u = 0$. If $p \in K$, then $(1-1/\lambda)p \in K$ and $(1-1/\lambda)p+(1/\lambda)\lambda v \in K$ by convexity, so $(1 - 1/\lambda)p \in I_v$. Given some $w \in \mathbb{S}^{d-1}$, there are points $p_1, p_2 \in K$ such that $p_2 - p_1 = w$. Then $(1 - 1/\lambda)p_i \in I_v$ for $i = 1, 2$, and therefore I_v contains two points q_1, q_2 with $q_2 - q_1 = (1 - 1/\lambda)w$. So we can pick $e_1, f_1, \ldots, e_d, f_d \in I_v$ such that $f_i - e_i$ is the vector with all coordinates zero, except the ith coordinate, which is $1 - 1/\lambda$. Let $c = \frac{1}{2d}\sum(e_i + f_i)$. By convexity of I_v, it is easy to see that $c \in I_v$, and whenever $|x_i| \leq \frac{1}{2d}(1 - 1/\lambda)$ for all i then $c + (x_1, \ldots, x_d) \in I_v$. So I_v contains a ball of radius $\frac{1}{2d}(1 - 1/\lambda)$ around c. □

The following useful lemma gives another condition for finding v such that I_v has non-empty interior, and it also gives some restrictions on what I_v can look like when I_v has empty interior: $I_v - I_v$ must be perpendicular to v.

Lemma 2.3.5. *Suppose that $p \in \mathbb{S}^{d-1}$, $x, q \in \mathbb{R}^d$ such that $\langle p, q\rangle > 1$ and $x, x + q \in K$. Then I_p has non-empty interior.*

In particular, if $v \in \mathbb{S}^{d-1}$ and $u, w \in \mathbb{R}^d$ such that $\langle v, w\rangle \neq 0$ and $u, u + w \in I_v$, then I_v has non-empty interior.

Proof. Let $\epsilon \in (0, 1)$ be small enough so that $0 \neq |p - \epsilon q| < 1 - \epsilon$. Note that such an ϵ exists, since $|p - \epsilon q|^2 = 1 - 2\langle p, q\rangle\epsilon + |q|^2\epsilon^2$ is less than $1 - 2\epsilon + \epsilon^2$ for ϵ small enough, as $\langle p, q\rangle > 1$. Let $p' = \frac{p-\epsilon q}{|p-\epsilon q|}$. Note that $|p'| = 1$. We know that there is some $y \in \mathbb{R}^d$ such that $y, y + p' \in K$. Let

$$z = \frac{\epsilon}{\epsilon + |p - \epsilon q|}x + \frac{|p - \epsilon q|}{\epsilon + |p - \epsilon q|}y,$$
$$z' = \frac{\epsilon}{\epsilon + |p - \epsilon q|}(x + q) + \frac{|p - \epsilon q|}{\epsilon + |p - \epsilon q|}(y + p').$$

Then $z, z' \in K$ by convexity. But

$$\begin{aligned} z' - z &= \frac{\epsilon}{\epsilon + |p - \epsilon q|} q + \frac{|p - \epsilon q|}{\epsilon + |p - \epsilon q|} p' \\ &= \frac{1}{\epsilon + |p - \epsilon q|} p \end{aligned}$$

But $\frac{1}{\epsilon + |p - \epsilon q|} > 1$, so I_p has non-empty interior by Lemma 2.3.4.

For the final part of the lemma, we may assume $\langle v, w \rangle > 0$ (otherwise replace u by $u + w$ and w by $-w$). But then $u, u + v + w \in K$, so we can apply the first part of the lemma with $p = v$, $q = v + w$, $x = u$. □

2.3.2 The main lemmas

The following lemma is one of the key observations. Essentially, it says that if I_v has empty interior but is not a single point, then for 'most' points p around v the set I_p has non-empty interior, and those p are reachable from v.

Lemma 2.3.6. *Suppose that $v \in \mathbb{S}^{d-1}$, and $u, w \in \mathbb{R}^d$ such that $u, u + w \in I_v$, $0 < |w| < 1$ and $\langle v, w \rangle = 0$. Let $P = \{p \in \mathbb{S}^{d-1} : \langle p, v + w \rangle > 1\} \cup \{p \in \mathbb{S}^{d-1} : \langle p, v - w \rangle > 1\}$. Then I_p has non-empty interior for all $p \in P$. Moreover, whenever $p \in P$, then p is reachable from v along a path which is $2|p - v|$-close to $\{v\}$.*

Note that the condition $\langle v, w \rangle = 0$ holds automatically when I_v has empty interior by the final part of Lemma 2.3.5. Figure 2.3 shows the set P in the case $d = 3$.

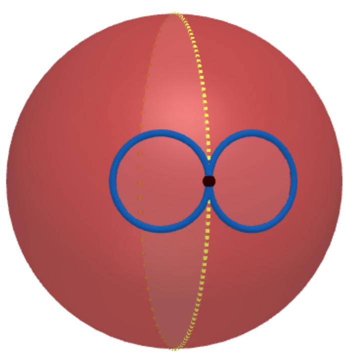

Figure 2.3: The set P in Lemma 2.3.6 is the region enclosed by the two blue circles ($d = 3$). The point v is the intersection of the two circles, and w is parallel to the line connecting the centres of the blue circles. The yellow dotted great circle gives the only direction (for $d = 3$) not pointing to the inside of the two circles.

Proof. The claim that I_p has non-empty interior for all $p \in P$ follows directly from Lemma 2.3.5. For the second claim, let $p \in P$ be arbitrary. We may assume $\langle p, v + w \rangle > 1$ (otherwise replace

u by $u+w$ and w by $-w$). Note that $\langle v,p\rangle = \langle v+w,p\rangle - \langle w,p\rangle > 1-1 = 0$, and similarly $\langle w,p\rangle > 0$. Pick some small $\lambda \in (0,1)$ (to be specified later). It is easy to see that if we write $\gamma(t) = \frac{v+2t\lambda w}{|v+2t\lambda w|}$ and $\delta(t) = u$ for all $t \in [0,1/2]$, then $\delta(t) \in I_{\gamma(t)}$ for all $t \in [0,1/2]$. Furthermore, if we write $q(s) = \frac{(1-s)(v+\lambda w)+sp}{|(1-s)(v+\lambda w)+sp|}$ for all $s \in [0,1]$, then $\langle q(s), v+w\rangle > 1$ for all $s \in [0,1]$. Indeed, it is easy to check that $\langle (1-s)(v+\lambda w)+sp, v+w\rangle > 0$, and we have

$$\begin{aligned}|(1-s)(v+\lambda w)+sp|^2 &= (1-s)^2(1+\lambda^2|w|^2) + s^2 + 2s(1-s)\langle p, v+\lambda w\rangle \\ &\leq (1-s)^2(1+\lambda^2|w|^2) + s^2 + 2s(1-s)\langle p, v+w\rangle\end{aligned}$$

and

$$\begin{aligned}\langle (1-s)(v+\lambda w)+sp, v+w\rangle^2 &= \big((1-s)(1+\lambda|w|^2) + s\langle p, v+w\rangle\big)^2 \\ &= (1-s)^2(1+\lambda|w|^2)^2 + s^2\langle p,v+w\rangle^2 + 2s(1-s)(1+\lambda|w|^2)\langle p, v+w\rangle \\ &> (1-s)^2(1+\lambda|w|^2) + s^2 + 2s(1-s)\langle p, v+w\rangle \\ &\geq (1-s)^2(1+\lambda^2|w|^2) + s^2 + 2s(1-s)\langle p, v+w\rangle,\end{aligned}$$

giving $\langle q(s), v+w\rangle^2 > 1$.

So $I_{q(s)}$ has non-empty interior for all s. Using Lemma 2.3.3 (and Lemma 2.3.2), it is easy to deduce that we can extend γ, δ to $[0,1]$ such that $\delta(t) \in I_{\gamma(t)}$ for all t and for all $t \geq 1/2$ there is some $s \in [0,1]$ such that $\gamma(t) = q(s)$.

Now, if $t \leq 1/2$, then

$$\begin{aligned}|v-\gamma(t)|^2 &= \left| v - \frac{v+2t\lambda w}{|v+2t\lambda w|}\right|^2 \\ &= 2 - 2\left\langle v, \frac{v+2t\lambda w}{|v+2t\lambda w|}\right\rangle \\ &= 2 - \frac{2}{|v+2t\lambda w|} \\ &= 2 - \frac{2}{(1+(2t\lambda|w|)^2)^{1/2}} \\ &\leq 2 - \frac{2}{(1+\lambda^2|w|^2)^{1/2}}.\end{aligned}$$

Furthermore, if $t \geq 1/2$ and $\gamma(t) = q(s)$, then

$$\begin{aligned}
|v - \gamma(t)|^2 &= \left| v - \frac{(1-s)(v+\lambda w) + sp}{|(1-s)(v+\lambda w) + sp|} \right|^2 \\
&= 2 - 2\left\langle v, \frac{(1-s)(v+\lambda w) + sp}{|(1-s)(v+\lambda w) + sp|} \right\rangle \\
&= 2 - 2\frac{(1-s) + s\langle v, p\rangle}{|(1-s)(v+\lambda w) + sp|} \\
&\leq 2 - 2\frac{\langle v, p\rangle}{|(1-s)(v+\lambda w) + sp|} \\
&\leq 2 - 2\frac{\langle v, p\rangle}{\max\{|v+\lambda w|, |p|\}} \\
&= 2 - \frac{2\langle v, p\rangle}{(1+\lambda^2|w|^2)^{1/2}}.
\end{aligned}$$

It follows that $|v - \gamma(t)| \leq \left(2 - \frac{2\langle v,p\rangle}{(1+\lambda^2|w|^2)^{1/2}}\right)^{1/2}$ for all t. But $\lambda \in (0,1)$ was arbitrary, and taking $\lambda \to 0^+$ we have $\left(2 - \frac{2\langle v,p\rangle}{(1+\lambda^2|w|^2)^{1/2}}\right)^{1/2} \to (2 - 2\langle v, p\rangle)^{1/2} = |v - p|$. It follows that we can choose λ such that $|v - \gamma(t)| \leq 2|v - p|$ for all t. $\square$

The next lemma is one of the main corollaries of Lemma 2.3.6.

Lemma 2.3.7. *Let $\epsilon > 0$, and suppose that $\gamma : [0,1] \to \mathbb{S}^{d-1}$ is continuously differentiable such that for all $t \in [0,1]$, one of the following holds.*

1. *$I_{\gamma(t)}$ has non-empty interior;*
2. *$I_{\gamma(t)}$ has empty interior, but there exist $u, w \in \mathbb{R}^d$ such that $u, u+w \in I_{\gamma(t)}$ and $\langle w, \gamma'(t)\rangle \neq 0$;*
3. *$I_{\gamma(t)}$ is a single point.*

Then $\gamma(1)$ is reachable from $\gamma(0)$ along a path which is ϵ-close to γ.

Note that in the second case we must have $\langle \gamma(t), w\rangle = 0$ by the final part of Lemma 2.3.5.

Proof. Let T_i be the set of $t \in [0,1]$ belonging to the ith case above ($i = 1, 2, 3$). We first claim that T_3 is closed. Indeed, it is easy to see that T_1 is open, and if $t \in T_2$ then by Lemma 2.3.6 there is some $\epsilon > 0$ such that $((t-\epsilon, t+\epsilon) \setminus \{t\}) \cap [0,1] \subseteq T_1$.

Since T_3 is closed, $[0,1] \setminus T_3$ is a union of disjoint (open) intervals: $[0,1] \setminus T_3 = \bigcup_{J \in \mathcal{J}} J$, where for all J either $J = (a_J, a_J + b_J)$ (with $0 \leq a_J < a_J + b_J \leq 1$), or $J = [0, b_J)$, or $J = (a_J, 1]$, or $J = [0,1]$ (and $J \cap J' = \emptyset$ if $J \neq J'$). For each $J \in \mathcal{J}$ (and positive integer m), define J_m as follows:

- If $J = (a_J, a_J + b_J)$, let $J_m = [a_J + \frac{1}{3m} b_J, a_J + (1 - \frac{1}{3m}) b_J]$;

- If $J = [0, b_J)$, let $J_m = [0, (1 - \frac{1}{3m})b_J]$;
- If $J = (a_J, 1]$, let $J_m = [a_J + \frac{1}{3m}(1 - a_J), 1]$;
- If $J = [0, 1]$, let $J_m = [0, 1]$.

Note that $\bigcup_{m \geq 1} J_m = J$ and $J_1 \subseteq J_2 \subseteq J_3 \subseteq \ldots$. Let us also write $J_0 = \emptyset$ for all J.

Observe that if $t \in T_1 \cup T_2$, then for any $\eta > 0$ there exists $\mu > 0$ such that if $t' \in (t - \mu, t + \mu) \cap [0, 1]$ then $\gamma(t')$ is reachable from $\gamma(t)$ along a path which is η-close to $\{\gamma(t)\}$. Indeed, this is easy to see (and follows from Lemma 2.3.3) when $t \in T_1$, and follows from Lemma 2.3.6 when $t \in T_2$.

Claim. We can recursively construct $\alpha_m : \bigcup_{J \in \mathcal{J}} J_m \to \mathbb{S}^{d-1}$ and $\beta_m : \bigcup_{J \in \mathcal{J}} J_m \to \mathbb{R}^d$ continuous functions such that

1. $\beta_m(t) \in I_{\alpha_m(t)}$ for all t (when defined);
2. If $m' > m$ then $\alpha_{m'}$ and $\beta_{m'}$ extend α_m and β_m, respectively;
3. For all $J \in \mathcal{J}$ and $m > 0$ we have $\alpha_m(\min J_m) = \gamma(\min J_m)$ and $\alpha_m(\max J_m) = \gamma(\max J_m)$;
4. If $t \in J_m \setminus J_{m-1}$, then there is some $t' \in J_m$ such that $|t - t'| < \text{length}(J)/m$ and $|\alpha_m(t) - \gamma(t')| < \min\{\epsilon, \text{length}(J)/m\}$.

Proof of Claim. It is enough to show that whenever $a < b$, $[a, b] \subseteq J$, $u \in I_{\gamma(a)}, v \in I_{\gamma(b)}$ and $\eta > 0$, then there exist $f : [a, b] \to \mathbb{S}^{d-1}$ and $g : [a, b] \to K$ continuous functions such that $g(t) \in I_{f(t)}$ for all t, $g(a) = u, g(b) = v, f(a) = \gamma(a), f(b) = \gamma(b)$, and for all t there is some $t' \in [a, b]$ with $|t - t'| < \eta$ such that $|f(t) - \gamma(t')| \leq \eta$. For each t, pick μ_t as in the observation above, we may assume $\mu_t < \eta$ for all t. Using the compactness of $[a, b]$, if N is large enough and we write $x(j) = a + j(b - a)/N$, then for all j we have $x(j) - x(j-1) < \eta$ and there is some t_j such that $|x(j) - t_j|, |x(j-1) - t_j| < \mu_{t_j}$. But then $\gamma(x(j)), \gamma(x(j+1))$ are both reachable from $\gamma(t_j)$ along a path which is η-close to $\{\gamma(t_j)\}$, and hence $x(j+1)$ is reachable from $x(j)$ along a path which is η-close to $\{\gamma(t_j)\}$. It follows that for any choice of $u_j \in I_{\gamma(x_j)}$ $(j = 0, \ldots, N)$ there exist $f_j : [x(j), x(j+1)] \to \mathbb{S}^{d-1}$ and $g_j : [x(j), x(j+1)] \to K$ such that $g_j(t) \in I_{f_j(t)}$ for all t, $g_j(x(j)) = u_j$, $g_j(x(j+1)) = u_{j+1}$, $f_j(x(j)) = \gamma(x(j))$, $f_j(x(j+1)) = \gamma(x(j+1))$, and for all t we have $|f_j(t) - \gamma(t_j)| \leq \eta$. Picking $u_0 = u$ and $u_N = v$ and then putting together these f_j, g_j gives the required functions f, g and finishes the proof of the claim. □

Define $\alpha : [0, 1] \to \mathbb{S}^{d-1}$ and $\beta : [0, 1] \to \mathbb{R}^d$ by setting $\alpha(t)$ to be $\alpha_m(t)$ and $\beta(t)$ to be $\beta_m(t)$ when $t \in T_1 \cup T_2$ (and m is large enough so that this exist), and when $t \in T_3$ then setting $\alpha(t) = \gamma(t)$ and $\beta(t)$ to be the unique point in $I_{\gamma(t)}$. It is clear that $\alpha(0) = \gamma(0)$, $\alpha(1) = \gamma(1)$, $\beta(t) \in I_{\alpha(t)}$ for all t, and α, β are continuous at all points in $T_1 \cup T_2$. Also, α is ϵ-close to γ. We show that α, β are continuous at all points in T_3 as well.

We first prove that if $t \in T_3$ then α is continuous at t. Take any sequence $(t_n) \to t$ in $[0,1]$, we show $(\alpha(t_n)) \to \alpha(t) = \gamma(t)$. If this is not true, then we can take a subsequence of $(\alpha(t_n))$ that converges to some $p \in \mathbb{S}^{d-1}$, $p \neq \gamma(t)$, so we may assume that $(\alpha(t_n))$ is convergent. Also, we may assume that $t_n \in T_1 \cup T_2$ for all n (since γ is continuous, and $\alpha(t') = \gamma(t')$ if $t' \in T_3$). We may also assume that (t_n) is either decreasing or increasing. Let $J(n) \in \mathcal{J}$ be such that $t_n \in J(n)$, and let $m(n)$ be the positive integer such that $t_n \in J(n)_{m(n)} \setminus J(n)_{m(n)-1}$. Furthermore, let t'_n be as given by point 4 above for $t_n \in J_{m(n)} \setminus J_{m(n)-1}$. Since (t_n) is either increasing or decreasing, either $J(n)$ is eventually constant and $m(n) \to \infty$, or $J(n)$ takes infinitely many different values and $\text{length}(J(n)) \to 0$. In either case, we have $\text{length}(J(n))/m(n) \to 0$. Hence $\alpha(t_n) - \gamma(t'_n) \to 0$ and $t_n - t'_n \to 0$. But then $t'_n \to t$ and hence $\gamma(t'_n) \to \gamma(t)$, which implies $\alpha(t_n) \to \gamma(t)$, as claimed.

We now show that β is also continuous at all $t \in T_3$. Assume that (t_n) is a sequence in $[0,1]$ converging to $t \in T_3$, we show $\beta(t_n) \to \beta(t)$. As before, by taking a subsequence we may assume that $\beta(t_n)$ converges to some $p \in K$. But $\beta(t_n) \in I_{\alpha(t_n)}$ for all n, i.e., $\beta(t_n), \beta(t_n) + \alpha(t_n) \in K$ for all n. Since K is closed and α is continuous, by taking limits we get $p, p + \alpha(t) \in K$, i.e., $p \in I_{\alpha(t)} = I_{\gamma(t)}$. But $I_{\gamma(t)} = \{\beta(t)\}$, hence $p = \beta(t)$, as claimed. □

We will attempt to find a 'good' path, i.e., one satisfying the conditions of Lemma 2.3.7. Note that the only case we need to avoid is having a point v on γ such that I_v has empty interior, is not a single point, and the tangent to γ at v is perpendicular to any $u - u'$ ($u, u' \in I_v$). To find such paths, it will be easier to work in $\mathbb{R}^{d-1}$ instead of on $\mathbb{S}^{d-1}$, using that locally they have the same structure. The next lemma captures the key property coming from Lemma 2.3.6 in terms of parametrizations.

While the formal statement is rather complicated, the lemma is intuitively quite simple, as we now explain. Let us focus on the case $d = 3$. Using Figure 2.3, we know that if γ is a path such that $\gamma(t)$ is a 'bad point', i.e., the conditions of Lemma 2.3.7 are not satisfied there, then we get the two blue circles touching at $v = \gamma(t)$ such that no point in the regions enclosed by the circles can be a bad point for any path. Moreover, we also know that γ must have tangent in the direction of the yellow dotted line at t. Our next lemma essentially states that if we take charts then we still get the blue circles whose interiors cannot contain bad points.

Lemma 2.3.8. *Let $\varphi : \mathbb{R}^{d-1} \to V$ be a smooth parametrization of some open set $V \subseteq \mathbb{S}^{d-1}$, and let $X \subseteq \mathbb{R}^{d-2}$ be an open neighbourhood of* 0. *Write $\gamma_x(t) = (t, x)$ for $x \in X, t \in [0,1]$ (so $\gamma_x : [0,1] \to \mathbb{R}^{d-1}$). Let Z be the set of all $(t,x) \in \mathbb{R}^{d-1}$ ($t \in [0,1], x \in X$) such that for $v = (\varphi \circ \gamma_x)(t) = \varphi(t,x)$ the set I_v has empty interior, but there exist $u, w \in \mathbb{R}^d$ such that $u, u+w \in I_v$, $w \neq 0$, $\langle v, w\rangle = 0$ and $\langle w, (\varphi \circ \gamma_x)'(t)\rangle = 0$. Let $X_Z = \{x \in X : (t,x) \in Z \text{ for some } t \in [0,1]\}$, and assume that $x \in X_Z$ and $t_x \in [0,1]$ are such that $(t_x, x) \in Z$. Then there is some $w_x \in \mathbb{R}^{d-2}$, $w_x \neq 0$ such that the open ball of radius $|w_x|$ centred at $(t_x, x + w_x)$ is disjoint from Z.*

The following lemma tells us that the conclusion of Lemma 2.3.8 guarantees that there are 'few' points we need to avoid.

Lemma 2.3.9. *Let $Z \subseteq \mathbb{R}^{d-1}$ and let X be an open neighbourhood of 0 in $\mathbb{R}^{d-2}$. Let $X_Z = \{x \in X : (t,x) \in Z \text{ for some } t \in [0,1]\}$, and for each $x \in X_Z$ let $t_x \in [0,1]$ be arbitrary such that $(t_x, x) \in Z$. Assume that for each $x \in X_Z$ there is some $w_x \in \mathbb{R}^{d-2}$, $w_x \neq 0$ such that the open ball of radius $|w_x|$ centred at $(t_x, x + w_x)$ is disjoint from Z. Then $X_Z \neq X$.*

Before we prove Lemma 2.3.8 and Lemma 2.3.9, let us first put them together to obtain the lemmas we will use later.

Lemma 2.3.10. *Let $\varphi : \mathbb{R}^{d-1} \to V$ be a smooth parametrization of some open set $V \subseteq \mathbb{S}^{d-1}$, and let $X \subseteq \mathbb{R}^{d-2}$ be an open neighbourhood of 0. Write $\gamma_x(t) = (t,x)$ for $x \in X, t \in [0,1]$ (so $\gamma_x : [0,1] \to \mathbb{R}^{d-1}$). Then there exists some $x \in X$ such that for all $\epsilon > 0$, $\varphi(\gamma_x(1))$ is reachable from $\varphi(\gamma_x(0))$ along a path which is ϵ-close to $\varphi \circ \gamma_x$.*

Proof. Define Z, X_Z and t_x (for $x \in X_Z$) as in Lemma 2.3.8. By Lemma 2.3.8, for each $x \in X_Z$ there is some $w_x \in \mathbb{R}^{d-2}$, $w_x \neq 0$ such that the open ball of radius $|w_x|$ centred at $(t_x, x + w_x)$ is disjoint from Z. So we can apply Lemma 2.3.9 to find some $x \in X$ such that $x \notin X_Z$. Then (using the final part of Lemma 2.3.5) we get that Lemma 2.3.7 applies for the path $\varphi \circ \gamma_x$ and hence $\varphi(\gamma_x(1))$ is reachable from $\varphi(\gamma_x(0))$ along a path which is ϵ-close to $\varphi \circ \gamma_x$. $\square$

For two points x and y in $\mathbb{R}^{d-1}$, let $\gamma_{x,y}$ denote the straight line segment from x to y (i.e., $\gamma_{x,y}(t) = (1-t)x + ty$ for $t \in [0,1]$). The following lemma is a more convenient version of Lemma 2.3.10.

Lemma 2.3.11. *Let $\varphi : \mathbb{R}^{d-1} \to V$ be a smooth parametrization of some open set $V \subseteq \mathbb{S}^{d-1}$. Let U_1, U_2 be non-empty open subsets of $\mathbb{R}^{d-1}$. Then there are some $x \in U_1, y \in U_2$ such that, for all $\epsilon > 0$, $\varphi(y)$ is reachable from $\varphi(x)$ along a path which is ϵ-close to $\varphi \circ \gamma_{x,y}$.*

Proof. We can take a bijective affine map $\psi : \mathbb{R}^{d-1} \to \mathbb{R}^{d-1}$ which maps U_1 to an open neighbourhood of 0 and U_2 to an open neighbourhood of $(1, 0, \dots, 0)$. Then the statement follows easily from Lemma 2.3.10 applied to the parametrization $\varphi \circ \psi^{-1}$. $\square$

We finish this subsection by giving the proofs of Lemmas 2.3.8 and 2.3.9.

Proof of Lemma 2.3.8. Let $v = \varphi(t_x, x)$. By the definition of Z, we may find $u, w \in \mathbb{R}^d$ such that $0 < |w| < 1$, $u, u + w \in I_v$, $\langle v, w \rangle = 0$ and $\langle w, (\varphi \circ \gamma_x)'(t) \rangle = 0$. By Lemma 2.3.6, the set $P = \{p \in \mathbb{S}^{d-1} : \langle p, v + w \rangle > 1\} \cup \{p \in \mathbb{S}^{d-1} : \langle p, v - w \rangle > 1\}$ has the property that for each $p \in P$, I_p has non-empty interior. In particular, $\varphi(Z)$ is disjoint from P.

By decreasing $|w|$ if necessary, we may assume that $P \subseteq V$. We want to show that for some $w' \in \mathbb{R}^{d-2}$ ($w' \neq 0$) the set $\varphi^{-1}(P)$ contains an open ball of radius $|w'|$ around $(t_x, x + w')$.

Let D be the derivative $D\varphi|_{(t_x,x)}$ of φ at (t_x, x), so D is a bijective linear map $\mathbb{R}^{d-1} \to \{v' \in \mathbb{R}^d : \langle v, v' \rangle = 0\}$. We can find an orthonormal basis $f_1, \ldots, f_{d-2}$ of $\mathbb{R}^{d-2}$ such that $f_1 = (1, 0, \ldots, 0)$, $\langle D(f_2), w \rangle > 0$ and $\langle D(f_i), w \rangle = 0$ for all $i \neq 2$. Consider the ball of radius ρ centred at $(t_x, x) + \rho f_2$. Any point of this open ball is of the form $q = (t_x, x) + \sum_{i=1}^{d-2} \lambda_i f_i$ with $(\lambda_2 - \rho)^2 + \sum_{i \neq 2} \lambda_i^2 < \rho^2$. But we have

$$\varphi(q) = v + \sum_{i=1}^{d-2} \lambda_i D(f_i) + O(\sum_{i=1}^{d-2} \lambda_i^2)$$

and hence

$$\varphi(q) = v + \sum_{i=1}^{d-2} \lambda_i D(f_i) + O(2\rho\lambda_2).$$

Using that $\langle v, D(f_i) \rangle = 0$ for all i and $\langle w, D(f_j) \rangle = 0$ for all $j \neq 2$,

$$\begin{aligned} \langle v + w, \varphi(q) \rangle &= \langle v + w, v + \sum_{i=1}^{d-2} \lambda_i D(f_i) + O(2\rho\lambda_2) \rangle \\ &= 1 + \lambda_2 \langle w, D(f_2) \rangle + O(2\rho\lambda_2). \end{aligned}$$

Since $\langle w, D(f_2) \rangle > 0$, we get that there is some $\rho_0 > 0$ such that if $\rho \leq \rho_0$ then $\langle v+w, \varphi(q) \rangle > 1$ (and hence $q \in \varphi^{-1}(P)$ and thus $q \notin Z$) for all such points q. Since f_1 is orthogonal to f_2, we have $f_2 = (0, y)$ for some $y \in \mathbb{R}^{d-2}$, $|y| = 1$. Then $w_x = \rho_0 y$ satisfies the conditions. □

Before we formally prove Lemma 2.3.9, let us give a sketch proof in the case when $d = 3$, $X = (-1, 1)$ and $w_x \in \mathbb{R}$ is the same for all x: $w_x = r \in (0, 1)$ for all $x \in X$. Assume that $X_Z = X$. Using that the circle of radius r centred at $(x, t_x + r)$ does not contain (y, t_y), it is easy to see that we must have $|t_y - t_x| \geq \Omega_r(\sqrt{y - x})$ whenever $0 < y - x < r$ (see Figure 2.4). So if we take $N + 1$ equally spaced points $x_0, \ldots, x_N$ between 0 and r ($x_j = jr/N$), then $|t_{x_i} - t_{x_j}| \geq \Omega_r(1/\sqrt{N})$ for all i, j. It is easy to see that this gives a contradiction as $N \to \infty$. We will use the Baire category theorem to reduce the general case to a case similar enough to the one discussed above.

Proof of Lemma 2.3.9. For each positive integer n, let $X_Z^n = \{x \in X_Z : |w_x| \geq 1/n\}$. Clearly, $X_Z = \bigcup_n X_Z^n$. By the Baire category theorem, it is enough to show that each X_Z^n is a finite union of nowhere dense sets. Assume, for contradiction, that X_Z^n cannot be written as such a finite union. Let $\eta = 1/4$, and for all $v \in \mathbb{S}^{d-3}$ let $U_v = \{u \in \mathbb{S}^{d-3} : \langle u, v \rangle > 1 - \eta\}$. Since $\mathbb{S}^{d-3}$ is compact, it is covered by finitely many such sets U_v. Write $Y_v = \{x \in X_Z^n : w_x/|w_x| \in U_v\}$. It follows that not every Y_v is nowhere dense, i.e., there exist $v \in \mathbb{S}^{d-3}$, $y \in X$ and $\epsilon > 0$ such that the closure of Y_v contains all $x \in X$ with $|x - y| \leq \epsilon$. We may assume $\epsilon < 1/n$. Write $x(j) = y + \frac{j}{N}\epsilon v$ for $j = 0, \ldots, N$, where N is some large positive integer (specified later). Note

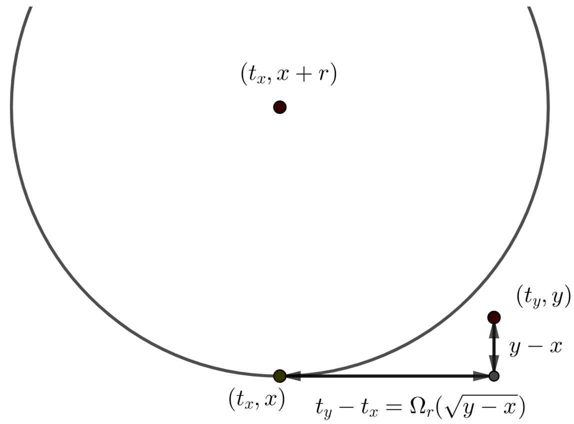

Figure 2.4: Since (t_y, y) is not contained in the ball of radius r centred at $(t_x, x+r)$, we have $|t_y - t_x| = \Omega_r(\sqrt{y-x})$.

that $|x(j) - y| \leq \epsilon$ for all j, so there are some $y(j) \in Y_v$ such that $|x(j) - y(j)| < \eta/N^2$.

Claim. If $0 \leq i < j \leq N$ then $|t_{y(i)} - t_{y(j)}| = \Omega_{n,\epsilon}(1/N^{1/2})$.

Note that if the claim holds, then $\max_i t_{y(i)} - \min_i t_{y(i)} = \Omega_{n,\epsilon}(N^{1/2})$. Then taking N large enough gives a contradiction. So the lemma follows from the claim above.

Proof of Claim. We will use that the open ball centred at $(t_{y(i)}, y(i)+w_{y(i)})$ of radius $|w_{y(i)}|$ does not contain $(t_{y(j)}, y(j))$. For simplicity, let us write t_i for $t_{y(i)}$, t_j for $t_{y(j)}$ and w for $w_{y(i)}$. We may assume $|w| = 1/n$. We have

$$|(t_i, y(i)+w) - (t_j, y(j))|^2 = |t_i - t_j|^2 + |y(i) + w - y(j)|^2.$$

But

$$\begin{aligned}
|y(i) + w - y(j)|^2 &= |w|^2 + |y(i) - y(j)|^2 - 2\langle w, y(j) - y(i)\rangle \\
&= |w|^2 + |y(i) - y(j)|^2 - 2\langle w, x(j) - x(i)\rangle - 2\langle w, y(j) - x(j)\rangle + 2\langle w, y(i) - x(i)\rangle \\
&\leq |w|^2 + |y(i) - y(j)|^2 - 2\langle w, \frac{j-i}{N}\epsilon v\rangle + 4\frac{\eta}{nN^2} \\
&\leq |w|^2 + (|x(i) - x(j)| + 2\eta/N^2)^2 - 2\frac{j-i}{N}\epsilon(1-\eta)/n + 4\frac{\eta}{nN^2} \\
&= |w|^2 + \left(\frac{j-i}{N}\epsilon + 2\eta/N^2\right)^2 - 2\frac{j-i}{N}\epsilon(1-\eta)/n + 4\frac{\eta}{nN^2}.
\end{aligned}$$

But we know $|w|^2 \le |(t_i, y(i)+w) - (t_j, y(j))|^2$, thus

$$\begin{aligned}|t_i - t_j|^2 &\ge 2\frac{j-i}{N}\epsilon(1-\eta)/n - \left(\frac{j-i}{N}\epsilon + 2\eta/N^2\right)^2 - 4\frac{\eta}{nN^2} \\ &= \frac{2(j-i)\epsilon(1-\eta)}{n}\frac{1}{N} - \frac{(j-i)^2\epsilon^2}{N^2} - O_{n,\epsilon}(1/N^2)\end{aligned}$$

Using $\frac{(j-i)^2\epsilon^2}{N^2} \le \frac{(j-i)\epsilon}{n}\frac{1}{N}$ (as $j-i \le N$ and $\epsilon \le 1/n$), we get

$$|t_i - t_j|^2 \ge \left(\frac{2(j-i)\epsilon(1-\eta)}{n} - \frac{(j-i)\epsilon}{n}\right)\frac{1}{N} - O_{n,\epsilon}(1/N^2).$$

As we picked $\eta = 1/4$, we get

$$|t_i - t_j|^2 \ge \frac{(j-i)\epsilon}{2n}\frac{1}{N}(1 - O_{n,\epsilon}(1/N)),$$

and hence

$$|t_i - t_j| \ge \left(\frac{\epsilon}{2n}\right)^{1/2}\frac{1}{N^{1/2}}(1 - O_{n,\epsilon}(1/N)),$$

proving the claim and hence the lemma. □

2.3.3 Finishing the proof

We now use our earlier lemmas (especially Lemma 2.3.11 and Lemma 2.3.6) to finish the proof of Theorem 2.3.1.

Lemma 2.3.12. *Let $\varphi : \mathbb{R}^{d-1} \to V$ be a smooth parametrization of some open set $V \subseteq \mathbb{S}^{d-1}$, and let $\epsilon > 0$. Assume that $v, v' \in V$ are such that $I_v, I_{v'}$ are non-empty. Then v' is reachable from v along a path which is ϵ-close to $\varphi \circ \gamma_{\varphi^{-1}(v),\varphi^{-1}(v')}$.*

Proof. Write u for $\varphi^{-1}(v)$ and u' for $\varphi^{-1}(v')$. As $I_v, I_{v'}$ are non-empty, there is an open set containing v and v' such that whenever p belongs to this set then I_p has non-empty interior. By Lemma 2.3.3, there are open balls $U_1, U_2 \subseteq \mathbb{R}^{d-1}$ around u and u' (respectively) such that for any $x \in U_1$, $\varphi(x)$ is reachable from v along a path which is 0-close to $\varphi \circ \gamma_{u,x}$, and similarly for any $y \in U_2$, $\varphi(y)$ is reachable from v' along a path which is 0-close to $\varphi \circ \gamma_{u',y}$. Pick $\eta > 0$ small (to be specified later). By Lemma 2.3.11, we can find $x \in U_1$, $|x-u| < \eta$ and $y \in U_2$, $|y-u'| < \eta$ such that $\varphi(y)$ is reachable from $\varphi(x)$ along a path which is $(\epsilon/2)$-close to $\varphi\circ\gamma_{x,y}$. It follows that u' is reachable from u along a path which is $(\epsilon/2)$-close to the union of the images of $\varphi \circ \gamma_{u,x}$, $\varphi \circ \gamma_{x,y}$, $\varphi \circ \gamma_{y,u'}$. However, by taking η small enough, we can guarantee that all points in these images are at most $\epsilon/2$ away from a point in the image of $\varphi \circ \gamma_{u,u'}$, proving the lemma. □

To extend Lemma 2.3.12 to all v, v', including when I_v or $I_{v'}$ is a single point, we will use the following lemma.

Lemma 2.3.13. *Let $\varphi : \mathbb{R}^{d-1} \to V$ be a smooth parametrization of some open set $V \subseteq \mathbb{S}^{d-1}$. Assume that $v \in V$ and I_v is a single point. Then one of the following statements hold.*

1. *For all $\eta > 0$ there is some $p \in V$ such that I_p has non-empty interior and p is reachable from v along a path which is η-close to $\{v\}$.*

2. *There is some open neighbourhood N of v such that whenever $p \in N$ then p is reachable from v along $\varphi \circ \gamma_{\varphi^{-1}(v),\varphi^{-1}(p)}$.*

Proof. First, assume that there is a sequence of points (p_n) in V converging to v such that for all n, I_{p_n} is not a single point. We will show that the first conclusion holds. By Lemma 2.3.6 (and the final part of Lemma 2.3.5), we may modify p_n slightly so that I_{p_n} has non-empty interior for all n. Let $\eta > 0$ be given. By Lemma 2.3.12, we can take $\gamma_n : [0,1] \to V$ and $\delta_n : [0,1] \to K$ continuous functions such that $\delta_n(t) \in I_{\gamma_n(t)}$ for all (n,t), $\gamma_n(0) = p_n$ for all n, $\gamma_n(1) = p_{n+1}$ for all n, and $\gamma_n(t)$ is at most $\eta/2^n$ away from some point on the image of $\varphi \circ \gamma_{\varphi^{-1}(p_n),\varphi^{-1}(p_{n+1})}$ for all (n,t). By taking a subsequence of the form $(p_n)_{n>N_0}$, we may assume that for all n, all points on the image of $\varphi \circ \gamma_{\varphi^{-1}(p_n),\varphi^{-1}(p_{n+1})}$ are at most $\eta/2$ away from v. So $|\gamma_n(t) - v| \leq \eta$ for all (n,t). Using Lemma 2.3.2, we may also assume that $\delta_n(1) = \delta_{n+1}(0)$ for all n.

Now define $\gamma : [0,1] \to V$ and $\delta : [0,1] \to K$ as follows. Let $\gamma(0) = v$ and let $\delta(0)$ be the unique point in I_v. For $t \in (0,1]$, let n be such that $\frac{1}{n+1} \leq t \leq \frac{1}{n}$, and set $\gamma(t) = \gamma_n(n(n+1)(\frac{1}{n} - t))$ and $\delta(t) = \delta_n(n(n+1)(\frac{1}{n} - t))$. It is easy to check that γ, δ are well-defined and continuous on $(0,1]$, and $|\gamma(t) - v| \leq \eta$ for all t. Moreover, using that $(p_n) \to v$ and $\gamma_n(t)$ is at most $\eta/2^n$ away from some point on the image of $\varphi \circ \gamma_{\varphi^{-1}(p_n),\varphi^{-1}(p_{n+1})}$ for all (n,t), we also get that γ is continuous at 0. To show continuity of δ at 0, assume that $(t_n) \to 0$ and $(\delta(t_n)) \to z$, we prove $z = \delta(0)$. We know $\delta(t_n), \delta(t_n) + \gamma(t_n) \in K$. Using that K is closed and γ is continuous, taking limits gives $z, z + \gamma(0) \in K$, i.e., $z, z + v \in K$, i.e., $z \in I_v$. Hence $z = \delta(0)$, as claimed. This proves the claim in the first case.

Now assume that such a sequence (p_n) does not exist. This means that there is an open neighbourhood of v consisting only of points p such that I_p is a single point. It follows that there is an open ball B around $u = \varphi^{-1}(v)$ such that whenever $x \in B$ then $I_{\varphi(x)}$ is a single point. Let $N = \varphi(B)$, so N is an open neighbourhood of v. Given $p \in N$, let $\varphi^{-1}(p) = q$. We show p is reachable from v along $\varphi \circ \gamma_{u,q}$. Indeed, let $\gamma(t) = \varphi((1-t)u + tq)$ and let $\delta(t)$ be the unique point in $I_{\gamma(t)}$. Then δ is continuous by an argument almost identical to the one above. Indeed, if $(t_n) \to t$ and $(\delta(t_n)) \to z$, then $\delta(t_n), \delta(t_n) + \gamma(t_n) \in K$. Taking limits gives $z, z + \gamma(t) \in K$, i.e., $z \in I_{\gamma(t)}$, i.e., $z = \delta(t)$, as required. This finishes the proof of the lemma. $\square$

Lemma 2.3.14. *Let $\varphi : \mathbb{R}^{d-1} \to V$ be a smooth parametrization of some open set $V \subseteq \mathbb{S}^{d-1}$, and let $\epsilon > 0$. Then for any $v, v' \in V$, v' is reachable from v along a path which is ϵ-close to $\varphi \circ \gamma_{\varphi^{-1}(v),\varphi^{-1}(v')}$.*

Proof. Write u for $\varphi^{-1}(v)$ and u' for $\varphi^{-1}(v')$. Let $\eta > 0$ be small (specified later). There is some open set $V_1 \subseteq V$ (not necessarily containing v) such that any $p \in V_1$ is reachable from v along a path which is η-close to $\{v\}$. Indeed, this follows from Lemma 2.3.6 if I_v is not a single point, and from Lemma 2.3.13 (together with Lemma 2.3.3) when I_v is a single point. Similarly, there is some $V_2 \subseteq V$ such that any $q \in V_2$ is reachable from v' along a path which is η-close to $\{v'\}$. In particular, $|v - p| \leq \eta$ and $|v' - q| \leq \eta$ for any such p, q.

But, by Lemma 2.3.11, there are some $p \in V_1, q \in V_2$ such that q is reachable from p along a path which is η-close to $\varphi \circ \gamma_{\varphi^{-1}(p),\varphi^{-1}(q)}$. Hence v' is reachable from v along a path which is η-close to $\{v, v'\} \cup \mathrm{Im}(\varphi \circ \gamma_{\varphi^{-1}(p),\varphi^{-1}(q)})$. By taking η small enough, we can guarantee that any point in $\mathrm{Im}(\varphi \circ \gamma_{\varphi^{-1}(p),\varphi^{-1}(q)})$ is at most $\epsilon/2$ away from some point in $\mathrm{Im}(\varphi \circ \gamma_{\varphi^{-1}(v),\varphi^{-1}(v')})$. The result follows. $\square$

Proof of Theorem 2.3.1. Using Lemma 2.3.14, it is easy to see that for any $t \in [0, 1]$ there is some $\delta_t > 0$ such that whenever $t' \in [0, 1]$ and $|t - t'| < \delta_t$, then $\gamma(t')$ is reachable from $\gamma(t)$ along a path which is ϵ-close to $\{\gamma(t)\}$. The result follows easily (using the compactness of $[0, 1]$ and Lemma 2.3.2). $\square$

Proof of Theorem 2.1.2. The result follows immediately from Theorem 2.3.1 when $d \geq 3$, and from Theorem 2.1.4 when $d = 2$ (using Lemma 2.3.2, which also holds for $d = 2$). $\square$

2.4 Counterexample for general bodies

In this section, our goal is to prove Theorem 2.1.3, restated below for convenience.

Theorem 2.1.3. *There exist convex bodies S and K in $\mathbb{R}^4$ such that K is S-Kakeya but there are two S-copies which cannot be rotated into each other within K.*

We will use similar ideas as for Theorem 2.1.5 (but this proof will be significantly more complicated). Note that it is sufficient to find a construction where S is compact but not necessarily convex, as the same set K will still provide a counterexample when S is replaced by its convex hull. The set S in our construction will be given by

$$S = \{(x, y, z, w) \in \mathbb{R}^4 : x^2 + y^2 + z^2 + w^2 = 1, x = \pm 1/2\},$$

see Figure 2.5.

In the proof of Theorem 2.1.5 we made sure that our set K lied inside the cylinder $\{(x, y, z) : y^2 + z^2 \leq 1/4\}$, and we controlled the intersection with the boundary of the cylinder. This control enabled us to prove discontinuity by observing that any segment in a direction of the yz plane had to intersect the boundary of the cylinder in a pair of points $(x, y, z), (x, -y, -z)$.

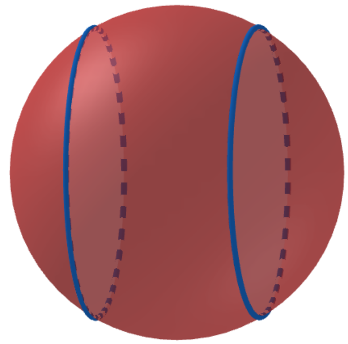

Figure 2.5: The set S in our construction is a 4-dimensional analogue of the blue set (or the convex hull of the blue set), which is a subset of the (red) unit sphere

We will attempt to do something similar here. Our construction will be contained inside the set $\{(x, y, z, w) : x^2 + y^2 \leq 1\}$, and we will control the intersection with the boundary $C = \{(x, y, z, w) \in \mathbb{R}^4 : x^2 + y^2 = 1\}$ of that set. Observe that any rotated copy of S is of the form

$$S_v = \{v' \in \mathbb{R}^4 : |v'| = 1, \langle v, v' \rangle = \pm 1/2\}$$

for some $v \in \mathbb{R}^4$ with $|v| = 1$. It is not difficult to deduce that if we only rotate S slightly, then the rotated copy intersects C in two pairs of antipodal points. (See Figure 2.5: great circles close to the one given by $x^2 + y^2 = 1, z = 0$ intersect the blue set in two pairs of antipodal points). We will have to make sure that K contains translated copies of any two such pairs of antipodal points (so that a translate of $\rho(S)$ is contained in K for all ρ), so for all $(x_1, y_1), (x_2, y_2) \in \mathbb{S}^1$ we will have some (z, w) such that $(\pm(x_1, y_1), z, w), (\pm(x_2, y_2), z, w) \in C \cap K$. Meanwhile, we will have restrictions on $C \cap K$ in such a way that we guarantee discontinuity.

Let us now turn to the formal proof of Theorem 2.1.3. As mentioned above, we will control the intersection of K with C, i.e., for all $(x, y) \in \mathbb{S}^1$ we will control the set $A_{x,y} = \{(z, w) : (x, y, z, w) \in K\}$. The following lemma lists all the properties that we will need – for now, we only show that such sets $A_{x,y}$ exist in $\mathbb{R}^2$, at this point they do not necessarily come from a body K in $\mathbb{R}^4$.

Lemma 2.4.1. *There exist compact convex sets $(A_p)_{p \in \mathbb{S}^1}$ in $\mathbb{R}^2$ such that the following properties hold.*

1. *For all $p, q \in \mathbb{S}^1$, $A_p \cap A_q \neq \emptyset$.*
2. *For all $p \in \mathbb{S}^1$, $A_p = A_{-p}$.*
3. *For all $p \in \mathbb{S}^1$ and all $t \in A_p$ we have $|t| \leq 1$.*
4. *The set $\{(p, t) : p \in \mathbb{S}^1, t \in A_p\}$ is closed, i.e, whenever $(p_n) \to p$ in $\mathbb{S}^1$ and $(t_n) \to t$ in $\mathbb{R}^2$ with $t_n \in A_{p_n}$ for all n, then $t \in A_p$.*

5. *For all $\epsilon > 0$ and $(z, w) \in \mathbb{R}^2$ there is some $r \in (0, \epsilon)$ such that whenever $p = (x, y) \in \mathbb{S}^1$ with $|x - 1/2| = r$ then all points of A_p are at least distance $1/100$ away from (z, w).*

Note that such sets A_p cannot exist in $\mathbb{R}$ instead of $\mathbb{R}^2$: each A_p would have to be a non-empty closed bounded interval, and then $\bigcap_{p\in\mathbb{S}^1} A_p$ would be non-empty by the first condition, so the last property could not be satisfied. This is the reason we need $\mathbb{R}^4$ for our construction instead of $\mathbb{R}^3$.

Before we prove Lemma 2.4.1, we state two lemmas which show why it is useful: Theorem 2.1.3 will follow immediately from Lemma 2.4.1 and these lemmas. Recall that $C = \mathbb{S}^1 \times \mathbb{R}^2$ and $S_v = \{v' \in \mathbb{R}^4 : |v'| = 1, \langle v, v' \rangle = \pm 1/2\}$.

Lemma 2.4.2. *Assume that we have compact convex sets $(A_p)_{p\in\mathbb{S}^1}$ in $\mathbb{R}^2$ such that the following properties hold.*

1. *For all $p, q \in \mathbb{S}^1$, $A_p \cap A_q \neq \emptyset$.*
2. *For all $p \in \mathbb{S}^1$, $A_p = A_{-p}$.*
3. *For all $p \in \mathbb{S}^1$ and all $t \in A_p$ we have $|t| \leq 1$.*
4. *The set $\{(p, t) : p \in \mathbb{S}^1, t \in A_p\}$ is closed, i.e, whenever $(p_n) \to p$ in $\mathbb{S}^1$ and $(t_n) \to t$ in $\mathbb{R}^2$ with $t_n \in A_{p_n}$ for all n, then $t \in A_p$.*

Then there exists a compact convex S-Kakeya set $K \subseteq \mathbb{R}^4$ such that $K \subseteq \{(x, y, z, w) \in \mathbb{R}^4 : x^2 + y^2 \leq 1\}$ and $K \cap C \subseteq \{(p, t) : p \in \mathbb{S}^1, t \in A_p\}$.

Lemma 2.4.3. *Assume that $(A_p)_{p\in\mathbb{S}^1}$ in $\mathbb{R}^2$ are compact convex sets such that the following property holds: for all $\epsilon > 0$ and $(z, w) \in \mathbb{R}^2$ there is some $r \in (0, \epsilon)$ such that whenever $p = (x, y) \in \mathbb{S}^1$ with $|x-1/2| = r$ then all points of A_p are at least distance $1/100$ away from (z, w). Assume furthermore that K is a compact convex set such that $K \subseteq \{(x, y, z, w) \in \mathbb{R}^4 : x^2+y^2 \leq 1\}$ and $K \cap C \subseteq \{(p, t) : p \in \mathbb{S}^1, t \in A_p\}$. Then whenever $\gamma : [0, 1] \to \mathbb{S}^3$ and $\delta : [0, 1] \to \mathbb{R}^4$ are continuous such that $\gamma(0) = (1, 0, 0, 0)$ and $S_{\gamma(t)} + \delta(t) \subseteq K$ for all t, then $\gamma(t) = (1, 0, 0, 0)$ for all t.*

We now prove Lemmas 2.4.1, 2.4.2 and 2.4.3.

Proof of Lemma 2.4.1. Consider the following four sets in $\mathbb{R}^2$:

$$\begin{aligned}
T_1 &= \{0\} \times [0, 1], \\
T_2 &= [0, 1] \times \{0\}, \\
T_3 &= \{(z, w) \in \mathbb{R}^2 : z + w = 1, 0 \leq z, w \leq 1\}, \\
T &= \{(z, w) \in \mathbb{R}^2 : 0 \leq z, w \leq 1, 0 \leq z + w \leq 1\}.
\end{aligned}$$

Given (x, y) with $x^2 + y^2 = 1$, we define $A_{x,y}$ as follows. Let $\min(|x - 1/2|, |x + 1/2|) = s$. If $s = 0$, then $A_{x,y} = T$. Otherwise, let k be the positive integer such that $1/2^k \geq s > 1/2^{k+1}$. If $s = 1/2^k$, then let $A_{x,y} = T$. Otherwise let $A_{x,y} = T_{k \bmod 3}$.

It is straightforward to check that each A_p is convex and compact, and that properties 1, 2 and 3 are satisfied. To see that property 4 holds, observe that if $(p_n) \to p$ and $(t_n) \to t$ as above, then either $A_p = T$, or A_{p_n} is eventually constant and equal to A_p. In either case, it is easy to deduce that $t \in A_p$.

Finally, we show that property 5 holds. Given such (z, w), we can find some $i \in \{1, 2, 3\}$ such that any point in T_i has distance at least $1/100$ from (z, w). Then we can find some $r \in (0, \epsilon)$ such that $1/2^k > r > 1/2^{k+1}$ for some positive integer k with $k \equiv i \bmod 3$. It is easy to see that this r satisfies the conditions. $\square$

Proof of Lemma 2.4.2. Observe that if $v \in \mathbb{S}^3$, the set S_v intersects C in 0, 2 or 4 points:

- S_v intersects C in 0 points if and only if $v_1^2 + v_2^2 < 1/4$;
- S_v intersects C in a pair of points $v', -v'$ if and only if $v_1^2 + v_2^2 = 1/4$;
- S_v intersects C in two pairs of (distinct) points $v', -v', v'', -v''$ if and only if $v_1^2 + v_2^2 > 1/4$.

Let

$$V_1 = \{v \in \mathbb{S}^3 : v_1^2 + v_2^2 \geq 1/4\},$$
$$V_2 = \{v \in \mathbb{S}^3 : 1/100 \leq v_1^2 + v_2^2 \leq 1/4\},$$
$$V_3 = \{v \in \mathbb{S}^3 : v_1^2 + v_2^2 \leq 1/100\}.$$

For all $v \in V_1$, let $T_v = \bigcap_{w \in C \cap S_v} A_{w_1, w_2} = \bigcap_{p \in \mathbb{S}^1 : (p,0,0) \in S_v} A_p$. Observe that we have $C \cap S_v = \{v', -v', v'', -v''\}$ for some $v', v'' \in \mathbb{S}^3$ (not necessarily distinct), so (using $A_{-p} = A_p$) we have $T_v = A_{v_1', v_2'} \cap A_{v_1'', v_2''}$. In particular, $T_v \neq \emptyset$. Let $K_{1,v} = S_v + \{(0, 0, t) : t \in T_v\}$ and

$$K_1 = \bigcup_{v \in V_1} K_{1,v}.$$

For all $v \in V_2$, let $p_v = \left(\frac{1/2}{\sqrt{v_1^2+v_2^2}} v_1, \frac{1/2}{\sqrt{v_1^2+v_2^2}} v_2, \frac{\sqrt{3}/2}{\sqrt{v_3^2+v_4^2}} v_3, \frac{\sqrt{3}/2}{\sqrt{v_3^2+v_4^2}} v_4 \right)$. So $(p_v)_1^2 + (p_v)_2^2 = 1/4$, $|p_v| = 1$, and we have $p_v = v$ if $v_1^2 + v_2^2 = 1/4$. Let $K_{2,v} = S_v + \{(0, 0, t) : t \in T_{p_v}\}$. (Note that $C \cap S_{p_v} = \{v', -v'\}$, where $v' = 2((p_v)_1, (p_v)_2, 0, 0)$ and hence $T_{p_v} = A_{2(p_v)_1, 2(p_v)_2}$.) Let

$$K_2 = \bigcup_{v \in V_2} K_{2,v}.$$

For all $v \in V_3$, let $K_{3,v} = S_v$, and let

$$K_3 = \bigcup_{v \in V_3} K_{3,v}.$$

Finally, let $K_0 = K_1 \cup K_2 \cup K_3$, and let K be the convex hull of K_0.

Claim. The set K_0 has the following properties.

1. For each $v \in \mathbb{S}^3$ there is some $w \in \mathbb{R}^4$ such that $S_v + w \subseteq K_0$.

2. We have $K_0 \cap C \subseteq \{(p,t) : p \in \mathbb{S}^1, t \in A_p\}$, and K_0 has no point (x,y,z,w) with $x^2+y^2 > 1$.

3. The set K_0 is compact.

Note that these properties are preserved when taking convex hull. So the claim above implies the statement of the lemma.

Proof of Claim. The first property holds because $K_{i,v}$ contains a translate of S_v if $v \in V_i$. To see that the second property holds, observe that K_0 is a union of sets of the form $S_v + (0,0,t)$ for some $t \in \mathbb{R}^2$. It follows that K_0 has no point (x,y,z,w) with $x^2+y^2 > 1$. Also, if $(p,t) \in K_0 \cap C$ ($p \in \mathbb{S}^1, t \in \mathbb{R}^2$), then $(p,t) \in S_v + (0,0,t)$ for some $v \in \mathbb{S}^3$ having $v_1^2 + v_2^2 \geq 1/4$, and $t \in T_v$ and $(p,0,0) \in C \cap S_v$. But $(p,0,0) \in C \cap S_v$ implies $T_v \subseteq A_p$, so $t \in A_p$, as claimed. It is easy to see that K_0 is bounded, so the only property left to check is that K_0 is closed. It is enough to show that K_1, K_2, K_3 are all closed.

We first show that K_3 is closed. Assume that (q_n) is a sequence of points in K_3 with $(q_n) \to q$, we show that $q \in K_3$. We know $q_n \in S_{v(n)}$ for some $v(n) \in V_3$. By taking an appropriate subsequence, we may assume that $v(n)$ converges to some $v \in V_3$. It is easy to see that $q \in S_v$ must hold, so then $q \in K_3$.

Next, we show that K_2 is closed. As before, assume that (q_n) is a sequence of points in K_2 with $(q_n) \to q$. We have $q_n \in S_{v(n)} + (0,0,t_n)$ for some $v(n) \in V_2$ and $t_n \in T_{p_{v(n)}} = A_{2(p_{v(n)})_1, 2(p_{v(n)})_2}$. By taking a subsequence, we may assume that $v(n)$ converges to some $v \in V_2$, and (t_n) converges to some $t \in \mathbb{R}^2$. Observe that $(p_{v(n)}) \to p_v$. But then $t \in A_{2(p_v)_1, 2(p_v)_2} = T_{p_v}$ and hence $q \in S_v + (0,0,t) \subseteq S_v + \{(0,0,t') : t' \in T_{p_v}\}$, so $q \in K_2$, as required.

Finally, we show that K_1 is also closed. Again, assume that (q_n) is a sequence of points in K_1 with $(q_n) \to q$. We have $q_n \in S_{v(n)} + (0,0,t_n)$ for some $v(n) \in V_1$ and $t_n \in T_{v(n)}$. As before, by taking a subsequence we may assume that $v(n)$ converges to some $v \in V_1$ and t_n converges to some $t \in \mathbb{R}^2$. We claim that this implies $t \in T_v$. Observe that $C \cap S_{v(n)}$ is of the form $\{v'(n), -v'(n), v''(n), -v''(n)\}$, where $v'(n) = \pm v''(n)$ if and only if $v(n)_1^2 + v(n)_2^2 = 1/4$. So we have

$$T_{v(n)} = A_{v'(n)_1, v'(n)_2} \cap A_{v''(n)_1, v''(n)_2}.$$

By taking an appropriate subsequence, we may assume that $v'(n)$ converges to v' and $v''(n)$ converges to v'', where $C \cap S_v = \{v', -v', v'', -v''\}$. But we have $t_n \in A_{v'(n)_1, v'(n)_2}$ for all n, and hence $t \in A_{v'_1, v'_2}$. Similarly, $t \in A_{v''_1, v''_2}$. Hence $t \in T_v$, as claimed. But then

$$q \in S_v + (0, 0, t) \subseteq S_v + \{(0, 0, t') : t' \in T_v\} = K_{1,v} \subseteq K_1,$$

as claimed. This finishes the proof of the claim and hence the lemma. □

Proof of Lemma 2.4.3. Assume, for contradiction, that $\gamma(t) \neq (1, 0, 0, 0)$ for some t. We may assume that $\gamma(t)_1 > 9/10$ for all t, and that for all $t > 0$ we have $\gamma(t) \neq (1, 0, 0, 0)$. There are some continuous functions $v', v'' : [0, 1] \to C$ such that $S_{\gamma(t)} \cap C = \{v'(t), -v'(t), v''(t), -v''(t)\}$, $\langle \gamma(t), v'(t) \rangle = \langle \gamma(t), v''(t) \rangle = 1/2$ and $v'(0), v''(0) = (1/2, \pm\sqrt{3}/2, 0, 0)$.

Observe that if $\gamma(t) \neq (1, 0, 0, 0)$ then $v'(t)_1 \neq 1/2$ or $v''(t)_1 \neq 1/2$. Indeed, we would have $v'(t), v''(t) = (1/2, \pm\sqrt{3}/2, 0, 0)$ and $1 = \langle \gamma(t), v'(t) + v''(t) \rangle = \langle \gamma(t), (1, 0, 0, 0) \rangle$, giving $\gamma(t) = (1, 0, 0, 0)$. It follows that for all $t > 0$, either $v'(t)_1 \neq 1/2$ or $v''(t)_1 \neq 1/2$.

By continuity, there is some $\epsilon > 0$ such that for all $t \leq \epsilon$ we have $|\delta(t) - \delta(0)| < 1/100$. We know $v'(\epsilon)_1 \neq 1/2$ or $v''(\epsilon)_1 \neq 1/2$, we may assume by symmetry that $v'(\epsilon)_1 \neq 1/2$. By assumption, there is an x_0 lying between $1/2$ and $v'(\epsilon)_1$ such that whenever $p \in \mathbb{S}^1$ is of the form $p = (x_0, y_0)$ (for some y_0) then any point of A_p is at least distance $1/100$ away from $(\delta(0)_3, \delta(0)_4)$. But, by continuity of v', there is some $t_0 \in [0, \epsilon]$ such that $v'(t_0)_1 = x_0$. Observe that

$$K \supseteq S_{\gamma(t_0)} + \delta(t_0) \supseteq \{v'(t_0), -v'(t_0)\} + \delta(t_0).$$

But if $u, u' \in K$ with $u - u' = 2(x, y, 0, 0)$ for some x, y with $x^2 + y^2 = 1$, then we must have $u, u' \in K \cap C$ and $u = (x, y, z, w)$, $u' = (-x, -y, z, w)$ for some $(z, w) \in A_{x,y}$. Hence $\delta(t_0) = (0, 0, z, w)$ for some $(z, w) \in A_{v'(t_0)_1, v'(t_0)_2}$. But then $|\delta(t_0) - \delta(0)| > 1/100$, giving a contradiction. □

Proof of Theorem 2.1.3. The result follows easily from Lemmas 2.4.1, 2.4.2 and 2.4.3. □

2.5 Concluding remarks

In this chapter we answered Question 2.1.1 and some related problems. However, there are still some open questions in this topic. For example, our counterexample in Theorem 2.1.3 requires $d \geq 4$, whereas we know that there can be no 2-dimensional counterexample (by Theorem 2.1.4). It would be interesting to see a counterexample in 3 dimensions (we believe that such a construction should exist).

Question 2.5.1. *Can we find convex bodies S and K in $\mathbb{R}^3$ such that S is K-Kakeya, but there are two S-copies in K which cannot be rotated into each other within K?*

Furthermore, we showed that if S is a unit segment, then any two S copies can be rotated into each other within a compact convex (S-)Kakeya set, but this fails for general bodies S. It would be interesting to determine if there are other sets S (or families of such) for which this property holds. (A trivial example is given by closed balls.)

Question 2.5.2. *Can we find (compact, convex) sets S in $\mathbb{R}^d$ with $d \geq 3$ such that S is not a segment or a ball, and whenever some convex body K is S-Kakeya then any two S copies can be rotated into each other within K?*

Chapter 3

Large hypergraphs without tight cycles

3.1 Introduction

A well-known basic fact about graphs states that a graph on n vertices containing no cycle of any length has at most $n-1$ edges, with this upper bound being tight. To find generalisations of this result (and other results concerning cycles) for r-uniform hypergraphs with $r \geq 3$, we need a corresponding notion of cycles in hypergraphs. There are several types of hypergraph cycles for which Turán-type problems have been widely studied, including Berge cycles and loose cycles [26, 68, 79, 80, 99, 109]. In this chapter we will consider tight cycles, for which it appears to be rather difficult to obtain extremal results.

Given positive integers $r \geq 2$ and $\ell > r$, an r-uniform tight cycle of length ℓ is a hypergraph with vertices $v_1, \dots, v_\ell$ and edges $\{v_i, v_{i+1}, \dots, v_{i+r-1}\}$ for $i = 1, \dots, \ell$, with the indices taken modulo ℓ. Observe that for $r = 2$ a tight cycle of length ℓ is just a cycle of length ℓ in the usual sense. Let $f_r(n)$ denote the maximal number of edges that an r-uniform hypergraph on n vertices can have if it has no subgraph isomorphic to a tight cycle of any length. So $f_2(n) = n-1$. It is easy to see that the hypergraph obtained by taking all edges containing a certain point is tight-cycle-free, giving a lower bound $f_r(n) \geq \binom{n-1}{r-1}$. Sós (see [117, 137]) and independently Verstraëte [137] raised the problem of estimating $f_r(n)$, and asked whether the lower bound $\binom{n-1}{r-1}$ is tight. This question was answered in the negative by Huang and Ma [86], who showed that for $r \geq 3$ there exists $c_r > 0$ such that if n is sufficiently large then $f_r(n) \geq (1+c_r)\binom{n-1}{r-1}$. Recently, Sudakov and Tomon [135] showed that $f_r(n) \leq n^{r-1+o(1)}$ for each fixed r, and commented[1] that it is widely believed that the correct order of magnitude is $\Theta(n^{r-1})$. The main result of this chapter is the following theorem, which shows that this is not the case.

[1] in the preprint version of [135]

Theorem 3.1.1. *For each fixed $r \geq 3$ we have $f_r(n) = \Omega(n^{r-1} \log n / \log\log n)$. In particular, $f_r(n)/n^{r-1} \to \infty$ as $n \to \infty$.*

The upper bound of Sudakov and Tomon [135] is $n^{r-1}e^{c_r\sqrt{\log n}}$. Very recently, Letzter [111] managed to improve their upper bound to $n^{r-1}\log^5 n$, so our lower bound in Theorem 3.1.1 is tight up to a factor of $(\log n)^{O(1)}$.

Concerning tight cycles of a given length, we mention the following interesting problem of Conlon (see [117]), which remains open.

Question 3.1.2 (Conlon)**.** *Given $r \geq 3$, does there exist some $c = c(r)$ constant such that whenever $\ell > r$ and ℓ is divisible by r then any r-uniform hypergraph on n vertices which does not contain a tight cycle of length ℓ has at most $O(n^{r-1+c/\ell})$ edges?*

Note that we need the assumption that ℓ is divisible by r, otherwise a (balanced, n-vertex) complete r-uniform r-partite hypergraph has no tight cycle of length ℓ and has $\Theta(n^r)$ edges.

3.2 Proof of our result

The key observation for our construction is the following lemma.

Lemma 3.2.1. *Assume that n, k, t are positive integers with $kt \leq n$, and $G_1, \dots, G_t$ are edge-disjoint subgraphs of $K_{n,n}$ such that no G_i contains a cycle of length at most $2k$. Then there is a tight-cycle-free 3-partite 3-uniform hypergraph on at most $3n$ vertices having $k\sum_{i=1}^{t}|E(G_i)|$ hyperedges.*

Proof. Let the two vertex classes of $K_{n,n}$ be X and Y, and let $Z = [t] \times [k]$. (As usual, $[m]$ denotes $\{1, \dots, m\}$.) Our 3-uniform hypergraph has vertex classes X, Y, Z and hyperedges

$$\{\{x,y,z\} : x \in X, y \in Y, z \in Z, z = (i,s) \text{ for some } i \in [t] \text{ and } s \in [k], \text{ and } \{x,y\} \in E(G_i)\}.$$

In other words, for each G_i we add k new vertices (denoted (i,s) for $s = 1, \dots, k$), and we replace each edge of G_i by the k hyperedges obtained by adding one of the new vertices corresponding to G_i to the edge.

We need to show that our hypergraph contains no tight cycles. Since our hypergraph is 3-partite, it is easy to see that any tight cycle is of the form $x_1y_1z_1x_2y_2z_2\dots x_\ell y_\ell z_\ell$ (for some $\ell \geq 2$ positive integer) with $x_j \in X, y_j \in Y, z_j \in Z$ for all j. Assume that $z_1 = (i, s_1)$. Then $\{x_1,y_1\},\{y_1,x_2\},\{x_2,y_2\} \in E(G_i)$. But $\{x_2,y_2\} \in E(G_i)$ implies that z_2 must be of the form (i, s_2) for some s_2. Repeating this argument, we deduce that there are $s_j \in [k]$ such that $z_j = (i, s_j)$ for all j, and $x_jy_j, y_jx_{j+1} \in E(G_i)$ for all j (with the indices taken mod ℓ). Hence $x_1y_1x_2y_2\dots x_\ell y_\ell$ is a cycle in G_i, giving $\ell > k$. But the vertices $z_j = (i, s_j)$ $(j = 1, \dots, \ell)$ must all be distinct, and there are k possible values for the second coordinate, giving $\ell \leq k$. We get a contradiction, giving the result. $\square$

We mention that Lemma 3.2.1 can be generalised to give $(r+r')$-uniform tight-cycle-free hypergraphs if we have edge-disjoint r-uniform hypergraphs $G_1, \ldots, G_t$ not containing tight cycles of length at most rk and edge-disjoint r'-uniform hypergraphs $H_1, \ldots, H_t$ not containing tight cycles of length more than $r'k$. Indeed, we can take all edges $e \cup f$ with $e \in E(G_i), f \in E(H_i)$ for some i. (Then Lemma 3.2.1 may be viewed as the special case $r = 2, r' = 1$.)

Lemma 3.2.2. *There exists $\alpha > 0$ such that whenever $k \leq \alpha \log n / \log \log n$ then we can find edge-disjoint subgraphs $G_1, \ldots, G_t$ of $K_{n,n}$ with $t = \lfloor n/k \rfloor$ such that no G_i contains a cycle of length at most $2k$, and $\sum_{i=1}^{t} |E(G_i)| = (1 - o(1))n^2$.*

Proof. It is well-known (and can be proved by a standard probabilistic argument) that there are constants $\beta, c > 0$ such that if n is sufficiently large and $k \leq \beta \log n$ then there exists a subgraph H of $K_{n,n}$ which has no cycle of length at most $2k$ and has $|E(H)| \geq n^{1+c/k}$. We randomly and independently pick copies $H_1, \ldots, H_t$ of H in $K_{n,n}$. Let $G_1 = H_1$ and $E(G_i) = E(H_i) \setminus \bigcup_{j=1}^{i-1} E(H_j)$ for $i \geq 2$. Then certainly the G_i are edge-disjoint and no G_i contains a cycle of length at most $2k$. Furthermore, the probability that a given edge is not contained in any H_i is

$$\begin{aligned}(1 - |E(H)|/n^2)^t &\leq \exp\left(-|E(H)|t/n^2\right) \\ &\leq \exp\left(-n^{1+c/k}\lfloor n/k \rfloor/n^2\right) = \exp\left(-n^{c/k}/k(1+o(1))\right).\end{aligned}$$

This is $o(1)$ as long as $k \leq \alpha \log n / \log \log n$ for some constant $\alpha > 0$. Therefore the expected value of $\left|\bigcup_{i=1}^{t} E(H_i)\right|$ is $(1-o(1))n^2$. Since $\sum_{i=1}^{t} |E(G_i)| = \left|\bigcup_{i=1}^{t} E(H_i)\right|$, the result follows. $\square$

Proof of Theorem 3.1.1. First consider the case $r = 3$. Lemma 3.2.2 and Lemma 3.2.1 together show that if $k \leq \alpha \log n / \log \log n$ then there is a tight-cycle-free 3-partite 3-uniform hypergraph on $3n$ vertices with $(1-o(1))kn^2$ edges. This shows $f_3(n) = \Omega(n^2 \log n / \log \log n)$, as claimed.

For $r \geq 4$, observe that $f_r(2n) \geq f_{r-1}(n)n$. Indeed, if H is an $(r-1)$-uniform tight-cycle-free hypergraph on n vertices, then we can construct a tight-cycle-free r-uniform hypergraph H' on $2n$ vertices with $n|E(H)|$ edges as follows. The vertex set of H' is the disjoint union of $[n]$ and the vertex set $V(H)$ of H, and the edges are $e \cup \{i\}$ with $e \in E(H)$ and $i \in [n]$. Then any tight cycle in H' must be of the form $v_1 v_2 \ldots v_{\ell r}$ with $v_i \in V(H)$ if i is not a multiple of r and $v_i \in [n]$ if i is a multiple of r. But then we get a tight cycle $v_1 v_2 \ldots v_{r-1} v_{r+1} v_{r+2} \ldots v_{2r-1} v_{2r+1} \ldots v_{\ell r-1}$ in H by removing each vertex from $[n]$ from this cycle. This is a contradiction, so H' contains no tight cycles. The result follows. $\square$

Chapter 4

Saturation for k-wise intersecting families

4.1 Introduction

Given positive integers $k \geq 2$ and n, we say that a family $\mathcal{F}$ of subsets of $[n] = \{1, 2, \ldots, n\}$ is *k-wise intersecting* if whenever $X_1, \ldots, X_k \in \mathcal{F}$ then $X_1 \cap \cdots \cap X_k \neq \emptyset$. It is well known (and easy to see) that any (2-wise) intersecting family over $[n]$ has size at most 2^{n-1}, so the largest possible size of a k-wise intersecting family is clearly 2^{n-1} for all k. This is achieved, for example, by taking $\mathcal{F} = \{A \in \mathcal{P}([n]) : 1 \in A\}$, where $\mathcal{P}(X)$ denotes the set of subsets of X.

However, the corresponding saturation problem of finding the smallest possible size of a maximal k-wise intersecting family is more interesting for $k \geq 3$. (A family $\mathcal{F}$ of subsets of $[n]$ is *maximal k-wise intersecting* if it is k-wise intersecting but no family $\mathcal{F}'$ over $[n]$ strictly containing $\mathcal{F}$ is k-wise intersecting. The $k = 2$ case is uninteresting, as any maximal 2-wise intersecting family has size 2^{n-1}.) This problem was briefly mentioned by Erdős and Kleitman [52] in 1974, and recently Hendrey, Lund, Tompkins and Tran studied this problem for $k = 3$. They determined the smallest possible size of a maximal 3-wise intersecting family exactly when n is sufficiently large and even. The case n odd was later resolved by Balogh, Chen and Luo, and the results of these two groups were published as a joint paper [14] giving a unified treatment for all n when $k = 3$. For general k, the authors of [14] showed that the smallest possible size $f_k(n)$ of a maximal k-wise intersecting family satisfies

$$c_k \cdot 2^{n/(k-1)} \leq f_k(n) \leq d_k \cdot 2^{n/\lceil k/2 \rceil}$$

(for some constants $c_k, d_k > 0$). They asked about closing the exponential gap between the lower and upper bounds.

In this chapter we prove the following result, which shows that the lower bound gives the right order of magnitude.

Theorem 4.1.1. *For each $k \geq 3$ there exists some constant C_k such that for all n there is a maximal k-wise intersecting family over $[n]$ of size at most $C_k \cdot 2^{n/(k-1)}$.*

In the case when $n \geq 2(k-1)$ is a multiple of $k-1$, the exact value of our upper bound will be $2^{n/(k-1)+k-3}(k-1)-(2^{k-1}-1)(k-2)$. In the special case $k=3$ this upper bound is $2^{n/2+1}-3$, which is tight for n sufficiently large (as shown by Balogh, Chen, Hendrey, Lund, Luo, Tompkins and Tran [14]), but for $k \geq 4$ the construction has more complicated structure than for $k=3$. In fact, in [14] it was shown that for $k=3$ (and n large), the unique maximal 3-wise intersecting families of smallest possible size are given by $\{A^c : A \in (\mathcal{P}(X) \setminus \{X\}) \cup (\mathcal{P}(Y) \setminus \{Y\})\}$ for some partition $X \cup Y$ of $[n]$ into two parts which are as close in size to each other as possible. This was proved by first obtaining a stability result stating that for any 'small' maximal 3-wise intersecting family $\mathcal{F}$, the family $\bar{\mathcal{F}} = \{A^c : A \in \mathcal{F}\}$ must be 'close' to the union of two cubes $\mathcal{P}(X) \cup \mathcal{P}(Y)$ (with X, Y as above).

However, the following result shows that, for $k \geq 4$, directly generalising this approach cannot work, and it is necessary to have more complicated structure.

Lemma 4.1.2. *Let $k \geq 4$, let $X_1 \cup \cdots \cup X_{k-1}$ be a partition of $[n]$ with each X_i having size $n/(k-1)+O(1)$, and let $Q = \mathcal{P}(X_1) \cup \cdots \cup \mathcal{P}(X_{k-1})$. If $|\mathcal{F} \setminus Q| = o(2^{n/(k-1)})$ and n is large enough, then $\bar{\mathcal{F}} = \{A^c : A \in \mathcal{F}\}$ cannot be maximal k-wise intersecting.*

We mention that many other saturation problems have already been studied in the context of set systems and intersection properties. For example, several authors gave bounds for the smallest possible size $m(r)$ of a set system which is maximal among (2-wise) intersecting families $\mathcal{F} \subseteq \binom{\mathbb{N}}{r}$ consisting of sets of size r – see, for example, [21, 31, 47]. A linear lower bound follows from a result of Erdős and Lovász [55], and Dow, Drake, Füredi, and Larson [47] showed that in fact $m(r) \geq 3r$ for $r \geq 4$. Blokhuis [21] proved a polynomial upper bound of $m(r) \leq r^5$, and for certain values of r quadratic upper bounds are also known – see, e.g., [21, 31]. Finding the order of magnitude of $m(r)$ is still an open problem. See the introduction and the references in [14] for other related saturation problems.

4.2 The construction

We now prove Theorem 4.1.1 by describing a family of size $O(2^{n/(k-1)})$ and showing that it is maximal k-wise intersecting over $[n]$. For simplicity, we will work with complements, using the observation that $\bar{\mathcal{G}} = \{X^c : X \in \mathcal{G}\}$ is k-wise intersecting if and only if no k elements of $\mathcal{G}$ have union $[n]$.

Fix some $k \geq 3$ and $n \geq 2(k-1)$ integers. Partition $[n]$ into $k-1$ sets $A_1, \ldots, A_{k-1}$ which are as close in size to each other as possible, and pick 'special' elements $a_i \in A_i$ for each i. Consider the following families. (All indices will be understood mod $k-1$, so, for example, $a_0 = a_{k-1}$.)

$$\begin{aligned}
\mathcal{F}_1(i) &= \mathcal{P}(A_i) \setminus \{A_i\} \\
\mathcal{F}_2(i) &= \{X \cup Y : X \subseteq A_i, Y \subseteq \{a_1, a_2, \ldots, a_{k-1}\} \setminus \{a_{i-1}, a_i\}, X \neq A_i \setminus \{a_i\}, X \neq A_i\} \\
\mathcal{F} &= \bigcup_{i=1}^{k-1} (\mathcal{F}_1(i) \cup \mathcal{F}_2(i)).
\end{aligned}$$

We will show that $\{A^c : A \in \mathcal{F}\}$ is maximal k-wise intersecting. Note that if $k = 3$ then $\mathcal{F}_2(i) \subseteq \mathcal{F}_1(i)$ for each $i = 1, 2$, so $\mathcal{F}$ is simply $\mathcal{F}_1(1) \cup \mathcal{F}_1(2) = (\mathcal{P}(A_1) \cup \mathcal{P}(A_2)) \setminus \{A_1, A_2\}$. This was shown to be (up to isomorphism) the unique minimal-sized construction when $k = 3$ and n is sufficiently large by Balogh, Chen, Hendrey, Lund, Luo, Tompkins and Tran [14]. Furthermore, note that

$$|\mathcal{F}| = 2^{n/(k-1)+k-3}(k-1) - (2^{k-1} - 1)(k-2)$$

if $k-1$ divides n. Indeed, the number of subsets of $\{a_1, \ldots, a_{k-1}\}$ appearing in $\mathcal{F}$ is $2^{k-1} - 1$, the number of sets which are not subsets of $\{a_1, \ldots, a_{k-1}\}$ but appear in some $\mathcal{F}_2(i)$ is $(k-1) \cdot (2^{n/(k-1)} - 4) \cdot 2^{k-3}$, and finally, the only elements of $\mathcal{F}$ we have not yet counted are the $(k-1)$ sets $A_i \setminus \{a_i\}$ for $i = 1, \ldots, k-1$. Summing these contributions gives the formula above.

Claim 4.2.1. The family $\mathcal{F}$ contains no k elements having union $[n]$.

Proof. Suppose we have sets $Y_1, \ldots, Y_k \in \mathcal{F}$ satisfying $Y_1 \cup \cdots \cup Y_k = [n]$. Then clearly at least one of them must come from some $\mathcal{F}_2(i)$, we may assume $Y_1 \in \mathcal{F}_2(1)$. If there is some $j \neq 1$ and $i \neq 1$ such that $Y_j \in \mathcal{F}_2(i)$, then for each t we can pick $b_t \in A_t \setminus \{a_t\}$ such that $b_t \notin Y_1 \cup Y_j$. Then no element of $\mathcal{F}$ contains more than one b_t, but $\{b_1, \ldots, b_{k-1}\} \subseteq \bigcup_{\ell \neq 1,j} Y_\ell$, giving a contradiction. On the other hand, if there is no such pair (i, j), then $\{Y_\ell : \ell \neq 1\}$ must contain at least one non-empty set from $\mathcal{F}_1(t)$ (to cover $A_t \setminus \{a_t\}$) for $t = 2, \ldots, k-2$, at least two different sets from $\mathcal{F}_1(k-1)$ (to cover A_{t-1}), and at least one set having an element in A_1, again giving a contradiction since these k sets must all be different. □

Claim 4.2.2. For any $X \in \mathcal{P}([n]) \setminus \mathcal{F}$ there exist $X_1, \ldots, X_{k-1} \in \mathcal{F}$ such that $X_1 \cup X_2 \cup \cdots \cup X_{k-1} \cup X = [n]$.

Proof. We first consider the following five cases, then check that each choice of X belongs to at least one of these cases.

- Case 1: $X \cap A_i \neq \emptyset$ for all i. Then let $X_i = A_i \setminus X$, so $X_i \in \mathcal{F}_1(i)$ and the X_i satisfy the conditions.
- Case 2: There is some i such that $X \cap A_i$ contains an element b_i with $b_i \neq a_i$ and $X \cap A_{i-1} \neq \emptyset$. We may assume that $i = 1$. Then let $X_1 = (A_1 \setminus \{b_1\}) \cup \{a_2, a_3, \ldots, a_{k-2}\}$ (so $X_1 \in \mathcal{F}_2(1)$),

let $X_j = A_j \setminus \{a_j\}$ for $j = 2, \ldots, k-2$ (so $X_j \in \mathcal{F}_1(j)$) and let $X_{k-1} = A_{k-1} \setminus X$ (so $X_{k-1} \in \mathcal{F}_1(k-1)$). These clearly satisfy the conditions.

- Case 3: There exist i, j, b_i, b_j such that $i \neq j$, $b_i \in X \cap (A_i \setminus \{a_i\})$, $b_j \in X \cap (A_j \setminus \{a_j\})$. Then let $X_i = (A_i \setminus X) \cup \{a_\ell : \ell \neq i, i-1\}$, $X_j = (A_j \setminus X) \cup \{a_\ell : \ell \neq j, j-1\}$, and $X_\ell = A_\ell \setminus \{a_\ell\}$ for $\ell \neq i, j$. Then $X_i \in \mathcal{F}_2(i), X_j \in \mathcal{F}_2(j)$, and $X_\ell \in \mathcal{F}_1(\ell)$ for $\ell \neq i, j$, so it is easy to see that the conditions are satisfied.

- Case 4: $X \supseteq (A_i \setminus \{a_i\}) \cup \{a_j\}$ for some i, j with $j \neq i$. Then let $X_i = \{a_t : t \neq j\}$ (so $X_i \in \mathcal{F}_2(j+1)$), and let $X_\ell = A_\ell \setminus \{a_\ell\}$ for $\ell \neq i$ (so $X_\ell \in \mathcal{F}_1(\ell)$). It is easy to see that the conditions are satisfied.

- Case 5: $X = A_i$ for some i. Then let $X_i = \{a_t : t \neq i\}$ (so $X_i \in \mathcal{F}_2(i+1)$) and $X_\ell = A_\ell \setminus \{a_\ell\}$ for $\ell \neq i$ (so $X_\ell \in \mathcal{F}_1(\ell)$). Then the conditions are again satisfied.

We now check that any $X \in \mathcal{P}([n]) \setminus \mathcal{F}$ belongs to at least one of these cases. If $X \subseteq \{a_1, \ldots, a_{k-1}\}$ then in fact $X = \{a_1, \ldots, a_{k-1}\}$ and we are in Case 1. Otherwise there is some i and some $b_i \in A_i$ such that $b_i \in X$ and $b_i \neq a_i$. If $X \cap A_{i-1} \neq \emptyset$ then we are in Case 2. Otherwise, if X is not a subset of $A_i \cup \{a_\ell : \ell \neq i, i-1\}$ then we are in Case 3. Finally, if X is a subset of $A_i \cup \{a_\ell : \ell \neq i, i-1\}$, then we are in Case 4 or Case 5 since $X \notin \mathcal{F}_1(i) \cup \mathcal{F}_2(i)$. □

Proof of Theorem 4.1.1. By Claims 4.2.1 and 4.2.2, the family $\bar{\mathcal{F}} = \{X^c : X \in \mathcal{F}\}$ is maximal k-wise intersecting for all $k \geq 3$ and $n \geq 2(k-1)$. However, $\mathcal{F}_1(i)$ and $\mathcal{F}_2(i)$ both have size $O_k(2^{n/(k-1)})$ for each i, so $|\mathcal{F}| = O_k(2^{n/(k-1)})$. The result follows. □

4.3 Non-existence of certain types of maximal families

Finally, we prove Lemma 4.1.2 stating that there can be no construction for $k \geq 4$ which is close to the union of $k-1$ cubes.

Proof of Lemma 4.1.2. Suppose that $\mathcal{F}$ is as in the statement of the lemma, and $\bar{\mathcal{F}}$ is maximal k-wise intersecting. Then at least one X_i does not appear in $\mathcal{F}$, we may assume $X_{k-1} \notin \mathcal{F}$. Let $\mathcal{G}$ be the family of sets S over $[n]$ satisfying the following conditions.

- The set S cannot be written as $S_1 \cup \cdots \cup S_{k-1}$ such that $S_i \in \mathcal{F}$ for all i and $S_1 \in \mathcal{F} \setminus Q$.

- For all $S' \in \mathcal{F} \setminus Q$ and $i \leq k-2$, we have $S' \cap X_i \neq S \cap X_i$.

- For all i, $S \cap X_i$ is non-empty.

- We have $S^c \cap (X_1 \cup \cdots \cup X_{k-2}) \notin \mathcal{F}$.

It is easy to see that each of these conditions fails for only $o(2^n)$ subsets of $[n]$, so $\mathcal{G}$ is non-empty. Pick any $S \in \mathcal{G}$, and let $T = S^c \cap (X_1 \cup \cdots \cup X_{k-2})$. Then $T \notin \mathcal{F}$, so (by maximality) there are $T_1, \ldots, T_{k-1} \in \mathcal{F}$ such that $T \cup T_1 \cup T_2 \cup \cdots \cup T_{k-1} = [n]$. Furthermore, by maximality, $\mathcal{F}$ must be a down-set, so we may assume that in fact $T^c = T_1 \cup \cdots \cup T_{k-1}$. So $S \cup X_{k-1} = T_1 \cup \cdots \cup T_{k-1}$. Write $S_i = T_i \setminus (X_{k-1} \setminus S)$, then $S = S_1 \cup \cdots \cup S_{k-1}$ and $S_i \in \mathcal{F}$ for all i (as $\mathcal{F}$ is a down-set). So we must have $S_1, \ldots, S_{k-1} \in Q \cap \mathcal{F}$. Since $S \cap X_i$ is non-empty for all i, we may assume that $S_i = S \cap X_i$ for all i. Then, for $i \leq k-2$, $T_i \cap X_i = S \cap X_i$, so $T_i \notin \mathcal{F} \setminus Q$. Hence $T_i \in Q$ and $T_i = S_i$. But then $T_{k-1} = X_{k-1}$, giving a contradiction. $\square$

Chapter 5

Generalizations of the Ruzsa–Szemerédi and rainbow Turán problems

5.1 Introduction

The famous Ruzsa–Szemerédi or $(6,3)$-problem is to determine how many edges there can be in a 3-uniform hypergraph on n vertices if no six vertices span three or more edges. This rather specific-sounding problem turns out to have several equivalent formulations and bounds in both directions have had many applications. It is not difficult to prove an upper bound of $O(n^2)$: one first observes that if two edges have two vertices in common, then neither of them can intersect any other edges, and after removing all such pairs of edges one is left with a linear hypergraph, for which the bound is trivial. Brown, Erdős and Sós [133] gave a construction achieving $\Omega(n^{3/2})$ edges and asked whether the maximum is $o(n^2)$.

The argument sketched in the previous paragraph shows that this question is equivalent to asking whether a graph on n vertices such that no edge is contained in more than one triangle must contain $o(n^2)$ triangles. A positive answer to this question was given by Ruzsa and Szemerédi [126], who obtained a bound of $O(n^2/\log^* n)$ with the help of Szemerédi's regularity lemma. This breakthrough result has been highly influential, as it is essentially the first appearance of the triangle removal lemma. Ruzsa and Szemerédi also gave a construction showing that the number of triangles can be as large as $n^2 e^{-O(\sqrt{\log n})} = n^{2-o(1)}$, so the exponent in their upper bound cannot be improved.

One of the applications they gave of their upper bound was an alternative proof of Roth's theorem. Indeed, let A be a subset of $\{1,\dots,N\}$ that contains no arithmetic progression of length 3. Define a tripartite graph G with vertex classes $X=\{1,2,\dots,N\}$, $Y=\{1,2,\dots,2N\}$ and $Z=\{1,2,\dots,3N\}$, where if $x\in X$, $y\in Y$ and $z\in Z$, then xy is an edge if and only if $y-x\in A$,

yz is an edge if and only if $z - y \in A$ and xz is an edge if and only if $(z-x)/2 \in A$. Note that these are the edges of the triangles with vertices belonging to triples of the form $(x, x+a, x+2a)$ with $x \in X$ and $a \in A$. If xyz is a triangle in this graph, then $a = y - x, b = z - y, c = (z-x)/2$ satisfy $a, b, c \in A$ and $a + b = 2c$, which gives us an arithmetic progression of length 3 in A unless $y - x = z - y$. Thus, the only triangles are the 'degenerate' ones of the form $(x, x+a, x+2a)$, which implies that each edge is contained in at most one triangle. Therefore, the number of triangles is $o(n^2)$ (where $n = 6N$). We also have that for each $a \in A$ there are N triangles of the form $(x, x+a, x+2a)$, so $|A| = o(N)$.

As Ruzsa and Szemerédi also observed, this argument can be turned round: it tells us that if A has density α, then there is a graph with $6N$ vertices and αN^2 triangles such that each edge is contained in at most one triangle. Since Behrend [16] proved that there exists a subset A of $\{1, \ldots, N\}$ of size $Ne^{-O(\sqrt{\log N})}$ that does not contain an arithmetic progression of length 3, this gives the lower bound mentioned above.

Several related questions have been studied, as well as applications and generalizations of the Ruzsa–Szemerédi problem: see, for example, [5, 8]. A natural generalization that we believe has not been considered is the following generalized Turán problem.

Question 5.1.1. *Let r and s be positive integers with $1 \leq r < s$. Let G be a graph on n vertices such that every subgraph of G isomorphic to K_r is contained in at most one subgraph of G isomorphic to K_s. What is the largest number of copies of K_s that G can contain?*

The Ruzsa–Szemerédi problem is the case $r = 2, s = 3$ of Question 5.1.1, and the answer is trivially $\Theta(n)$ if $r = 1$. Because of the influence of the Ruzsa–Szemerédi theorem, proving an upper bound of $o(n^r)$ if $r \geq 2$ is now standard using the graph removal lemma. Indeed, if every K_r is in at most one K_s, then there are $O(n^r) = o(n^s)$ copies of K_s, therefore, by the graph removal lemma (see, e.g., [38]), we can delete $o(n^2)$ edges to remove all copies of K_s. But each edge is contained in at most $O(n^{r-2})$ copies of K_r and hence $O(n^{r-2})$ copies of K_s, so we removed $o(n^r)$ copies of K_s in total, as claimed.

In the case $r = 2$, the construction of Ruzsa and Szemerédi for the lower bound can be generalized (for example, by using h-sum-free sets from [4]) to get a lower bound of $n^2 e^{-O(\sqrt{\log n})}$. However, there is no obvious way of generalizing the algebraic construction for $r \geq 3$. We shall present a geometric construction instead, in order to prove the following result, which is the first of the two main results of this chapter. The idea behind the construction is similar to a famous construction of Bollobás and Erdős [29] (which was introduced in the context of a Ramsey–Turán problem).

Theorem 5.1.2. *For each $1 \leq r < s$ and positive integer n there is a graph on n vertices with $n^r e^{-O(\sqrt{\log n})} = n^{r-o(1)}$ copies of K_s such that every K_r is contained in at most one K_s.*

We shall also use a modification of our construction to answer a question about rainbow colourings. Given an edge-colouring of a graph G, we say that a subgraph H is *rainbow* if all of its edges have different colours. We denote by $\mathrm{ex}^*(n, H)$ the maximal number of edges that a graph on n vertices can contain if it can be properly edge-coloured (that is, no two edges of the same colour meet at a vertex) in such a way that it contains no rainbow copy of H. The rainbow Turán problem (i.e., the problem of estimating $\mathrm{ex}^*(n, H)$) was introduced by Keevash, Mubayi, Sudakov and Verstraëte [107], and was studied for several different families of graphs H, such as complete bipartite graphs [107], even cycles [45, 97, 107] and paths [57, 100]. Gerbner, Mészáros, Methuku and Palmer [72] considered the following generalized rainbow Turán problem (analogous to the generalized Turán problem introduced by Alon and Shikhelman [6]). Given two graphs H and F, let $\mathrm{ex}(n, H, \text{rainbow-}F)$ denote the maximal number of copies of H that a properly edge-coloured graph on n vertices can contain if it has no rainbow copy of F. Note that $\mathrm{ex}^*(n, H)$ is the special case $\mathrm{ex}(n, K_2, \text{rainbow-}H)$. The authors of [72] focused on the case $H = F$ and obtained several results, for example when H is a path, cycle or a tree, and also gave some general bounds. One of their concluding questions was the following.

Question 5.1.3 (Gerbner, Mészáros, Methuku and Palmer [72])**.** *What is the order of magnitude of* $\mathrm{ex}(n, K_r, \text{rainbow-}K_r)$ *for* $r \geq 4$*?*

For fixed r, a straightforward double-counting argument shows that if H has r vertices, then $\mathrm{ex}(n, H, \text{rainbow-}H) = O(n^{r-1})$. Indeed, if G is a graph with n vertices that contains no rainbow copy of H, then every copy of H contains two edges of the same colour. But the number of such pairs of edges is at most $\binom{n}{2}\frac{n-2}{2} = O(n^3)$ (since there are at most $\frac{n-2}{2}$ edges with the same colour as any given edge), and each such pair can be extended to at most $O(n^{r-4})$ copies of H.

The authors above improved this bound to $o(n^{r-1})$, and gave an example that shows that $\mathrm{ex}(n, K_r, \text{rainbow-}K_r) = \Omega(n^{r-2})$. They also asked whether there is a graph H for which the exponent $r-1$ in the upper bound is sharp. Our next result shows that $H = K_r$ is such a graph.

Theorem 5.1.4. *For each* $r \geq 4$ *we have* $\mathrm{ex}(n, K_r, \text{rainbow-}K_r) = n^{r-1-o(1)}$.

Note that a triangle is always rainbow in a proper edge-colouring, so we trivially have $\mathrm{ex}(n, K_r, \text{rainbow-}K_r) = 0$ for $r < 4$.

In fact, our method can be used to prove the following more general result.

Theorem 5.1.5. *Let* $r \geq 4$*, let* H *be a graph, and let* H *have a proper edge-colouring with no rainbow* K_r*. Suppose that for each vertex* v *of* H *there is a* $p_v \in \mathbb{R}^m$*, and for each colour* κ *in the colouring there is a non-zero vector* z_κ *such that for every edge* vw *of colour* κ*,* z_κ *is a linear combination of* p_v *and* p_w *with non-zero coefficients. Then* $\mathrm{ex}(n, H, \text{rainbow-}K_r) \geq n^{m_0-o(1)}$*, where* m_0 *is the dimension of the subspace of* $\mathbb{R}^m$ *spanned by the points* p_v*.*

It is easy to see that Theorem 5.1.4 is a special case of Theorem 5.1.5, but Theorem 5.1.5 also allows us to determine the behaviour of $\mathrm{ex}(n, H, \text{rainbow-}K_r)$ for several other natural choices of H. We give some examples in Section 5.5.

Theorem 5.1.5 is 'almost equivalent' to the following, slightly weakened, alternative version.

Theorem 5.1.5'. *Let $r \geq 4$, let H be a graph, and let c be a proper edge-colouring of H without a rainbow K_r. Suppose that for each vertex $v \in V(H)$ we have a vector $p_v \in \mathbb{R}^{m-1}$, and for each colour κ of c the lines through the pairs p_v, p_w with $c(vw) = \kappa$ are either all parallel, or all go through the same point and that point is different from p_v, p_w unless $p_v = p_w$. Assume that no $(m-2)$-dimensional affine subspace contains all the points p_v. Then $\mathrm{ex}(n, H, \text{rainbow-}K_r) \geq n^{m-o(1)}$.*

It is easy to see that Theorem 5.1.5' is equivalent to the weakened version of Theorem 5.1.5 where we make the additional assumption that each p_v is non-zero. Indeed, given a configuration of points p_v as in Theorem 5.1.5 (with $m = m_0$), we can project it from the origin to an appropriate affine $(m-1)$-dimensional subspace not going through the origin to get a configuration as in Theorem 5.1.5'. Conversely, a configuration of points p_v as in Theorem 5.1.5' gives a configuration as in Theorem 5.1.5 by taking the points $p_v \times \{1\} \in \mathbb{R}^m$.

5.2 The idea of the construction, and a preliminary lemma

We now briefly describe the construction used in our proof of Theorem 5.1.2. As mentioned before, it is similar to the construction of Bollobás and Erdős introduced in [29]. For simplicity, we focus on the case $r = 2, s = 3$, i.e., the Ruzsa–Szemerédi problem.

Consider the d-dimensional sphere $S^d = \{x \in \mathbb{R}^{d+1} : \|x\| = 1\}$. (We will choose d to be about $\sqrt{\log n}$.) Join two points of the sphere by an edge if the angle between the corresponding vectors is between $2\pi/3 - \delta$ and $2\pi/3 + \delta$, where δ is some appropriately chosen small number (roughly $e^{-\sqrt{\log n}}$). Then there are 'few' triangles containing any given edge, since if xy is an edge then any point z such that xyz is a triangle is restricted to lie in a small neighbourhood around the point $-(x+y)$. However, there are 'many' edges, since the edge-neighbourhood of a point is a set of points around a codimension-1 surface, which is much larger then the neighbourhood of a single point. Choosing the parameters appropriately, we can achieve that if we pick n random points then any two of them form an edge with probability $n^{-o(1)}$, and any three of them form a triangle with probability $n^{-1-o(1)}$. Then any edge is expected to be in $n^{-o(1)}$ triangles and there are $n^{2-o(1)}$ edges. After some modification, we get a graph with $n^{2-o(1)}$ triangles in which any edge extends to at most one triangle.

The general construction is quite similar. We want to define the edges in such a way that knowing the position of any r of the vertices of a K_s restricts the remaining $s - r$ vertices to small neighbourhoods around certain points, but knowing the position of i points with $i < r$ only restricts the remaining points to a neighbourhood of a codimension-i surface. For example, when $(r, s) = (3, 4)$, we can define our graph by joining two points if the angle between the corresponding vectors is close to the angle given by two vertices of a regular tetrahedron (centred at the origin).

In fact, our construction and the construction of Ruzsa and Szemerédi based on the Behrend set are more similar than they might at first appear, which also explains why they give similar bounds (namely $n^2 e^{-O(\sqrt{\log n})}$ for the case $r = 2, s = 3$). Behrend's construction [16] of a large set with no arithmetic progression of length 3 starts by observing that for any positive integers k, d there is some m such that the grid $\{1, \ldots, k\}^d$ intersects the sphere $\{x \in \mathbb{R}^d : \|x\|^2 = m\}$ in a set A consisting of at least $k^d/(dk^2)$ points. This set A has no arithmetic progression of length 3. (In Behrend's construction, this is transformed into a subset of $\mathbb{Z}$ using an appropriate map, but this is unnecessary for our purposes.) Repeating the construction from Section 5.1, we define a tripartite graph G on vertex set $X \cup Y \cup Z$ where $X = \{1, \ldots, k\}^d, Y = \{1, \ldots, 2k\}^d, Z = \{1, \ldots, 3k\}^d$, and edges given by the edges of the triangles $(x, x+a, x+2a) \in X \times Y \times Z$ for $x \in X, a \in A$. Explicitly, for $x \in X, y \in Y, z \in Z$, we join x and y if $\|x - y\| = m^{1/2}$ (and $y_i \geq x_i$ for all i), we join y and z if $\|z - y\| = m^{1/2}$ (and $z_i \geq y_i$), and we join x and z if $\|x - z\| = 2m^{1/2}$ (and $z_i \geq x_i$). This gives the same phenomenon as our construction: the neighbourhood of a point x is given by a codimension-1 condition, but the joint neighbourhood of two points is a single point, since y must be the midpoint of x and z.

We conclude this section with the following technical fact, whose proof we include for completeness. Given unit vectors v, w, we write $\angle(v, w)$ for the angle between v and w – that is, for $\cos^{-1}(\langle v, w \rangle)$.

Lemma 5.2.1. *There exist constants $0 < \alpha < B$ such that the following holds. Let d be a positive integer, let $0 < \rho \leq 2$ and let $v \in S^d$. Let $X_\rho = \{w \in S^d : \|v - w\| < \rho\}$. Let μ denote the usual probability measure on S^d. Then*

$$\alpha^d \rho^d \leq \mu(X_\rho) \leq B^d \rho^d.$$

Furthermore, for any $-1 < \xi < 1$ there exists $\beta > 0$ such that for every positive integer d, every point $v \in S^d$, and every $0 \leq \delta \leq 2$, the set $Y_{\xi,\delta} = \{w \in S^d : |\langle v, w \rangle - \xi| < \delta\}$ has

$$\mu(Y_{\xi,\delta}) \geq \beta^d \delta.$$

Proof. Using the usual spherical coordinate system, we see that for $0 \leq \varphi \leq \pi$ the set $Z_\varphi = \{w \in S^d : \angle(v, w) < \varphi\}$ satisfies

$$\mu(Z_\varphi) = \frac{\int_0^\varphi \sin^{d-1}\theta \, \mathrm{d}\theta}{\int_0^\pi \sin^{d-1}\theta \, \mathrm{d}\theta}. \tag{5.1}$$

But we have $\theta \geq \sin\theta \geq \frac{2}{\pi}\theta$ for $0 \leq \theta \leq \pi/2$. Thus, $\int_0^t \sin^{d-1}\theta \, \mathrm{d}\theta$ is between $\frac{c_1^{d-1}}{d} t^d$ and $\frac{1}{d} t^d$ for all $0 \leq t \leq \pi/2$ (for some constant $0 < c_1 < 1$). Using this bound for both the numerator and the denominator in (5.1), we deduce that $\alpha_0^d \varphi^d \leq \mu(Z_\varphi) \leq B_0^d \varphi^d$ for some absolute constants

$0 < \alpha_0 < B_0$. But if $\angle(v,w) = \varphi$ and $\|v-w\| = \rho \le 2$, then $\rho \le \varphi \le \rho\pi/2$, so $Z_\rho \subseteq X_\rho \subseteq Z_{\rho\pi/2}$. The first claim follows.

For the second claim, let $0 < \varphi < \pi$ such that $\xi = \cos\varphi$ and let $\epsilon_0 = \min\{\varphi/2, (\pi-\varphi)/2\}$. Write $W_{\varphi,\epsilon} = \{w \in S^d : \varphi - \epsilon < \angle(v,w) < \varphi + \epsilon\}$. For $0 < \epsilon < \epsilon_0$ we have

$$\mu(W_{\varphi,\epsilon}) = \frac{\int_{\varphi-\epsilon}^{\varphi+\epsilon} \sin^{d-1}\theta \, \mathrm{d}\theta}{\int_0^{\pi/2} \sin^{d-1}\theta \, \mathrm{d}\theta}.$$

But also $\sin\theta \ge \min\{\sin(\varphi-\epsilon_0), \sin(\varphi+\epsilon_0)\} = \sin(\epsilon_0)$ when $\varphi - \epsilon_0 \le \theta \le \varphi + \epsilon_0$. Writing $\beta_0 = \sin(\epsilon_0) > 0$, it follows that whenever $\epsilon < \epsilon_0$, then $\mu(W_{\varphi,\epsilon}) \ge \frac{2\epsilon\beta_0^{d-1}}{\pi/2}$. However, we have $|\cos(\theta) - \cos(\varphi)| \le |\theta - \varphi|$, so $Y_{\xi,\delta} \supseteq W_{\varphi,\delta}$. Choosing some sufficiently small β, the second claim follows. □

5.3 The generalized Ruzsa–Szemerédi problem

In this section we prove the first of our main results, Theorem 5.1.2. In the case $r = 2, s = 3$, the construction is based, as we saw in Section 5.2, on the observation that if we wish to find three vectors in $S^d = \{x \in \mathbb{R}^{d+1} : \|x\| = 1\}$ in such a way that the angle between any two of them is $120°$, and if we choose the vertices one by one, then there are d degrees of freedom for the first vertex and $d-1$ for the second, but the third is then uniquely determined. This gives us an example of a 'continuous graph' with 'many' edges, such that each edge is in exactly one triangle, and a suitable perturbation and discretization of this graph gives us a finite graph with $n^{2-o(1)}$ triangles such that each edge belongs to at most one triangle.

To generalize this to arbitrary (r,s) we need to find a configuration of s unit vectors (where by 'configuration' we mean an $s \times s$ symmetric matrix that specifies the angles, or equivalently inner products, between the unit vectors) with the property that if we choose the points of the configuration one by one, then for $i \le r$ the i^{th} point can be chosen with $d+1-i$ degrees of freedom, but from the $(r+1)^{\text{st}}$ point onwards all points are uniquely determined. It turns out that all we have to do is choose an arbitrary collection of s points $p_1, \dots, p_s$ in general position from the sphere S^{r-1} and take the angles $\angle(p_i, p_j)$. To see that this works, suppose we wish to choose $x_1, \dots, x_s \in S^d$ one by one in such a way that $\langle x_i, x_j\rangle = \langle p_i, p_j\rangle$ for every i,j. Suppose that we have chosen $x_1, \dots, x_r$ and let V be the r-dimensional subspace that they generate. Let u_{r+1} be the orthogonal projection of x_{r+1} to V. Then $\langle u_{r+1}, x_i\rangle = \langle x_{r+1}, x_i\rangle$ for each $i \le r$, and $u_{r+1} \in V$, so u_{r+1} is uniquely determined. Furthermore, since the angles $\langle p_i, p_j\rangle$ are equal to the angles $\langle x_i, x_j\rangle$ when $i,j \le r$ and to the angles $\langle x_i, u_{r+1}\rangle$ when $i \le r, j = r+1$, and p_{r+1} is a unit vector, it must be that u_{r+1} is a unit vector, which implies that $x_{r+1} = u_{r+1}$. Since this argument made no use of the ordering of the vectors, it follows that any r vectors in a configuration determine the rest, as claimed.

We shall now use this observation as a guide for constructing a finite graph with many copies of K_s such that each K_r is contained in at most one K_s.

As above, pick s 'reference' points $p_1, \ldots, p_s$ in general position on the sphere S^{r-1}. Since for any set $B \subseteq \{1, \ldots, s\}$ of size r the points p_b ($b \in B$) form a basis of $\mathbb{R}^r$, we may write, for any a,

$$p_a = \sum_{b \in B} \lambda_{B,a,b} p_b$$

for some real constants $\lambda_{B,a,b}$.

For any $c > 0$ and positive integers N, d we define an s-partite random graph $G_{N,d,c}$ as follows. (The graph will also depend on $r, s, p_1, \ldots, p_s$, but for readability we drop these dependencies from the notation.) Consider the usual probability measure on the d-sphere S^d. Pick, independently and uniformly at random, sN points $x_{a,i}$ ($1 \leq a \leq s, 1 \leq i \leq N$) on S^d: these points form the vertex set. Join two points $x_{a,i}$ and $x_{b,j}$ by an edge if $a \neq b$ and $|\langle x_{a,i}, x_{b,j} \rangle - \langle p_a, p_b \rangle| < c$. Write $V_a = \{x_{a,i} : 1 \leq i \leq N\}$ so that $G_{N,d,c}$ is s-partite with classes $V_1, \ldots, V_s$.

We also define a graph $G'_{N,d,c}$ as follows. Let M_0 be the maximum among all values of $|\lambda_{B,a,b}|$ and $\lambda^2_{B,a,b}$, and let $M = 2(r+1)\sqrt{M_0}$. Then $G'_{N,d,c}$ is obtained from $G_{N,d,c}$ by deleting all vertices $x_{a,i}$ for which there is another vertex $x_{a,j}$ ($i \neq j$) such that $\|x_{a,i} - x_{a,j}\| < M\sqrt{c}$.

This graph is designed to be finite and to have the property that any copy of K_s must be close to a configuration with angles determined by the points $p_1, \ldots, p_s$. The vertex deletions are there to ensure that the vertices are reasonably well separated. This will imply that no K_r is contained in more than one K_s, since once r vertices of a K_s are chosen, the remaining vertices are constrained to lie in small neighbourhoods.

Lemma 5.3.1. *The graph $G'_{N,d,c}$ has the property that any of its subgraphs isomorphic to K_r is contained in at most one subgraph isomorphic to K_s (for any choices of $r, s, p_1, \ldots, p_s, N, d, c$).*

Proof. Let $x_{a_1,i_1}, \ldots, x_{a_r,i_r}$ be points that form a K_r. Then necessarily all a_t are distinct. Suppose that we have two extensions H_1, H_2 of this K_r to a K_s. Then both H_1 and H_2 intersect each class V_a in exactly one point. We now show that for each a this point must be the same for H_1 and H_2, which will imply the lemma.

Suppose that H_1 intersects V_a in point x. Write $B = \{a_1, \ldots, a_r\}$ and write x_{a_t} for x_{a_t,i_t}

$(t = 1, \ldots, r)$. Then

$$\begin{aligned}
\left\| x - \sum_{t=1}^{r} \lambda_{B,a,a_t} x_{a_t} \right\|^2 &= \left\langle x - \sum_{t=1}^{r} \lambda_{B,a,a_t} x_{a_t}, x - \sum_{t=1}^{r} \lambda_{B,a,a_t} x_{a_t} \right\rangle \\
&= \langle x, x \rangle - 2 \sum_{t=1}^{r} \lambda_{B,a,a_t} \langle x, x_{a_t} \rangle + \sum_{t,t'=1}^{r} \lambda_{B,a,a_t} \lambda_{B,a,a_{t'}} \langle x_{a_t}, x_{a_{t'}} \rangle \\
&\leq 1 - 2 \sum_{t=1}^{r} \lambda_{B,a,a_t} \langle p_a, p_{a_t} \rangle + 2c \sum_{t=1}^{r} |\lambda_{B,a,a_t}| \\
&\qquad + \sum_{t,t'=1}^{r} \lambda_{B,a,a_t} \lambda_{B,a,a_{t'}} \langle p_{a_t}, p_{a_{t'}} \rangle + c \sum_{t,t'=1}^{r} |\lambda_{B,a,a_t} \lambda_{B,a,a_{t'}}| \\
&\leq \left\langle p_a - \sum_{t=1}^{r} \lambda_{B,a,a_t} p_{a_t}, p_a - \sum_{t=1}^{r} \lambda_{B,a,a_t} p_{a_t} \right\rangle + 2rcM_0 + r^2 cM_0 \\
&= (r^2 + 2r) cM_0.
\end{aligned}$$

It follows that $\|x - \sum_{t=1}^{r} \lambda_{B,a,a_t} x_{a_t}\| < (r+1)\sqrt{M_0 c}$. Similarly, if H_2 intersects V_a in point y then $\|y - \sum_{t=1}^{r} \lambda_{B,a,a_t} x_{a_t}\| < (r+1)\sqrt{M_0 c}$. Hence $\|x - y\| < 2(r+1)\sqrt{M_0 c} = M\sqrt{c}$. By the definition of $G'_{N,d,c}$, we must have $x = y$. $\square$

To prove Theorem 5.1.2, it suffices to show that the expected number of copies of K_s in $G'_{N,d,c}$ is at least $N^r e^{-O(\sqrt{\log N})}$ for suitable choices of d and c. For this purpose we shall use the following technical lemma. For later convenience (in Section 5.4), we state it in a slightly more general form than required here, to allow the possibility that $r = s$ and the possibility that $p_1, \ldots, p_s$ are not in general position (but still span $\mathbb{R}^r$).

Lemma 5.3.2. *Let $1 \leq r \leq s$ be positive integers and let $p_1, \ldots, p_s$ be points on S^{r-1} such that $p_1, \ldots, p_r$ form a basis of $\mathbb{R}^r$. Then there exist constants $\alpha > 0$ and h such that for any $d \geq r$ and $0 < c < 1$ the probability that a set $\{x_a : 1 \leq a \leq s\}$ of random unit vectors (chosen independently and uniformly) on S^d satisfies $|\langle x_a, x_b \rangle - \langle p_a, p_b \rangle| < c$ for all a, b is at least $\alpha^d c^{d(s-r)/2+h}$.*

We may think of the conclusion of Lemma 5.3.2 as follows. The dominant (smallest) factor in the probability above is the factor $c^{d(s-r)/2}$. The probability should be close to this because if we imagine placing the s points one by one and we have already picked $x_1, \ldots, x_i$ joined to each other, then

- if $i < r$, then x_{i+1} is restricted to a neighbourhood of a codimension-i surface, so with reasonably large probability (comparable to c^i) it is connected to all previous vertices;
- if $i \geq r$, then the linear dependencies between the points restrict x_{i+1} to be in a ball of radius about $c^{1/2}$ around a certain point, which has measure about $c^{d/2}$ (which is much smaller than c^r).

The proof of Lemma 5.3.2 is given in an appendix at the end of this chapter.

Proof of Theorem 5.1.2. By Lemma 5.2.1, there are constants c_0, B, C such that if $c < c_0$ then the probability that a given vertex $x_{a,i}$ is removed from $G_{N,d,c}$ when forming $G'_{N,d,c}$ is at most $NB^d(M\sqrt{c})^d \leq NC^d c^{d/2}$. Here $B > 0$ is an absolute constant and the constants $C, c_0 > 0$ depend on $r, s, p_1, \ldots, p_s$ only. Moreover, the event '$x_{a,i}$ is removed' is independent of any event of the form '$x_{1,i_1}, \ldots, x_{s,i_s}$ form a K_s in $G_{N,d,c}$'. Using Lemma 5.3.2, we deduce that the probability that $x_{1,i_i}, \ldots, x_{s,i_s}$ is contained in $G'_{N,d,c}$ and forms a K_s is at least $(1 - sNC^d c^{d/2})\alpha^d c^{d(s-r)/2+h}$ (where α, h depend on $r, s, p_1, \ldots, p_s$ only). So the expected number of copies of K_s in $G'_{N,d,c}$ is at least

$$N^s(1 - sNC^d c^{d/2})\alpha^d c^{d(s-r)/2+h}.$$

If $c = (2sNC^d)^{-2/d}$, then this is at least

$$\frac{1}{2}N^s\alpha^d \frac{1}{(2sC^d)^{s-r}} N^{r-s}(2sNC^d)^{-2h/d} \geq \eta^d N^{r-2h/d} = N^r e^{-Ed-(2h/d)\log N} \tag{5.2}$$

for some constants $\eta > 0$ and E not depending on N, d.

Choosing $d = \lfloor\sqrt{\log N}\rfloor$, this is $N^r e^{-O(\sqrt{\log N})}$, and $c < c_0$ when N is sufficiently large. The result follows, as $G'_{N,d,c}$ has at most Ns vertices. □

Note that our proof in fact also gives the correct (and trivial) lower bound $\Theta(n)$ in the case $r = 1$, since if $r = 1$ then $h = 0$ so we may choose d to be a constant and get $\Theta(N)$ in (5.2).

5.4 Generalized rainbow Turán numbers for complete graphs

We now turn to the proofs of our results about generalized rainbow Turán numbers (Theorems 5.1.4 and 5.1.5). First we recall a general result of Gerbner, Mészáros, Methuku and Palmer [72], which can be proved using the graph removal lemma.

Proposition 5.4.1 (Gerbner, Mészáros, Methuku and Palmer [72])**.** *For any graph H on r vertices, we have* $\mathrm{ex}(n, H, \text{rainbow-}H) = o(n^{r-1})$.

In particular, we know that $\mathrm{ex}(n, K_r, \text{rainbow-}K_r) = o(n^{r-1})$. We would like to match this with a lower bound of the form $n^{r-1-o(1)}$.

Before we prove such a bound, let us briefly discuss the ideas that underlie the proof. It is easy to show that a lower bound $\mathrm{ex}(n, K_4, \text{rainbow-}K_4) \geq n^{3-o(1)}$ would imply that we have $\mathrm{ex}(n, K_r, \text{rainbow-}K_r) \geq n^{r-1-o(1)}$ for all $r \geq 4$, so it suffices to consider the case $r = 4$. However, when $r = 4$ and G is a properly edge-coloured graph with no rainbow K_4, then every triangle of G is contained in at most three copies of K_4. (Indeed, if the vertices of the K_3 are x, y, z, then the only way that adding a further vertex w can lead to a non-rainbow K_4 is if wx has the same colour as yz, wy has the same colour as xz or wz has the same colour as xy. But since the edge-colouring is proper, we cannot find more than one w such that the same one of these three events occurs.) So it is natural to expect that our construction for Theorem 5.1.2 is relevant here.

To see how a similar construction gives the desired result, it is helpful, as earlier, to look at a simpler continuous example that serves as a guide to the construction. Consider the graph where the vertex set is S^d and two unit vectors v, w are joined if and only if $\langle v, w\rangle = -1/3$ (the angle between vectors that go through the origin and two distinct vertices of a regular tetrahedron). Then any K_4 in this graph must be given by the vertices of a regular tetrahedron. We colour an edge by the line that joins the origin to the midpoint of that edge. This is a proper colouring with the property that opposite edges have the same colour, so each K_4 is 3-coloured in this colouring. The construction we are about to describe is a suitable perturbation and discretization of this one.

For the discretized graph, we will again have 'near-regular' tetrahedra forming K_4s. To ensure that each copy of K_4 is still rainbow, we shall have to modify the colouring slightly. We shall take only certain 'allowed lines' as colours, and we shall colour an edge by the allowed line that is closest to the line through the midpoint (if that line is not very far – otherwise we delete the edge). We need to choose the allowed lines in such a way that no two allowed lines are close (so that near-regular tetrahedra are still 3-coloured), but a large proportion of lines are close to an allowed line (so that not too many edges are deleted). This can be achieved using the following lemma.

Lemma 5.4.2. *There exists $\delta > 0$ with the following property. For any positive integer d and any $0 < c_1 < 1$, we can choose $L \geq (\delta/c_1)^d$ points $q_1, \ldots, q_L$ on S^d such that $\|q_i - \epsilon q_j\| \geq 3c_1$ for any $i \neq j$ and any $\epsilon \in \{1, -1\}$.*

Proof. Take a maximal set of points satisfying the condition above. Then the balls of radius $3c_1$ around the points $\pm q_1, \ldots, \pm q_L$ cover the entire sphere. But any such ball covers a proportion of surface area at most $(Bc_1)^d$ for some constant B (by Lemma 5.2.1). Therefore $2L(Bc_1)^d \geq 1$, which gives the result. □

One can prove Theorem 5.1.4 using the method described above. However, the proof naturally yields the more general Theorem 5.1.5 (which is restated below), so that is what we shall do. Essentially, we can prove a lower bound of $n^{m-o(1)}$ for a graph H whenever we can draw H in $\mathbb{R}^m$ in such a way that for each colour there is a line through the origin meeting (the line of) each edge of that colour, and the vertices of the graph span $\mathbb{R}^m$.

Theorem 5.1.5. *Let $r \geq 4$, let H be a graph, and let H have a proper edge-colouring with no rainbow K_r. Suppose that for each vertex v of H there is a $p_v \in \mathbb{R}^m$, and for each colour κ in the colouring there is a non-zero vector z_κ such that for every edge vw of colour κ, z_κ is a linear* *combination of p_v and p_w with non-zero coefficients. Then* $\mathrm{ex}(n, H, \text{rainbow-}K_r) \geq n^{m_0 - o(1)}$, *where m_0 is the dimension of the subspace of $\mathbb{R}^m$ spanned by the points p_v.*

Proof. Passing to a subspace, we may assume that $m = m_0$ and $\{p_v : v \in V(H)\}$ spans $\mathbb{R}^m$. Furthermore, by rescaling we may assume that each z_κ and each non-zero p_v has unit length. Write $V_0 = \{v \in V(H) : p_v = 0\}$ and $V_1 = \{v \in V(H) : p_v \neq 0\}$. For each $c > 0$ and any two positive integers N, d, we define a (random) graph $F_{N,d,c}$ as follows. The vertex set of $F_{N,d,c}$ has $|H|$ vertex classes labelled by the vertices of H. If $v \in V_0$ then there is a single point $x_{v,1} = 0$ in the class labelled by v. If $v \in V_1$, then we pick (uniformly and independently at random) N points $x_{v,1}, \dots, x_{v,N}$ on S^d: these will be the vertices in the class labelled by v. We join two vertices $x_{v,a}$ and $x_{w,b}$ by an edge if and only if $vw \in E(H)$ and $|\langle x_{v,a}, x_{w,b}\rangle - \langle p_v, p_w\rangle| < c$.

By assumption, we know that for each edge vw of colour κ there exist $\lambda_{\kappa,v}, \lambda_{\kappa,w}$ non-zero real coefficients such that $z_\kappa = \lambda_{\kappa,v} p_v + \lambda_{\kappa,w} p_w$. Let λ be the minimum and M_0 the maximum over all values of $|\lambda_{\kappa,v}|$. Write $c_1 = (12 M_0^2 c)^{1/2}$. Form a new graph $F'_{N,d,c}$ out of $F_{N,d,c}$ by removing any vertex $x_{v,i}$ for which there is another vertex $x_{v,j}$ (with $j \neq i$) such that $\|x_{v,i} - x_{v,j}\| \leq \frac{2}{\lambda} c_1$. (The exact values of the constants are not particularly important – they were chosen so that the graph described below will be properly coloured with no rainbow K_r. That is, we could replace $12M_0^2$ and $2/\lambda$ by other sufficiently large constants.)

Let $q_1, \dots, q_L$ be points on S^d with $L \geq (\delta/c_1)^d$ such that $\|q_i - q_j\| \geq 3c_1$ for all $i \neq j$. Here δ is some positive (absolute) constant, and the existence of such a set follows from Lemma 5.4.2. Also, pick independently and uniformly at random a rotation $R_\kappa \in \mathrm{SO}(d+1)$ for each colour κ used in the edge-colouring of H. The probability measure we use on $\mathrm{SO}(d+1)$ is the usual (Haar) measure, so for any $q \in S^d$ the points $R_\kappa q$ are independently and uniformly distributed on S^d. We think of the points $R_\kappa q_\ell$ ($\ell = 1, \dots, L$) as the allowed colours for the edges $x_{v,i} x_{w,j}$ when $vw \in E(H)$ has colour κ (and we take different rotations for different colours to have independence).

We form an edge-coloured graph $F''_{N,d,c}$ from $F'_{N,d,c}$ as follows. For any edge $x_{v,i} x_{w,j}$ of $F'_{N,d,c}$, we perform the following modification. Let κ be the colour of vw in $E(H)$, and let $\lambda_{\kappa,v}, \lambda_{\kappa,w} \neq 0$ be as before, so that $z_\kappa = \lambda_{\kappa,v} p_v + \lambda_{\kappa,w} p_w$.

- If there is some ℓ with $\|\lambda_{\kappa,v} x_{v,i} + \lambda_{\kappa,w} x_{w,j} - R_\kappa q_\ell\| < c_1$, then we colour the edge $x_{v,i} x_{w,j}$ with colour (κ, ℓ). Note that such an ℓ must be unique since $\|R_\kappa q_\ell - R_\kappa q_{\ell'}\| \geq 3c_1$ if $\ell' \neq \ell$.

- Otherwise we delete the edge $x_{v,i} x_{w,j}$.

Claim 1. The edge-colouring of $F''_{N,d,c}$ is proper.

Proof. Suppose that $x_{v,i} x_{w,j}$ and $x_{v,i} x_{w',j'}$ are both edges with colour (κ, ℓ). Then vw and

vw' both have colour κ in $E(H)$, thus $w = w'$. Also,

$$\begin{aligned}\|x_{w,j'} - x_{w,j}\| &\leq \frac{1}{|\lambda_{\kappa,w}|}\left(\|\lambda_{\kappa,v}x_{v,i} + \lambda_{\kappa,w}x_{w,j'} - R_\kappa q_\ell\| + \|\lambda_{\kappa,v}x_{v,i} + \lambda_{\kappa,w}x_{w,j} - R_\kappa q_\ell\|\right)\\ &\leq \frac{1}{|\lambda_{\kappa,w}|}2c_1\\ &\leq \frac{2}{\lambda}c_1.\end{aligned}$$

But then $j = j'$ by the definition of $F'_{N,d,c}$. So the edge-colouring of $F''_{N,d,c}$ is indeed proper. $\square$

Claim 2. There is no rainbow copy of K_r in $F''_{N,d,c}$.

Proof. Suppose that the vertices $x_{v_1,i_1}, \ldots, x_{v_r,i_r}$ form a K_r in $F''_{N,d,c}$. Then $v_1, \ldots, v_r$ form a K_r in H. This K_r is not rainbow (by assumption). By symmetry, we may assume that the edges v_1v_2 and v_3v_4 both have colour κ. Write x_a for x_{v_a,i_a} and λ_a for λ_{κ,v_a} for $a = 1, 2, 3, 4$. Then we have (recalling that $M_0 = \max_{\kappa',v}|\lambda_{\kappa',v}|$)

$$\begin{aligned}&\|\lambda_1x_1 + \lambda_2x_2 - \lambda_3x_3 - \lambda_4x_4\|^2\\ &\quad= \langle\lambda_1x_1 + \lambda_2x_2 - \lambda_3x_3 - \lambda_4x_4, \lambda_1x_1 + \lambda_2x_2 - \lambda_3x_3 - \lambda_4x_4\rangle\\ &\quad= \sum_{a=1}^4 \lambda_a^2\|x_a\|^2 + 2\lambda_1\lambda_2\langle x_1, x_2\rangle - 2\lambda_1\lambda_3\langle x_1, x_3\rangle - 2\lambda_1\lambda_4\langle x_1, x_4\rangle\\ &\qquad\qquad - 2\lambda_2\lambda_3\langle x_2, x_3\rangle - 2\lambda_2\lambda_4\langle x_2, x_4\rangle + 2\lambda_3\lambda_4\langle x_3, x_4\rangle\\ &\quad\leq \sum_{a=1}^4 \lambda_a^2\|p_{v_a}\|^2 + 2\lambda_1\lambda_2\langle p_{v_1}, p_{v_2}\rangle - 2\lambda_1\lambda_3\langle p_{v_1}, p_{v_3}\rangle - 2\lambda_1\lambda_4\langle p_{v_1}, p_{v_4}\rangle\\ &\qquad\qquad - 2\lambda_2\lambda_3\langle p_{v_2}, p_{v_3}\rangle - 2\lambda_2\lambda_4\langle p_{v_2}, p_{v_4}\rangle + 2\lambda_3\lambda_4\langle p_{v_3}, p_{v_4}\rangle + 12M_0^2c\\ &\quad= \langle\lambda_1p_{v_1} + \lambda_2p_{v_2} - \lambda_3p_{v_3} - \lambda_4p_{v_4}, \lambda_1p_{v_1} + \lambda_2p_{v_2} - \lambda_3p_{v_3} - \lambda_4p_{v_4}\rangle + 12M_0^2c\\ &\quad= 12M_0^2c.\end{aligned}$$

Since $c_1 = (12M_0^2c)^{1/2}$, we get that $\|\lambda_1x_1 + \lambda_2x_2 - \lambda_3x_3 - \lambda_4x_4\| \leq c_1$. But if x_1x_2 has colour (κ, ℓ) and x_3x_4 has colour (κ, ℓ'), then

$$\|q_\ell - q_{\ell'}\| \leq \|\lambda_1x_1 + \lambda_2x_2 - R_\kappa q_\ell\| + \|\lambda_3x_3 + \lambda_4x_4 - R_\kappa q_{\ell'}\| + \|\lambda_1x_1 + \lambda_2x_2 - \lambda_3x_3 - \lambda_4x_4\| < 3c_1.$$

It follows that $\ell = \ell'$ and hence the K_r with vertices $x_{v_1,i_1}, \ldots, x_{v_r,i_r}$ is not rainbow. $\square$

Claim 3. The expected number of copies of H in $F''_{N,d,c}$ is at least $N^{m-o(1)}$ if d and c are chosen appropriately.

Proof. Pick arbitrary vertices x_{v,i_v} in the classes (with $i_v = 1$ if $v \in V_0$ and $1 \leq i_v \leq N$ otherwise). We consider the probability that they form a copy of H in $F''_{N,d,c}$. Write x_v for x_{v,i_v}.

Let $\epsilon > 0$ be a small constant to be specified later. By Lemma 5.3.2, we have

$$\mathbb{P}[|\langle x_v, x_w\rangle - \langle p_v, p_w\rangle| < \epsilon c \text{ for all } v, w \in V(H)] \geq \alpha^d (\epsilon c)^{d(|V_1|-m)/2+h} \tag{5.3}$$

for some constants $\alpha > 0$ and h.

Let $v \in V_1$. By Lemma 5.2.1, the probability that x_v is removed when we form $F'_{N,d,c}$ is at most $NB_1^d c^{d/2}$, for some constant $B_1 > 0$ that does not depend on N, d, c. By independence, if $NB_1^d c^{d/2} < 1$ then

$$\mathbb{P}[\text{none of the } x_v \text{ are removed when we form } F'_{N,d,c}] \geq (1 - NB_1^d c^{d/2})^{|V_1|}. \tag{5.4}$$

Finally, for each colour κ in the colouring of $E(H)$, pick an edge $v_\kappa w_\kappa$ of that colour in H. Write $y_\kappa = \lambda_{\kappa,v_\kappa} x_{v_\kappa} + \lambda_{\kappa,w_\kappa} x_{w_\kappa}$ and $y'_\kappa = \frac{y_\kappa}{\|y_\kappa\|}$. Note that $\|y_\kappa\| \neq 0$ with probability 1, since all $\lambda_{\kappa,v}$ are non-zero and at least one of p_{v_κ} and p_{w_κ} is non-zero. For each κ, if ϵ is sufficiently small then by Lemma 5.2.1 we have

$$\mathbb{P}[\text{there is some } \ell_\kappa \text{ such that } \|y'_\kappa - R_\kappa q_{\ell_\kappa}\| < \epsilon c^{1/2}] \geq L\eta^d (\epsilon c^{1/2})^d \geq \eta_1^d \epsilon^d \tag{5.5}$$

for some constants $\eta, \eta_1 > 0$.

Observe that the events in (5.3), (5.4) and (5.5) (for all κ) are independent. It follows that

$$\mathbb{P}[\text{the events in (5.3), (5.4), and, for all } \kappa, \text{ (5.5) hold}] \geq \gamma_\epsilon^d c^{d(|V_1|-m)/2+h} (1 - NB_1^d c^{d/2})^{|V_1|} \tag{5.6}$$

where γ_ϵ is some constant depending on ϵ (but not on N, d, c). We show that these events together imply that the x_v form a copy of H, if ϵ is sufficiently small. The only property that we need to check is that no edge is removed when $F''_{N,d,c}$ is formed out of $F'_{N,d,c}$. Consider then an edge uu' of H. Let κ be its colour and write $v = v_\kappa, w = w_\kappa, y = y_\kappa, y' = y'_\kappa, \lambda_v = \lambda_{\kappa,v}, \lambda_w = \lambda_{\kappa,w}, \lambda_u = \lambda_{\kappa,u}$, and $\lambda_{u'} = \lambda_{\kappa,u'}$. We have

$$\begin{aligned}
\langle y, y\rangle &= \langle \lambda_v x_v + \lambda_w x_w, \lambda_v x_v + \lambda_w x_w\rangle \\
&= \langle \lambda_v p_v + \lambda_w p_w, \lambda_v p_v + \lambda_w p_w\rangle + O(\epsilon c) \\
&= \langle z_\kappa, z_\kappa\rangle + O(\epsilon c) \\
&= 1 + O(\epsilon c).
\end{aligned}$$

So

$$\|y - y'\| = |\|y\| - 1| = O(\epsilon c).$$

Furthermore, if we write $y'' = \lambda_u x_u + \lambda_{u'} x_{u'}$, then

$$\begin{aligned}\langle y - y'', y - y''\rangle &= \langle \lambda_v x_v + \lambda_w x_w - \lambda_u x_u - \lambda_{u'} x_{u'}, \lambda_v x_v + \lambda_w x_w - \lambda_u x_u - \lambda_{u'} x_{u'}\rangle \\ &= \langle \lambda_v p_v + \lambda_w p_w - \lambda_u p_u - \lambda_{u'} p_{u'}, \lambda_v p_v + \lambda_w p_w - \lambda_u p_u - \lambda_{u'} p_{u'}\rangle + O(\epsilon c) \\ &= \langle z_\kappa - z_\kappa, z_\kappa - z_\kappa\rangle + O(\epsilon c) \\ &= O(\epsilon c).\end{aligned}$$

It follows that

$$\begin{aligned}\|y'' - R_\kappa q_{\ell_\kappa}\| &\leq \|y'' - y\| + \|y - y'\| + \|y' - R_\kappa q_{\ell_\kappa}\| \\ &\leq O((\epsilon c)^{1/2}) + O(\epsilon c) + \epsilon c^{1/2}.\end{aligned}$$

This is indeed less than $c_1 = (12M_0^2 c)^{1/2}$ if ϵ is sufficiently small.

Choosing ϵ appropriately, (5.6) gives that the expected number of copies of H in $F''_{N,d,c}$ is at least

$$N^{|V_1|}\gamma^d c^{(|V_1|-m)d/2+h}(1 - NB_1^d c^{d/2})^{|V_1|}$$

for some constant γ. Letting $c = \left(\frac{1}{2NB_1^d}\right)^{2/d}$ and $d = \lfloor\sqrt{\log N}\rfloor$, we get that the expected number of copies of H in $F''_{N,d,c}$ is at least $N^{m-o(1)}$, which proves the claim. $\square$

The theorem follows from Claims 1, 2 and 3. $\square$

Deduction of Theorem 5.1.4 from Theorem 5.1.5. Given a complete graph K_r on vertex set $\{1, \dots, r\}$, we can properly edge-colour it by giving the edges 12 and 34 the same colour κ, and giving arbitrary different colours to the remaining edges. Pick $r-1$ linearly independent points $p_2, p_3, \dots, p_r$ in $\mathbb{R}^{r-1}$, and let $p_1 = p_2 + p_3 + p_4$. Let $z_\kappa = p_3 + p_4 = p_1 - p_2$ and let $z_{\kappa'} = p_i + p_j$ when ij is an edge of colour $\kappa' \neq \kappa$. Theorem 5.1.5 gives that $\mathrm{ex}(n, K_r, \text{rainbow-}K_r) \geq n^{r-1-o(1)}$, and we have a matching upper bound by Proposition 5.4.1. $\square$

5.5 Some applications of Theorem 5.1.5

We have already seen that Theorem 5.1.5 can be used to answer the question of Gerbner, Mészáros, Methuku and Palmer about the order of magnitude of $\mathrm{ex}(n, K_r, \text{rainbow-}K_r)$. In this section we give some other examples of applications of the theorem.

To show that our lower bounds are sharp, we shall use a simple proposition to give matching upper bounds. This will require the following definition. Given a graph H and a proper edge-colouring c of H, we say that a subset $V_0 \subseteq V(H)$ is a *c-spanning set* if there is an ordering $v_1, \dots, v_k$ of the vertices in $V(H)\backslash V_0$ such that for all i there are some $u, u', w \in V_0 \cup \{v_1, \dots, v_{i-1}\}$ such that $uu' \in E(H)$, $v_i w \in E(H)$ and $c(uu') = c(v_i w)$. In other words, we can add the remaining vertices to V_0 one by one in such a way that new vertices are joined to some vertex in the set by a colour already used.

Proposition 5.5.1. *Let H and F be graphs and let r be a positive integer. Assume that for every proper edge-colouring c of H that does not contain a rainbow copy of F there is a c-spanning set of size at most r. Then* $\mathrm{ex}(n, H, \text{rainbow-}F) = O(n^r)$*. If we also have $r < |V(H)|$, and if for every such c and every edge e of H there is a c-spanning set of size at most r containing e, then* $\mathrm{ex}(n, H, \text{rainbow-}F) = o(n^r)$.

Proof. Let G be a graph on n vertices and let κ be a proper edge-colouring of G without a rainbow copy of F. Let G contain M copies of H. Then we can partition the vertices into classes X_v for $v \in V(H)$ in such a way that there are $\Omega(M)$ choices of $\mathbf{x} = (x_v)_{v\in V(H)}$ such that $x_v \in X_v$ and $v \mapsto x_v$ is a graph homomorphism from H. (To see this, place each vertex independently, uniformly at random into one of the classes. If $\{x_v : v \in V(H)\}$ is an isomorphic copy of H in G (such that $v \mapsto x_v$ is the corresponding isomorphism), then we have $\mathbb{P}[x_v \in X_v \text{ for all } v] = 1/|V(H)|^{|V(H)|}$, so the expected number of such tuples $\mathbf{x}$ is $M/|V(H)|^{|V(H)|} = \Omega(M)$.)

For each $\mathbf{x}$ as above pick an isomorphic proper edge-colouring $c_{\mathbf{x}} : E(H) \to \{1, \dots, |E(H)|\}$, that is, $c_{\mathbf{x}}(vw) = c_{\mathbf{x}}(v'w')$ if and only if $\kappa(x_v x_w) = \kappa(x_{v'} x_{w'})$ for all edges $vw, v'w'$ of H. Note that $c_{\mathbf{x}}$ cannot contain a rainbow copy of F. Then there is a colouring $c : E(H) \to \{1, \dots, |E(H)|\}$ that appears for $\Omega(M)$ choices of $\mathbf{x}$. Let V_0 be a c-spanning set of size at most r.

Note that any $\mathbf{x}$ with $c_{\mathbf{x}} = c$ is determined by $(x_v)_{v\in V_0}$, since the edge-colouring is proper. But there are $O(n^r)$ choices for $(x_v)_{v\in V_0}$, hence $M = O(n^r)$.

Now assume that $r < |V(H)|$ and that for every proper edge-colouring c' of H without a rainbow F and every edge e of H there is a c'-spanning set of size at most r that contains e. By the graph removal lemma (see, e.g., [38]) and the first part of our proposition, we can remove $o(n^2)$ edges from G so that the new graph G' contains no copy of H. So it suffices to show that each edge appeared in at most $O(n^{r-2})$ tuples $\mathbf{x}$ with $c_{\mathbf{x}} = c$. Given an edge $e = y_v y_w$ with $y_v \in X_v, y_w \in X_w, vw \in E(H)$ we can pick in H a c-spanning set $V_{0,e}$ of size at most r containing vw. Then any $\mathbf{x}$ with $c_{\mathbf{x}} = c$ and $x_v = y_v, x_w = y_w$ is determined by $(x_u)_{u\in V_0\setminus\{v,w\}}$, which gives the result. □

Now we give some sample applications of Theorem 5.1.5 and Proposition 5.5.1. We shall give two illustrations, but it is quite easy to generate additional examples.

5.5.1 Complete graphs

Perhaps the most natural extension of Question 5.1.3 is to determine the behaviour of the function $\mathrm{ex}(n, K_r, \text{rainbow-}K_s)$. Note that trivially $\mathrm{ex}(n, K_r, \text{rainbow-}K_s) = \Theta(n^r)$ when $s > r$ (by taking a complete r-partite graph), and we have seen that $\mathrm{ex}(n, K_s, \text{rainbow-}K_s) = n^{s-1-o(1)}$ (when $s \geq 4$). We also have $\mathrm{ex}(n, K_r, \text{rainbow-}K_s) = 0$ whenever $r \geq r_s$ for some integer r_s depending on s. Indeed, if we have a K_r with no rainbow copy of K_s, and the largest rainbow subgraph has order $t \leq s$, then any of the remaining $(r-t)$ vertices must be joined to this K_t by one of the $\binom{t}{2}$ colours appearing in the K_t. But each such colour appears at most once at each vertex, giving $r = O(s^3)$. In fact, Alon, Lefmann and Rödl [7] showed that $r_s = \Theta(s^3/\log s)$.

However, the question is non-trivial for $s < r < r_s$. First note that $\mathrm{ex}(n, K_r, \text{rainbow-}K_s) = o(n^{s-1})$ whenever $r \geq s$ by Proposition 5.5.1 (since any maximal rainbow subgraph is a c-spanning set). The simplest case for the lower bound is $(r, s) = (5, 4)$. In this case Theorem 5.1.5 gives a matching lower bound $n^{3-o(1)}$. Indeed, take an arbitrary proper edge-colouring of K_5 with no rainbow K_4, and take points $p_1, \ldots, p_5$ in general position in $\mathbb{R}^3$. The existence of appropriate values of z_κ follows from the fact that any four of the p_i are linearly dependent (but any three are independent), and each colour is used at most twice. It is easy to deduce that $\mathrm{ex}(n, K_{s+1}, \text{rainbow-}K_s) = n^{s-1-o(1)}$ for all $s \geq 4$.

When $s = 4$ then $r_s = 7$, leaving the case $(r, s) = (6, 4)$. Unfortunately, in this case Theorem 5.1.5 does not give a lower bound of $n^{3-o(1)}$. (To see this, observe that to get such a bound the corresponding points p_v would all have to be non-zero. Then we can use the alternative formulation Theorem 5.1.5' to see that we would have to be able to draw a properly edge-coloured K_6 in the plane such that there is no rainbow K_4 and lines of edges of the same colour are either all parallel or go through the same point. Applying an appropriate projection and affine transformation, we may assume that we have two colour classes where the edges are all parallel, and these two parallel directions are perpendicular. This leaves essentially two cases to be checked, and neither of them yields an appropriate configuration.)

However, we can still deduce a lower bound of $\mathrm{ex}(n, K_6, \text{rainbow-}K_4) \geq n^{12/5-o(1)}$, as sketched below. We can take 6 points $p_0 = 0$ and $p_a = e^{2\pi i a/5}$ (for $a = 1, \ldots, 5$), that is, the vertices of a regular pentagon together with its centre. We define a colouring c as follows. Give parallel lines between vertices of the pentagon the same colour, and also give the same colour to the edge incident at the centre which is perpendicular to these lines (see Figure 5.1). This gives a proper edge-colouring of K_6 and corresponding points in 2 dimensions for which the conditions of Theorem 5.1.5 are satisfied, giving a lower bound of $n^{2-o(1)}$. (The point z_κ is chosen to be p_a when $p_0 p_a$ has colour κ.) This can be improved to $n^{12/5-o(1)}$ by a product argument as follows. Looking at the construction, we see that our graph G is 6-partite with classes $V_0, \ldots, V_5$, at most n vertices, and a proper edge-colouring κ such that the following hold.

- There are (at least) $n^{2-o(1)}$ copies of K_6 in G.
- The class $|V_0|$ has size 1.
- There is a 5-colouring c of the edges of K_6 (on vertex set $\{0, \ldots, 5\}$) with no rainbow K_4 such that whenever $v_{i_1}, v_{i_2}, v_{i_3}, v_{i_4}$ form a K_4 in G with $v_{i_j} \in V_{i_j}$, then $i_j \mapsto v_{i_j}$ gives an isomorphism of colourings between the restrictions of c and κ to the appropriate four-vertex graphs (i.e., $\kappa(v_{i_j} v_{i_\ell}) = \kappa(v_{i_{j'}} v_{i_{\ell'}})$ if and only if $c(i_j i_{j'}) = c(i_{j'} i_{\ell'})$). Moreover, this

5-colouring c has the property that for all $i, j \in \{1, \dots, 6\}$ there is a permutation of the vertices $\{0, \dots, 5\}$ which is an automorphism of colourings and maps i to j. (Indeed, we can take rotations of the pentagon when $i, j \neq 0$, and we can take the permutation $(01)(34)$ when $i = 0$, $j = 1$.)

We construct a new graph as follows. For each $i \in \{0, \dots, 5\}$, pick a permutation π_i of $\{0, \dots, 5\}$ which gives a colouring automorphism of c and sends i to 0. Define a 6-partite graph G_i obtained from G by permuting the vertex classes: G_i has classes $V_0^i, \dots, V_5^i$ given by $V_a^i = V_{\pi_i(a)}$ and same edge set as G. Let G' be the product of these 6-partite graphs, that is, it is 6-partite with vertex classes $W_a = V_a^0 \times V_a^1 \times \cdots \times V_a^5$, and two vertices $(v_0, \dots, v_5) \in W_a$ and $(w_0, \dots, w_5) \in W_b$ are joined by an edge if $v_i w_i \in E(G)$ for all i. Moreover, colour such an edge by colour $(\kappa(v_0 w_0), \dots, \kappa(v_5 w_5))$. It is easy to check that the colouring is proper, G' contains no rainbow K_4, G' has at most n^5 vertices in each class, and G' contains at least $n^{12-o(1)}$ copies of K_6, giving the bound stated.

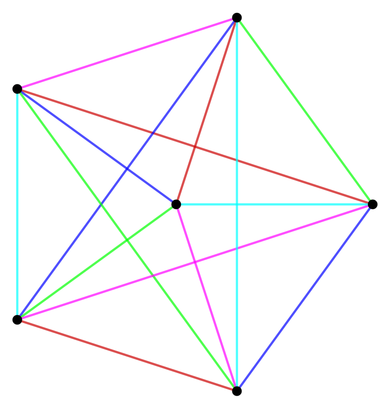

Figure 5.1: The colouring and points used for $(r, s) = (6, 4)$ to get a lower bound.

This leaves some open questions about $\mathrm{ex}(n, K_r, \text{rainbow-}K_s)$. It would be interesting to determine its order of magnitude for $(r, s) = (6, 4)$, or the magnitude for other pairs with $s < r < r_s$.

5.5.2 King's graphs

Given positive integers $k, \ell \geq 2$, write $H_{k,\ell}$ for the graph with vertex set $\{1, \dots, k\} \times \{1, \dots, \ell\}$ where (a, b) and (a', b') are joined by an edge if and only if they are distinct and $|a - a'|, |b - b'| \leq 1$. In other words, $H_{k,\ell}$ is the strong product of a path with k points and a path with ℓ points, sometimes called the $k \times \ell$ king's graph. We can use our results to show that $\mathrm{ex}(n, H_{k,\ell}, \text{rainbow-}K_4) = n^{k+\ell-1-o(1)}$.

First consider the upper bound. It is easy to see that any sequence of vertices $p_1, \dots, p_{k+\ell-1}$

is a c-spanning set (for all proper edge-colourings c of $H_{k,\ell}$ without a rainbow K_4) if either of the following statements holds.

1. We have $p_1 = (1, 1)$, $p_{k+\ell-1} = (k, \ell)$ and $p_{i+1} - p_i \in \{(0, 1), (1, 0)\}$ for all i.
2. We have $p_1 = (1, \ell)$, $p_{k+\ell-1} = (k, 1)$ and $p_{i+1} - p_i \in \{(0, 1), (-1, 0)\}$ for all i.

(Indeed, this follows from the fact that we can add the other vertices one by one, creating a new copy of K_4 in our set in each step.) Since any edge is contained in such a sequence, Proposition 5.5.1 gives $\mathrm{ex}(n, H_{k,\ell}, \text{rainbow-}K_4) = o(n^{k+\ell-1})$.

For the lower bound, consider an edge-colouring c of $H_{k,\ell}$ with $c((a, b)(a + 1, b)) = a$, where the other edges are given arbitrary distinct colours. This gives a proper edge-colouring of $H_{k,\ell}$ with no rainbow K_4. For each vertex (a, b) of H, define $p_{a,b} \in \mathbb{R}^{k+\ell}$ to be the vector with i^{th} coordinate

$$(p_{a,b})_i = \begin{cases} 0 & \text{if } i \neq a, k+b \\ 1 & \text{if } i = a \\ (-1)^a & \text{if } i = k+b \end{cases}$$

For each $1 \leq a \leq k-1$ we let $z_a \in \mathbb{R}^{k+\ell}$ be the vector with all entries zero except the a^{th} and $(a+1)^{\text{th}}$ coordinates which are 1, and for each other colour κ used in the colouring of $H_{k,\ell}$ we take $z_\kappa = p_v + p_w$, where vw is the unique edge of colour κ. Then we have $p_{a,b} + p_{a+1,b} = z_a$, so the conditions of Theorem 5.1.5 are satisfied. The dimension of the subspace of $\mathbb{R}^{k+\ell}$ spanned by the vectors $p_{a,b}$ is at least $k + \ell - 1$, since $p_{1,\ell}, p_{1,\ell-1}, \ldots, p_{1,1}, p_{2,1}, p_{3,1}, \ldots, p_{k,1}$ are linearly independent. We get the required lower bound $n^{k+\ell-1-o(1)}$.

5.6 Appendix

In this appendix, we prove Lemma 5.3.2, which is recalled below.

Lemma 5.3.2. *Let $1 \leq r \leq s$ be positive integers and let $p_1, \ldots, p_s$ be points on S^{r-1} such that $p_1, \ldots, p_r$ form a basis of $\mathbb{R}^r$. Then there exist constants $\alpha > 0$ and h such that for any $d \geq r$ and $0 < c < 1$ the probability that a set $\{x_a : 1 \leq a \leq s\}$ of random unit vectors (chosen independently and uniformly) on S^d satisfies $|\langle x_a, x_b\rangle - \langle p_a, p_b\rangle| < c$ for all a, b is at least $\alpha^d c^{d(s-r)/2+h}$.*

Lemma 5.6.1. *Let r be a positive integer. Let $p_1, \ldots, p_{r+1}$ be points on S^{r-1} such that $p_1, \ldots, p_r$ are linearly independent. Then there exist real numbers $\delta > 0$, $\alpha_0 > 0$ and h_0 such that whenever $d \geq r$ is a positive integer, $0 < c < 1$, and $x_1, \ldots, x_r$ are points on S^d with $|\langle x_i, x_j\rangle - \langle p_i, p_j\rangle| < \delta c$ for all $1 \leq i, j \leq r$, then the probability that a random point x_{r+1} on S^d satisfies $|\langle x_i, x_{r+1}\rangle - \langle p_i, p_{r+1}\rangle| < c$ for all $i \leq r$ is at least $\alpha_0^d c^{d/2+h_0}$.*

Proof. We prove the statement by induction on r. If $r = 1$, then $p_1, p_2 \in \{-1, 1\}$ and the condition $|\langle x_i, x_{r+1}\rangle - \langle p_i, p_{r+1}\rangle| < c$ becomes $|\langle p_1p_2x_1, x_2\rangle - 1| < c$, which is equivalent to

$\|p_1p_2x_1 - x_2\| < \sqrt{2c}$. By Lemma 5.2.1, this happens with probability at least $\alpha_0^d c^{d/2}$, giving the claim. (Here $\delta = 1$ and $h_0 = 0$.)

Now assume that $r \geq 2$ and the result holds for smaller values of r. We may assume that $p_{r+1} \neq \pm p_1$ (otherwise swap p_1 and p_2). By symmetry, we may assume that $x_1 = (0, 0, \ldots, 0, 1) \in S^d$ and $p_1 = (0, 0, \ldots, 0, 1) \in S^{r-1}$. Write x_a^i for the i^{th} coordinate of x_a. For each $2 \leq a \leq r+1$, define a normalized projected vector

$$x'_a = \frac{(x_a^1, x_a^2, \ldots, x_a^d)}{\|(x_a^1, x_a^2, \ldots, x_a^d)\|} \in S^{d-1}.$$

Note that the denominator is non-zero for $a = 2, \ldots, r$ if δ is sufficiently small, and it is non-zero with probability 1 for $a = r+1$. Also, x_{r+1} is uniformly distributed on S^{d-1}. Similarly, for each $2 \leq a \leq r+1$, define

$$p'_a = \frac{(p_a^1, p_a^2, \ldots, p_a^{r-1})}{\|(p_a^1, p_a^2, \ldots, p_a^{r-1})\|} \in S^{r-2}.$$

Note that $p'_2, \ldots, p'_r$ are linearly independent in $\mathbb{R}^{r-1}$.

Note that for $2 \leq a, b \leq r$ we have

$$\begin{aligned}
\langle x'_a, x'_b \rangle &= \frac{\langle x_a, x_b \rangle - \langle x_1, x_a \rangle \langle x_1, x_b \rangle}{(1 - \langle x_1, x_a \rangle^2)^{1/2} (1 - \langle x_1, x_b \rangle^2)^{1/2}} = \frac{\langle p_a, p_b \rangle - \langle p_1, p_a \rangle \langle p_1, p_b \rangle}{(1 - \langle p_1, p_a \rangle^2)^{1/2} (1 - \langle p_1, p_b \rangle^2)^{1/2}} + O(\delta c) \\
&= \langle p'_a, p'_b \rangle + O(\delta c).
\end{aligned}$$

Let $\epsilon > 0$ be a small constant to be specified later. Applying the induction hypothesis for $r' = r-1$ and points $p'_2, \ldots, p'_{r+1}$, we have that

$$\mathbb{P}[|\langle x'_a, x'_{r+1} \rangle - \langle p'_a, p'_{r+1} \rangle| < \epsilon c \text{ for all } 2 \leq a \leq r] \geq \alpha_0^{d-1} (\epsilon c)^{(d-1)/2 + h_0} \tag{5.7}$$

whenever $\delta < \delta_0 \epsilon$, for some constants $\alpha_0, \delta_0 > 0$ and h_0 depending on $p_1, \ldots, p_{r+1}$ only.

By Lemma 5.2.1, there is a constant β depending on p_1, p_{r+1} only such that

$$\mathbb{P}[|\langle x_1, x_{r+1} \rangle - \langle p_1, p_{r+1} \rangle| < \epsilon c] \geq \beta^d \epsilon c. \tag{5.8}$$

Note that the events in (5.7) and (5.8) are independent, since conditioning on the second event we still have an independent uniform distribution for the vector x'_{r+1}. It follows that

$$\mathbb{P}[|\langle x'_a, x'_{r+1} \rangle - \langle p'_a, p'_{r+1} \rangle| < \epsilon c \text{ for all } 2 \leq a \leq r \text{ and } |\langle x_1, x_{r+1} \rangle - \langle p_1, p_{r+1} \rangle| < \epsilon c] \geq \gamma^d (\epsilon c)^{d/2 + h_1} \tag{5.9}$$

whenever $\delta < \delta_0 \epsilon$, for some constants $\gamma > 0$ and h_1 (with the constants depending on $p_1, \ldots, p_{r+1}$ only).

So it suffices to show that if ϵ and δ are sufficiently small (depending on $p_1, \ldots, p_{r+1}$ only), then the event above implies that $|\langle x_a, x_{r+1}\rangle - \langle p_a, p_{r+1}\rangle| < c$ for all $2 \le a \le r$. But we have

$$\begin{aligned}\langle x_a, x_{r+1}\rangle &= \langle x'_a, x'_{r+1}\rangle(1-\langle x_1, x_a\rangle^2)^{1/2}(1-\langle x_1, x_{r+1}\rangle^2)^{1/2} + \langle x_1, x_a\rangle\langle x_1, x_{r+1}\rangle \\ &= (\langle p'_a, p'_{r+1}\rangle + O(\epsilon c))((1-\langle p_1, p_a\rangle^2)^{1/2} + O(\delta c))((1-\langle p_1, p_{r+1}\rangle^2)^{1/2} + O(\epsilon c)) \\ &\quad + (\langle p_1, p_a\rangle + O(\delta c))(\langle p_1, p_{r+1}\rangle + O(\epsilon c)) \\ &= \langle p_a, p_{r+1}\rangle + O((\epsilon+\delta)c),\end{aligned}$$

which gives the result. □

Lemma 5.6.2. *Let r be a positive integer and let $p_1, \ldots, p_r$ be linearly independent points in S^{r-1}. Then there exist real numbers $\alpha_1 > 0$ and h_1 such that whenever $d \ge r$ is a positive integer and $0 < c < 1$ then the probability that r points $x_1, \ldots, x_r$ chosen independently and uniformly at random on S^d satisfy $|\langle x_i, x_j\rangle - \langle p_i, p_j\rangle| < c$ for all $1 \le i, j \le r$ is at least $\alpha_1^d c^{h_1}$.*

Proof. The proof is essentially the same as for the previous lemma. We prove the statement by induction on r. The case $r = 1$ is trivial. Now assume that $r \ge 2$ and that the statement holds for smaller values of r. By symmetry, we may assume that $x_1 = (0, 0, \ldots, 0, 1) \in S^d$ and $p_1 = (0, 0, \ldots, 0, 1) \in S^{r-1}$. Define p'_a and x'_a for $a \ge 2$ as in the proof of Lemma 5.6.1. Let $\epsilon > 0$ be some small constant to be determined later.

By induction, we have

$$\mathbb{P}[|\langle x'_a, x'_b\rangle - \langle p'_a, p'_b\rangle| < \epsilon c \text{ for all } 2 \le a, b \le r] \ge \alpha_1^{d-1}(\epsilon c)^{h_1}$$

for some $\alpha_1 > 0$ and h_1 (where the constants depend only on $p_1, \ldots, p_{r+1}$).

By Lemma 5.2.1, there are constants $\beta_2, \ldots, \beta_r$ depending on $p_1, \ldots, p_r$ only such that for each $2 \le a \le r$

$$\mathbb{P}[|\langle x_1, x_a\rangle - \langle p_1, p_a\rangle| < \epsilon c] \ge \beta_a^d \epsilon c.$$

By independence,

$$\mathbb{P}[|\langle x'_a, x'_b\rangle - \langle p'_a, p'_b\rangle| < \epsilon c \text{ for all } 2 \le a, b \le r \text{ and } |\langle x_1, x_a\rangle - \langle p_1, p_a\rangle| < \epsilon c \text{ for all } a] \ge \gamma^d(\epsilon c)^{h_2}$$

for some real numbers $\gamma > 0$ and h_2.

However, if the event above holds then

$$\begin{aligned}\langle x_a, x_b\rangle &= \langle x'_a, x'_b\rangle(1-\langle x_1, x_a\rangle^2)^{1/2}(1-\langle x_1, x_b\rangle^2)^{1/2} + \langle x_1, x_a\rangle\langle x_1, x_b\rangle \\ &= \langle p'_a, p'_b\rangle(1-\langle p_1, p_a\rangle^2)^{1/2}(1-\langle p_1, p_b\rangle^2)^{1/2} + \langle p_1, p_a\rangle\langle p_1, p_b\rangle + O(\epsilon c) \\ &= \langle p_a, p_b\rangle + O(\epsilon c).\end{aligned}$$

The result follows by taking a sufficiently small ϵ. □

Proof of Lemma 5.3.2. By Lemma 5.6.1, we can choose constants $0 < \delta < 1$, $\alpha_0 > 0$ and h_0 such that whenever $d \geq r$ is a positive integer, $0 < c < 1$ and $x_1, \ldots, x_r$ are points on S^d with $|\langle x_i, x_j\rangle - \langle p_i, p_j\rangle| < \delta c$ for all $1 \leq i, j \leq r$, then for all $a > r$ the probability that a random point x_a on S^d satisfies $|\langle x_i, x_a\rangle - \langle p_i, p_a\rangle| < c$ for all $i \leq r$ is at least $\alpha_0^d c^{d/2+h_0}$.

Now let ϵ be a small constant to be specified later. Using Lemma 5.6.2, the observation above and independence of $x_{r+1}, \ldots, x_s$ conditional on $x_1, \ldots, x_r$, we have that

$$\begin{aligned}
\mathbb{P}[|\langle x_i, x_j\rangle - \langle p_i, p_j\rangle| < \delta\epsilon c \text{ whenever } i, j \leq r \text{ and } |\langle x_i, x_a\rangle - \langle p_i, p_a\rangle| < \epsilon c \text{ for all } i \leq r < a \leq s]& \\
\geq \alpha_1^d (\delta\epsilon c)^{h_1} \alpha_0^{(s-r)d} (\epsilon c)^{(s-r)(d/2+h_0)}& \\
\geq \alpha^d (\epsilon c)^{d(s-r)/2+h}&
\end{aligned}$$

for some constants $\alpha > 0$ and h. We show that the event above implies that $|\langle x_a, x_b\rangle - \langle p_a, p_b\rangle| < c$ even if $a, b > r$ (if ϵ is sufficiently small). Given $b > r$, we can find coefficients $\lambda_{b,a}$ such that $p_b = \sum_{a=1}^r \lambda_{b,a} p_a$. Write $y_b = \sum_{a=1}^r \lambda_{b,a} x_a$. Then

$$\begin{aligned}
\|y_b - x_b\|^2 &= \left\langle \sum_{a=1}^r \lambda_{b,a} x_a - x_b, \sum_{a=1}^r \lambda_{b,a} x_a - x_b \right\rangle \\
&= \left\langle \sum_{a=1}^r \lambda_{b,a} p_a - p_b, \sum_{a=1}^r \lambda_{b,a} p_a - p_b \right\rangle + O(\epsilon c) \\
&= O(\epsilon c).
\end{aligned}$$

Thus $\|y_b - x_b\| = O\left((\epsilon c)^{1/2}\right)$. Furthermore, we have, for each $1 \leq i \leq r$,

$$\begin{aligned}
\langle x_i, y_b - x_b\rangle &= \left\langle x_i, \sum_{a=1}^r \lambda_{b,a} x_a - x_b \right\rangle \\
&= \left\langle p_i, \sum_{a=1}^r \lambda_{b,a} p_a - p_b \right\rangle + O(\epsilon c) \\
&= O(\epsilon c).
\end{aligned}$$

It follows that whenever $b, b' > r$ then

$$\begin{aligned}
\langle x_b, x_{b'} \rangle &= \langle y_b + (x_b - y_b), y_{b'} + (x_{b'} - y_{b'}) \rangle \\
&= \langle y_b, y_{b'} \rangle + \langle x_b - y_b, y_{b'} \rangle + \langle y_b, x_{b'} - y_{b'} \rangle + O(\epsilon c) \\
&= \left\langle \sum_{a=1}^{r} \lambda_{b,a} x_a, \sum_{a=1}^{r} \lambda_{b',a} x_a \right\rangle + \left\langle x_b - y_b, \sum_{a=1}^{r} \lambda_{b',a} x_a \right\rangle + \left\langle \sum_{a=1}^{r} \lambda_{b,a} x_a, x_{b'} - y_{b'} \right\rangle + O(\epsilon c) \\
&= \left\langle \sum_{a=1}^{r} \lambda_{b,a} p_a, \sum_{a=1}^{r} \lambda_{b',a} p_a \right\rangle + O(\epsilon c) \\
&= \langle p_b, p_{b'} \rangle + O(\epsilon c).
\end{aligned}$$

Choosing a sufficiently small $\epsilon > 0$ gives the result. $\square$

Chapter 6

The generalized rainbow Turán problem for cycles

6.1 Introduction

In this chapter we continue the study of the generalised rainbow Turán problem. To make this chapter self-contained, we will repeat some definitions from the previous chapter.

The problem of estimating the maximal possible size $\mathrm{ex}(n, F)$ of an F-free graph on n vertices is one of the most fundamental problems in extremal graph theory. It is a well known fact that $\mathrm{ex}(n, F)/\binom{n}{2} \to 1 - 1/(r-1)$ as $n \to \infty$ if F has chromatic number r, determining the asymptotic behaviour of this function when F is not bipartite. However, much less is known in the bipartite case. See [69] for a survey on the topic.

Alon and Shikhelman [6] initiated the systematic study of the following generalisation of the problem above. Given two graphs H and F, let $\mathrm{ex}(n, H, F)$ denote the maximal number of copies of H that an F-free graph on n vertices can contain. Note that the usual Turán number $\mathrm{ex}(n, F)$ is the special case $\mathrm{ex}(n, K_2, F)$. This problem has been studied for several different choices of H and F, see, e.g., [6, 71, 73].

Another generalisation of the Turán problem was introduced by Keevash, Mubayi, Sudakov and Verstraëte [107]. Given an edge-coloured graph, we say that a subgraph is rainbow if all of its edges have different colours. Let $\mathrm{ex}^*(n, F)$ denote the maximal number of edges that a properly edge-coloured graph on n vertices can have if it contains no rainbow copy of F. Note that clearly $\mathrm{ex}(n, F) \le \mathrm{ex}^*(n, F)$, and in fact it was shown in [107] that $\mathrm{ex}^*(n, F) = \mathrm{ex}(n, F) + o(n^2)$, giving the asymptotic behaviour when F is not bipartite. This rainbow Turán problem has been studied for graphs F including paths [57, 100], cycles [45, 97, 107] and complete bipartite graphs [107], and for several graphs exact results are also known [107].

As discussed in Chapter 5, a common generalisation to the problems above was studied by Gerbner, Mészáros, Methuku and Palmer [72]. Let $\mathrm{ex}(n, H, \text{rainbow-}F)$ denote the maximal number of copies of H that a properly edge-coloured graph on n vertices can contain if it has no rainbow subgraph isomorphic to F. The authors of [72] focused mainly on the case $H = F$, and obtained several results, for example, when F is a path, cycle or a tree. Concerning cycles, they proved the following theorem.

Theorem 6.1.1 (Gerbner, Mészáros, Methuku and Palmer [72]). *If $k \geq 2$ is an integer, then*

$$\mathrm{ex}(n, C_{2k+1}, \text{rainbow-}C_{2k+1}) = \Theta(n^{2k-1})$$

and

$$\Omega(n^{k-1}) \leq \mathrm{ex}(n, C_{2k}, \text{rainbow-}C_{2k}) \leq O(n^k).$$

Moreover, if $\ell \geq 2$ is an integer with $\ell \neq k$, then

$$\mathrm{ex}(n, C_{2\ell}, \text{rainbow-}C_{2k}) = \Theta(n^\ell).$$

(Throughout this chapter, whenever we use the Ω, Θ or O notation, the implied constants may depend, as usual, on the other parameters present, such as k and ℓ above.) The authors of [72] asked what the correct order of magnitude is for the generalised rainbow Turán number $\mathrm{ex}(n, C_{2k}, \text{rainbow-}C_{2k})$. (They were able to improve the lower bound to $\Omega(n^{3/2})$ when $k = 2$ and the upper bound to $O(n^{8/3})$ when $k = 3$.) The main aim of this chapter is to obtain the following extension of Theorem 6.1.1.

Theorem 6.1.2. *If $s \geq 4$ and $t \geq 3$ are positive integers, then*

$$\mathrm{ex}(n, C_s, \text{rainbow-}C_t) = \begin{cases} \Theta(n^{s/2}) & \text{if } t = 4 \\ \Theta(n^{s/2}) & \text{if } s, t \text{ are even with } s \neq t \\ \Theta(n^{s/2-1}) & \text{if } s = t \geq 6 \text{ and } t \text{ is even} \\ \Theta(n^{(s-1)/2}) & \text{if } t \geq 6 \text{ is even and } s \text{ is odd} \\ \Theta(n^{s-2}) & \text{if } s, t \text{ are odd with } s \leq t \\ \Theta(n^s) & \text{if } t \text{ is odd, and } s > t \text{ or } s \text{ is even.} \end{cases}$$

In particular, this shows that $\mathrm{ex}(n, C_{2k}, \text{rainbow-}C_{2k}) = \Theta(n^{k-1})$ for all $k \geq 3$, whereas $\mathrm{ex}(n, C_4, \text{rainbow-}C_4) = \Theta(n^2)$, answering the question of Gerbner, Mészáros, Methuku and Palmer [72].

For comparison, we mention the order of magnitude of this function in the non-rainbow setting. We note that in many cases more precise bounds are known than the ones given below.

Theorem 6.1.3 (Gishboliner and Shapira [73], Gerbner, Győri, Methuku and Vizer [71]). *If*

$s \geq 4$ and $t \geq 3$ are distinct positive integers, then

$$\mathrm{ex}(n, C_s, C_t) = \begin{cases} \Theta(n^{s/2}) & \text{if } t = 4 \\ \Theta(n^{s/2}) & \text{if } s, t \text{ are even} \\ \Theta(n^{(s-1)/2}) & \text{if } t \geq 6 \text{ is even and } s \text{ is odd} \\ \Theta(n^{(s-1)/2}) & \text{if } s, t \text{ are odd with } s < t \\ \Theta(n^{s}) & \text{if } t \text{ is odd, and } s > t \text{ or } s \text{ is even.} \end{cases}$$

As part of our proof, we will also determine the order of magnitude of the maximal number of paths of length ℓ if there is no rainbow copy of C_{2k} whenever $k, \ell \geq 2$. (By the path P_ℓ of length ℓ we mean the path with ℓ edges and $\ell + 1$ vertices.) This result is given in the following theorem. Note that the answer is of the same order of magnitude as in the case of the corresponding (non-rainbow) generalised Turán problem [73], although our proof is rather different. Also, we trivially have $\mathrm{ex}(n, P_\ell, \text{rainbow-}C_t) = \Theta(n^{\ell+1})$ if t is odd.

Theorem 6.1.4. *If $k, \ell \geq 2$ are integers, then*

$$\mathrm{ex}(n, P_\ell, \text{rainbow-}C_{2k}) = \begin{cases} \Theta(n^{\lceil (\ell+1)/2 \rceil}) & \text{if } k \geq 3 \\ \Theta(n^{\ell/2+1}) & \text{if } k = 2. \end{cases}$$

Note that a path of length $\ell = 1$ is just an edge, so the corresponding generalised rainbow Turán number $\mathrm{ex}(n, P_1, \text{rainbow-}C_{2k})$ is $\mathrm{ex}^*(n, C_{2k})$. Very recently, this was shown to be $\Theta(n^{1+1/k})$ by O. Janzer [97]. We mention that we believe that the most difficult (new) results in this chapter are Theorem 6.1.4 and the closely related $s = t = 2k$ case of Theorem 6.1.2.

Theorem 6.1.2 deals with all cases except when $s = 3$. In that case the correct order of magnitude is unknown in general even in the non-rainbow setting, where the following bounds are known.

Theorem 6.1.5 (Győri and Li [81], Alon and Shikhelman [6], Gishboliner and Shapira [73])**.** *For every $k \geq 2$, we have*

$$\Omega(\mathrm{ex}(n, \{C_4, C_6, \ldots, C_{2k}\})) \leq \mathrm{ex}(n, C_3, C_{2k}) \leq O(\mathrm{ex}(n, C_{2k}))$$

and

$$\Omega(\mathrm{ex}(n, \{C_4, C_6, \ldots, C_{2k}\})) \leq \mathrm{ex}(n, C_3, C_{2k+1}) \leq O(\mathrm{ex}(n, C_{2k})).$$

Note that the lower and upper bounds are only known to be of the same order of magnitude when $k \in \{2, 3, 5\}$, in which case both bounds are $\Theta(n^{1+1/k})$. For the rainbow version, we have the following.

Theorem 6.1.6. *For every $k \geq 2$ integer, we have*

$$\mathrm{ex}(n, C_3, \text{rainbow-}C_{2k}) = O(n^{1+1/k}).$$

Moreover, if $k \geq 2$ is odd then $\mathrm{ex}(n, C_3, \text{rainbow-}C_{2k}) = \Omega(n^{1+1/k})$, and if k is even then $\mathrm{ex}(n, C_3, \text{rainbow-}C_{2k+1}) = \Omega(n^{1+1/k})$. Furthermore, for every $k \geq 2$ integer, we have

$$\mathrm{ex}(n, C_3, \text{rainbow-}C_{2k}) \geq \mathrm{ex}(n, C_3, C_{2k}) = \Omega(\mathrm{ex}(n, \{C_4, C_6, \ldots, C_{2k}\})),$$
$$\mathrm{ex}(n, C_3, \text{rainbow-}C_{2k+1}) \geq \mathrm{ex}(n, C_3, C_{2k+1}) = \Omega(\mathrm{ex}(n, \{C_4, C_6, \ldots, C_{2k}\})).$$

Note that $\mathrm{ex}(n, \{C_4, C_6, \ldots, C_{2k}\})$ is only known to be $\Omega(n^{1+1/k})$ when $k = 2, 3, 5$.

After the results in this chapter were published, Balogh, Delcourt, Heath and Li [13] complemented Theorem 6.1.6 by proving that $\mathrm{ex}(n, C_3, \text{rainbow-}C_{2k+1}) = O(n^{1+1/k})$ for any $k \geq 2$.

6.2 Forbidden rainbow C_{2k}

In this section we consider graphs having no rainbow C_{2k} subgraph, and prove the corresponding cases of Theorem 6.1.2, as well as Theorem 6.1.4 concerning the number of paths. We will use the following lemma of Gerbner, Mészáros, Methuku and Palmer [72]. We also include its proof below for completeness.

Lemma 6.2.1 (Gerbner, Mészáros, Methuku, Palmer [72])**.** *Let G be a properly edge-coloured graph on n vertices containing no rainbow C_{2k}. Then for every $a \in V(G)$, the number of paths axy of length 2 starting at a is $O(n)$.*

Proof. We may assume that G is bipartite, since a random bipartition is expected to preserve a quarter of all paths of length 2 starting at a. Let $X = N(a)$ and $Y = N(N(a)) \setminus \{a\}$. Observe that the number of paths axy is $e(X, Y)$, that is, the number of edges between X and Y. Using the well known fact that $\mathrm{ex}(n, T) = O(n)$ for any tree T, it suffices to show that the induced subgraph $G[X \cup Y]$ does not contain a $(100k)$-ary tree of depth $2k$.

Assume that it does contain such a tree. Then it also contains a $(100k)$-ary tree of depth $2k-1$ rooted at some $x_1 \in X$. Then we can recursively find distinct vertices $y_1, x_2, y_2, \ldots, y_{k-1}, x_k$ (with $x_i \in X, y_j \in Y$) such that for all i, $x_i y_i, y_i x_{i+1} \in E(G)$, and the $3k-2$ colours of the form $c(x_i y_i), c(y_i x_{i+1})$ or $c(a x_i)$ are all distinct. (Here c denotes the edge-colouring.) But then $a x_1 y_1 x_2 y_2 \ldots y_{k-1} x_k a$ is a rainbow cycle of length $2k$, giving a contradiction. $\square$

We now state explicitly the cases of Theorem 6.1.2 we deal with in the next two subsections.

Theorem 6.2.2. *Let $k \geq 2$ be an integer. Then*

$$\mathrm{ex}(n, C_{2k}, \text{rainbow-}C_{2k}) = \begin{cases} \Theta(n^{k-1}) & \text{if } k \geq 3 \\ \Theta(n^2) & \text{if } k = 2. \end{cases}$$

Theorem 6.2.3. *If $k, \ell \geq 2$ are integers, then*

$$\mathrm{ex}(n, C_{2\ell+1}, \text{rainbow-}C_{2k}) = \begin{cases} \Theta(n^{\ell}) & \text{if } k \geq 3 \\ \Theta(n^{\ell+1/2}) & \text{if } k = 2. \end{cases}$$

For the remainder of this section, unless otherwise stated, we will assume that $k \geq 2$ is an integer, G is a properly edge-coloured graph on n vertices with no rainbow copy of C_{2k}, and $c : E(G) \to \mathbb{Z}$ denotes the edge-colouring.

6.2.1 Paths and even cycles

In this subsection, we will prove Theorems 6.1.4 and 6.2.2. Note that for the upper bounds in Theorems 6.1.4 and 6.2.2 it suffices to consider bipartite graphs G, since a random bipartition is expected to preserve a fixed positive proportion of subgraphs isomorphic to a given bipartite graph, so from now on we assume that G is bipartite.

In light of Lemma 6.2.1, to prove the upper bound in Theorem 6.1.4 for $k \geq 3$, it is sufficient to show that the number of paths of length 3 is $O(n^2)$. Let us say that a pair x, y of vertices of G is bad if x and y have at least $100k$ common neighbours, and it is good otherwise. Then there are three types of paths $axyz$ of length 3: either ay and xz are both good, or both bad, or one of them is good and the other one is bad. We will treat these cases in three separate lemmas, as follows.

Lemma 6.2.4. *Let $k \geq 3$. For every $a \in V(G)$, the number of paths $axyz$ such that ay and xz are both bad is $O(n)$.*

Lemma 6.2.5. *Let $k \geq 3$. For every $a \in V(G)$, the number of paths $axyz$ such that ay is good and xz is bad is $O(n)$.*

Lemma 6.2.6. *Let $k \geq 3$. The number of paths $axyz$ such that ay and xz are both good is $O(n^2)$.*

It will be important later that for two of these cases we prove not only that the number of P_3s of that type is $O(n^2)$, but also the stronger statement that for any vertex a the number of paths $axyz$ with xz bad is $O(n)$. However, it is not true that for any vertex a the number of paths $axyz$ of length 3 starting at a has to be $O(n)$. To see this, take a C_{2k}-free bipartite graph G_0 on vertex classes X, Y with $|X| = |Y| = n/4$ and $|E(G_0)| = \omega(n)$. For each $x \in X$ add a new vertex x', and join each pair xx' by an edge of the same colour. Finally, add a vertex a and join it to all vertices x'. Then the (bipartite) graph we get contains no rainbow C_{2k}, and the number of paths of length 3 starting at a is $|E(G_0)|$.

Deducing Theorem 6.1.4. For $k \geq 3$, Lemma 6.2.1 shows that there are $O(n^2)$ copies of P_2, and Lemmas 6.2.4, 6.2.5 and 6.2.6 show that there are $O(n^2)$ copies of P_3. The required upper bound then follows by repeated application of Lemma 6.2.1. For the lower bound, take an $(\ell+1)$-partite graph with vertex classes $X_1, \dots, X_{\ell+1}$ such that $|X_i| = 1$ if i is even and $|X_i| = \Theta(n)$ if i is odd, and join vertices x and y if and only if $x \in X_i$ and $y \in X_j$ with $i - j = \pm 1$. (The edge-colouring is arbitrary.)

When $k = 2$, the number of paths of length 2 is $O(n^2)$ by Lemma 6.2.1, and the number of paths of length 1 is at most $\mathrm{ex}^*(n, C_4) = \Theta(n^{3/2})$ (see [107]). The required upper bound then follows by repeated application of Lemma 6.2.1. For the lower bound, we can take a C_4-free d-regular graph on $\Theta(n)$ vertices with $d = \Theta(n^{1/2})$. (It is well known that such graphs exist. For example, such a bipartite graph can be constructed by taking a projective plane P with $\Theta(n)$ points, adding a vertex for each point and line in P, and adding an edge between the vertex corresponding to a point p and a vertex corresponding to a line L if L is incident with p.) $\square$

We now prove Lemmas 6.2.4, 6.2.5 and 6.2.6.

Proof of Lemma 6.2.4. Let $Y = \{y \in N(N(a)) \setminus \{a\} : ay \text{ is bad}\}$, and let $Z = N(Y)$. Observe that $G[Y \cup Z]$ cannot contain a rainbow path of length $2k-3$. Indeed, if there is such a rainbow path, then there is a rainbow path $y_1 z_1 \dots y_{k-2} z_{k-2} y_{k-1}$ of length $2k-4$ with $y_i \in Y, z_j \in Z$. Since ay_1 and ay_{k-1} are bad, we can choose $b \in N(a) \cap N(y_1)$ and $b' \in N(a) \cap N(y_{k-1})$ such that $aby_1 z_1 \dots y_{k-2} z_{k-2} y_{k-1} b' a$ is a rainbow $2k$-cycle, giving a contradiction. It follows that $e(Y, Z) = O(n)$, i.e., $\sum_{y \in Y} \deg_G(y) = O(n)$. (We are using the fact that for any ℓ we have $\mathrm{ex}^*(n, P_\ell) = O(n)$. See [57] for the best known upper bound.)

For each $y \in Y$, define an auxiliary graph H_y on vertex set $N(y)$ by letting zz' be an edge if and only if zz' is bad. Note that H_y cannot contain a path of length $k-1$. Indeed, if $z_1 \dots z_k$ is such a path, then we can choose $b_i \in N_G(z_i) \cap N_G(z_{i+1})$ in such a way that $yz_1 b_1 z_2 \dots b_{k-1} z_k y$ is a rainbow $2k$-cycle in G, giving a contradiction. It follows that $|E(H_y)| \leq k|H_y| = k \deg_G(y)$. But the number of triples (x, y, z) such that xyz is a path, xz is bad and $y \in Y$ is $2 \sum_{y \in Y} |E(H_y)| \leq 2k \sum_{y \in Y} \deg_G(y) = O(n)$. The statement of the lemma follows. $\square$

Proof of Lemma 6.2.5. Let $Y = \{y \in N(N(a)) \setminus \{a\} : ay \text{ is good}\}$, and let

$$Z = \{z \in V(G) : \text{for any set } S \subseteq V(G) \text{ with } |S| \leq 100k \text{ there is a path } axyz \text{ of length 3 such that } x \notin S,\ ay \text{ is good and } xz \text{ is bad}\}.$$

Consider first the number of paths $axyz$ with $z \in Z$ such that ay is good (and xz is bad). The number of these is at most $100k \cdot e(Y, Z)$, as after picking yz there are at most $100k$ possible choices for x.

Claim. $G[Y \cup Z]$ cannot contain a rainbow path of length $2k-5$.

Proof of Claim. Suppose it contains such a rainbow path. Then it also contains a rainbow path P: $z_1 y_1 \dots z_{k-3} y_{k-3} z_{k-2}$ of length $2k-6$ such that $z_i \in Z, y_j \in Y$. Let

$$S_1 = V(P) \cup \{x \in N(a) : c(ax) = c(z_i y_i) \text{ or } c(ax) = c(y_i z_{i+1}) \text{ for some } i\}.$$

Then $|S_1| < 100k$, so we can pick a P_3 $axyz_1$ from a to z_1 such that $x \notin S_1$, ay is good and xz_1 is bad. Let $S_2 = S_1 \cup \{x\}$ and pick a path $ax'y'z_{k-2}$ such that $x' \notin S_2$, ay' is good and $x'z_{k-2}$ is bad. Then we can pick $y'' \in N(x) \cap N(z_1)$ such that $c(xy'')$ and $c(y''z_1)$ are distinct from all $c(z_i y_i), c(y_i z_{i+1}), c(ax), c(ax')$, and y'' is distinct from a and each y_i. Similarly, we can pick y''' such that $c(xy''')$ and $c(y'''z_1)$ are distinct from all $c(z_i y_i), c(y_i z_{i+1}), c(ax), c(ax'), c(xy''), c(y''z_1)$, and y''' is distinct from a, y'' and each y_i. Then $axy''z_1 y_1 z_2 \dots y_{k-3} z_{k-2} y''' x' a$ is a rainbow C_{2k}, giving a contradiction. The claim follows. □

So $G[Y \cup Z]$ contains no rainbow P_{2k-5}, so $e(Y, Z) = O(n)$. So there are $O(n)$ $P_3 s$ $axyz$ with $z \in Z$ such that ay is good (and xz is bad).

Now consider the number of $P_3 s$ $axyz$ with $z \notin Z$ such that ay is good and xz is bad. Given $z \notin Z$, there is a set S with $|S| \le 100k$ such that any P_3 $axyz$ such that ay is good and xz is bad must have $x \in S$. So for each $z \in Z$ we can pick $x_z \in N(a)$ such that at least a proportion of $1/(100k)$ of all such P_3s from a to z go through x_z. For each $x \in N(a)$ let $Z_x = \{z \notin Z : x_z = x\}$. Also let $Y_x = Y \cap N(x)$. Then the number of such P_3s starting at a and ending outside Z is at most

$$\begin{aligned}\sum_{z \notin Z} 100k \cdot |N(x_z) \cap N(z) \cap Y| &= \sum_{z \notin Z} 100k \cdot e(Y_{x_z}, \{z\}) \\ &= \sum_{x \in N(a)} 100k \cdot e(Y_x, Z_x).\end{aligned}$$

Note that $e(Y_x, Z_x)$ is the number of paths of length 2 starting at x in the graph $G[\{x\} \cup Y_x \cup Z_x]$. Since that graph contains no rainbow C_{2k}, Lemma 6.2.1 gives that $e(Y_x, Z_x) = O(|Y_x| + |Z_x| + 1)$. Note, however, that

$$\sum_{x \in N(a)} |Y_x| = \sum_{y \in Y} |N(y) \cap N(a)| \le 100k|Y| = O(n)$$

and

$$\sum_{x \in N(a)} |Z_x| = \sum_{z \notin Z} 1 = O(n).$$

Putting these bounds together, we get that the number of such P_3s starting at a and ending outside Z is $O(n)$. The statement of the lemma follows. $\square$

Some parts of the next proof will be similar to the proof of the fact $\mathrm{ex}^*(n, C_6) = O(n^{4/3})$ in [107].

Proof of Lemma 6.2.6. We start similarly as in the proof of Lemma 6.2.5. Let

$$W = \{(a,z) \in V(G) \times V(G) : \text{ for any set } S \text{ of at most } (100k)^2 \text{ colours there is a rainbow path } axyz \text{ such that } c(ax), c(xy), c(yz) \notin S \text{ and } ay, xz \text{ are good.}\}$$

Given $a \in V(G)$, let $Z_a = \{z : (a,z) \in W\}$, and let $Y_a = \{y \in N(N(a)) \setminus \{a\} : \ ay \text{ is good}\}$.

Claim. $G[Y_a \cup Z_a]$ contains no rainbow path of length $2k-5$.

Proof of Claim. Suppose it does. Then it also contains a rainbow path P of length $2k-6$: $z_1y_1 \dots z_{k-3}y_{k-3}z_{k-2}$, with $z_i \in Z_a, y_j \in Y_a$. Let S_1 be the set consisting of colours coming from the following sets.

- The colours appearing on the path P.
- The colours appearing on a path of length 2 of the form axy_i for some i and some vertex $x \in N(a) \cap N(y_i)$.
- The colours of edges of the form az_i for some i with $z_i \in N(a)$.

Note that $|S_1| \le 2k + k \cdot 100k \cdot 2 + k < (100k)^2$, so we can pick a rainbow path $axyz_1$ such that ay, xz_1 are good and $c(ax), c(xy), c(yz_1) \notin S_1$. Note that $y \neq y_i$ for all i and $x \neq z_j$ for all j. Let S_2 be the set

$$S_1 \cup \{c(ax), c(xy), c(yz_1)\} \cup \{c(aw) : w \in N(a) \cap N(y)\} \cup \{c(wy) : w \in N(a) \cap N(y)\}.$$

We have $|S_2| < (100k)^2$, so we can pick a rainbow path $ax'y'z_{k-2}$ such that $ay', x'z_{k-2}$ are good and $c(ax'), c(x'y'), c(x'z_{k-2}) \notin S_2$. Note that $y' \neq y_i, y$ and $x' \neq z_j, x$. But then the cycle $axyz_1y_1 \dots z_{k-3}y_{k-3}z_{k-2}y'x'a$ is a rainbow C_{2k}, giving a contradiction. The claim follows. $\square$

By the Claim, we have $e(Y_a, Z_a) = O(n)$ for all a. Hence the number of paths $axyz$ such that ay and xz are good and $(a,z) \in W$ is $O(n^2)$ (since for any a, each edge yz extends to at most $100k$ such paths $axyz$).

Now consider P_3s $axyz$ with $(a,z) \notin W$. For any a and z, let $f(a,z)$ denote the number of rainbow P_3s $axyz$ from a to z such that ay and xz are both good. If $(a,z) \notin W$, we can pick a colour c_{az} such that there are at least $\lceil f(a,z)/(100k)^2 \rceil$ P_3s $axyz$ such that ay, xz are good and $c_{az} \in \{c(ax), c(xy), c(yz)\}$. Note that at most $100k$ of these P_3s have $c(ax) = c_{az}$, since the colouring is proper and xz is good. Similarly, at most $100k$ of these P_3s have $c(yz) = c_{az}$. We deduce that there are at least $N_{az} = \lceil f(a,z)/(100k)^2 \rceil - 200k$ P_3s $axyz$ such that $c(xy) = c_{az}$

and ay, xz are good. Note that these paths must be internally vertex-disjoint. So we can list N_{az} such paths as ax_iy_iz for $i = 1, 2, \ldots, N_{az}$ such that if $i \neq j$ then $x_i \neq x_j$ and $y_i \neq y_j$.

Using the observations above, we now show that there are 'many' 6-cycles of the form $ax_iy_izy_jx_ja$ such that $c(x_iy_i) = c(x_jy_j) = c_{az}$ and each pair (of distance 2) in the 6-cycle is good. (Note that if we did not require that x_ix_j and y_iy_j are good then we would immediately get at least $\binom{N_{az}}{2}$ such 6-cycles if $N_{az} > 0$). Write $N = N_{az}$. Define an auxiliary graph H on vertex set $\{x_1, \ldots, x_N\}$ such that x_ix_j is an edge if and only if x_ix_j is bad. Observe that H contains no path of length $k-1$. Indeed, if $x_{i_1}x_{i_2}\ldots x_{i_k}$ is such a path in H, then we can choose some vertices $b_1, \ldots, b_{k-1}$ in G such that $ax_{i_1}b_1x_{i_2}b_2\ldots x_{i_{k-1}}b_{k-1}x_{i_k}a$ is a rainbow cycle of length $2k$, giving a contradiction. It follows that $|E(H)| \leq kN$. So there are at most kN pairs $\{i, j\}$ such that x_ix_j is bad. Similarly, there are at most kN pairs $\{i, j\}$ such that y_iy_j is bad. It follows that if $N \geq 1$ then there are at least $\binom{N}{2} - 2kN$ 6-cycles $ax_iy_izy_jx_ja$ in which each pair of vertices of distance 2 is good.

Write $T = \{(a, z) \notin W : f(a, z) > (100k)^2 \cdot 200k\}$. By the argument above, the number of 6-cycles $axyzy'x'a$ in which $c(xy) = c(x'y')$ and each pair of vertices of distance 2 is good is at least

$$\frac{1}{6} \sum_{(a,z) \in T} \left[\binom{N_{az}}{2} - 2kN_{az} \right],$$

which is at least

$$\sum_{(a,z) \in T} (\alpha f(a, z)^2 - \beta f(a, z))$$

for some positive constants α, β.

On the other hand, if L denotes the number of paths $axyz$ in which ay, xz are both good, then the number of such 6-cycles is at most $100kL$. Indeed, there are L ways to choose $xyzy'$, then x' is uniquely determined by the condition $c(xy) = c(x'y')$, and then there are at most $100k$ possible choices for a, since we need xx' to be good. Hence

$$\sum_{(a,z) \in T} (\alpha f(a, z)^2 - \beta f(a, z)) \leq 100kL.$$

But we have

$$L \leq \sum_{(a,z) \in T} f(a, z) + O(n^2). \tag{6.1}$$

Indeed, we know that the number of P_3s $axyz$ (such that ay and xz are good) having $(a, z) \in W$ is $O(n^2)$, the number of such rainbow P_3s $axyz$ with $(a, z) \in T$ is $\sum_{(a,z) \in T} f(a, z)$, the number of such rainbow P_3s $axyz$ with $(a, z) \notin T, (a, z) \notin W$ is at most $((100k)^2 \cdot 200k)n^2$, and finally, the number of such non-rainbow P_3s is at most the number of P_2s xyz with xz good, which is

$O(n^2)$. It follows that

$$\sum_{(a,z)\in T} (\alpha f(a,z)^2 - \beta f(a,z)) \leq 100k \sum_{(a,z)\in T} f(a,z) + O(n^2),$$

and hence

$$\sum_{(a,z)\in T} f(a,z)^2 \leq A \sum_{(a,z)\in T} f(a,z) + Bn^2$$

for some positive constants $A, B > 0$. But we have

$$\sum_{(a,z)\in T} f(a,z)^2 \geq \left[\sum_{(a,z)\in T} f(a,z)\right]^2 \cdot \frac{1}{|T|} \geq \left[\sum_{(a,z)\in T} f(a,z)\right]^2 \cdot \frac{1}{n^2}.$$

We get

$$\left[\sum_{(a,z)\in T} f(a,z)\right]^2 \leq An^2 \sum_{(a,z)\in T} f(a,z) + Bn^4,$$

which gives $\sum_{(a,z)\in T} f(a,z) = O(n^2)$. The statement of the lemma then follows using (6.1). $\square$

We now prove Theorem 6.2.2. Although the upper bound is proved for $k = 2$ and the lower bound is proved for $k \geq 3$ in [72], we include proofs of these for completeness.

Proof of Theorem 6.2.2. Consider first the case $k = 2$. For the upper bound, observe that there can be no bad pair if there is no rainbow C_4, thus any two vertices x and z are contained in $O(1)$ 4-cycles of the form $xyzw$. The upper bound $\mathrm{ex}(n, C_4, \text{rainbow-}C_4) = O(n^2)$ follows. For the lower bound when $k = 2$, let A be a Sidon set in $\mathbb{Z}_n$ of size $\Theta(\sqrt{n})$, i.e., a set such that whenever $a, b, a', b' \in A$ with $a + b = a' + b'$ then $(a, b) = (a', b')$ or $(a, b) = (b', a')$. (See, e.g., [56] for the construction of such sets.) Partition A into two subsets A_1, A_2 of size $\Theta(\sqrt{n})$ each. Let G be a 4-partite graph with vertex classes $X_{00}, X_{01}, X_{10}, X_{11}$ each being copies of $\mathbb{Z}_n$, and edges given as follows. If $x_{00} \in X_{00}, x_{01} \in X_{01}, x_{10} \in X_{10}, x_{11} \in X_{11}$, then we join:

- x_{00} to x_{10} by an edge of colour a_1 if $x_{10} - x_{00} = a_1 \in A_1$;
- x_{00} to x_{01} by an edge of colour a_2 if $x_{01} - x_{00} = a_2 \in A_2$;
- x_{10} to x_{11} by an edge of colour a_2 if $x_{11} - x_{10} = a_2 \in A_2$;
- x_{01} to x_{11} by an edge of colour a_1 if $x_{11} - x_{01} = a_1 \in A_1$.

Clearly, the graph G we get is properly edge-coloured. We claim that the 4-cycles in G are exactly those of the form $x_{00}x_{10}x_{11}x_{01}$ with $x_{ij} \in X_{ij}$ for all i, j and

$$(x_{00}, x_{10}, x_{11}, x_{01}) = (x, x + a_1, x + a_1 + a_2, x + a_2)$$

for some $x \in \mathbb{Z}_n, a_1 \in A_1, a_2 \in A_2$. (Then there are $n|A_1||A_2| = \Theta(n^2)$ 4-cycles, none of which are rainbow, giving the result.) To see this, first consider 4-cycles of the form $x_{00}x_{10}x_{11}x_{01}$ with $x_{ij} \in X_{ij}$ for all i, j. Then, writing $a_1 = x_{10} - x_{00}, a_2 = x_{01} - x_{00}, a_1' = x_{11} - x_{01}$ and $a_2' = x_{11} - x_{10}$, we get $a_1 + a_2' = a_1' + a_2$. Since A is a Sidon set partitioned into A_1 and A_2, and $a_j, a_j' \in A_j$ for $j = 1, 2$, we get $a_1 = a_1'$ and $a_2 = a_2'$, as required. On the other hand, cycles $x_1x_2x_3x_4$ not of this form must have two vertices, say x_1 and x_3, in the same vertex class X_{ij}. Writing a_{ij} for the colour of the edge x_ix_j in the cycle, we get that $x_1 - x_3 = (x_1 - x_2) - (x_3 - x_2) = \pm(a_{12} - a_{23})$, and similarly $x_1 - x_3 = \pm(a_{41} - a_{34})$. Thus $a_{12} - a_{23} = \pm(a_{41} - a_{34})$. Using that A is a Sidon set, we see that a_{12} agrees with one of a_{23}, a_{34}, a_{41}. Since the four vertices $x_1, \ldots, x_4$ are distinct, we must have $a_{12} = a_{34}$, and similarly, $a_{23} = a_{14}$. But then x_2 and x_4 must belong to the same vertex class $X_{i'j'}$ of G, and expressing $x_1 - x_3$ in two ways again gives $a_{12} - a_{23} = a_{41} - a_{34}$, implying $a_{12} = a_{23}$, a contradiction.

Now consider the lower bound for $k \geq 3$. Take a $(2k)$-partite graph with vertex classes $X_1, \ldots, X_{2k}$, where $|X_1| = |X_2| = |X_4| = |X_5| = 1$, $|X_6| = |X_8| = |X_{10}| = \cdots = |X_{2k}| = n$, $|X_3| = n$ and $|X_7| = |X_9| = \cdots = |X_{2k-1}| = 1$. Join two vertices x and y by an edge if and only if $x \in X_i$, $y \in X_j$ with $i - j \equiv \pm 1 \bmod 2k$. Give the unique edge X_1 to X_2 and the unique edge X_4 to X_5 colour 1, and arbitrary distinct colours to the remaining edges. It is easy to see that any $2k$-cycle must contain both of the edges of colour 1, there are $\Theta(n)$ vertices and $\Theta(n^{k-1})$ copies of C_{2k}.

It remains to prove the upper bound for $k \geq 3$. Given a $2k$-cycle $x_1 \ldots x_{2k}x_1$, define its pattern to be the list of i such that x_ix_{i+2} is good (indices understood mod $2k$), together with the list of pairs (i, j) such that $c(x_ix_{i+1}) = c(x_jx_{j+1})$. Note that there are finitely many patterns, so it suffices to show that for each pattern the number of $2k$-cycles of that pattern is $O(n^{k-1})$.

Consider first the case $k \geq 4$. Assume that we have a pattern and an i such that $x_{i-1}x_{i+1}$ is good but $x_{i-3}x_{i-1}$ is bad in the pattern. Then, by Theorem 6.1.4, we can choose vertices $x_{i+1}x_{i+2} \ldots x_{i+2k-4}$ in $O(n^{k-2})$ ways, since we have to pick a path of length $2k - 5$. (Note that $x_{i+2k-4} = x_{i-4}$.) Then, by Lemmas 6.2.4 and 6.2.5, there are at most $O(n)$ ways of choosing the path $x_{i-4}x_{i-3}x_{i-2}x_{i-1}$ according to the pattern (since $x_{i-3}x_{i-1}$ has to be bad). Then there are at most $100k$ possible ways of choosing x_i, since $x_{i-1}x_{i+1}$ is good. So we get $O(n^{k-1})$ $2k$-cycles for these patterns.

So (when $k \geq 4$) it remains to consider the case when there is no i such that $x_{i-1}x_{i+1}$ is good but $x_{i-3}x_{i-1}$ is bad. Observe that for any $2k$-cycle $x_1 \ldots x_{2k}x_1$, at least one (in fact, at least two) of the pairs $x_2x_4, x_4x_6, \ldots, x_{2k}x_2$ has to be good (otherwise we can find a rainbow C_{2k}). So it remains to consider patterns such that each of these pairs is good. Similarly, we may assume that each of $x_1x_3, \ldots, x_{2k-1}x_1$ is a good pair.

Now consider the colours for the pattern. We must have a pair of different edges with the same colour. We may assume that we have $c(x_1x_2) = c(x_ix_{i+1})$ for some i with $3 \leq i \leq k+1$.

Then we can choose $x_2x_3\dots x_{2k-1}$ in $O(n^{k-1})$ ways (since it is a path of length $2k-3$). Then x_1 is uniquely determined by the condition $c(x_1x_2) = c(x_ix_{i+1})$, and then there are at most $100k$ possible choices for x_{2k} (according to the pattern), since x_1x_{2k-1} is good. This gives $O(n^{k-1})$ $2k$-cycles of this pattern, as required.

It remains to consider the case $k = 3$. Observe that if $k = 3$, then for any edge ab there is at most one way to extend this edge to a path abc such that ac is bad. Indeed, if we have two different extensions abc and abc' then there is a rainbow 6-cycle of the form $axcbc'x'a$. Consider any pattern, we show that there are $O(n^2)$ 6-cycles of that pattern. We may assume that $c(x_1x_2) = c(x_ix_{i+1})$ for some $i \in \{3,4\}$. If x_5x_1 is good in the pattern, then we are done exactly as above: we can choose $x_2x_3x_4x_5$ in $O(n^2)$ ways, then x_1 is determined by the condition $c(x_1x_2) = c(x_ix_{i+1})$, and there are at most $100k$ choices for x_6. So we may assume that x_5x_1 is bad. But there are $O(n^2)$ ways of choosing $x_3x_2x_1x_6$, and then there is at most one way of extending x_1x_6 to a path $x_1x_6x_5$ such that x_1x_5 is bad, and there is at most one way of picking x_4 such that $c(x_ix_{i+1}) = c(x_1x_2)$, since $i \in \{3,4\}$. So we get $O(n^2)$ copies of C_6, as required. □

6.2.2 Odd cycles

We now turn to the case of odd cycles. Once we have established Theorem 6.1.4, the proof of Theorem 6.2.3 is essentially the same as the proof of Gishboliner and Shapira [73] for the non-rainbow version of the problem.

Proof of Theorem 6.2.3. The lower bounds follow from the fact

$$\mathrm{ex}(n, F, \text{rainbow-}H) \geq \mathrm{ex}(n, F, H)$$

and the corresponding results for the non-rainbow problem, see [73]. (Note that the only difficult case is when $k = 2$.)

For the upper bound when $k = 2$, observe that there can be no bad pair of vertices if there is no rainbow C_4, hence the number of $(2\ell+1)$-cycles is at most $100k = 200$ times the number of paths of length $2\ell - 1$, which is $O(n^{\ell+1/2})$ by Theorem 6.1.4.

Now consider the case $k \geq 3$. Given a path $P : x_1x_2\dots x_{2\ell-1}$ of length $2\ell - 2$ in G, write $X_P = N(x_1) \setminus V(P)$ and $Y_P = N(x_{2\ell-1}) \setminus V(P)$. Then the number of ways of extending path P to a cycle $x_1x_2\dots x_{2\ell+1}x_1$ is $e(X_P, Y_P)$. But this is at most the number of paths of length 2 starting at x_1 in the graph $G[\{x_1\} \cup X_P \cup Y_P]$, which is $O(1 + |X_P| + |Y_P|)$ by Lemma 6.2.1. It follows that P extends to at most $O(1 + |X_P| + |Y_P|)$ cycles of length $2\ell + 1$. But $|X_P|$ is the number of ways of extending P to a path $x_0x_1x_2\dots x_{2\ell-1}$, and similarly, $|Y_P|$ is the number of ways of extending P to a path $x_1\dots x_{2\ell}$. It follows that if the number of paths of length s is p_s, then $\sum_P |X_P| = O(p_{2\ell-1})$, and similarly for Y_P. Hence the number of cycles of length $2\ell + 1$ is at most

$$\sum_P e(X_P, Y_P) = O\left(\sum_P (1 + |X_P| + |Y_P|)\right) = O(p_{2\ell-2}) + O(p_{2\ell-1}) + O(p_{2\ell-1}),$$

which is $O(n^\ell)$ by Theorem 6.1.4. □

6.3 Forbidden rainbow C_{2k+1}

In this section we prove the following result, which is the only non-trivial case of Theorem 6.1.2 with t odd.

Theorem 6.3.1. *If $k \geq \ell \geq 2$ are integers, then*

$$\mathrm{ex}(n, C_{2\ell+1}, \text{rainbow-}C_{2k+1}) = \Theta(n^{2\ell-1}).$$

From now on, unless otherwise stated, we will assume that G is a properly edge-coloured graph of order n, and c denotes the edge-colouring. Also, we will say (as before) that a pair x, y of vertices is bad if $|N(x) \cap N(y)| \geq 100k$, and good otherwise.

We will deduce Theorem 6.3.1 from the following two lemmas.

Lemma 6.3.2. *Let G be any properly edge-coloured graph, and let $\ell \geq 2$ be an integer. Then the number of non-rainbow copies of $C_{2\ell+1}$ in G is*

$$O(n^{2\ell-1}) + O(\text{number of rainbow } C_{2\ell+1}\text{s in } G).$$

Lemma 6.3.3. *Let $k \geq \ell \geq 2$ be integers and let G be a properly edge-coloured graph with no rainbow C_{2k+1}. Assume that every edge of G is contained in a rainbow $C_{2\ell+1}$. Then for every $a \in V(G)$ the number of paths axy of length 2 starting at a in G is $O(n)$.*

Deducing Theorem 6.3.1. For the lower bound, we take the following $(2\ell+1)$-partite graph. Take vertex classes $X_1, \dots, X_{2\ell+1}$ all having size n. Add a perfect matching between X_1 and X_2 and a perfect matching between X_3 and X_4, and colour all these edges with colour 1. Furthermore, for all $i \neq 1, 3$, join each pair of vertices $x \in X_i$ and $y \in X_{i+1}$ by an edge of arbitrary unused colour (with indices understood mod $2\ell+1$). It is clear that the graph we get is properly edge-coloured, there are $\Theta(n)$ vertices and $\Theta(n^{2\ell-1})$ copies of $C_{2\ell+1}$. Furthermore, no copy of C_{2k+1} is rainbow, since any C_{2k+1} must contain an edge between each pair of X_i, X_{i+1} (otherwise it would be a subgraph of a bipartite graph). The lower bound follows.

Now consider the upper bound. By Lemma 6.3.2, it suffices to show that if G contains no rainbow C_{2k+1} then the number of rainbow $C_{2\ell+1}$s is $O(n^{2\ell-1})$. For this, we may assume that any edge is contained in a rainbow copy of $C_{2\ell+1}$. But then, by Lemma 6.3.3, for any vertex

$a \in V(G)$ there are $O(n)$ paths of length 2 starting at a. By repeated application of this fact, it follows that for any a there are $O(n^{\ell})$ paths of length 2ℓ starting at a, and hence there are $O(n^{\ell+1}) \leq O(n^{2\ell-1})$ copies of $C_{2\ell+1}$. $\square$

Proof of Lemma 6.3.2. We will consider patterns of $(2\ell+1)$-cycles. Recall that the pattern $\mathcal{P}$ of a $(2\ell+1)$-cycle $x_1 \dots x_{2\ell+1}x_1$ is the list of i such that $x_i x_{i+2}$ is good, together with the list of pairs (i,j) such that $c(x_i x_{i+1}) = c(x_j x_{j+1})$ (with the indices understood mod $2\ell+1$). Since there are finitely many patterns, it suffices to show that for any non-rainbow pattern the required bound holds for cycles of that pattern.

Consider first the case when there are three edges with the same colour in a pattern $\mathcal{P}$, say $x_p x_{p+1}$, $x_q x_{q+1}$, $x_r x_{r+1}$. Then we can pick $(x_i)_{i \neq p,q}$ in $O(n^{2\ell-1})$ ways, and there is at most one way of extending those points to a $(2\ell+1)$-cycle of the appropriate pattern. This shows that there are $O(n^{2\ell-1})$ cycles with this pattern.

Now consider the case when there are two different colours such that each of them appears at least twice as the colour of an edge. For both of these colours, pick two edges of the appropriate colour. So we have $c(e) = c(e')$ and $c(f) = c(f')$ in our pattern for four different edges e, e', f, f'. Note that we must have $e \cup e' \neq f \cup f'$. So we can pick i, j such that $x_i \in (e \cup e') \setminus (f \cup f')$ and $x_j \in (f \cup f') \setminus (e \cup e')$. Then picking the vertices $(x_a)_{a \neq i,j}$ determines the $(2\ell+1)$-cycle uniquely by the colour conditions. It follows that there are $O(n^{2\ell-1})$ cycles of this pattern.

It remains to consider patterns $\mathcal{P}$ in which there is only one pair of edges of the same colour, say $c(x_i x_{i+1}) = c(x_j x_{j+1})$, with $i \neq j-1, j, j+1$. Given a choice $X = \{x_a : a \neq i, j\}$ of all vertices except x_i, x_j, consider the number of ways of extending X to a $(2\ell+1)$-cycle. Write $d_1 = |N(x_{i-1}) \cap N(x_{i+1}) \setminus X|$ and $d_2 = |N(x_{j-1}) \cap N(x_{j+1}) \setminus X|$. Then the number of ways of extending X to a $(2\ell+1)$-cycle of pattern $\mathcal{P}$ is at most $\min\{d_1, d_2\}$, whereas the number of ways of extending X to a rainbow $C_{2\ell+1}$ is at least $(d_1 - 5\ell)(d_2 - 5\ell)$. But we have $\min\{d_1, d_2\} \leq 10\ell + \max\{0, (d_1 - 5\ell)(d_2 - 5\ell)\}$, so the number of extensions of pattern $\mathcal{P}$ is at most $O(1)$ plus the number of rainbow extensions. Summing over all possible choices of X, we get the required bound. $\square$

Lemma 6.3.3 is proved similarly to Lemma 6.2.1.

Proof of Lemma 6.3.3. Given a bipartition $V(G) = X \cup Y$ of the vertex set of G, let $G_{X,Y}$ be the corresponding bipartite graph obtained from G (i.e., $G_{X,Y}$ is obtained by deleting all edges inside X and inside Y). Since a random bipartition is expected to preserve a quarter of all paths of length 2 starting at a, it suffices to show that for every bipartition $V(G) = X \cup Y$ with $a \in Y$, the number of paths of length 2 starting at a in $G_{X,Y}$ is $O(n)$, where the implied constant is independent of the bipartition. So let $V(G) = X \cup Y$ be any bipartition. Write $X_1 = N_G(a) \cap X$ and $Y_1 = N_G(X_1) \cap Y \setminus \{a\}$, so that we would like to show $e_{G_{X,Y}}(X_1, Y_1) = O(n)$. It suffices to show that $G_{X,Y}[X_1 \cup Y_1]$ does not contain a $(100k)$-ary tree of depth $2k$.

Suppose it contains such a tree, then it also contains a $(100k)$-ary tree T of depth $2k-1$ rooted at some $x \in X_1$. Since $ax \in E(G)$, the edge ax of G is contained in a rainbow cycle of length $2\ell+1$ in G. Hence we can find a rainbow path $P : az_1z_2 \dots z_{2\ell-1}x$ of length 2ℓ from a to x in G. Then we can recursively find distinct vertices $x = x_1, x_2, \dots, x_{2(k-\ell)+1}$ on our tree T such that

- for all i we have $x_ix_{i+1} \in E(G_{X,Y})$;
- for all i even we have $x_i \in Y_1 \setminus V(P)$;
- for all $i \geq 3$ odd we have $x_i \in X_1 \setminus V(P)$;
- for all i, $c(x_ix_{i+1})$ does not appear on the path $az_1z_2 \dots z_{2\ell-1}x_1 \dots x_i$;
- the colour $c(ax_{2(k-\ell)+1})$ is distinct from all the colours appearing on the path given by $az_1z_2 \dots z_{2\ell-1}x_1 \dots x_{2(k-\ell)}$.

But then $az_1z_2 \dots z_{2\ell-1}x_1x_2 \dots x_{2(k-\ell)+1}a$ is a rainbow cycle of length $2k+1$ in G, giving a contradiction. □

6.4 Deducing Theorem 6.1.2 and Theorem 6.1.6

We now summarise how we deduce each case in Theorem 6.1.2.

Proof of Theorem 6.1.2. We have the following cases.

- If $s = t = 4$, then the result follows from Theorem 6.2.2. If $t = 4, s \neq 4$ and s is even, then it follows from Theorem 6.1.1. If $t = 4$ and s is odd, it follows from Theorem 6.2.3.
- If s, t are even with $s \neq t$, then the result follows from Theorem 6.1.1.
- If $s = t \geq 6$ is even, then the result follows from Theorem 6.2.2.
- If $t \geq 6$ is even and s is odd, then the result follows from Theorem 6.2.3.
- If s, t are odd with $s \leq t$, then the result follows from Theorem 6.3.1.
- If t is odd, and s is even or $s > t$, then the upper bound is trivial, and for the lower bound we can take a blowup of C_s (i.e., we replace each vertex of C_s by n vertices and each edge by a complete bipartite graph. The edge-colouring is arbitrary.)

□

Finally, we prove Theorem 6.1.6 concerning triangles.

Proof of Theorem 6.1.6. We first show $\mathrm{ex}(n, C_3, \text{rainbow-}C_{2k}) = O(\mathrm{ex}^*(n, C_{2k}))$. (Using the result $\mathrm{ex}^*(n, C_{2k}) = O(n^{1+1/k})$ of O. Janzer [97], this implies the required upper bound.) Let G be a properly edge-coloured graph on n vertices with no rainbow C_{2k}, and define good and bad pairs as before. Observe that the number of triangles containing a good pair is at most $100k|E(G)|$, since we can pick the good pair in at most $|E(G)|$ ways. So it suffices to show that the number of paths xyz with xz bad is $O(|E(G)|)$. But for any $y \in V(G)$, if we define an auxiliary graph H_y with vertex set $N(y)$ and edges being the bad pairs, then there can be no path $x_1 \dots x_k$ of length $k-1$ in H_y (otherwise we can find a rainbow cycle $yx_1b_1x_2b_2 \dots x_ky$). It follows that H_y has at most $k|V(H_y)| = k \deg_G(y)$ edges, so each y is contained in at most $k \deg(y)$ paths xyz with xz bad. But $\sum_y \deg(y) = 2|E(G)|$, giving the required bound.

For the lower bound, the two inequalities

$$\mathrm{ex}(n, C_3, \text{rainbow-}C_{2k}) \geq \mathrm{ex}(n, C_3, C_{2k}),$$
$$\mathrm{ex}(n, C_3, \text{rainbow-}C_{2k+1}) \geq \mathrm{ex}(n, C_3, C_{2k+1})$$

are clear, and the lower bounds

$$\mathrm{ex}(n, C_3, C_{2k}) = \Omega\left(\mathrm{ex}(n, \{C_4, C_6, \dots, C_{2k}\})\right),$$
$$\mathrm{ex}(n, C_3, C_{2k+1}) = \Omega\left(\mathrm{ex}(n, \{C_4, C_6, \dots, C_{2k}\})\right)$$

follow from Theorem 6.1.5.

Finally, we prove that we have $\mathrm{ex}(n, C_3, \text{rainbow-}C_{2k}) = \Omega(n^{1+1/k})$ when k is odd and $\mathrm{ex}(n, C_3, \text{rainbow-}C_{2k+1}) = \Omega(n^{1+1/k})$ when k is even. Take a B_k-set A of size $\Theta(n^{1/k})$ in $\mathbb{Z}_n$, that is, a set such that any $m \in \mathbb{Z}_n$ can be written as $a_1 + \cdots + a_k$ with $a_i \in A$ in at most one way (ignoring permutations of the summands). (See [32] for the construction of such 'dense' B_k-sets.) Then we take a tripartite graph G with vertex classes X_1, X_2, Y all being copies of $\mathbb{Z}_n$ and edges given as follows. For each $x \in \mathbb{Z}_n$, we join the copy of x in X_1 to the copy of x in X_2 by an edge of colour 0, and we join the copy $x \in X_i$ to $x + a \in Y$ by an edge of colour (a, i) for each $a \in A$ and $i = 1, 2$. Clearly, G has $\Theta(n^{1+1/k})$ triangles. We claim that this graph contains no rainbow C_{2k} if k is odd and no rainbow C_{2k+1} if k is even. Indeed, assume that k is odd an there is a rainbow C_{2k}. Then it must be of the form $x_1y_1x_2y_2 \dots x_ky_kx_1$ with $y_j \in Y$ and $x_i \in X_1 \cup X_2$. Then we get a representation $0 = a_1 - b_1 + a_2 - b_2 + \cdots + a_k - b_k$ with $a_i, b_j \in A$ by letting $a_i = y_i - x_i$, $b_i = y_i - x_{i+1}$ (where $x_{k+1} = x_1$). So the a_i must be a permutation of the b_j. But k is odd, so we have $|\{x_1, \dots, x_k\} \cap X_1| \neq |\{x_1, \dots, x_k\} \cap X_2|$, and hence there exist i and j such that $a_i = b_j$ and x_i, x_{j+1} are in the same vertex class X_ℓ. But then $c(x_iy_i) = c(y_jx_{j+1})$, so the cycle is not rainbow, giving a contradiction. The case when k is even and G contains a rainbow $(2k+1)$-cycle is similar. □

Chapter 7

Local rainbow colourings

7.1 Introduction

7.1.1 Local rainbow colourings

Estimating the minimum possible size of a program that computes specific Boolean functions is a major research area in Theoretical Computer Science. In 1993, Karchmer [103] introduced the so-called fusion method for finding circuit lower bounds. This technique unifies and generalizes the topological method of Sipser [131] and the approximation method of Razborov [124]. Karchmer and Wigderson [104, 138] demonstrated that proving lower bounds for circuit sizes can be reduced to extremal combinatorics problems, and Wigderson [138] presented three problems that arise this way. One of them is as follows.

Problem 7.1.1 (Karchmer and Wigderson [138])**.** Given a positive integer n, estimate the smallest k for which the following is true. There exist colourings $c_1, \dots, c_n$ of the n-dimensional cube $\{0,1\}^n$ with k colours such that for any three distinct $x, y, z \in \{0,1\}^n$ there is a coordinate $i \in [n]$ such that x_i, y_i and z_i are not all equal and the three colours $c_i(x)$, $c_i(y)$ and $c_i(z)$ are pairwise distinct.

Karchmer and Wigderson [104] proved that k has to grow with n; more precisely that k needs to be at least $\Omega(\frac{\log\log^* n}{\log\log\log^* n})$, where $\log^* n$ is the smallest integer m such that applying the function $\log_2(x)$ iteratively m times, starting with input n, one obtains a number not exceeding 1.

Alon and Ben-Eliezer [3] improved this significantly by showing that k needs to be at least $\Omega((\frac{\log n}{\log\log n})^{1/4})$. As part of their approach, they introduced the following problem, which will be our main focus in this chapter.

Definition 7.1.2. Let n be a positive integer and let H be a graph. For each vertex v of a given clique K_n, let f_v be a (not necessarily proper) colouring of the edges of the same K_n. We say that the collection of these n colourings is (n, H)-local if for any copy T of H in K_n, there exists some $u \in V(T)$ such that all edges of T receive different colours in f_u.

Problem 7.1.3 (Alon and Ben-Eliezer [3])**.** Let $g(n, H)$ be the smallest k for which there is a collection of colourings, each using k colours, which is (n, H)-local. Estimate the growth of $g(n, H)$ as $n \to \infty$.

To see the connection to Problem 7.1.1, note that $g(n, P_3)$ is a lower bound for the smallest possible k in the Karchmer–Wigderson problem. (Here and below, P_ℓ denotes the path with ℓ edges.) Indeed, we can think of the n coordinates of $\{0,1\}^n$ as the n vertices of K_n and the elements of Hamming weight 2 in $\{0,1\}^n$ as edges in K_n. A valid collection of colourings in Problem 7.1.1 is then necessarily an (n, P_3)-local colouring, since we may choose x, y, z to be three sets in $\{0,1\}^n$ which correspond to the three edges of some P_3. Alon and Ben-Eliezer showed that $g(n, P_3) = \Omega((\frac{\log n}{\log \log n})^{1/4})$, implying the same lower bound for the problem of Karchmer and Wigderson.

Alon and Ben-Eliezer also studied Problem 7.1.3 for general graphs H. They characterized the family of graphs for which $g(n, H)$ is bounded.

Theorem 7.1.4 (Alon and Ben-Eliezer [3])**.** *For a fixed graph H, there is a constant $c(H)$ such that $g(n, H) \leq c(H)$ for every n if and only if H contains at most 3 edges and H is neither P_3 nor P_3 together with any number of isolated vertices. Moreover, in all these cases $g(n, H) \leq 5$ for every n.*

They used the local lemma to obtain the following general upper bound.

Theorem 7.1.5 (Alon and Ben-Eliezer [3])**.** *Let H be a fixed graph with r vertices. Then $g(n, H) = O(r^4 n^{1-\frac{2}{r}})$.*

They also proved polynomial lower bounds for various small graphs and used this to obtain the following result.

Theorem 7.1.6 (Alon and Ben-Eliezer [3])**.** *For any graph H with at least 13 edges, there is a constant $b = b(H) > 0$ such that $g(n, H) = \Omega(n^b)$.*

They posed three concrete open problems in their paper.

1. Improve the bounds for $g(n, P_3)$.

2. Is it true that if $g(n, H)$ is unbounded, then it grows polynomially?

3. Is it true that for every $\varepsilon > 0$ there is some $r = r(\varepsilon)$ such that $g(n, K_r) \geq n^{1-\varepsilon}$ for every sufficiently large n?

The first problem is well motivated by its connection to Problem 7.1.1. The second one is motivated by Theorem 7.1.6. The third one is motivated by Theorem 7.1.5 and the fact that if H' is a subgraph of H on the same set of vertices, then $g(n, H') \leq g(n, H)$.

Some progress on these questions was made by Cheng and Xu [37]. They showed that $g(n, P_\ell)$ is polynomial in n for every $\ell \geq 4$. Combined with other new bounds for small graphs, they used this to prove that if H is a graph with at least 6 edges, then $g(n, H)$ is polynomial, improving Theorem 7.1.6. Finally, they showed that $g(n, K_r) = \Omega(n^{2/3})$ holds for all $r \geq 8$, which can be seen as progress towards answering the third question above.

In this chapter, we answer the second and third question of Alon and Ben-Eliezer, and also show that $g(n, P_3)$ grows subpolynomially, which essentially answers the first question as well.

Theorem 7.1.7. *We have $g(n, P_3) = n^{o(1)}$.*

Together with the lower bound $g(n, P_3) = \Omega((\frac{\log n}{\log \log n})^{1/4})$ of Alon and Ben-Eliezer, Theorem 7.1.7 answers the second question of theirs in the negative. Our next result answers their third question affirmatively.

Theorem 7.1.8. *For each $\ell \geq 2$, we have $g(n, C_{2\ell}) = \Omega(n^{1-\frac{2}{\ell+1}})$. Consequently, for any even $r \geq 4$, $g(n, K_r) = \Omega\left(n^{1-\frac{4}{r+2}}\right)$.*

Remark 7.1.9. Using a variant of the proof of Theorem 7.1.8, we can also prove that for each sufficiently large odd r, we have $g(n, K_r) = \Omega\left(n^{1-\frac{10}{r-3}}\right)$, showing that the exponent tends to 1 in this case as well.

We also obtain a near-complete characterization of the family of graphs H for which $g(n, H)$ is polynomial. The only graph H for which we cannot decide whether $g(n, H)$ is polynomial is the disjoint union of a P_3 and a P_1 (together with an arbitrary number of isolated vertices).

Theorem 7.1.10. *Let H be a graph which is not the disjoint union of P_3 and P_1 together with an arbitrary number of isolated vertices. Then there exists some $b = b(H) > 0$ such that $g(n, H) = \Omega(n^b)$ if and only if H has at least 5 edges or H has precisely 4 edges and is triangle-free.*

In particular, we obtain a full characterization of the family of connected graphs H for which $g(n, H)$ is polynomial.

Corollary 7.1.11. *Let H be a connected graph. Then there exists some $b = b(H) > 0$ such that $g(n, H) = \Omega(n^b)$ if and only if H has at least 4 edges and H is different from the triangle with a pendant edge.*

7.1.2 The Erdős–Gyárfás function

We will see (in Section 7.2) that local rainbow colourings for certain graphs H are related to the Erdős–Gyárfás function in Ramsey Theory (especially to Theorem 7.1.14 below). In this subsection we describe the Erdős–Gyárfás problem, and state a new result resolving a family of special cases of a conjecture of Conlon, Fox, Lee and Sudakov [40, 41].

Definition 7.1.12. Let $p, q, r, n \geq 2$ be positive integers with $q \leq \binom{p}{r}$. An edge-colouring of the r-uniform complete hypergraph $K_n^{(r)}$ is a (p,q)-colouring if at least q distinct colours appear among any p vertices. Let $f_r(n,p,q)$ be the smallest positive integer k such that there exists a k-colouring of the edges of $K_n^{(r)}$ forming a (p,q)-colouring.

The function $f_r(n,p,q)$ was introduced by Erdős and Shelah [53], and first studied in more detail by Erdős and Gyárfás [54] (for $r=2$). Let us first consider the graph case $r=2$. When $q=2$, then a (p,q)-colouring is simply a colouring which avoids monochromatic sets of size p, so as a special case we get the classical multicolour Ramsey problem. In particular, $f_2(n,3,2)$ (and hence $f_2(n,p,2)$) is at most logarithmic in n. On the other extreme, when $q=\binom{p}{2}$, we trivially have $f_2(n,p,\binom{p}{2}) = \binom{n}{2}$ (as long as $p \geq 4$). The function $f_r(n,p,q)$ is clearly increasing in q, and Erdős and Gyárfás [54] investigated how the behaviour of $f_2(n,p,q)$ changes as q increases from 2 to $\binom{p}{2}$. Among other results, they proved that when $p=q$, it is polynomial in n, i.e., $f_2(n,p,p) = \Omega(n^{\alpha_p})$ for some $\alpha_p > 0$. They asked if this is the smallest value of q for which $f_2(n,p,q)$ is polynomial in n, i.e., whether or not $f_2(n,p,p-1)$ is subpolynomial in n.

The first difficult case $p=4$ was settled by Mubayi [116], who gave a construction showing that $f_2(n,4,3) = n^{o(1)}$. This was first extended to $p=5$ as well by Eichhorn and Mubayi [50], and later Conlon, Fox, Lee and Sudakov [39] proved that $f_2(n,p,p-1)$ is subpolynomial for all p, fully answering this question of Erdős and Gyárfás.

Theorem 7.1.13 (Conlon, Fox, Lee and Sudakov [39])**.** *For any fixed $p \geq 4$, we have*

$$f_2(n,p,p-1) \leq e^{(\log n)^{1-1/(p-2)+o(1)}} = n^{o(1)}.$$

Consider now the Erdős–Gyárfás function for general uniformity r (this is the setting in which Erdős and Shelah [53] originally introduced the problem). Answering a question of Graham, Rothschild and Spencer [75], Conlon, Fox, Lee and Sudakov [40] proved that $f_3(n,4,3)$ is subpolynomial in n. (This has close connections to the proof of Shelah [130] of primitive recursive bounds for the Hales–Jewett theorem, see [40] for details.)

Theorem 7.1.14 (Conlon, Fox, Lee and Sudakov [40])**.** *We have*

$$f_3(n,4,3) \leq e^{(\log n)^{2/5+o(1)}} = n^{o(1)}.$$

Conlon, Fox, Lee and Sudakov [40] also proved that $f_r(n,p,\binom{p-1}{r-1}+1)$ is at least polynomial in n. In light of this result, as well as Theorem 7.1.13 and Theorem 7.1.14, they proposed the following conjecture.

Conjecture 7.1.15 (Conlon, Fox, Lee and Sudakov [40, 41])**.** *For any positive integers p and r with $2 \leq r < p$,*

$$f_r\left(n,p,\binom{p-1}{r-1}\right) = n^{o(1)}.$$

Note that the case $r = 2$ holds by Theorem 7.1.13, and Theorem 7.1.14 is the case $r = 3$, $p = 4$. As further evidence towards Conjecture 7.1.15, Conlon, Fox, Lee and Sudakov [40] proved that its statement holds for $r = 3$, $p = 5$ as well. These results (i.e., $r = 2$; or $r = 3$ and $p \in \{4, 5\}$) are the only known cases of Conjecture 7.1.15. We show that Conjecture 7.1.15 holds whenever $p = r + 1$.

Theorem 7.1.16. *For any $r \geq 3$, we have*

$$f_r(n, r + 1, r) \leq e^{(\log n)^{2/5+o(1)}} = n^{o(1)}.$$

In addition to proving a family of special cases of Conjecture 7.1.15, Theorem 7.1.16 is also directly related to the problem of Karchmer and Wigderson (Problem 7.1.1): in Subsection 7.4.2 we briefly describe how Theorem 7.1.16 implies that a natural approach to finding polynomial lower bounds for Problem 7.1.1 cannot work.

The rest of the chapter is organized as follows. In Section 7.2, we prove Theorem 7.1.7 and the "only if" part of Theorem 7.1.10, that is, a subpolynomial upper bound for $g(n, H)$ whenever H has at most 3 edges or has precisely 4 edges and contains a triangle. In Section 7.3, we prove Theorem 7.1.8 and the "if" part of Theorem 7.1.10. In Section 7.4 we prove Theorem 7.1.16. We finish the chapter by giving some brief concluding remarks in Section 7.5.

7.2 Upper bounds

In this section we give constructions providing subpolynomial upper bounds for $g(n, H)$ when H is one of the following graphs:

- P_3, the path with 3 edges;
- T_p, a triangle with a pendant edge; or
- T_e, the disjoint union of a triangle and an edge.

Note that if H' is formed from H by adding some isolated vertices, then any collection of colourings which is (n, H)-local is also (n, H')-local, hence $g(n, H') \leq g(n, H)$. So subpolynomiality for the 3 graphs above, together with Theorem 7.1.4, deals with all cases of the "only if" part of Theorem 7.1.10. Furthermore, since P_3 is a subgraph of T_p (on the same vertex set), the result that $g(n, T_p)$ is subpolynomial easily implies that $g(n, P_3)$ is subpolynomial. Nevertheless, we will first focus on the proof for P_3, as it is slightly simpler and motivates the construction for T_p (and also interesting on its own due to its connection to Problem 7.1.1).

We will pick our colourings in such a way that in each colouring f_v, the edges which contain v receive different colours from the ones that do not contain v. Note that if this property holds, then the only way a P_3 *abcd* can be non-rainbow (i.e., have some colour appearing more than once) for each colouring corresponding to its vertices if

$$f_a(bc) = f_a(cd), f_b(ab) = f_b(bc), f_c(bc) = f_c(cd) \text{ and } f_d(ab) = f_d(bc).$$

It is not difficult to prove that for any collection of colourings with $n^{o(1)}$ colours we can find a P_3 in which the first three of the four equalities above hold, i.e., the colourings f_a, f_b and f_c of the P_3 are all non-rainbow. However, we will show that we can construct a collection of colourings with $n^{o(1)}$ colours such that any P_3 *abcd* is rainbow in either f_a or f_d.

The main idea is the following. Assume that the colourings f_v are defined in such a way that any edge $\{x, y\}$ not containing v is coloured by the colour $\gamma(\{v, x, y\})$, where γ is some colouring of the edges of the complete 3-uniform hypergraph formed by our n vertices. Then the condition $f_a(bc) = f_a(cd)$ becomes $\gamma(abc) = \gamma(acd)$, and the condition $f_d(ab) = f_d(bc)$ becomes $\gamma(abd) = \gamma(bcd)$. Thus, if both of these conditions hold, then γ uses at most 2 colours on the 4 vertices a, b, c, d. However, recall from Subsection 7.1.2 the following result of Conlon, Fox, Lee and Sudakov [40] about the Erdős–Gyárfás problem.

Theorem 7.2.1 (Conlon, Fox, Lee and Sudakov [40])**.** *The edges of the complete* 3*-uniform hypergraph* $K_n^{(3)}$ *on* n *vertices can be coloured using* $e^{(\log n)^{2/5+o(1)}} = n^{o(1)}$ *colours in such a way that at least* 3 *different colours appear among the four edges spanned by any* 4 *vertices.*

We are ready to prove that $g(n, P_3)$ is subpolynomial, which follows easily from the discussion above.

Proof of Theorem 7.1.7. Using Theorem 7.2.1, we can take a colouring γ of the complete 3-uniform hypergraph formed on our n vertices such that γ uses at most $e^{(\log n)^{2/5+o(1)}}$ colours and at least 3 colours appear among any 4 vertices. We now define our collection of colourings as follows. Let z_0 be a colour not used by γ (i.e., $z_0 \notin \mathrm{Im}(\gamma)$). Define, for any vertex v and edge $e = \{x, y\}$,

$$f_v(e) = \begin{cases} z_0 & \text{if } v \in e \\ \gamma(e \cup \{v\}) & \text{if } v \notin e. \end{cases}$$

Let *abcd* be any P_3 in our K_n. As noted in the discussion above, if the edges of this P_3 do not all receive different colours in f_a, then $\gamma(abc) = \gamma(acd)$. Similarly, if the edges of the P_3 do not all receive different colours in f_d, then $\gamma(abd) = \gamma(bcd)$. But we know that at least 3 different colours appear among $\gamma(abc), \gamma(abd), \gamma(acd)$ and $\gamma(bcd)$, so $\gamma(abc) = \gamma(acd)$ and $\gamma(abd) = \gamma(bcd)$ cannot simultaneously hold. Hence our collection of colourings is (n, P_3)-local with $1 + e^{(\log n)^{2/5+o(1)}} = n^{o(1)}$ colours. □

We now turn to the proof that $g(n, T_p)$ is subpolynomial, where T_p is the triangle with a pendant edge. As noted before, this result is stronger than Theorem 7.1.7, and correspondingly the proof will be a refinement of the one above.

Theorem 7.2.2. *We have $g(n, T_p) = n^{o(1)}$.*

Proof. As before, take an edge-colouring γ of the 3-uniform complete hypergraph $K_n^{(3)}$ using $n^{o(1)}$ colours such that among any 4 vertices at least 3 colours appear. Furthermore, take an edge-colouring δ of the clique K_n using $O(\log n)$ colours such that there is no monochromatic triangle. (It is well-known that this is possible – for example, label the vertices by elements of $\{0,1\}^m$ and colour the edge xy by the minimal i for which $x_i \neq y_i$.) We may assume that γ and δ have disjoint images.

We define our collection of colourings f_v by setting

$$f_v(e) = \begin{cases} \delta(e) & \text{if } v \in e \\ \gamma(e \cup \{v\}) & \text{if } v \notin e. \end{cases}$$

Take any copy $abcd$ of T_p: the vertices b, c, d form a triangle and a is joined to b. We show that this copy must be rainbow in one of f_a, f_c or f_d.

If this copy of T_p is not rainbow under f_c, then we must have either $\delta(bc) = \delta(cd)$ or $\gamma(abc) = \gamma(bcd)$. Similarly, if the copy is not rainbow under f_d, then either $\delta(bd) = \delta(cd)$ or $\gamma(abd) = \gamma(bcd)$. Note that, by the definition of δ, we cannot have both $\delta(bc) = \delta(cd)$ and $\delta(bd) = \delta(cd)$. Thus, without loss of generality, we have $\gamma(abc) = \gamma(bcd)$.

But if our copy of T_p is not rainbow under f_a, then at least two of abc, abd, acd have the same colour in γ. Together with $\gamma(abc) = \gamma(bcd)$, this would imply that at most two colours appear among a, b, c, d in γ, giving a contradiction. So our collection of colourings is (n, T_p)-local with $n^{o(1)} + O(\log n) = n^{o(1)}$ colours. □

Finally, we prove that $g(n, T_e)$ is subpolynomial, where T_e denotes the disjoint union of a triangle and an edge. In fact, we will prove a logarithmic bound here.

Theorem 7.2.3. *We have $g(n, T_e) = O(\log n)$.*

Proof. We may assume that the vertices of our clique K_n are elements of $\{0,1\}^m$, where $m = \lceil \log_2 n \rceil$. Given vertices x and y, let $\delta(xy) = \min\{i : x_i \neq y_i\}$ be the first coordinate where x and y differ. Observe that δ has the property that for any three vertices x, y, z, exactly two of $\delta(xy), \delta(xz), \delta(yz)$ are equal. Moreover, if $\delta(xz) = \delta(yz)$ then $\delta(xy) > \delta(xz) = \delta(yz)$.

Define our collection of colourings f_v as follows. For any edge xy, let

$$f_v(xy) = \begin{cases} -\delta(xy) & \text{if } v \in \{x, y\} \\ \max\{\delta(vx), \delta(vy)\} & \text{if } v \notin \{x, y\}. \end{cases}$$

(The only purpose of the minus sign in the first case is to ensure that f_v takes different values on edges that contain v and on edges that do not.) Consider any copy of T_e formed by a triangle

abc and an edge xy; we show that this copy is rainbow under one of f_a, f_b or f_c. Without loss of generality, we may assume that $\delta(ab) > \delta(ac) = \delta(bc)$.

Note that $f_a(ab) \neq f_a(ac)$. So if this copy of T_e is non-rainbow under f_a, then we must have $f_a(xy) = f_a(bc)$, i.e., $\max\{\delta(ax), \delta(ay)\} = \max\{\delta(ab), \delta(ac)\} = \delta(ab)$. Without loss of generality, we have $\delta(ab) = \delta(ax)$. But then $\delta(bx) > \delta(ab)$ and hence $\max\{\delta(bx), \delta(by)\} > \delta(ab)$. This implies that $f_b(ac) \neq f_b(xy)$. As $f_b(ab) \neq f_b(bc)$, our copy of T_e must be rainbow under f_b, finishing the proof. $\square$

7.3 Lower bounds

7.3.1 The proof of Theorem 7.1.8

In this subsection, we prove Theorem 7.1.8. The proof uses the following lemma of O. Janzer [97, Theorem 3.1], which is a significant generalization of the Bondy–Simonovits theorem [30].

Lemma 7.3.1. *Let $\ell \geq 2$ and s be positive integers. Then there exists a constant $C = C(\ell, s)$ with the following property. Suppose that $G = (V, E)$ is a graph with N vertices and at least $CN^{1+1/\ell}$ edges. Let $\sim$ be a symmetric binary relation on V such that for every $u \in V$ and $v \in V$, v has at most s neighbours $w \in V$ which satisfy $u \sim w$. Then G contains a 2ℓ-cycle $x_1x_2 \dots x_{2\ell}$ such that $x_i \not\sim x_j$ for every $i \neq j$.*

Proof of Theorem 7.1.8. Let $C = C(\ell, 1)$ be provided by Lemma 7.3.1 and let $c = (4C)^{-\frac{\ell}{\ell+1}}$. For each vertex $v \in V(K_n)$, let f_v be an edge-colouring of $E(K_n)$ which uses $[k]$ as colours, where $k \leq cn^{1-\frac{2}{\ell+1}}$. Define an auxiliary graph G whose vertex set is $V = V(K_n) \times [k]$ and in which (u, i) and (v, j) are joined by an edge if and only if $u \neq v$, $f_u(uv) = i$ and $f_v(uv) = j$. Observe that there is a natural bijection between the edges of G and the edges of K_n, so $e(G) = e(K_n) = \binom{n}{2}$. Moreover, clearly, the number of vertices in G is $N = nk$. Hence,

$$CN^{1+1/\ell} \leq C(cn^{2-\frac{2}{\ell+1}})^{1+1/\ell} = Cc^{1+1/\ell}n^2 = n^2/4 \leq \binom{n}{2}.$$

For vertices $(u, i), (v, j) \in V(G)$, let us write $(u, i) \sim (v, j)$ if $u = v$. We claim that if $x, y \in V(G)$, then y has at most one neighbour z in G such that $x \sim z$. Indeed, let u be the first coordinate of x and let v be the first coordinate of y. Then (as $x \sim z$) the first coordinate of z must be u, and (as z is a neighbour of y in G) the second coordinate of z has to be the colour of the edge uv in the colouring f_u. Hence, by Lemma 7.3.1 applied with $s = 1$, it follows that there is a 2ℓ-cycle $x_1x_2 \dots x_{2\ell}$ in G such that the first coordinate of each x_i is different. Let $x_i = (u_i, \alpha_i)$. Then, for each i, we have $f_{u_i}(u_{i-1}u_i) = f_{u_i}(u_i, u_{i+1}) = \alpha_i$, where indices are considered mod 2ℓ. Therefore, the 2ℓ-cycle $u_1u_2 \dots u_{2\ell}$ in K_n witnesses that the collection of colourings $\{f_v : v \in V(K_n)\}$ is not $(n, C_{2\ell})$-local. Hence, any $(n, C_{2\ell})$-local colouring must use more than $cn^{1-\frac{2}{\ell+1}}$ colours, which means that $g(n, C_{2\ell}) > cn^{1-\frac{2}{\ell+1}}$.

The second assertion of Theorem 7.1.8 follows trivially since K_r contains C_r as a subgraph on the same vertex set, so $g(n, K_r) \geq g(n, C_r)$. $\square$

When r is odd, we can use a variant of the above method to prove the bound stated in Remark 7.1.9. Let $\theta_{\ell,t}$ be the union of t paths of length ℓ which share the same endpoints but are pairwise internally vertex-disjoint. Note that $\theta_{\ell,2} = C_{2\ell}$. A result of O. Janzer [97, Theorem 3.7] shows that Lemma 7.3.1 can be generalized to find, under the same conditions (with a constant C that depends in addition on t) a copy of $\theta_{\ell,t}$ without a pair of vertices related by $\sim$. Using this result, an argument very similar to the proof of Theorem 7.1.8 shows that if f_v are k-colourings of the edge set of K_n for $k \leq cn^{1-\frac{2}{\ell+1}}$ where c is a sufficiently small constant, then we can find a copy T of $\theta_{\ell,3}$ in K_n with the property that for each $v \in V(T)$, the colour of every edge of T incident to v is the same in f_v. It follows that $g(n, \theta_{\ell,3} \cup C_{2q}) = \Omega(n^{1-\frac{2}{\min(\ell,q)+1}})$, where $\theta_{\ell,3} \cup C_{2q}$ is the disjoint union of a $\theta_{\ell,3}$ and a C_{2q}. Choosing $\ell \in \{\lfloor \frac{r+1}{5} \rfloor, \lfloor \frac{r+1}{5} \rfloor - 1\}$ such that ℓ is even and setting $q = \frac{r+1-3\ell}{2}$, we obtain $q \geq \ell \geq \frac{r-8}{5}$ and $|V(\theta_{\ell,3} \cup C_{2q})| = 3\ell - 1 + 2q = r$, so $g(n, K_r) \geq g(n, \theta_{\ell,3} \cup C_{2q}) \geq \Omega(n^{1-\frac{2}{\ell+1}}) \geq \Omega(n^{1-\frac{10}{r-3}})$.

7.3.2 The proof of Theorem 7.1.10

In this subsection, we complete the proof of Theorem 7.1.10. Our results from Section 7.2 together with Theorem 7.1.4 already prove the "only if" part of Theorem 7.1.10, so it suffices to prove the "if" part, which amounts to giving a polynomial lower bound for $g(n, H)$ in the remaining cases.

We remark that if $g(n, H)$ is polynomial, then so is $g(n, H^+)$, where H^+ is the graph obtained from H by adding an isolated vertex. Indeed, assume that there are k-colourings f_v of the edges of K_n for each $v \in V(K_n)$ which form an (n, H^+)-local collection. Let $s = |V(H)| + 1$ and let $u_1, \ldots, u_s$ be arbitrary distinct vertices in K_n. Now for each $v \in V(K_n)$, define the colouring f'_v of $E(K_n)$ by setting $f'_v(e) = (f_v(e), f_{u_1}(e), \ldots, f_{u_s}(e))$. This is a colouring which uses at most k^{s+1} colours. We claim that these colourings form an (n, H)-local colouring. Indeed, otherwise we could find a copy T of H in K_n such that for each $v \in V(T)$ there are at least two edges in T which have the same colour in f'_v. By definition, any such pair of edges have the same colour in all of $f_v, f_{u_1}, \ldots, f_{u_s}$, so, as $s > |V(H)|$, we may find a copy T^+ of H^+ (obtained by adding a vertex u_i to T) such that for each $v \in V(T^+)$ there are two edges in T^+ which have the same colour in f_v. This contradicts the assumption that the f_v are (n, H^+)-local. Hence, the f'_v are indeed (n, H)-local, so $g(n, H) \leq g(n, H^+)^{s+1}$.

Together with our earlier observation, it also follows that if $g(n, H)$ is polynomial and F contains H as a (not necessarily spanning) subgraph, then $g(n, F)$ is also polynomial.

By the above discussion, it suffices to consider graphs with no isolated vertices. We first focus on graphs H with precisely 4 edges and no isolated vertices; the list of such graphs can be found in Table 7.1. It was shown in [37] that $g(n, C_4) = \Omega(n^{1/3})$ and that $g(n, P_4) = \Omega(n^{1/5})$. When H contains a triangle, then $g(n, H)$ is subpolynomial, and when H is the union of P_3 and P_1, we do not know whether $g(n, H)$ is polynomial or not. The remaining 6 cases are all covered by the following definition and theorem.

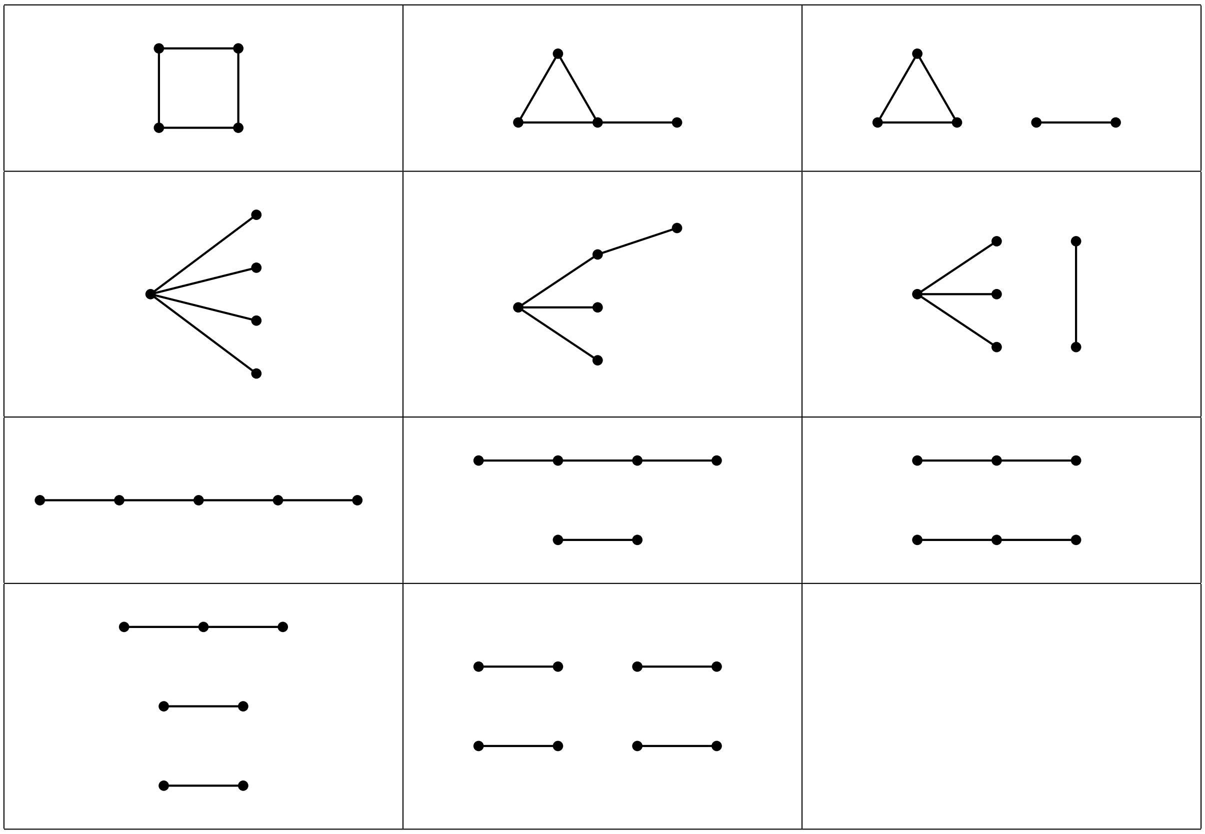

Table 7.1: The list of graphs with 4 edges and no isolated vertices

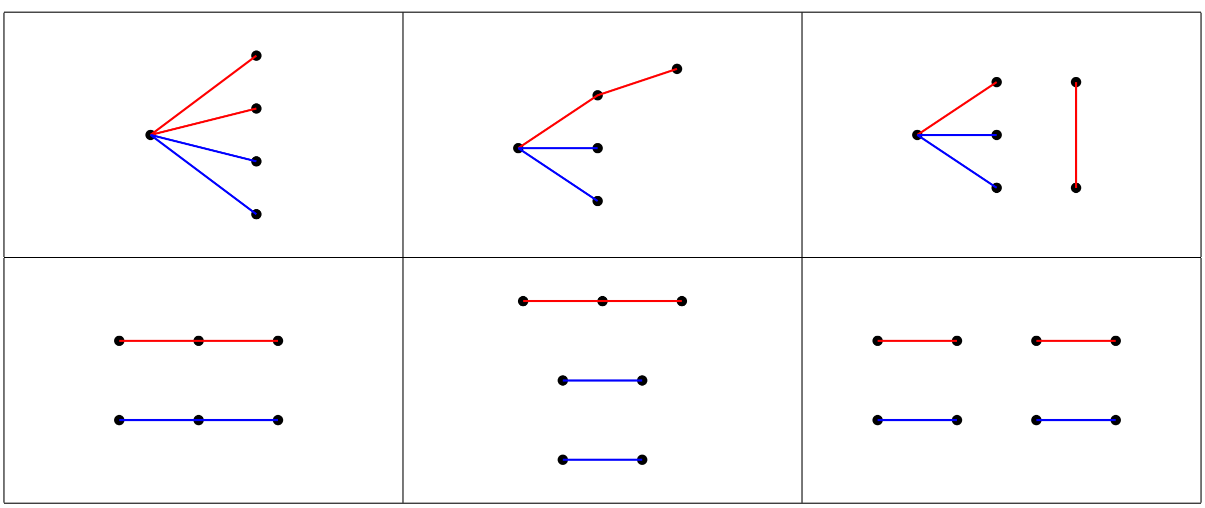

Table 7.2: The list of nice graphs with 4 edges and no isolated vertices

Definition 7.3.2. Let us call a graph H *nice* if it contains distinct edges e_1, e_2, f_1, f_2 such that $(e_1 \cup e_2) \cap (f_1 \cup f_2) \subset f_1 \cap f_2$.

Remark 7.3.3. That the remaining 6 graphs from Table 7.1 are all nice is demonstrated in Table 7.2. For each graph, suitable edges e_1 and e_2 are coloured red, while suitable f_1 and f_2 are coloured blue.

Theorem 7.3.4. *If H is a nice graph, then $g(n, H) = \Omega(n^{1/6})$.*

Proof. Choose distinct edges e_1, e_2, f_1, f_2 in H such that $(e_1 \cup e_2) \cap (f_1 \cup f_2) \subset f_1 \cap f_2$. For each $v \in V(K_n)$, let f_v be a colouring of the edges of K_n which uses $k \leq cn^{1/6}$ colours, where c is a sufficiently small constant which depends on H. It suffices to prove that the collection of these colourings is not (n, H)-local. We need the following claim, which follows from a simple counting argument.

Claim. The following two statements hold.

(a) There exist disjoint edges p and q in K_n such that the number of vertices v in K_n with $f_v(p) = f_v(q)$ is $\Omega(n/k)$.

(b) There exist distinct, intersecting edges s and t in K_n such that the number of vertices v in K_n with $f_v(s) = f_v(t)$ is $\Omega(n/k)$.

Proof of Claim. (a) It is easy to see by convexity that for any vertex $v \in V(K_n)$, there are $\Omega(n^4/k)$ pairs of disjoint edges p and q in K_n such that $f_v(p) = f_v(q)$. Hence, the number of triples (v, p, q) where p and q are disjoint edges in K_n and $f_v(p) = f_v(q)$ is $\Omega(n^5/k)$. It follows from the pigeonhole principle that there exist p and q for which the number of suitable choices for v is $\Omega(n/k)$.

(b) It is easy to see by convexity that for any vertex $v \in V(K_n)$, there are $\Omega(n^3/k)$ pairs of distinct, intersecting edges s and t in K_n such that $f_v(s) = f_v(t)$. Hence, the number of triples (v, s, t) where s and t are distinct, intersecting edges in K_n and $f_v(s) = f_v(t)$ is $\Omega(n^4/k)$. It follows from the pigeonhole principle that there exist s and t for which the number of suitable choices for v is $\Omega(n/k)$. □

Now if e_1 and e_2 are disjoint in H, let us use part (a) of the claim to find disjoint edges p and q in K_n such that there is a set A of $\Omega(n/k)$ vertices in $V(K_n) \setminus (p \cup q)$ such that each $v \in A$ satisfies $f_v(p) = f_v(q)$. Similarly, if e_1 and e_2 intersect each other in H, then let us use part (b) of the claim to find distinct, intersecting edges p and q in K_n such that there is a set A of $\Omega(n/k)$ vertices in $V(K_n) \setminus (p \cup q)$ such that each $v \in A$ satisfies $f_v(p) = f_v(q)$. Our aim is to find a copy of H in K_n in which e_1 and e_2 are mapped to p and q (in an arbitrary way), and all the remaining vertices of H are mapped to A. By the definition of A, for any such embedding T of H we have that T is not rainbow with respect to f_v whenever $v \in V(T) \setminus (p \cup q)$ (since we have

$f_v(p) = f_v(q)$ for any such vertex). We will embed the vertices in $(f_1 \cup f_2) \setminus (e_1 \cup e_2)$ into K_n in a way that for every $v \in p \cup q$ the images of the edges f_1 and f_2 will have the same colour with respect to the colouring f_v. If we can do this, we obtain an embedding T of H such that for each $v \in V(T)$ the graph T is not rainbow with respect to f_v, showing that the collection of colourings is not (n, H)-local.

We consider two cases. First, assume that f_1 and f_2 are disjoint in H. In particular, $f_1 \cup f_2$ is disjoint from $e_1 \cup e_2$. Label each edge between two vertices of A by its colours with respect to the colourings f_v for all $v \in p \cup q$. Depending on whether p and q intersect or not, this labels the edges by a triple or quadruple of colours. In particular, there are at most k^4 possible labels. Since there are $\Omega(n^2/k^2)$ edges between two vertices of A, there will be a label that appears on $\Omega(n^2/k^6)$ different edges. If c is sufficiently small, then we obtain two non-intersecting edges within A with the same label. Choosing these two edges as the image of f_1 and f_2, and mapping all remaining vertices of H arbitrarily to A, we obtain the desired embedding of H.

The second case is when f_1 and f_2 intersect each other in H. If the common vertex of f_1 and f_2 belongs to $e_1 \cup e_2$, then we have already mapped it to some vertex x; else let us choose an arbitrary vertex $x \in A$ as its image in the embedding. Labelling each $y \in A \setminus \{x\}$ by the colours of $f_v(xy)$ for all $v \in p \cup q$, there are at most k^4 possible labels, so if c is sufficiently small, then there exist $y \neq z$ in $A \setminus \{x\}$ such that $f_v(xy) = f_v(xz)$ holds for all $v \in p \cup q$. Mapping f_1 to xy and f_2 to xz, and mapping all remaining vertices of H to arbitrary vertices in A, we obtain a suitable embedding of H. □

We are now in a position to complete the proof of Theorem 7.1.10. Our Theorem 7.3.4 shows that whenever H contains one of the graphs in Table 7.2 as a subgraph, we have $g(n, H) = \Omega(n^{1/6})$. As mentioned above, Cheng and Xu [37] proved that $g(n, C_4) = \Omega(n^{1/3})$ and $g(n, P_4) = \Omega(n^{1/5})$. Observe that any graph with 4 edges which is triangle-free and not the disjoint union of P_3 and P_1 is equal to C_4, P_4 or one of the graphs in Table 7.2. This proves Theorem 7.1.10 for graphs H with 4 edges. Finally, observe that any graph with at least 5 edges contains C_4, P_4 or one of the graphs in Table 7.2 as a subgraph (clearly, it suffices to verify this for graphs obtained by adding an edge to one of the following three graphs: the triangle with a pendant edge, the triangle with an isolated edge and the union of a P_3 and a P_1). This completes the proof of Theorem 7.1.10.

7.4 Bounds for the Erdős–Gyárfás function

7.4.1 The proof of Theorem 7.1.16

In this subsection we prove Theorem 7.1.16 about subpolynomial values for the Erdős–Gyárfás function. We begin by briefly discussing our approach. We will use Theorem 7.1.14, i.e., an appropriate colouring for $r = 3$, to construct our colouring for $r = 4$ (and larger values of r). Let c denote the colouring of the triples provided by Theorem 7.1.14. Let us first try to colour the edges of $K_n^{(4)}$ by simply ignoring one of the vertices: e will receive colour $c(e - v(e))$, where $v(e)$ is some special vertex of e (and $e - v(e)$ is the set $e \setminus \{v(e)\}$). Note that we will need to describe a rule to choose the special vertex $v(e)$ that we ignore.

Now let us see when this approach provides a colouring satisfying the conditions. Observe that a copy of $K_5^{(4)}$ receives at least 4 colours if and only if there is at most one colour repetition, i.e., at most one pair of edges of the $K_5^{(4)}$ share the same colour. Assume instead that in our colouring we have two pairs of edges e, e' and f, f' in some $K_5^{(4)}$ sharing the same colour. Since any edge misses exactly one vertex of this $K_5^{(4)}$, there is a vertex p contained in the intersection $e \cap e' \cap f \cap f'$. If we can make sure that p is the special vertex that we ignore in each of e, e', f, f', then we get a contradiction from $c(e - p) = c(e' - p)$ and $c(f - p) = c(f' - p)$ by the definition of c.

To choose the special vertex $v(e)$, one simple method is to take an ordering of all of the vertices, and pick $v(e)$ to be the largest element of e. Of course, the point $p \in e \cap e' \cap f \cap f'$ does not in general have to be the largest element of e, e', f, f'. So, instead of taking just one ordering, we will take many total orders, and our final colouring will be the product of the colourings corresponding to each of these total orders. We will have to pick our list of total orders in a special way; we will make use of the following result of Hajnal (see [134]).

Theorem 7.4.1 ([134]). *Let k be a fixed positive integer. For any positive integer n and any set V of size n, we can find $M = O(\log \log n)$ total orders $<_1, \dots, <_M$ on V such that whenever $a_1, \dots, a_k$ are distinct elements of V, then there is some j (with $1 \leq j \leq M$) such that $a_i <_j a_1$ for $i = 2, \dots, k$. In other words, a_1 is maximal among $a_1, \dots, a_k$ in at least one of our total orders.*

We are now ready to prove Theorem 7.1.16.

Proof of Theorem 7.1.16. We show the statement by induction on r. The $r = 3$ case is exactly Theorem 7.1.14, so assume that $r \geq 4$ and the statement holds for smaller values of r. Let V be our set of n vertices. By the induction hypothesis, we can pick a colouring c of the $(r-1)$-element subsets of V such that at least $r - 1$ colours appear among any r vertices and c uses $e^{(\log n)^{2/5+o(1)}}$ colours. Furthermore, by Theorem 7.4.1 applied for $k = r + 1$, we can pick $M = O(\log \log n)$ total orders $<_1, \dots, <_M$ on V such that for any distinct vertices $a_1, \dots, a_{r+1} \in V$, a_1 is maximal among these $r + 1$ vertices in one of the total orders $<_j$. Given a (non-empty) set $W \subseteq V$, let us write $\max_j W$ for the element of W which is maximal in W in the ordering $<_j$. We define a colouring c' of r-element subsets of V by setting

$$c'(e) = (c(e - \max_1(e)), c(e - \max_2(e)), \dots, c(e - \max_M(e))).$$

In other words, c_j is the product of all the colourings $c'_j(e) = c(e - \max_j(e))$ formed by using the colouring c after ignoring the largest element of e in the ordering $<_j$. Note that the number of colours used is at most

$$\left(e^{(\log n)^{2/5+o(1)}}\right)^{O(\log\log n)} = e^{(\log n)^{2/5+o(1)}}.$$

We claim that in any $K_{r+1}^{(r)}$, at least r colours appear under the colouring c'. Assume, for contradiction, that this is not the case. Then there exist a set W of $r+1$ vertices and subsets $e, e', f, f' \subseteq W$ of size r (with $e \neq e', f \neq f'$) such that $c'(e) = c'(e')$, $c'(f) = c'(f')$ and $\{e, e'\} \neq \{f, f'\}$. Since any r-edge inside W misses exactly one element of W, we have $|e \cap e' \cap f \cap f'| \geq (r+1) - 4 \geq 1$. Pick any element $p \in e \cap e' \cap f \cap f'$, and let $j \in \{1, \dots, M\}$ be such that $p = \max_j(W)$.

Since $c'(e) = c'(e')$, we have (by taking jth coordinates) $c(e-p) = c(e'-p)$. Similarly, $c(f-p) = c(f'-p)$. But this means that, under the colouring c, at most $r-2$ different colours appear in the $K_r^{(r-1)}$ induced by $W-p$. This contradicts our choice of c and finishes the proof. □

7.4.2 Connections to Problem 7.1.1

Recall that lower bounds on $g(n, P_3)$ also imply the same lower bounds for Problem 7.1.1. However, we have seen (Theorem 7.1.7) that $g(n, P_3)$ is subpolynomial, so this method cannot give a polynomial lower bound for the problem of Karchmer and Wigderson.

It is very natural to try to give lower bounds to Problem 7.1.1 by considering only specific forms of triples x, y, z; in particular, it is natural to take $(x \cup y \cup z) \setminus (x \cap y \cap z)$ to be small to make sure that only a few conditions need to be satisfied. Thus, instead of taking x, y, z to be edges of a P_3, we could add to each of them the same set of vertices, i.e., take $x = x_0 \cup S, y = y_0 \cup S, z = z_0 \cup S$, where x_0, y_0, z_0 are sets of size 2 forming a P_3.

However, Theorem 7.1.16 implies that this cannot work for sets S of some given bounded size ℓ. Indeed, we can proceed similarly as for P_3s. Let γ be a colouring coming from Theorem 7.1.16 for $r = \ell + 3$, and define $f_v(x)$ to be $\gamma(\{v\} \cup x)$ if $v \notin x$ and some arbitrary different colour if $v \in x$. Then, similarly to the proof of Theorem 7.1.7, we find that whenever $x = x_0 \cup S$, $y = y_0 \cup S$ and $z = z_0 \cup S$ such that $|S| = \ell$ and x_0, y_0, z_0 form a P_3 *abcd* (disjoint from S), then x, y, z receive distinct colours in either f_a or f_d.

More generally, similar arguments can be used to construct colourings f_v for each $v \in V(K_n)$ using a subpolynomial number of colours such that whenever x, y, z are distinct subsets of $V(K_n)$ of bounded size, then there exists some $v \in (x \cup y \cup z) \setminus (x \cap y \cap z)$ for which $f_v(x)$, $f_v(y)$ and $f_v(z)$ are distinct.

Note, however, that the argument above does not work when $|S|$ is large (say, at least logarithmic in n); perhaps considering such S might give better bounds.

7.5 Concluding remarks

In this chapter we have determined, for each graph H other than the disjoint union of P_3 and P_1 with an arbitrary number of isolated vertices, whether the growth of $g(n, H)$ is polynomial. The natural question left open by our investigations is whether $g(n, H)$ is polynomial when H is $P_3 \cup P_1$.

Problem 7.1.1 also remains open, with the best known lower bound coming from the result $g(n, P_3) = \Omega((\frac{\log n}{\log \log n})^{1/4})$ due to Alon and Ben-Eliezer [3]. Our Theorem 7.1.7 shows that this approach cannot give a polynomial lower bound. Indeed, as we mentioned in Subsection 7.4.2, one cannot obtain a polynomial lower bound by considering only vectors of bounded Hamming weight. On the other hand, one can prove a polynomial upper bound using the asymmetric Lovász Local Lemma. It remains an interesting open question whether the answer to Problem 7.1.1 is polynomial.

Chapter 8

Long running times for hypergraph bootstrap percolation

8.1 Introduction

The *hypergraph bootstrap percolation process* is an infection process on hypergraphs which was introduced by Bollobás [23] in 1968 under the name of *weak saturation*. For an integer $r \geq 2$ and a set S, denote by $\binom{S}{r}$ the set of all subsets of S of size r. Given an r-uniform hypergraph H and a positive integer n, the *H-bootstrap percolation process* is a deterministic process defined as follows. We start with a given r-uniform hypergraph G_0 on vertex set $[n] = \{1, \ldots, n\}$. For each time step $t \geq 1$, we define the hypergraph G_t on the same vertex set $[n]$ by letting

$$E(G_t) = E(G_{t-1}) \cup \left\{ e \in \binom{[n]}{r} : \exists \text{ an } H\text{-copy } H' \text{ s.t. } e \in E(H') \subseteq E(G_{t-1}) \cup \{e\} \right\},$$

that is, G_t is an r-uniform hypergraph on $[n]$ defined by including all edges of G_{t-1} together with all edges $e \in \binom{[n]}{r}$ which create a new copy of H with the edges of G_{t-1}. The hypergraph G_0 is called the *initial infection*, and the edges $E(G_t) \setminus E(G_{t-1})$ are said to be *infected* at time t. If there exists some $T \geq 0$ such that $G_T = K_n^{(r)}$, we say that G_0 *percolates* under this process. In the weak saturation interpretation, we say that the hypergraph G_0 is weakly H-saturated if G_0 is H-free and percolates under H-bootstrap percolation, that is, if there exists an ordering of $E(K_n^{(r)}) \setminus E(G_0) = \{e_1, \ldots, e_t\}$ such that the addition of e_i to $G_0 \cup \{e_1, \ldots, e_{i-1}\}$ will create a new copy of H, for every $i \in [t]$.

Given a fixed hypergraph H, one of the most studied extremal problems in this setting is establishing the minimum size of an n-vertex hypergraph which is weakly H-saturated. For the most basic case, where $r = 2$ and $H = K_k$, it was conjectured by Bollobás [23] that the minimum size of a weakly K_k-saturated n-vertex graph is $(k-2)n - \binom{k-1}{2}$. About a decade later, Lovász [113] was the first to confirm this conjecture (using a generalisation of the Bollobás Two

Families Theorem [22]). This was later independently reproved by Alon [2], Frankl [66] and Kalai [101, 102] – all using methods from linear algebra. For the hypergraph case, the work of Frankl [66] also settles this problem for $K_k^{(r)}$ with $r \geq 3$. This problem has also been studied for other graphs H, and for host graphs other than the complete graph, and other related settings; see, e.g., [2, 35, 110, 115, 121, 122, 128].

Even though the initial infection graphs which are solutions to the weak saturation problem have the smallest possible number of edges, it is interesting to note that, in many of the known examples, they require only very few steps until the infection process stabilises. For example, the standard construction of a weakly K_k-saturated graph achieving the minimum size is given by removing the edges of a clique of size $n-k+2$ from K_n, which means that only one step is needed in order to complete the infection process. In this direction, Bollobás raised the problem of finding the initial infection for which the running time of the H-bootstrap percolation process is maximised. This was then also studied in the related setting of neighbourhood percolation by Benevides and Przykucki [17, 18, 123], and for a random initial infection by Gunderson, Koch and Przykucki [78].

Here we consider this problem in the hypergraph bootstrap percolation setting. Given a fixed r-uniform hypergraph H and an r-uniform initial infection G_0, we define the *running time* of the H-bootstrap percolation process on G_0 to be

$$M_H(G_0) = \min\{t \geq 0 : G_t = G_{t+1}\}.$$

We denote the *maximum running time* over all r-uniform hypergraphs G_0 on n vertices as $M_H(n)$. We shall simplify these notations to $M_k^r(G_0)$ and $M_k^r(n)$ when $H = K_k^{(r)}$ is the complete r-uniform hypergraph on k vertices, and drop the superscript to $M_k(n)$ in the graph setting ($r=2$). Note that a trivial upper bound for $M_H(n)$ is given by $\binom{n}{r}$, the total number of edges of $K_n^{(r)}$.

The simplest setting to consider is for graph bootstrap percolation and $H = K_k$. For $k=3$, it is not hard to see that $M_3(n) = \lceil \log_2(n-1) \rceil$, where an extremal example is given by a path of length $n-1$ (see, e.g., [28] for details). Bollobás, Przykucki, Riordan, and Sahasrabudhe [28] and independently Matzke [114] considered this problem for higher values of k. By carefully analysing the growth of cliques during the percolation process, both groups of authors showed that $M_4(n) = n-3$. Moreover, for $k \geq 5$, Bollobás, Przykucki, Riordan, and Sahasrabudhe [28] obtained the lower bound $M_k(n) \geq n^{2-\alpha_k - o(1)}$, where $\alpha_k = (k-2)/(\binom{k}{2}-2)$, using a probabilistic argument. The authors of [28] conjectured that $M_k(n) = o(n^2)$ for all $k \geq 5$. However, in a subsequent paper, Balogh, Kronenberg, Pokrovskiy, and Szabó [15] disproved this conjecture for $k \geq 6$, showing that the natural upper bound is tight up to a constant factor. The authors of [15] also improved the lower bound for $k=5$ to $M_5(n) \geq n^{2-O(1/\sqrt{\log n})}$, using Behrend's construction of 'dense' 3-AP-free sets [16], and conjectured that $M_5(n) = o(n^2)$. It remains an open problem to determine whether this is the case.

In this chapter we consider the question of the maximum running time when H is an r-uniform hypergraph with $r \geq 3$. This was recently investigated by Noel and Ranganathan [119]. By providing an explicit construction to establish the lower bound (noting the trivial upper bound of $\binom{n}{r}$), they proved the following theorem for the case $k \geq r+2$.

Theorem 8.1.1 (Noel and Ranganathan [119])**.** *Let $r \geq 3$. If $k \geq r+2$, then $M_k^r(n) = \Theta(n^r)$.*

For the case $k = r+1$, they established the following lower bound.

Theorem 8.1.2 (Noel and Ranganathan [119])**.** *Let $r \geq 3$. If $k = r+1$, then $M_k^r(n) = \Omega(n^{r-1})$.*

This theorem leaves a gap between the lower bound and the trivial upper bound $M_{r+1}^r(n) = O(n^r)$. Noel and Ranganathan conjectured that $M_4^3(n) = O(n^2)$ [119, Conjecture 5.1], but suggested that, for sufficiently large r, it is indeed true that the maximum running time achieves $M_{r+1}^r(n) = \Theta(n^r)$ [119, Question 5.2].

In this chapter, we show the conjecture to be false and prove that the trivial upper bound is in fact tight, up to a constant factor, for all $r \geq 3$. This also gives a positive answer to their question, in a strong sense.

Theorem 8.1.3. *For any fixed integer $r \geq 3$, we have $M_{r+1}^r(n) = \Theta(n^r)$.*

Another proof for Theorem 8.1.3 was independently obtained by Hartarsky and Lichev [82].

We note that Theorem 8.1.3 establishes a clear difference with respect to the graph case $r = 2$, where $M_k(n) = o(n^r)$ for $k \in \{r+1, r+2\}$ (and possibly also $r+3$). It may therefore seem that the behaviour of hypergraph bootstrap percolation is less rich than its graph counterpart. We propose a modification of the problem above that shows this is not the case, and that different (and very interesting) asymptotic running times may still occur in the hypergraph setting.

Indeed, recall that we may think of H-bootstrap percolation as an infection process where the infection spreads to a new copy of H if only one edge of said copy was not infected in the previous step. It is reasonable then to consider models where the infection is more powerful, in the sense that it will extend to copies of H which are missing at most m edges, for some fixed integer m. We consider here in particular the case $m = 2$. Note that if $m = 2$ and H is a complete hypergraph (which is the case we will focus on), then this modified model is equivalent to the original hypergraph percolation process for the hypergraph H' obtained by deleting an arbitrary edge from H.

Formally, let H be a given r-uniform hypergraph, and let G be an r-uniform hypergraph on $[n]$. For each copy H' of H on $[n]$, if $|E(H') \setminus E(G)| \leq m$, we say that H' is *m-completable* in G. We define the (H, m)-bootstrap percolation process on an initial infection G_0 on $[n]$ to be the sequence of hypergraphs $G_0, G_1, \ldots$ on $[n]$ given by setting, for each $t \geq 1$,

$$E(G_t) = E(G_{t-1}) \cup \bigcup_{\substack{H' \text{ copy of } H \text{ on } [n] \\ H' \ m\text{-completable in } G_{t-1}}} E(H').$$

Note that the $(H,1)$-bootstrap percolation process simply corresponds to the usual H-bootstrap percolation process. Let us denote the running time of this hypergraph percolation process as $M_{(H,m)}(G_0) = \min\{t \geq 0 : G_t = G_{t+1}\}$, and the maximum running time over all r-uniform n-vertex hypergraphs G_0 as $M_{(H,m)}(n)$. The next result shows that we get interesting new behaviour when $m = 2$ and $H = K_4^{(3)}$ (which is probably the most natural first case to consider).

Theorem 8.1.4. *For all $n \geq 4$, we have $M_{(K_4^{(3)},2)}(n) = 2n - \lfloor \log_2(n-2) \rfloor - 6$.*

It is worth remarking here that this is the first nontrivial exact result about running times of hypergraph bootstrap percolation. The only nontrivial exact results in graph bootstrap percolation were those for K_3- and K_4-bootstrap percolation [28], until very recently Fabian, Morris, and Szabó [61] determined the maximal possible running time for all cycles C_k (when n is sufficiently large).

We also prove that in the next case, $H = K_5^{(3)}$, the running time can once again be cubic (i.e., as large as possible).

Theorem 8.1.5. *We have $M_{(K_5^{(3)},2)}(n) = \Theta(n^3)$.*

Let $K_s^{(r)} - e$ denote the hypergraph obtained by deleting an edge from $K_s^{(r)}$. As mentioned above, the $(K_s^{(r)}, 2)$-process is the same as the usual bootstrap percolation process for $K_s^{(r)} - e$, so the results above can be reformulated as follows.

Theorem 8.1.4'. *For all $n \geq 4$, we have $M_{K_4^{(3)}-e}(n) = 2n - \lfloor \log_2(n-2) \rfloor - 6$.*

Theorem 8.1.5'. *We have $M_{K_5^{(3)}-e}(n) = \Theta(n^3)$.*

We present our proof of Theorem 8.1.3 in Section 8.2. We defer the proofs of Theorems 8.1.4 and 8.1.5 to Section 8.3. We also propose some open problems in our concluding remarks.

8.2 Long running times for simple infections

In order to prove Theorem 8.1.3, we will use a result of Noel and Ranganathan [119] that allows us to focus on the case $r = 3$. To state their result, we need to recall some definitions from [119]. Let G_0 be an r-uniform hypergraph, let G_t be the hypergraph at time t for the $K_{r+1}^{(r)}$-bootstrap process starting with G_0 as initial infection, and let $T = M_{r+1}^r(G_0)$ be the time the process stabilises. We say that G_0 is $K_{r+1}^{(r)}$-*civilised* if the following conditions are satisfied for some edge e_0 of G_0.

1. For each $t \in [T]$, G_t contains only one more edge e_t than G_{t-1}, and one more copy H_t of $K_{r+1}^{(r)}$.

2. For all $t \in [T]$ we have $E(H_t) \cap \{e_0, e_1, \ldots, e_T\} = \{e_{t-1}, e_t\}$.

3. The $K_{r+1}^{(r)}$-bootstrap percolation process starting with $G_0 - e_0$ infects no edge.

Lemma 8.2.1 (Noel and Ranganathan [119, Lemma 2.11])**.** *If for all n there exists a $K_4^{(3)}$-civilised hypergraph G_0 on $\Theta(n)$ vertices such that $M_4^3(G_0) = \Theta(n^3)$, then for all $r \geq 3$ we have $M_{r+1}^r(n) = \Theta(n^r)$.*

Before we give the formal proof of Theorem 8.1.3, let us give an informal description of the construction that gives a lower bound for the number of steps of the percolation process. As noted above, by Lemma 8.2.1 it is enough to consider the case $r = 3$. The main part of the construction consists of three layers of vertices: 'top' vertices labelled t_i, 'bottom' vertices labelled b_j, and 'middle' vertices labelled m_ℓ. In each time step, just one new edge will become infected. That infection will happen because one copy of $K_4^{(3)}$, which had only two edges present in the initial infection, has a third edge infected in the previous step of the process.

The process will consist mainly of chains of infections, where we move from one chain to another by using special gadgets. The chains will have the format of the so-called 'beachball hypergraph'. The vertex set of this hypergraph consists of one top and one bottom vertex, and some ordered vertices in the middle; the edges are the triples consisting of two consecutive middle vertices, and either the top or the bottom vertex. See Figure 8.1 for an illustration.

It will be convenient to think of the process as having n phases, each phase having $\Theta(n)$ stages, and each stage having $\Theta(n)$ infection steps. A phase will represent the infection process that occurs when we fix a top vertex t_i. In each phase, we have $\Theta(n)$ stages, where each stage is the process that occurs when we fix b_j (for the fixed t_i of this phase). At a specified phase and stage, the initial infected set will be the above mentioned beachball hypergraph, and the infection will spread through the middle vertices. This gives $\Theta(n)$ infection steps for each stage.

The challenge will then be to move to another top or bottom vertex without infecting more than one edge in each step of the process. For this purpose, we will introduce, at the end of each stage, a new middle vertex and a special 'switching' gadget. Each stage of the process will be represented by a tuple of a top vertex t_i, a bottom vertex b_j, and consecutive middle vertices starting from m_s and ending in m_ℓ, where $-(n-1) \leq s \leq 0$ and $n \leq \ell \leq 2n$. For moving between phases, we will introduce a different type of gadget.

Let us first describe the first few stages of the process to give a better intuition. The first n infection steps will come from a 'path' on the middle layer. The edges $t_1 m_\ell m_{\ell+1}$ and $b_1 m_\ell m_{\ell+1}$ will be present at time zero for all $0 \leq \ell \leq n-1$, as well as the edge $t_1 b_1 m_0$. Once the edge $t_1 b_1 m_\ell$ becomes infected, it propagates the infection in the next step to $t_1 b_1 m_{\ell+1}$. See Figure 8.1 for an illustration.

After $\Theta(n)$ such infections, we want to swap out b_1 to another bottom vertex (labelled b_{-1}). We do this by making sure that the last infected edge using b_1 (namely, the edge $t_1 b_1 m_n$) makes the middle path longer, that is, it makes the edge $t_1 m_n m_{n+1}$ infected in the next step. To achieve this, we will have $b_1 m_n m_{n+1}$ and $t_1 b_1 m_{n+1}$ present in the original hypergraph G_0; see Figure 8.2.

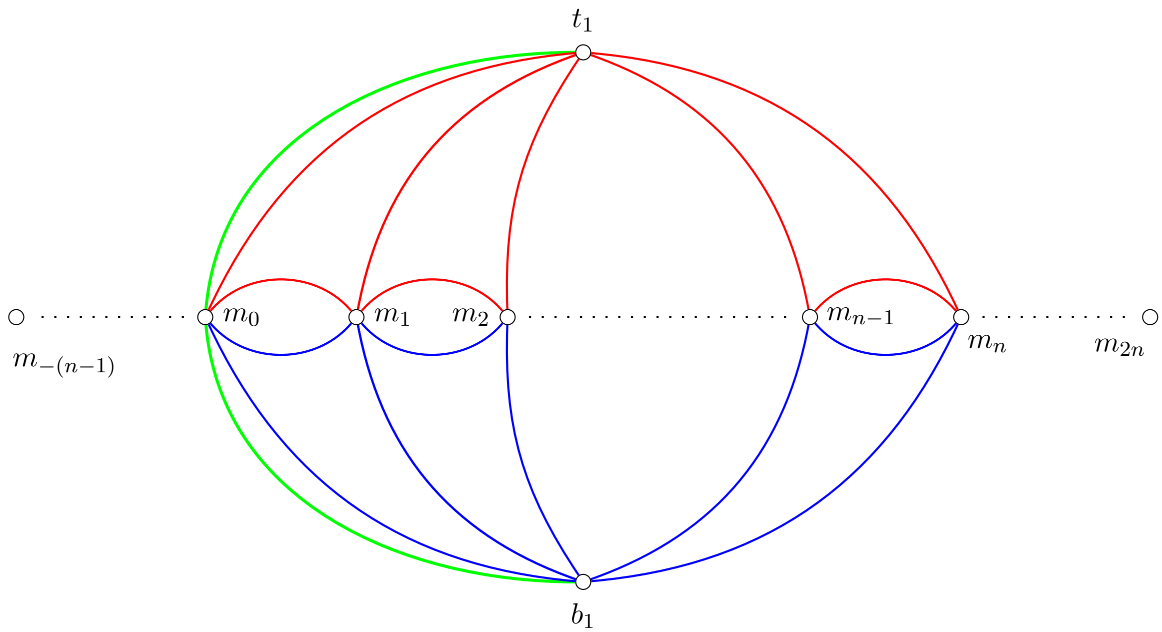

Figure 8.1: Initial infection G_0 showing only the first top and bottom vertices, t_1 and b_1. Each red or blue triangle represents an edge of G_0, and together they form the first beachball hypergraph in our process. The green arc represents an edge containing the vertices it passes through. To form G_1, the edge $t_1m_1b_1$ is added, as this completes a copy of $K_4^{(3)}$ on $\{t_1, m_0, m_1, b_1\}$. It is clear to see that subsequently all edges of the form $t_1m_\ell b_1$ for ℓ increasing from 2 to n are added in turn.

Once $t_1m_nm_{n+1}$ is infected, it can start a chain of infections using the new bottom vertex b_{-1}. However, this time the chain of infections will go in the opposite direction on the middle path: we will first infect $t_1b_{-1}m_n$ (for this we will need the edges $t_1b_{-1}m_{n+1}$ and $b_{-1}m_nm_{n+1}$ to be present initially, as in Figure 8.2), then we infect $t_1b_{-1}m_{n-1}$, and so on, until $t_1b_{-1}m_0$.

At this point we again swap out the bottom vertex to a different one (labelled b_2) — we do this using the same trick as above, i.e., making the middle path one longer, and then changing the direction we traverse the path. We keep repeating the steps above for $\Theta(n)$ bottom vertices to get $\Theta(n^2)$ infections which all use the same top vertex t_1.

Once we have the $\Theta(n^2)$ infections using t_1, we wish to swap out the top vertex t_1 to a different one (labelled t_2). We could do this similarly to how we swapped the b_j's, but it is more convenient to simply introduce a gadget using three 'dummy' vertices d_1, d_2, d_3 to do this swap. The last infection using t_1 (namely, $t_1m_{2n-1}m_{2n}$) will start a short chain of infections using the $K_4^{(3)}$'s given by $t_1m_{2n-1}m_{2n}d_1$, $m_{2n-1}m_{2n}d_1d_2$, $m_{2n}d_1d_2d_3$, $d_1d_2d_3t_2$, $d_2d_3t_2m_0$, and $d_3t_2m_0m_1$. The last one of these allows us to start a repeat of the previous infection process, using t_2 instead of t_1. We will use three such dummy vertices d_i for each of the $n-1$ swaps at the top — so only $3n-3 = \Theta(n)$ dummy vertices in total. See Figure 8.3 for an illustration.

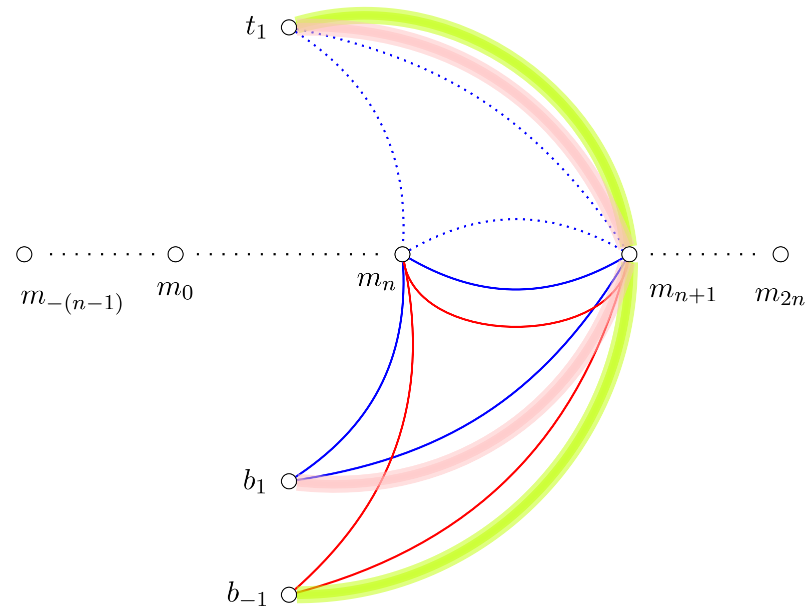

Figure 8.2: Switching gadget to change from $K_4^{(3)}$ copies containing b_1 to those containing b_{-1}. The edges $b_1m_nm_{n+1}$ and $b_{-1}m_nm_{n+1}$ are present in the initial infection G_0. After the edge $t_1m_nb_1$ is created by the percolation process, the copy of $K_4^{(3)}$ induced by the vertices $\{t_1, m_n, m_{n+1}, b_1\}$ is present, except for the missing edge $t_1m_nm_{n+1}$ shown in the dotted blue line. Thus this edge is added, followed by $t_1m_nb_{-1}$. This triggers the process to run backwards and create all edges of form $t_1m_ib_{-1}$, for i decreasing from $i = n-1$ to $i = 0$, in turn.

Let us now turn to the formal proof of our theorem.

Proof of Theorem 8.1.3. By Lemma 8.2.1, it suffices to consider the case $r = 3$ and show that there are $K_4^{(3)}$-civilised 3-uniform hypergraphs on $\Theta(n)$ vertices such that the $K_4^{(3)}$-bootstrap process takes $\Theta(n^3)$ steps to stabilise. We now describe a construction achieving this.

The initial infection hypergraph G_0 has $9n-4 = \Theta(n)$ vertices, which are labelled as follows: $t_1, \dots, t_n$, $b_1, \dots, b_n, b_{-1}, \dots, b_{-(n-1)}$, $m_{-(n-1)}, m_{-(n-2)} \dots, m_{2n}$, and $d_{i,1}, d_{i,2}, d_{i,3}$ for $i \in [n-1]$. The edges of G_0 are given below:

(a) $t_1m_0m_1$;

(b) $t_im_\ell m_{\ell+1}$ for all $i \in [n]$ and $\ell \in [n-1]$;

(c) $b_jm_\ell m_{\ell+1}$ for all $j \in [n]$ and $\ell \in [-(j-1), n+j-1]$;

(d) $b_{-j}m_\ell m_{\ell+1}$ for all $j \in [n-1]$ and $\ell \in [-j, n+j-1]$;

(e) $t_ib_jm_{-(j-1)}$ and $t_ib_jm_{n+j}$ for all $i, j \in [n]$;

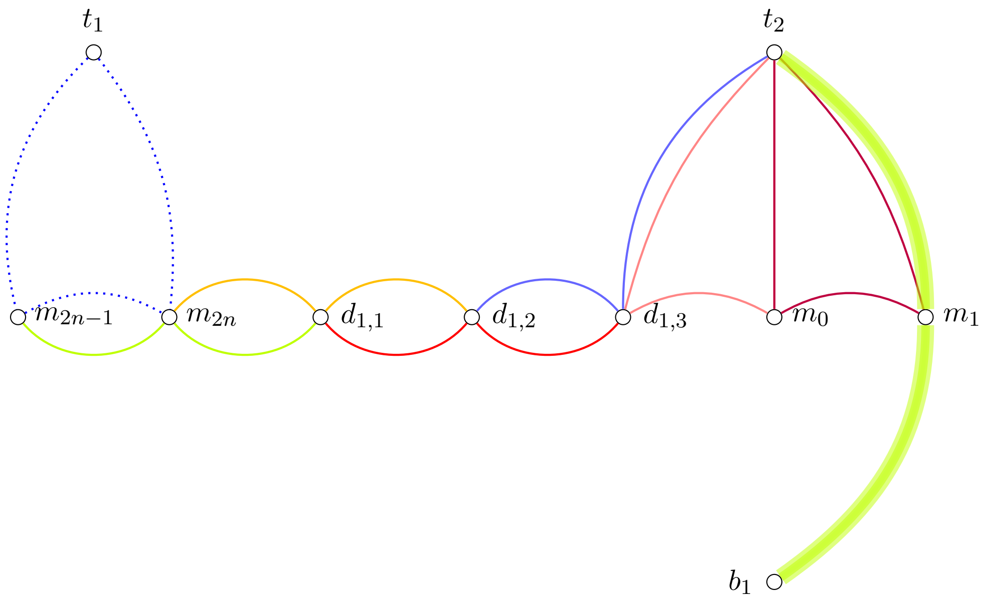

Figure 8.3: Switching gadget for changing the top vertex t_1 to t_2. Edges present in the initial infection G_0 are omitted for clarity. When the dotted blue edge $t_1 m_{2n-1} m_{2n}$ is infected, this causes the edges along the chain to become infected, ending in $t_2 m_0 m_1$. This triggers the infection of $t_2 m_1 b_1$, and in turn the process from the stage as shown in Figure 8.1, with t_1 replaced with t_2.

(f) $t_i b_{-j} m_{n+j}$ and $t_i b_{-j} m_{-j}$ for all $i \in [n]$ and $j \in [n-1]$;

(g) $t_i m_{2n-1} d_{i,1}$, $t_i m_{2n} d_{i,1}$, $m_{2n-1} m_{2n} d_{i,2}$, $m_{2n-1} d_{i,1} d_{i,2}$, $m_{2n} d_{i,1} d_{i,3}$, $m_{2n} d_{i,2} d_{i,3}$, $d_{i,1} d_{i,2} t_{i+1}$, $d_{i,1} d_{i,3} t_{i+1}$, $d_{i,2} d_{i,3} m_0$, $d_{i,2} t_{i+1} m_0$, $d_{i,3} t_{i+1} m_1$, and $d_{i,3} m_0 m_1$, for all $i \in [n-1]$.

As mentioned in the informal discussion, it will be easier to think about the initial infected hypergraph as a set of beachball hypergraphs, and gadgets connecting between them. For this purpose, we note the following.

- The edges from (b) and (c), as well as those from (b) and (d), (nearly) form beachball hypergraphs. For the beachballs with edges from (b) and (c) the infection process increases with the indices of the middle vertices, whereas for those from (b) and (d) it decreases with the indices of the middle vertices. These hypergraphs are used as the main ingredients of the infection process.

- The second type of edges from (e) together with the first type of edges from (f) form the gadgets that help us swap from b_j to b_{-j}, where t_i is fixed; that is, they help us move between the beachball with t_i, b_j as top and bottom and the beachball with t_i, b_{-j}. Figure 8.2 illustrates the gadget swapping from b_1 to b_{-1}.

- The first type of edges from (e) and second type of edges from (f) create the gadgets that help us swap from b_{-j} to b_{j+1}, where t_i is fixed; that is, these help us move from the beachball with t_i, b_{-j} as top and bottom to the beachball with t_i, b_{j+1}.

- The edges in (g) form the gadgets swapping between top vertices, from t_i to t_{i+1}, using the dummy vertices $d_{i,s}$. In other words, these gadgets move us from the beachball with t_i, b_n as top and bottom to the beachball with t_{i+1}, b_1. Figure 8.3 illustrates the gadget swapping from t_1 to t_2.

We will show that there are three types of edges that are being infected during the process:

(I) missing edges of the beachballs, that is, edges of the form $t_i b_j m_\ell$;

(II) edges from the gadgets swapping bottom vertices, of the form $t_i m_\ell m_{\ell+1}$, and

(III) edges from the gadgets swapping top vertices (these have several different forms).

We will now name the edges being infected during the process. For each $i, j \in [n]$, let $A_{i,j}$ denote the following sequence of edges:

$$A_{i,j} = (t_i b_j m_{-(j-2)},\ t_i b_j m_{-(j-3)}, \dots,\ t_i b_j m_{n+j-1},\ t_i m_{n+j-1} m_{n+j}). \tag{8.1}$$

These will be the edges infected in the stage of phase i corresponding to the bottom vertex b_j. Similarly, for each $i \in [n]$ and $j \in [n-1]$, let

$$A_{i,-j} = (t_i b_{-j} m_{n+j-1},\ t_i b_{-j} m_{n+j-2}, \dots,\ t_i b_{-j} m_{-(j-1)},\ t_i m_{-(j-1)} m_{-j}). \tag{8.2}$$

These will correspond to the stage with b_{-j} as bottom vertex. Concatenating these, we get the sequence A_i of edges corresponding to phase i (these are edges of types (I) and (II) above):

$$A_i = A_{i,1} A_{i,-1} A_{i,2} A_{i,-2} \dots A_{i,n-1} A_{i,-(n-1)} A_{i,n}.$$

For the phase change using the dummy vertices $d_{i,j}$, let us write, for each $i \in [n-1]$,

$$D_i = (m_{2n-1} m_{2n} d_{i,1},\ m_{2n} d_{i,1} d_{i,2},\ d_{i,1} d_{i,2} d_{i,3},\ d_{i,2} d_{i,3} t_{i+1},\ d_{i,3} t_{i+1} m_0,\ t_{i+1} m_0 m_1). \tag{8.3}$$

These are the edges of type (III). Finally, let us write A for the concatenation

$$A = A_1 D_1 A_2 D_2 \dots A_{n-1} D_{n-1} A_n.$$

We will show that during the infection process, edges become infected one-by-one, according to the sequence A.

Let T be the number of triples in A, and let $A = (e_1, e_2, \dots, e_T)$. Note that $T = \Theta(n^3)$. Let us also write e_0 for the edge $t_1 m_0 m_1$, and for all $a \in [T-1]$ let H_a be the copy of $K_4^{(3)}$ with vertex set $e_{a-1} \cup e_a$ (note that $|e_{a-1} \cap e_a| = 2$ for all a). For each $s \in [T]$, let us write G_s for the hypergraph with edge set $E(G_0) \cup \{e_1, e_2, \dots, e_s\}$. (Note that we do not yet know that these coincide with the hypergraphs obtained during the $K_4^{(3)}$-bootstrap process, but we will see that they do.) Let us also write $G_{-1} = G_0 - e_0$.

Claim 8.2.2. Assume that $s \in [-1, T]$ is an integer and $e = x_1x_2x_3$ is a triple not contained in $E(G_s)$. Suppose that adding e to G_s completes a copy of $K_4^{(3)}$ whose fourth vertex is x_4. Then, $s \in [0, T-1]$, $e = e_{s+1}$ and $\{x_1, x_2, x_3, x_4\} = V(H_{s+1})$.

We note here that the case $s = -1$ is needed to formally justify that G_0 is $K_4^{(3)}$-civilised below.

Proof. We consider the following two cases.

Case 1: a vertex $d_{i,c}$ appears among $x_1, \dots, x_4$ (for some $i \in [n]$ and $c \in [3]$). Let us temporarily write $d_{i,-2} = t_i$, $d_{i,-1} = m_{2n-1}$, $d_{i,0} = m_{2n}$, $d_{i,4} = t_{i+1}$, $d_{i,5} = m_0$ and $d_{i,6} = m_1$, so the edges $d_{i,a}d_{i,a+1}d_{i,a+3}$ and $d_{i,a}d_{i,a+2}d_{i,a+3}$ are present in G_0 for all $-2 \le a \le 3$. Observe that the only vertices appearing in an edge of G_s together with $d_{i,c}$ (recall $1 \le c \le 3$) are of the form $d_{i,a}$ with $|c - a| \le 3$ (see (g) as well as (8.3)). Hence, $x_1, \dots, x_4$ are all of the form $d_{i,a}$ for some $-2 \le a \le 6$. Observe furthermore that every edge of G_s of the form $d_{i,p}d_{i,q}d_{i,r}$ ($-2 \le p < q < r \le 6$) satisfies $|r - p| \le 3$, or $(p, q, r) = (-2, 5, 6)$ or $(p, q, r) = (-1, 0, 4)$. It is easy to deduce that the only possible quadruples of vertices $d_{i,a}$ forming a $K_4^{(3)}$ minus an edge are of the form $\{x_1, x_2, x_3, x_4\} = \{d_{i,a}, d_{i,a+1}, d_{i,a+2}, d_{i,a+3}\}$ (for some $-2 \le a \le 3$). So exactly one of $d_{i,a}d_{i,a+1}d_{i,a+2}$ and $d_{i,a+1}d_{i,a+2}d_{i,a+3}$ appears in G_s, as the other two triples appear in G_0 (recall (g)). Since these are the edges e_N and e_{N+1}, respectively, for some $N \in [T-1]$, we must have $e = e_{N+1}$ and $e_N \in G_s$. So $N = s$, $e = e_{s+1}$, and $\{x_1, x_2, x_3, x_4\} = \{d_{i,a}, d_{i,a+1}, d_{i,a+2}, d_{i,a+3}\} = e_s \cup e_{s+1} = V(H_{s+1})$, as claimed.

Case 2: no vertex of the form $d_{i,c}$ ($c \in [3]$) appears among $x_1, \dots, x_4$. Then, $x_1, \dots, x_4$ are all of the form t_i, b_j or m_ℓ. Observe that no pair of the form $t_it_{i'}$ or $b_jb_{j'}$ appears simultaneously in an edge of G_s ($i \neq i', j \neq j'$), so $X = \{x_1, \dots, x_4\}$ contains at most one vertex of the form t_i and at most one vertex of the form b_j. So it must contain at least two vertices of the form m_ℓ. But m_ℓ and $m_{\ell'}$ appear simultaneously in an edge only if $|\ell - \ell'| \le 1$. It follows that X must be of the form $\{t_i, b_j, m_\ell, m_{\ell+1}\}$ for some i, j, ℓ. Assume that $j > 0$ (the case $j < 0$ is similar). If $\ell \le -j$, then neither $t_ib_jm_\ell$ nor $b_jm_\ell m_{\ell+1}$ appear in G_s (see (c), (e) and (8.1)), giving a contradiction. Similarly, if $\ell \ge n + j$, then neither $t_ib_jm_{\ell+1}$ nor $b_jm_\ell m_{\ell+1}$ appear in G_s, again giving a contradiction. Hence, we have $-(j-1) \le \ell \le n + j - 1$. It follows that $b_jm_\ell m_{\ell+1}$ is an edge of $G_0 - e_0$. So e is one of $t_ib_jm_\ell$, $t_ib_jm_{\ell+1}$ and $t_im_\ell m_{\ell+1}$.

First, consider the case $e = t_ib_jm_\ell$. Since $t_ib_jm_{\ell+1}$ is already present, we must have $\ell = n+j-1$ (see (e)). But if $t_im_\ell m_{\ell+1} = t_im_{n+j-1}m_{n+j} = e_N$ appears in G_s, then so does its preceding edge $e_{N-1} = t_ib_jm_{n+j-1} = t_ib_jm_\ell$ (see (8.1)), giving a contradiction.

If the new edge is $e = t_i b_j m_{\ell+1}$, then $-(j-1) \leq \ell \leq n+j-2$. So we have $e = e_N$ for some $N \in [T]$, and the edge e_{N-1} is either $t_i b_j m_\ell$ (if $\ell \neq -(j-1)$) or $t_i m_{-(j-2)} m_{-(j-1)} = t_i m_\ell m_{\ell+1}$ (if $\ell = -(j-1)$). In either case, we have $e_{N-1} \in E(G_s)$ and $e_N \notin E(G_s)$, giving $s = N-1$, $e = e_{s+1}$, $\{x_1, x_2, x_3, x_4\} = e_s \cup e_{s+1} = V(H_{s+1})$, as claimed.

Finally, consider the case when the new edge is $e = t_i m_\ell m_{\ell+1}$. So $t_i b_j m_\ell$ and $t_i b_j m_{\ell+1}$ are edges of G_s. Note that $t_i b_j m_\ell$ or $t_i b_j m_{\ell+1}$ is of the form e_N for some $N \in [T]$. It follows that all edges $e_{N'}$ with $N' < N$ appear in G_s, so, in particular, $t_i m_{\ell'} m_{\ell'+1}$ is in G_s for all $-(j-1) \leq \ell' \leq n+j-2$. Hence $\ell = n+j-1$. But then there is some $M \in [T]$ such that $e_M = t_i m_\ell m_{\ell+1}$, and we have $e_{M-1} = t_i b_j m_{n+j-1} = t_i b_j m_\ell$, which appears in G_s. If follows that $s = M-1$, $e = e_{s+1}$ and $\{x_1, x_2, x_3, x_4\} = e_s \cup e_{s+1} = V(H_{s+1})$, as claimed. □

It is straightforward to check that for all $s \in [T]$ we have $E(H_s) \setminus E(G_{s-1}) = \{e_s\}$ and $E(H_s) \cap \{e_0, e_1, \ldots, e_s\} = \{e_{s-1}, e_s\}$. Using these observations and the claim above, we see that all conditions of being $K_4^{(3)}$-civilised are satisfied for G_0, and the result follows from Lemma 8.2.1. □

Remark 8.2.3. It immediately follows from the construction and the proof above that our proposed initial infection has $9n + O(1)$ vertices and that the infection process takes $4n^3 + O(n^2)$ steps. It therefore follows that $M_4^3(n) \geq 4n^3/9^3 + O(n^2)$. We note that we have made no effort to optimise the leading constant.

8.3 Long running times for double infections

8.3.1 Double infections for $K_4^{(3)}$

We now move on to the proof of our results about the variant where we allow two edges to be infected at the same time if they together complete a copy of H. We begin with Theorem 8.1.4, giving tight bounds in the case $H = K_4^{(3)}$. Our approach is motivated by the proof of Bollobás, Przykucki, Riordan, and Sahasrabudhe [28] of the fact that $M_4(3) = n - 3$, but both the construction and the proof of the upper bound are significantly more complicated here. We start with an informal description of the infection process for the extremal construction.

We will construct the initially infected hypergraph inductively. Assume that for some n we have already constructed a hypergraph G_0 on n vertices $\{x_1, \ldots, x_n\}$ for which the process runs for T steps. Furthermore, assume that G_T is complete, but there exist two vertices $u, v \in \{x_1, \ldots, x_n\}$ such that no edge of G_{T-1} contains u and v simultaneously. (These conditions might at first seem arbitrary, but they are satisfied in the obvious construction when $n = 4$.) Then, we can add another vertex x_{n+1} and another edge $x_{n+1}uv$ to G_0 without changing the first T steps of the infection process. (Indeed, the process will only be altered if we create a new 2-completable copy of $K_4^{(3)}$, and this requires having two edges sharing two vertices.) Moreover, at time $T+1$, the new edges that become infected are all those of the form $x_{n+1}ws$ with $w \in W_1 = \{u, v\}$ and

$s \in \{x_1, \ldots, x_n\} \setminus W_1$. Finally, at time $T+2$, the vertices $x_1, \ldots, x_{n+1}$ will form a complete hypergraph. This gives a construction on $n+1$ vertices with running time $T+2$.

To obtain our construction for $n+2$ vertices, notice that if we pick some $u_2 \in \{x_1, \ldots, x_n\} \setminus W_1$, then no edge of G_T contains both x_{n+1} and u_2. So if we add a new vertex x_{n+2} and a new edge $x_{n+2}x_{n+1}u_2$, then the first $T+1$ steps of the infection will remain unaffected by this change. Furthermore, one can check that at time $T+2$ the edges containing x_{n+2} are given as $x_{n+2}x_{n+1}w$ and $x_{n+2}u_2w$ with $w \in W_1$. Moreover, at time $T+3$ the edges $x_{n+2}ws$ with $w \in W_2 = W_1 \cup \{u_2, x_{n+1}\}$ and $s \in \{x_1, \ldots, x_{n+1}\} \setminus W_2$ will become infected, and at time $T+4$ we get a complete hypergraph.

We can keep repeating these steps: take some $u_j \in \{x_1, \ldots, x_n\} \setminus W_{j-1}$, add a new vertex x_{n+j} and a new edge $x_{n+j}x_{n+j-1}u_j$. This will extend the process by 2 steps, and at time $T+2j-1$ the edges containing the new vertex x_{n+j} will be of the form $x_{n+j}ws$ with $w \in W_j = W_{j-1} \cup \{u_j, x_{n+j-1}\}$ and $s \in \{x_1, \ldots, x_{n+j-1}\}$. Moreover, at time $T+2j$ our hypergraph will contain all edges on $\{x_1, \ldots, x_{n+j}\}$. This means that we can keep adding a vertex and extending the process by 2 steps each time. However, the set W_j is growing, and at some point it will contain all of our vertices. When this happens, we will no longer be able to pick an appropriate u_j, and we will 'lose' 1 step of the infection process (i.e., by adding a new vertex we can only extend the process by 1 step at this point). So 'usually' adding a vertex extends the infection by 2 steps, giving the leading term $2n$ for the running time, but sometimes (when W becomes everything) we only gain one extra step, and this will contribute the term $-\lfloor \log_2(n-2) \rfloor$.

Let us now start the formal construction. Let G_0 be any 3-uniform hypergraph on some vertex set V, and let $G_0, G_1, \ldots$ be the corresponding $(K_4^{(3)}, 2)$-process. Let T be the running time of this process. We say that G_0 is *nice* if $T \neq 0$, G_T is complete, and there exist distinct vertices $u, v \in V$ such that no edge of G_{T-1} contains both u and v. The following lemma will be used to obtain the lower bound.

Lemma 8.3.1. *Suppose that there is a nice hypergraph on $k \geq 4$ vertices such that the corresponding $(K_4^{(3)}, 2)$-process has running time T. Then, for all $\ell \in [k+1, 2k-3]$ we have*

$$M_{(K_4^{(3)},2)}(\ell) \geq T + 2(\ell - k).$$

Furthermore,

$$M_{(K_4^{(3)},2)}(2k-2) \geq T + 2k - 5,$$

and there exists a nice hypergraph on $2k-2$ vertices whose corresponding $(K_4^{(3)}, 2)$-process has running time $T + 2k - 5$.

Proof. Let G_0 be a nice 3-uniform hypergraph on k vertices $x_1, \ldots, x_k$ such that the corresponding $(K_4^{(3)}, 2)$-process $G_0, G_1, \ldots$ has running time T, G_T is complete, and x_1, x_k do not appear in any edge of G_{T-1} simultaneously. Let $\ell \in [k+1, 2k-2]$ be arbitrary. We define a hypergraph G'_0 on a vertex set $\{x_1, \ldots, x_\ell\}$ of size ℓ as follows. For any $i \in [\ell - k]$, let

$$e_i = x_{k+i}x_{k+i-1}x_i,$$

and let $\mathcal{E} = \{e_i : i \in [\ell - k]\}$. Then, set

$$E(G'_0) = E(G_0) \cup \mathcal{E}.$$

Let $G'_0, G'_1, \ldots$ be the corresponding $(K_4^{(3)}, 2)$-bootstrap percolation process with initial infection G'_0, and let us write $W_j = \{x_1, \ldots, x_j, x_k, x_{k+1}, \ldots, x_{k+j-1}\}$ for all $j \in [\ell - k]$.

Claim 8.3.2. We have

$$E(G'_t) = \begin{cases} E(G_t) \cup \mathcal{E} & \text{if } t \in [T], \\ \binom{\{x_1, \ldots, x_{k+j-1}\}}{3} \cup \mathcal{E} \cup F_j^{\text{odd}} & \text{if } t = T + 2j - 1 \text{ with } j \in [\ell - k], \\ \binom{\{x_1, \ldots, x_{k+j}\}}{3} \cup \mathcal{E} \cup F_j^{\text{even}} & \text{if } t = T + 2j \text{ with } j \in [\ell - k], \end{cases}$$

where

$$F_j^{\text{odd}} = \{x_{k+j}wx_a : w \in W_j, a \in [k+j-1], x_a \neq w\}$$

and

$$F_j^{\text{even}} = \{x_{k+j+1}wx_a : w \in W_j, a \in \{j+1, k+j\}\},$$

unless $j = \ell - k$, in which case $F_{\ell-k}^{\text{even}} = \emptyset$.

Proof. We show this statement by induction on t. The case $t \in [T]$ is straightforward. Indeed, note that $|e_i \cap e_j| < 2$ for all distinct $i, j \in [\ell - k]$, and $|e_i \cap f| < 2$ whenever $f \in E(G_{T-1})$ (by our assumption that G_{T-1} does not contain any triple containing both x_1 and x_k). Recall that, if a copy H of $K_4^{(3)}$ is 2-completable in a hypergraph G, then G contains two edges of H, which must share two vertices. Thus, in the first T steps of the process, the addition of the edges in $\mathcal{E}$ does not result in any infections that did not occur for G_0. Now assume that $t > T$ and the statement above holds for $t - 1$. For notational purposes, set $F_0^{\text{even}} = \emptyset$. We split the analysis into two cases.

Case 1. Consider first the case $t = T + 2j - 1$ (with $j \in [\ell - k]$). By the induction hypothesis,

$$E(G'_{t-1}) = \binom{\{x_1, \ldots, x_{k+j-1}\}}{3} \cup \mathcal{E} \cup F_{j-1}^{\text{even}}.$$

We first verify that $F_j^{\text{odd}} \subseteq E(G'_t)$. Let $w \in W_j$ and $a \in [k+j-1]$ with $x_a \neq w$, so that $x_{k+j}wx_a \in F_j^{\text{odd}}$. Then, there exists some $w' \in W_j$ with $x_{k+j}ww' \in E(G'_{t-1})$. (Indeed, if we pick

w' such that $w' \in \{x_j, x_{k+j-1}\}$ and $w' \neq w$, then $x_{k+j}ww' \in F^{\text{even}}_{j-1} \cup \{e_j\}$.) If $x_a = w'$, then trivially $x_{k+j}wx_a \in E(G'_t)$. If $x_a \neq w'$, then we also have $ww'x_a \in E(G'_{t-1})$. It follows that the copy of $K_4^{(3)}$ with vertex set $\{x_{k+j}, w, w', x_a\}$ is 2-completable in G'_{t-1}, and so $x_{k+j}wx_a \in E(G'_t)$.

We next show that any edge infected at time t must belong to F^{odd}_j. Indeed, if $h \in [j+1, \ell-k]$ then $|e_h \cap f| < 2$ for all $f \in E(G'_{t-1}) \setminus \{e_h\}$, so e_h cannot appear in a copy of $K_4^{(3)}$ completed in this step. It follows that any added edge must be of the form $e = x_{k+j}x_ax_b$ with $a, b \in [k+j-1]$ distinct. Furthermore, either x_a or x_b must appear together with x_{k+j} in an edge of $E(G'_{t-1}) \setminus \{e_i : i \in [j+1, \ell-k]\}$, so one of x_a, x_b must belong to $W_{j-1} \cup \{x_j, x_{k+j-1}\} = W_j$. But then $e \in F^{\text{odd}}_j$, as claimed.

Case 2. Consider now the case $t = T + 2j$ (with $j \in [\ell - k]$). By induction, we know

$$E(G'_{t-1}) = \binom{\{x_1, \ldots, x_{k+j-1}\}}{3} \cup \mathcal{E} \cup F^{\text{odd}}_j.$$

Observe first that, whenever $a \in [k + j - 2]$, we have $x_{k+j}x_{k+j-1}x_a \in F^{\text{odd}}_j \subseteq E(G'_{t-1})$. It follows that, whenever $a, b \in [k + j - 2]$ are distinct, the copy of $K_4^{(3)}$ with vertex set $\{x_{k+j}, x_{k+j-1}, x_a, x_b\}$ is 2-completable in G'_{t-1}. Hence, $x_{k+j}x_cx_d \in E(G'_t)$ whenever $c, d \in [k + j-1]$ (distinct). Thus, $\binom{\{x_1,\ldots,x_{k+j}\}}{3} \subseteq E(G'_t)$. Furthermore, assume that $j \neq \ell - k$ and $e \in F^{\text{even}}_j$, so $e = x_{k+j+1}wx_a$ with $w \in W_j$ and $a \in \{j+1, k+j\}$. Then, $e_{j+1} = x_{k+j+1}x_{k+j}x_{j+1} \in E(G'_{t-1})$ and $x_{k+j}wx_{j+1} \in F^{\text{odd}}_j \subseteq E(G'_{t-1})$, so the copy of $K_4^{(3)}$ with vertex set $\{x_{k+j+1}, x_{k+j}, x_{j+1}, w\}$ is 2-completable in G'_{t-1}, which implies $x_{k+j+1}wx_a \in E(G'_t)$. So $F^{\text{even}}_j \subseteq E(G'_t)$.

It remains to show that any edge added in this step must belong to $\binom{\{x_1,\ldots,x_{k+j}\}}{3} \cup F^{\text{even}}_j$. Indeed, as in the previous case, we see that any copy of $K_4^{(3)}$ which is 2-completable in G'_{t-1} must have vertex set $\{x_a, x_b, x_c, x_d\}$ with $a, b, c, d \in [k+j+1]$ (distinct). So any edge which is infected at time t is of the form $e = x_ax_bx_c$ with $a, b, c \in [k + j + 1]$. If $a, b, c \in [k + j]$, the containment holds trivially, so we may assume that $c = k + j + 1$. In order for e to become infected at time t, we must have that $x_{k+j+1}x_ax_d$ or $x_{k+j+1}x_bx_d$ appears in $E(G'_{t-1})$; we may assume that it is the former. But this edge must be $e_{j+1} = x_{k+j+1}x_{k+j}x_{j+1}$, and hence $\{a, d\} = \{j + 1, k + j\}$. It also follows that $x_{k+j+1}x_bx_d \notin E(G'_{t-1})$, and hence $x_ax_bx_d \in E(G'_{t-1})$, i.e., $x_{k+j}x_bx_{j+1} \in E(G'_{t-1})$. This implies $x_{k+j}x_bx_{j+1} \in F^{\text{odd}}_j$ and, therefore, $x_b \in W_j$. So $c = k + j + 1$, $a \in \{j + 1, k + j\}$ and $x_b \in W_j$, hence $x_ax_bx_c \in F^{\text{even}}_j$, as claimed. □

By the claim above, $G'_{T+2(\ell-k)}$ is complete, but $G'_{T+2(\ell-k)-1}$ is not unless $\ell = 2k - 2$ (indeed, if $\ell \neq 2k - 2$, then $x_{k-2}x_{k-1}x_\ell \notin E(G'_{T+2(\ell-k)-1})$). Moreover, if $\ell = 2k - 2$, then $G'_{T+2(\ell-k-1)} = G'_{T+2k-6}$ does not contain an edge in which both x_{2k-2} and x_{k-1} appear. The statement of the lemma follows. □

We are ready to deduce the lower bound.

Lemma 8.3.3. *For all $n \geq 4$ we have $M_{(K_4^{(3)},2)}(n) \geq 2n - \lfloor \log_2(n - 2) \rfloor - 6$.*

Proof. Observe first that there is a nice 3-uniform hypergraph G_0 on 4 vertices $\{x_1, x_2, x_3, x_4\}$, given by $E(G_0) = \{x_1x_2x_3, x_2x_3x_4\}$, for which the running time of the $(K_4^{(3)}, 2)$-process is $T = 1$. A straightforward induction using Lemma 8.3.1 shows that for all $m \geq 1$ there exists a nice hypergraph on 2^m+2 vertices for which the running time of the $(K_4^{(3)}, 2)$-process is $2^{m+1}-(m+2)$. Furthermore, also by Lemma 8.3.1, whenever $2^m + 2 \leq n < 2^{m+1} + 2$ we have

$$M_{(K_4^{(3)},2)}(n) \geq 2^{m+1} - (m+2) + 2(n - 2^m - 2) = 2n - m - 6 = 2n - \lfloor \log_2(n-2) \rfloor - 6,$$

as claimed. $\square$

We now turn to the proof of the upper bound. For any $t \geq 1$, let $m = m(t)$ denote the unique positive integer which satisfies

$$2^{m+1} - (m+2) \leq t < 2^{m+2} - (m+3).$$

The following key lemma essentially shows that the infections must contain a substructure similar to the one in our construction.

Lemma 8.3.4. *Let G_0 be a 3-uniform hypergraph on $n \geq 4$ vertices, and consider the $(K_4^{(3)}, 2)$-process $G_0, G_1, \ldots$ with G_0 as initial infection. Assume that $a \geq 1$ and $e \in E(G_a) \setminus E(G_{a-1})$. Then, there exist some $t_e \geq a$, $S_e \subseteq V(G_0)$, $v_e \in S_e$ and $W_e \subseteq S_e \setminus \{v_e\}$ such that, for $m = m(t_e)$,*

(P1) $e \subseteq S_e$,

(P2) $|S_e| = (t_e + m + 6)/2$,

(P3) $|W_e| = t_e - (2^{m+1} - (m+2))$,

(P4) $G_{t_e}[S_e]$ *is complete, and*

(P5) $\binom{S_e \setminus \{v_e\}}{3} \cup \{v_e w s : w \in W_e, s \in S_e \setminus \{v_e, w\}\} \subseteq E(G_{t_e - 1})$.

Proof. We prove the statement by induction on a. If $a = 1$, then we know e is in some copy H of $K_4^{(3)}$ in G_1. We can set $t_e = 1$ (so $m = 1$), $S_e = V(H)$, $v_e \in V(H)$ such that $V(H) \setminus \{v_e\} \in E(G_0)$, and $W_e = \emptyset$; the properties (P1)–(P5) are then satisfied.

Now assume that $a \geq 2$ and the statement holds for smaller values of a. We know there is some copy H of $K_4^{(3)}$ in G_a such that $e \in E(H)$, $|E(H) \cap E(G_{a-1})| \geq 2$ and $|E(H) \cap E(G_{a-2})| < 2$. It follows that there is some $f \in E(H)$ such that $f \in E(G_{a-1}) \setminus E(G_{a-2})$, i.e., f is infected at time $a-1$. Furthermore, there is another edge $f' \in E(H) \cap E(G_{a-1})$ (with $f' \neq f$). Let us write $t = t_f$, $S = S_f$, $W = W_f$ and $v = v_f$, and let $m = m(t)$. We consider several cases according to how e, f and f' overlap with S (and v, W). Note that, if $e \not\subseteq S$, then there is some $p \in V(G_0)$ such that $e \setminus f = e \setminus S = \{p\}$ and $p \in f'$.

Case 1: $e \subseteq S$. Since $e \notin E(G_{a-1})$ and $G_t[S]$ is complete, we have $t \geq a$. It follows that $t_e = t$, $S_e = S$, $v_e = v$, $W_e = W$ satisfy properties (P1)–(P5).

Case 2: $e \not\subseteq S$, $f' \in E(G_{t-1})$, and $f' = pss'$ for some $s, s' \in S \setminus W$ (recall $\{p\} = e \setminus S$). By (P5) for f, we know that, whenever $w \in W$, we have $wss' \in E(G_{t-1})$. This, together with the fact that $f' = pss' \in E(G_{t-1})$, guarantees that

$$pws, pws' \in E(G_t). \tag{8.4}$$

Let $u \in W \cup \{s, s'\}$ and $u' \in \{s, s'\}$ with $u' \neq u$, and let $z \in S \setminus \{s, s'\}$ with $z \neq u$. Then, (8.4) and our assumption on f' tell us that $puu' \in E(G_t)$, and by (P4) for f we have that $uu'z \in E(G_t)$. This implies $puz \in E(G_{t+1})$. That is, for all $u \in W \cup \{s, s'\}$ and all $z \in S \setminus \{u\}$ we have $puz \in E(G_{t+1})$. This in turn implies that $G_{t+2}[S \cup \{p\}]$ is complete (as $psz, psz' \in E(G_{t+1})$ implies $pzz' \in E(G_{t+2})$).

If $t = 2^{m+2} - (m+4)$, then $|W| = 2^{m+1} - 2$ and $|S| = 2^{m+1} + 1$, so we have $|W \cup \{s, s'\}| = |S| - 1$ and hence $G_{t+1}[S \cup \{p\}]$ is complete by the observation above. Hence, $t_e = t+1$, $S_e = S \cup \{p\}$, $v_e = p$, $W_e = \emptyset$ satisfy properties (P1)–(P5) (note that in this case $m(t_e) = m+1$).

On the other hand, if $t \neq 2^{m+2} - (m+4)$, then $t \leq 2^{m+2} - (m+6)$ (since $t + m = 2|S| - 6$ is even by (P2)). So $m(t+2) = m(t)$. It follows that $t_e = t+2$, $S_e = S \cup \{p\}$, $v_e = p$, $W_e = W \cup \{s, s'\}$ satisfy the properties.

Case 3: $e \not\subseteq S$, $f' \in E(G_{t-1})$, and $f' = pws$ for some $w \in W$, $s \in S$, $\{p\} = e \setminus S$. Then, by (P5) for f, whenever $z \in S \setminus \{w, s\}$ we have $wsz \in E(G_{t-1})$. Since $f' = pws \in E(G_{t-1})$, it follows that $pwz \in E(G_t)$ for all $z \in S \setminus \{w\}$. Therefore, whenever $z, z' \in S \setminus \{w\}$ are distinct, we have $pwz, pwz' \in E(G_t)$, and hence $pzz' \in E(G_{t+1})$. Thus, $G_{t+1}[S \cup \{p\}]$ is complete.

If $t = 2^{m+2} - (m+4)$, then $t_e = t+1$, $S_e = S \cup \{p\}$, $v_e = p$, $W_e = \emptyset$ satisfy properties (P1)–(P5).

On the other hand, assume that $t < 2^{m+2} - (m+4)$ (as in case 2, we then have $m(t+2) = m(t)$). Then, (P2) and (P3) imply that $|W| < |S| - 3$. Let W' be an arbitrary subset of S of size $|W| + 2$. Then, $t_e = t+2$, $S_e = S \cup \{p\}$, $v_e = p$, $W_e = W'$ satisfy properties (P1)–(P5).

Case 4: $e \not\subseteq S$, $f' \notin E(G_{t-1})$. Then, we must have $t = a - 1$ and $f' \in E(G_t) \setminus E(G_{t-1})$. We may assume that $t_{f'} = t$, since otherwise we can swap the roles of f and f' and we are done by the previous cases. Let us write S' for $S_{f'}$. Note that $|S| = |S'|$, and $S \cap S' \supseteq f \cap f'$ has size at least 2.

Assume first that $S' \setminus S = \{p\}$ for some $p \in V(G_0)$ (where necessarily $\{p\} = e \setminus f$). Then, $S \setminus S' = \{q\}$ for some $q \in V(G_0)$ (with $\{q\} = e \setminus f'$). Observe that, by (P4), whenever $s, s' \in S \cap S'$ are distinct, we have $pss' \in E(G_t)$ and $qss' \in E(G_t)$. This implies that $pqs \in E(G_{t+1})$ for every $s \in S \cap S'$. Hence, $G_{t+1}[S \cup S']$ is complete, where $|S \cup S'| = |S| + 1$. If $t = 2^{m+2} - (m+4)$, then properties (P1)–(P5) are satisfied for $t_e = t+1$, $S_e = S \cup S'$, $v_e = p$ and $W_e = \emptyset$. Otherwise, $t \leq 2^{m+2} - (m+6)$ (as $t + m$ is even), so (P1)–(P5) are satisfied for $t_e = t+2$, $S_e = S \cup S'$,

$v_e = p$, and W_e an arbitrary subset of S of size $|W| + 2$.

Observe that, whenever $x, y \in S \cap S'$ (distinct), $s \in S \setminus S'$ and $s' \in S' \setminus S$, by (P4) we know $xys, xys' \in E(G_t)$, and therefore $xss' \in E(G_{t+1})$. By the same argument, if $\bar{s}' \in S' \setminus S$ (with $\bar{s}' \neq s'$), then $xs\bar{s}' \in E(G_{t+1})$. This then implies that $ss'\bar{s}' \in E(G_{t+2})$. Similarly, if $s' \in S' \setminus S$ and $s, \bar{s} \in S \setminus S'$ (distinct), then $s\bar{s}s' \in E(G_{t+2})$. Hence, $G_{t+2}[S \cup S']$ is complete, where $|S \cup S'| \geq |S| + 2$. Now pick any two vertices $p, p' \in S' \setminus S$ with $e \subseteq S \cup \{p, p'\}$. If $t = 2^{m+2} - (m+6)$, then let $t_e = t+3$, $S_e = S \cup \{p, p'\}$, $v_e = p$ and $W_e = \emptyset$. If $t = 2^{m+2} - (m+4)$, then let $t_e = t+3$, $S_e = S \cup \{p, p'\}$, $v_e = p$ and W_e an arbitrary subset of $S \cap S'$ of size 2. Finally, assume $t < 2^{m+2} - (m+6)$. Since $t+m$ is even, we have $t \leq 2^{m+2} - (m+8)$. Then, let $t_e = t+4$, $S_e = S \cup \{p, p'\}$, $v_e = p$ and W_e an arbitrary subset of S of size $|W| + 4$. These choices satisfy properties (P1)–(P5). This finishes the proof of the lemma. $\square$

Proof of Theorem 8.1.4. The lower bound follows from Lemma 8.3.3. For the upper bound, given an arbitrary hypergraph G_0 on a vertex set V of size $n \geq 4$, consider the $(K_4^{(3)}, 2)$-process $G_0, G_1, \dots$ with G_0 as initial infection. Let $T = M_{(K_4^{(3)},2)}(G_0)$ be the running time of the process, and let $e \in E(G_T) \setminus E(G_{T-1})$ be arbitrary. By Lemma 8.3.4, there is some $t \geq T$ such that G_t contains a clique of size $\frac{t+m(t)+6}{2} \geq \frac{T+m(T)+6}{2}$. Hence,

$$\frac{T + m(T) + 6}{2} \leq n.$$

It follows that

$$T \leq 2n - \lfloor \log_2(n-2) \rfloor - 6,$$

as we wanted to prove. Indeed, if $\lfloor \log_2(n-2) \rfloor = \alpha$, then $2^\alpha + 2 \leq n \leq 2^{\alpha+1} + 1$ and hence $2^{\alpha+1} - (\alpha+2) \leq 2n - \alpha - 6 < 2^{\alpha+2} - (\alpha+3)$, giving $m(2n-\alpha-6) = \alpha$ and $\frac{(2n-\alpha-6)+m(2n-\alpha-6)+6}{2} = n$. $\square$

8.3.2 Double infections for $K_5^{(3)}$

We now turn our attention to Theorem 8.1.5. In order to prove it, observe that, for an r-uniform hypergraph H, the trivial upper bound $M_{(H,k)}(n) \leq \binom{n}{r}$ still holds, so it suffices to provide a lower bound. We will proceed by constructing an initial infection for which the $(K_5^{(3)}, 2)$-bootstrap percolation process runs for a cubic number of steps. At most two edges will become infected in each step of the infection process, which will make it easier to analyse the number of steps. Our construction is intuitively similar to the one we constructed for the proof of Theorem 8.1.3, albeit a bit more convoluted. Let us begin with an intuitive description.

Our vertex set will again be split into three layers, with vertices t_i playing the role of 'top' vertices, vertices b_j playing the role of 'bottom' vertices, and vertices m_ℓ conforming the 'middle' layer. We will also have a number of 'dummy' vertices. For fixed top and bottom vertices, the vertices in the middle layer will allow us to infect two edges at a time, while traversing this layer, for a linear total number of steps. For each fixed bottom vertex, there will be some extra edges at the end of the middle layer which will allow us to swap the bottom vertex for the next one and continue the process. Finally, the dummy vertices will allow us to swap the top vertex and start the process anew.

To be more precise, let us describe the first few stages of the infection process. For each $0 \leq \ell \leq 2n-2$, our initial infection will contain all edges of the copy of $K_5^{(3)}$ with vertex set $t_1, b_1, m_\ell, m_{\ell+1}, m_{\ell+2}$ except for $t_1 m_\ell b_1$, $t_1 m_{\ell+1} b_1$, and $t_1 m_{\ell+2} b_1$. It will also contain $t_1 m_0 b_1$. This edge will trigger the infection of $t_1 m_1 b_1$ and $t_1 m_2 b_1$ in the first step of the process, then of $t_1 m_3 b_1$ and $t_1 m_4 b_1$, and the infection will keep propagating towards higher values of ℓ, until finally, after n steps, the edges $t_1 m_{2n-1} b_1$ and $t_1 m_{2n} b_1$ become infected.

At this point, we will swap out b_1 to b_{-1}. This can be achieved in two steps of the infection process. Our initial infection will already contain all edges of the copy of $K_5^{(3)}$ with vertex set $t_1, b_1, m_{2n}, m_{2n+1}, m_{2n+2}$ except for $t_1 m_{2n} m_{2n+1}$, $t_1 m_{2n+1} m_{2n+2}$, and the edge $t_1 m_{2n} b_1$ which was just added in the previous step; $t_1 m_{2n} m_{2n+1}$ and $t_1 m_{2n+1} m_{2n+2}$ will therefore become infected in the next step. The initial infection also contains all the edges of the copy of $K_5^{(3)}$ with vertex set $t_1, b_{-1}, m_{2n}, m_{2n+1}, m_{2n+2}$ except for $t_1 m_{2n} b_{-1}$, $t_1 m_{2n+1} b_{-1}$, and the two that were just added. These two edges now become infected as well, and start a new infection process where now the indices decrease through the middle layer.

Finally, suppose we have reached a point where the copy of $K_5^{(3)}$ defined on the vertices t_1, b_n, m_{4n-4}, m_{4n-3}, m_{4n-2} has been completely infected, with the edges infected in the last step of the process being $t_1 m_{4n-3} b_n$ and $t_1 m_{4n-2} b_n$. We now want to swap out the top vertex to t_2, using for this purpose four dummy vertices. Similarly as in the proof of Theorem 8.1.3, these will simply create a short chain of infections that allows us to restart the process.

We now give a formal proof.

Proof of Theorem 8.1.5. Consider an initial infection hypergraph G_0 whose vertex set consists of $12n-5$ vertices, labelled as $t_1, \ldots, t_n$, $b_1, \ldots, b_n$, $b_{-1}, \ldots, b_{-(n-1)}$, $m_{-2(n-1)}, \ldots, m_{4n}$, and $d_{i,1}, d_{i,2}, d_{i,3}$ for $i \in [n-1]$. For notational purposes, for each $i \in [n]$ let $d_{i,-2} = t_i$, $d_{i,-1} = b_n$, $d_{i,0} = m_{4n-2}$, $d_{i,4} = m_0$, $d_{i,5} = b_1$, and $d_{i,6} = t_{i+1}$. The edges of G_0 appear in the following list:

(a) $t_1 m_0 b_1$;

(b) $m_\ell m_{\ell+1} m_{\ell+2}$, for all $-2(n-1) \leq \ell \leq 4n-4$;

(c) $t_i m_\ell m_{\ell+2}$, for all $i \in [n]$ and $-2(n-1) \leq \ell \leq 4n-4$;

(d) $b_j m_\ell m_{\ell+2}$, for all $j \in [n]$ and $-2(j-1) \leq \ell \leq 2(n+j-1)$;

(e) $b_{-j} m_\ell m_{\ell+2}$, for all $j \in [n-1]$ and $-2j \leq \ell \leq 2(n+j-1)$;

(f) $t_i m_\ell m_{\ell+1}$, for all $i \in [n]$ and $0 \leq \ell \leq 2n-1$;

(g) $b_j m_\ell m_{\ell+1}$, for all $j \in [n]$ and $-2(j-1) \leq \ell \leq 2(n+j)-1$;

(h) $b_{-j} m_\ell m_{\ell+1}$, for all $j \in [n-1]$ and $-2j \leq \ell \leq 2(n+j)-1$;

(i) $t_i m_{2(n+j)-1} b_j$, $t_i m_{2(n+j)} b_j$ and $t_i m_{2(n+j)} b_{-j}$, for all $i \in [n]$ and $j \in [n-1]$;

(j) $t_i m_{-2j+1} b_{-j}$, $t_i m_{-2j} b_{-j}$ and $t_i m_{-2j} b_{j+1}$, for all $i \in [n]$ and $j \in [n-1]$;

(k) $d_{i,j}d_{i,j+1}d_{i,j+3}$, $d_{i,j}d_{i,j+1}d_{i,j+4}$, $d_{i,j}d_{i,j+2}d_{i,j+3}$, $d_{i,j}d_{i,j+2}d_{i,j+4}$, $d_{i,j}d_{i,j+3}d_{i,j+4}$, $d_{i,j+1}d_{i,j+2}d_{i,j+3}$, and $d_{i,j+1}d_{i,j+3}d_{i,j+4}$, for all $i \in [n-1]$ and $j \in \{-2, 0, 2\}$.

To compare this with the construction hinted at before the proof, consider the following. The edge in (a) is an edge e_0 which starts the whole infection process. The edges in (b)–(h) are there to ensure the correct propagation of the infection through the middle layer, where the edges in (d) and (g) will be used to propagate the infection towards larger values of ℓ, using some bottom vertex of the form b_j, and those in (e) and (h) will be used to propagate the infection towards smaller values of ℓ, using some bottom vertex of the form b_{-j}. The edges in (i) and (j) are needed for swapping the bottom vertices. Finally, the edges which appear in (k) are used to swap the top vertices.

For each pair (i, j) with $i, j \in [n]$, let $A_{i,j}$ be the sequence of edges

$$A_{i,j} = (t_i m_{-2(j-1)+\ell} b_j)_{\ell=1}^{2n+4(j-1)}. \tag{8.5}$$

Similarly, for each pair (i, j) with $i \in [n]$ and $j \in [n-1]$, we define

$$A_{i,-j} = (t_i m_{2(n+j)-\ell} b_{-j})_{\ell=1}^{2n+4(j-1)+2}. \tag{8.6}$$

Note that each of these has an even number of elements. Using $\bigtimes$ to denote concatenation, for each phase $i \in [n]$ we define the sequence

$$A_i = A_{i,1} \bigtimes_{j=1}^{n-1} (t_i m_{2(n+j-1)} m_{2(n+j)-1}, t_i m_{2(n+j)-1} m_{2(n+j)}) A_{i,-j} (t_i m_{-2(j-1)} m_{-2j+1}, t_i m_{-2j+1} m_{-2j}) A_{i,j+1}. \tag{8.7}$$

Finally, we set

$$\begin{aligned} A &= A_1 \bigtimes_{i=1}^{n-1} (d_{i,-1}d_{i,0}d_{i,2}, d_{i,0}d_{i,1}d_{i,2}, d_{i,1}d_{i,2}d_{i,4}, d_{i,2}d_{i,3}d_{i,4}, d_{i,3}d_{i,4}d_{i,6}, d_{i,4}d_{i,5}d_{i,6}) A_{i+1} \\ &= A_1 \bigtimes_{i=1}^{n-1} (b_n m_{4n-2} d_{i,2}, m_{4n-2} d_{i,1} d_{i,2}, d_{i,1} d_{i,2} m_0, d_{i,2} d_{i,3} m_0, d_{i,3} m_0 t_{i+1}, m_0 b_1 t_{i+1}) A_{i+1}. \end{aligned} \tag{8.8}$$

We will sometimes abuse notation and treat each of the above sequences as sets. Observe that A has an even number of elements and that none appear repeatedly. We may label these elements as $(e_1, e'_1, e_2, e'_2, \dots, e_T, e'_T)$, for some $T > 0$. Note that $T = 4n^3 + O(n^2)$ by construction. Note, moreover, that by construction we are guaranteed that $|\{e_t, e'_t\} \cap A_{i,j}| \in \{0, 2\}$ for all i and j. Additionally, any two consecutive triples in A share exactly two vertices, thus, it is easy to check that any three consecutive triples span five vertices.

Let $e'_0 = t_1 m_0 b_1$. For each $t \in [T-1]$, let H_t denote the copy of $K_5^{(3)}$ with vertex set $e'_{t-1} \cup e_t \cup e'_t$, and let H'_t denote the copy of $K_5^{(3)}$ with vertex set $e_{t-1} \cup e'_{t-1} \cup e_t$ (if $t > 1$). For each $t \in [T]$, let G_t be the hypergraph with edge-set $E(G_{t-1}) \cup \{e_t, e'_t\}$. We will show that these hypergraphs indeed coincide with those obtained by the $K_5^{(3)}$-bootstrap percolation process with initial infection G_0.

Claim 8.3.5. Let H be a copy of $K_5^{(3)}$ on the vertex set of G_0. Assume that, for some $t \in [0, T-1]$, we have that $H \nsubseteq G_t$ but H is 2-completable in G_t. Then, the following hold.

- If H is 1-completable in G_t, suppose that adding e to G_t completes H. Then, $t \geq 1$, $e = e_{t+1}$ and $H = H'_{t+1}$.
- If H is not 1-completable in G_t, suppose that adding e and e' to G_t completes H. Then, $\{e, e'\} = \{e_{t+1}, e'_{t+1}\}$ and $H = H_{t+1}$.

Proof. Consider any copy H of $K_5^{(3)}$ on $V(G_0)$. If H contains two vertices of the form t_i and $t_{i'}$ with $1 \leq i < i' \leq n$, since G_T does not contain any edge with two 'top' vertices (see (a)–(k) as well as (8.5)–(8.8)), H must be missing at least three edges in G_T. As $G_0 \subseteq G_1 \subseteq \ldots \subseteq G_T$, it follows that H is not 2-completable in G_t for any $t \in [0, T-1]$. The same argument holds if H contains two vertices of the form b_j and $b_{j'}$. Hence, we may assume that H contains at most one vertex t_i and one vertex b_j. Similarly, if H contains two vertices m_ℓ and $m_{\ell'}$ with $|\ell - \ell'| \geq 3$, then G_T does not contain any edge containing both m_ℓ and $m_{\ell'}$, so H is not 2-completable in G_T. Therefore, H contains at most three vertices of the form m_ℓ, and their indices must be within distance two of each other.

Assume first that H contains some vertex of the form $d_{i,c}$ with $i \in [n-1]$ and $c \in [3]$. All triples in G_T containing one such vertex are of the form $d_{i,r} d_{i,p} d_{i,q}$ with $-2 \leq r < p < q \leq 6$ and $q - r \leq 4$ (see (k) and (8.8)). It follows easily that, if $V(H)$ does not consist of five consecutive vertices $d_{i,p}, d_{i,p+1}, \dots, d_{i,p+4}$ with $-2 \leq p \leq 2$, then H cannot be 2-completable in G_T. Moreover, if we assume $V(H) = \{d_{i,p+h} : 0 \leq h \leq 4\}$ for some $i \in [n-1]$ and $p \in \{-1, 1\}$, then we also know that the triples $d_{i,p} d_{i,p+1} d_{i,p+4}$, $d_{i,p} d_{i,p+2} d_{i,p+4}$ and $d_{i,p} d_{i,p+3} d_{i,p+4}$ do not appear in G_T, so again H cannot be 2-completable. So we must have $V(H) = \{d_{i,p+h} : 0 \leq h \leq 4\}$ with $p \in \{-2, 0, 2\}$. But then, by (k) and (8.8), the only three edges missing from H in G_0 are e'_{N-1}, e_N and e'_N, for some $N \in [T]$. Let $t \in [0, T-1]$ be such that H is 2-completable in G_t but $E(H) \nsubseteq E(G_t)$. It is easy to see that we must have $t = N - 1$; furthermore, $H = H_{t+1}$, H is not 1-completable in G_t, and $E(G_{t+1}) \setminus E(H) = \{e_N, e'_N\}$, as desired.

Assume next that H does not contain any vertex of the form $d_{i,c}$ with $i \in [n-1]$ and $c \in [3]$, so it must contain one top vertex t_i, one bottom vertex b_j, and three consecutive middle vertices m_ℓ, $m_{\ell+1}$ and $m_{\ell+2}$. Assume $j > 0$ (the other case can be argued analogously). If $\ell \geq 2(n+j)-1$, then G_T is missing the edges $t_i m_{\ell+2} b_j$, $b_j m_{\ell+1} m_{\ell+2}$ and $b_j m_\ell m_{\ell+2}$ (see (d), (g), (i), (8.5) and (8.7)), so H cannot be 2-completable at any stage of the process. Similarly, if $\ell < -2(j-1)$, then G_T is missing the edges $t_i m_\ell b_j$, $b_j m_\ell m_{\ell+1}$ and $b_j m_\ell m_{\ell+2}$ (see (d), (g), (j), (8.5) and (8.7)), hence H is not 2-completable. Thus, we must have $-2(j-1) \leq \ell \leq 2(n+j-1)$. However, for the case when $j = n$ and $\ell \in \{4n-3, 4n-2\}$, it follows from (b), (c), (f) and (8.7) that G_T is missing the triples $m_\ell m_{\ell+1} m_{\ell+2}$, $t_i m_\ell m_{\ell+2}$ and $t_i m_{\ell+1} m_{\ell+2}$, hence H cannot be 2-completable. So we must have $-2(j-1) \leq \ell \leq 2(n+j-1)$ when $j \in [n-1]$ and $-2(j-1) \leq \ell \leq 2(n+j-2)$ when $j = n$. Let $t \in [0, T-1]$ be such that H is 2-completable in G_t but $E(H_t) \not\subseteq E(G_t)$. We now split the analysis into further cases.

Assume first that $j < n$ and $\ell = 2(n+j-1)$. It follows from (b)–(j) that the only triples of H missing from G_0 are $t_i m_\ell b_j$, $t_i m_\ell m_{\ell+1}$ and $t_i m_{\ell+1} m_{\ell+2}$. These three are added throughout the sequence of edges defined above, as follows from (8.5) and (8.7), as e'_{N-1}, e_N and e'_N, respectively, for some $N \in [T]$. Then, in order for H to be 2-completable in G_t, we must have $e'_{N-1} \in E(G_t)$, and hence $t = N$, H=H_{t+1}, H is not 1-completable, and $E(G_{t+1}) \setminus E(H) = \{e_N, e'_N\}$, as desired.

Consider next the case that $j < n$ and $\ell = 2(n+j-1)-1$. The triples of H missing from G_0 are $t_i m_\ell b_j$, $t_i m_{\ell+1} b_j$ and $t_i m_{\ell+1} m_{\ell+2}$ (see (b)–(j)), which are e_{N-1}, e'_{N-1} and e_N, respectively, for some $N \in [T]$ (see (8.5) and (8.7)). Thus, in order for H to be 2-completable in G_t, this hypergraph must contain at least one of the missing triples; however, since e_{N-1} and e'_{N-1} are added simultaneously in the sequence of hypergraphs, we must have $e_{N-1}, e'_{N-1} \in E(G_t)$, and so H is 1-completable. Then, the only edge that can complete H is $e = e_N$, and so it follows that $N = t+1$ and $H = H'_{t+1}$.

Assume now that $\ell = -2(j-1)$. Here we have two further subcases. If $j = 1$, then the only triples of H missing in G_0 are precisely e_1 and e'_1 (see (a)–(j) as well as (8.5)). Therefore, we have $\{e, e'\} = \{e_1, e'_1\}$ and $H = H_1$. So suppose that $j \geq 2$. Then, the triples of H missing in G_0 are $t_i m_\ell m_{\ell+1}$, $t_i m_{\ell+1} m_{\ell+2}$, $t_i m_{\ell+1} b_j$ and $t_i m_{\ell+2} b_j$ (see (b)–(j)). But then it follows from (8.5) and (8.7) that these triples take the form e_{N-1}, e'_{N-1}, e_N, e'_N, for some $N \in [T]$. In order for H to be 2-completable in G_t, we must have $e_{N-1}, e'_{N-1} \in E(G_t)$. Then, it follows that $t = N-1$, $H = H_{t+1}$, H is not 1-completable, and $E(G_{t+1}) \setminus E(H) = \{e_N, e'_N\}$.

Suppose now that $\ell = -2(j-1)+1$. Again, we must consider two subcases. If $j = 1$, the edges of H missing in G_0 are $t_i m_\ell b_j$, $t_i m_{\ell+1} b_j$ and $t_i m_{\ell+2} b_j$, which correspond to e_1, e'_1 and e_2 (see (b)–(j) as well as (8.5)). Thus, in order for H to be 2-completable in G_t we must have $e_1, e'_1 \in E(G_t)$, so $t = 1$ and H is 1-completable in G_1. It then follows that $e = e_2$, and $H = H'_2$. So suppose $j \geq 2$. Then, the triples of H missing in G_0 are $t_i m_\ell m_{\ell+1}$, $t_i m_\ell b_j$, $t_i m_{\ell+1} b_j$ and $t_i m_{\ell+2} b_j$ (see (b)–(j)). By (8.5) and (8.7), these triples take the form e'_{N-2}, e_{N-1}, e'_{N-1}, e_N, for some $N \in [2, T]$. In order for H to be 2-completable in G_t, at least two of these edges must be added. But e_{N-1}

and e'_{N-1} are added simultaneously, so we conclude that $e'_{N-2}, e_{N-1}, e'_{N-1} \in E(G_t)$ and H is 1-completable in G_t. Therefore, $E(G_{t+1}) \setminus E(H) = \{e_N\}$, $N = t+1$ and $H = H'_{t+1}$.

Suppose finally that $-2(j-2) \leq \ell \leq 2(n+j-2)$. Then, the only edges of H missing in G_0 are $t_i m_\ell b_j$, $t_i m_{\ell+1} b_j$ and $t_i m_{\ell+2} b_j$ (see (b)–(j)). If ℓ is even, it follows from (8.5) that these edges take the form e'_{N-1}, e_N and e'_N, respectively, for some $N \in [T]$; on the contrary, if ℓ is odd, then they take the form e_{N-1}, e'_{N-1} and e_N. In the former case, in order for H to be 2-completable in G_t, we must have $e'_{N-1} \in E(G_t)$, and it follows that $E(G_{t+1}) \setminus E(H) = \{e_N, e'_N\}$, $N = t+1$ and $H = H_{t+1}$. In the latter case, we must have $e_{N-1}, e'_{N-1} \in E(G_t)$, so H is 1-completable, and it follows that $E(G_{t+1}) \setminus E(H) = \{e_N\}$, $N = t+1$ and $H = H'_{t+1}$. $\square$

By applying Claim 8.3.5 iteratively, we conclude that the $(K_5^{(3)}, 2)$-bootstrap percolation process with initial infection G_0 indeed generates the sequence of hypergraphs $G_0, G_1, \ldots, G_T, \ldots$, so its running time is at least $T = 4n^3 + O(n^2)$. By taking into account the number of vertices of the hypergraphs we are considering, we conclude that $M_{(K_5^{(3)},2)}(n) \geq 4n^3/12^3 + O(n^2)$. $\square$

8.4 Concluding remarks

Graph and hypergraph bootstrap percolation have seen a lot of research in recent years, with many intriguing questions remaining open and many possible avenues for further research. We have focused particularly on understanding the maximum running time of these processes. Our first main result, Theorem 8.1.3, building on the previous work of Noel and Ranganathan [119], has allowed us to conclude that the maximum running time of $K_k^{(r)}$-bootstrap percolation is of order $\Theta(n^r)$ for any $k > r \geq 3$. A first very natural (although possibly very difficult) problem is to determine the leading constant in this asymptotic behaviour.

Problem 8.4.1. For each $k > r \geq 2$, determine the limit of $M_k^r(n)/n^r$ (if it exists).

In particular, all results in this hypergraph context have relied on the trivial upper bound that $M_H^r(n) \leq \binom{n}{r}$; obtaining better upper bounds should be the first step towards this problem. We also note that the lower bound arising from our construction (see Remark 8.2.3) is not tight.

Another very natural direction is to study the asymptotic growth of $M_H(n)$ when H is an r-uniform hypergraph which is not complete. We have made the first progress in this direction by addressing two particular cases, see Theorems 8.1.4' and 8.1.5'. A more general study of this problem for different instances of H is crucial towards a unified understanding of hypergraph bootstrap percolation.

More generally, the notion of more 'powerful' infections that we proposed when considering (H, m)-bootstrap percolation leads to many new open problems. Here we have only addressed two particular instances to showcase that this notion leads to interesting results. In the case of $(K_5^{(3)}, 2)$-bootstrap percolation, Theorem 8.1.5 shows that the maximum running time is cubic, that is, as large as it could possibly be (up to constant factors). In the case of $(K_4^{(3)}, 2)$-bootstrap percolation, however, the maximum running time is only linear, and in Theorem 8.1.4 we have determined the exact value of this maximum running time for all values of n. Remarkably, this is the only nontrivial exact result in the area other than those for K_3- and K_4-bootstrap percolation [28] and C_k-bootstrap percolation [61]. It would certainly be desirable to understand the behaviour of the maximum running time of (H, m)-bootstrap percolation more generally. To begin, we propose the following problem.

Problem 8.4.2. Given $k > r \geq 2$ and $m \in [\binom{k}{r}]$, determine the asymptotic behaviour of the maximum running time of the $(K_k^{(r)}, m)$-bootstrap percolation process.

It would also be interesting to consider this more general (H, m)-bootstrap percolation process in other contexts where (hyper)graph bootstrap percolation has been studied. In particular, one may consider the extremal problem, i.e., what is the minimum number of edges an initial r-uniform infection G_0 on n vertices can have if we know the (H, m)-percolation process $G_0, \ldots, G_T$ satisfies $G_T = K_n^{(r)}$?

Chapter 9

Partial shuffles by lazy swaps

9.1 Introduction

Let S_n denote the symmetric group on n elements, i.e., the set of bijections $[n] \to [n]$. A *lazy transposition* with parameters (a, b, p) is a random permutation T such that $T = (a, b)$ with probability p and T is the identity with probability $1 - p$, where (a, b) denotes the transposition swapping a and b. We say that the independent lazy transpositions $T_1, \dots, T_\ell$ form a *transposition shuffle* (of order n and length ℓ) if their product $T_1 \dots T_\ell$, which is a random permutation of S_n, is uniformly distributed among all the $n!$ elements of S_n. What is the shortest possible length $U(n)$ of a transposition shuffle of order n? This problem was first raised by Fitzsimons [64], and also independently studied by Angel and Holroyd [12].

It is not difficult to show that $U(n) \leq \binom{n}{2}$ – there are many constructions achieving this. One example is obtained as follows. We may inductively take $\binom{n-1}{2}$ lazy transpositions $T_1, \dots, T_{\binom{n-1}{2}}$, only permuting $\{1, \dots, n-1\}$, which form a transposition shuffle of order $n - 1$. We may also construct lazy transpositions $T'_1, \dots, T'_{n-1}$ such that $T'_1 \dots T'_{n-1}$ maps the element n to each of $1, 2, \dots, n$ with probability $1/n$. Then the random permutation $T'_1 \dots T'_{n-1} T_1 \dots T_{\binom{n-1}{2}}$ is easily seen to be uniform. Fitzsimons [64] and Angel and Holroyd [12] asked whether the upper bound $\binom{n}{2}$ is tight. Very recently, Groenland, Johnston, Radcliffe and Scott [76] proved the following theorem, answering this question in the negative.

Theorem 9.1.1 (Groenland, Johnston, Radcliffe and Scott [76])**.** *For all $n \geq 6$ we have $U(n) < \binom{n}{2}$. In fact,*

$$U(n) \leq \frac{2}{3}\binom{n}{2} + O(n \log n).$$

A simple lower bound for $U(n)$ can be obtained from the observation that the product of ℓ lazy transpositions takes at most 2^ℓ possible values in S_n, giving $U(n) \geq \log_2(n!) = \Theta(n \log n)$. Surprisingly, this is the best known lower bound, even though the argument above ignores uniformity and only uses that each permutation can be reached with positive probability. In fact, if we only want to achieve that the final permutation is 'close' to uniform, but not necessarily exactly uniform, then it is enough to take $O(n \log n)$ lazy transpositions, as shown by Czumaj [44].

While the result of Groenland, Johnston, Radcliffe and Scott [76] answered the question of Fitzsimons [64] and Angel and Holroyd [12], and showed that the natural upper bound $\binom{n}{2}$ is not tight, there is still a large gap between the best known upper and lower bounds for $U(n)$. The authors of [76] conjectured that in fact $U(n) = o(n^2)$, and also asked about improving the lower bound $\Omega(n \log n)$.

A natural approach to better understand the numbers $U(n)$ is to consider, for some integer $k \leq n$, the smallest possible number of (independent) lazy transpositions $T_1, \dots, T_\ell$ such that the product $T_1 \dots T_\ell$ maps the elements $1, \dots, k$ uniformly to the $n(n-1)\dots(n-k+1)$ possible ordered k-tuples. Such a sequence is called a (k, n)-*shuffle*, and the shortest possible length of a (k, n)-shuffle is denoted $U_k(n)$.

Note that we have $U_1(n) = n - 1$ for all n. Applying this fact repeatedly (similarly to how we derived the bound $U(n) \leq \binom{n}{2}$), we easily get

$$U_k(n) \leq kn - \binom{k+1}{2}. \tag{9.1}$$

Moreover, any improvement on this upper bound for $U_k(n)$ gives an upper bound better than $\binom{n}{2}$ for $U(n)$. Groenland, Johnston, Radcliffe and Scott [76] managed to improve (9.1) for all $k \geq 3$ (and used this to obtain their upper bound for $U(n)$ in Theorem 9.1.1). For the case $k = 2$, they conjectured that (9.1) cannot be improved.

Conjecture 9.1.2 (Groenland, Johnston, Radcliffe and Scott [76, 77])**.** *For all $n \geq 2$ we have*

$$U_2(n) = 2n - 3.$$

Towards Conjecture 9.1.2, Groenland, Johnston, Radcliffe and Scott [77] proved that if we only use lazy transpositions of the form $(1, b, p)$, then we need at least $1.6n - O(1)$. Moreover, they proved that if instead of uniformity we only ask for *reachability*, i.e., if we only want to achieve that the elements $1, 2$ may end up at any other pair (i, j), then the minimal number of lazy transpositions required is $\lceil 3n/2 \rceil - 2$ – demonstrating a gap between the reachability problem and the uniformity problem for pairs.

Our first result in this chapter is a proof of Conjecture 9.1.2.

Theorem 9.1.3. *For all $n \geq 2$ we have*

$$U_2(n) = 2n - 3.$$

Perhaps an even more natural question than determining $U_k(n)$ is as follows: rather than asking that one fixed k-tuple is mapped to each k-tuple with the same probability, we ask that this should be true for *every* k-tuple. More precisely, given $1 \leq k \leq n$, let $U_k^{\text{all}}(n)$ denote the shortest possible sequence of (independent) lazy transpositions $T_1, \ldots, T_\ell$ such that whenever $1 \leq x_1 < x_2 < \cdots < x_k \leq n$, then the image of the k-tuple $(x_1, \ldots, x_k)$ under $T_1 \ldots T_\ell$ is uniformly distributed among the $n(n-1)\ldots(n-k+1)$ possible values. If $T_1, \ldots, T_\ell$ are as above, let us say that they form a *strong* (k,n)*-shuffle*. It is clear that $U_k(n) \leq U_k^{\text{all}}(n) \leq U(n)$ and $U_n^{\text{all}}(n) = U(n)$.

We start with the case $k = 1$: thus we want to ensure that, for each i and j, the probability that i maps to j is $1/n$. Our next result gives an essentially tight bound for $U_1^{\text{all}}(n)$. Note that this bound, being $\Theta(n \log n)$, already matches the order of magnitude of the best known lower bound for $U(n)$.

Theorem 9.1.4. *For any positive integer n, we have*

$$\frac{1}{2} n \log_2 n \leq U_1^{\text{all}}(n) \leq \frac{1}{2} n \log_2 n + 2n.$$

Moreover, if n is a power of 2, then

$$U_1^{\text{all}}(n) = \frac{1}{2} n \log_2 n.$$

We digress to mention an appealing geometric reformulation of this result. For the problem above, the natural 'matrix of probabilities' has i,j entry given by $\mathbb{P}(i \text{ maps to } j)$. If we consider the columns of this matrix as vectors in $\mathbb{R}^n$, then it is easy to see that the problem is equivalent to the following. We start with n independent vectors in $\mathbb{R}^n$, and at any step we replace two of the vectors, say u and v, with the convex combinations $tu + (1-t)v$ and $(1-t)u + tv$ respectively. How many such steps are needed before we have mapped all the vectors to their centroid? The above theorem shows that the answer is about $\frac{1}{2} n \log_2 n$. Interestingly, we do not see any 'directly geometric' argument to establish this result, even approximately.

It is interesting to compare Theorem 9.1.4 with the corresponding reachability problem, i.e., the minimal number of lazy transpositions required so that their product maps any i to any j with positive probability. This is equivalent to a well-known problem, sometimes called 'gossiping dons' (see, e.g., [24, 36]), and the minimal number of lazy transpositions required is $2n - 4$ for $n \geq 4$.

The next case is $k = 2$: thus we are asking that every ordered pair maps to every other ordered pair with probability $\frac{1}{n(n-1)}$. Here one might expect that a quadratic number of lazy transpositions is needed. We prove that, surprisingly, $n(\log n)^{O(1)}$ swaps still suffice.

Theorem 9.1.5. *We have*

$$U_2^{\text{all}}(n) = O(n \log^2 n).$$

Note that we clearly have $U_2^{\text{all}}(n) \geq U_1^{\text{all}}(n)$ for all n, so Theorem 9.1.4 implies a lower bound

$U_2^{\mathrm{all}}(n) \geq \frac{1}{2}n\log_2 n$. We have not managed to obtain any non-trivial improvement on this lower bound.

We can once again contrast Theorem 9.1.5 with the corresponding reachability problem. As mentioned before, there exist lazy transpositions $T_1, \ldots, T_\ell$ such that $\ell = \frac{3}{2}n + O(1)$ and $T_1 \ldots T_\ell$ maps $(1, 2)$ to any pair (i, j) (where $j \neq i$) with positive probability. Thus $T_\ell T_{\ell-1} \ldots T_1$ maps each (i, j) to $(1, 2)$ with positive probability. So if we take an independent copy $T'_1, \ldots, T'_\ell$ of our sequence of lazy transpositions, we get that $T_1 T_2 \ldots T_\ell T'_\ell T'_{\ell-1} \ldots T'_1$ maps each (i, j) to each (i', j') with positive probability, giving an upper bound of $3n + O(1)$ for the reachability problem. It would be interesting to know what the asymptotic behaviour is.

9.2 Tight bounds for $(2, n)$-shuffles

In this section we prove Theorem 9.1.3 concerning lazy transpositions mapping the pair $(1, 2)$ uniformly to all pairs. Our proof is motivated by a new proof of the lower bound for the corresponding reachability problem – our short proof is different from the proof of Groenland, Johnston, Radcliffe and Scott [77], so we find it helpful to include it here. Let $R_2(n)$ denote the minimal number ℓ of transpositions $T_1, \ldots, T_\ell$ such that for any i and j (distinct), there is a subsequence of $T_1, \ldots, T_\ell$ whose product maps $(1, 2)$ to (i, j). (Equivalently, it is the minimal number of lazy transpositions such that their product maps $(1, 2)$ to any (i, j) with positive probability.)

Theorem 9.2.1 (Groenland, Johnston, Radcliffe and Scott [77]). *For all $n \geq 2$ we have*

$$R_2(n) = \lceil 3n/2 \rceil - 2.$$

For reasons that will soon become clear, our proof of the lower bound below is somewhat simpler and more intuitive if we assume that the first possible swap T_ℓ is $(1, 2)$ – equivalently, if we work with the modified problem where we consider unordered pairs $\{i, j\}$ instead of ordered ones. We recommend that the reader focuses on this case.

Let $T_1, \ldots, T_\ell$ be a sequence of transpositions such that for any (i, j), the pair $(1, 2)$ is mapped to (i, j) under some subsequence. Observe that all the relevant information about the first t possible swaps $T_{\ell-t+1}, T_{\ell-t+2}, \ldots, T_\ell$ is carried by a directed graph G_t on vertex set $[n]$ with edges

$$E(G_t) = \{(i, j) : \text{there is a subsequence of } T_{\ell-t+1}, \ldots, T_\ell \text{ mapping } (1, 2) \text{ to } (i, j)\}.$$

(In the simplified setting, this becomes an undirected graph.) We can describe the result of adding an additional transposition $T_{\ell-t} = (a, b)$ (to the beginning of this sequence) in terms of these graphs. Indeed, $E(G_{t+1})$ consists of the following edges.

- All edges in $E(G_t)$;

- All pairs (a, j) such that $(b, j) \in E(G_t)$, $j \neq a$;
- All pairs (j, a) such that $(j, b) \in E(G_t)$, $j \neq a$;
- All pairs (b, j) such that $(a, j) \in E(G_t)$, $j \neq b$;
- All pairs (j, b) such that $(j, a) \in E(G_t)$, $j \neq b$;
- The edges (a, b) and (b, a), provided that at least one of them appear in G_t.

Our proof relies on finding an appropriate invariant based on these digraphs. For any digraph G, let $f_1(G)$ denote the number of vertices in G that appear in an edge (i.e., have positive in-degree or positive out-degree). Also, let us say that $X \subseteq V(G)$ is a *nice clique* if for all $x, x' \in X$ (with $x' \neq x$), at least one of (x, x') and (x', x) appear in $E(G)$, and furthermore every $x \in X$ has both positive in-degree and positive out-degree in G (but not necessarily in X). (In the simplified setting, this is just the usual notion of a clique in a simple graph.) Let us write $f_2(G)$ for the maximal size of a nice clique in G. The next lemma shows that $F(t) = f_1(G_t) + \frac{1}{2} f_2(G_t)$ increases by at most 1 in each step. Note that this immediately gives the tight lower bound $\lceil 3n/2 \rceil - 2$, as G_0 is a single directed edge, and G_ℓ is a complete directed graph on n vertices.

Lemma 9.2.2. *For all* $0 \leq t \leq \ell - 1$, *we have*

$$F(t+1) \leq F(t) + 1.$$

Proof. Let $T_{\ell - t}$ be the transposition (a, b). For all i, let S_i denote the set of vertices in G_i appearing in an edge. Observe that

$$S_{t+1} = \begin{cases} S_t & \text{if } a, b \text{ are both elements or both non-elements of } S_t \\ S_t \cup \{a\} & \text{if } a \notin S_t,\ b \in S_t \\ S_t \cup \{b\} & \text{if } b \notin S_t,\ a \in S_t. \end{cases}$$

In particular, $f_1(G_{t+1}) \leq f_1(G_t)+1$. Let us take a nice clique $X \subseteq [n]$ of maximal possible size $f_2(G_{t+1})$ in G_{t+1}.

Assume first that $f_1(G_{t+1}) = f_1(G_t) + 1$. Then, without loss of generality, $a \notin S_t$ and $b \in S_t$. Hence $(a, b), (b, a) \notin E(G_{t+1})$ and so $\{a, b\} \not\subseteq X$. If $a \notin X$, then X is a nice clique in G_t, thus $f_2(G_t) = f_2(G_{t+1})$, giving the result. However, if $a \in X$, then $(X \setminus \{a\}) \cup \{b\}$ is a nice clique in G_t, once again giving $f_2(G_t) = f_2(G_{t+1})$ and hence $F_2(t+1) = F_2(t) + 1$.

Now assume that $f_1(G_{t+1}) = f_1(G_t)$. Observe that $X \setminus \{a, b\}$ is a nice clique in G_t, so we immediately get $f_2(G_t) \geq f_2(G_{t+1}) - 2$ and the result follows. $\square$

Proof of Theorem 9.2.1. The lower bound $R_2(n) \geq \lceil 3n/2 \rceil - 2$ follows immediately from Lemma 9.2.2, since $F(0) = 2$ and $F(\ell) = 3n/2$.

For the upper bound, it is easy to see that $R_2(n+1) \leq R_2(n) + 2$ for all n, so it suffices to show that $R_2(n) \leq 3n/2 - 2$ when n is even. Let $X = \{x_1, \ldots, x_{n/2-1}\}$ and $Y = \{y_1, \ldots, y_{n/2-1}\}$ partition $[n] \setminus \{1, 2\}$. Let T be the swap $(1, 2)$, furthermore, let T_i denote the transposition $(1, x_i)$, and similarly let T'_i denote $(2, y_i)$. Finally, let W_i denote (x_i, y_i). It is easy to check that

$$W_1, W_2, \ldots W_{n/2-1}, T'_1, T'_2, \ldots T'_{n/2-1}, T_1, T_2, \ldots T_{n/2-1}, T$$

has a subsequence mapping $(1, 2)$ to (i, j), for any pair (i, j). The result follows. $\square$

Let us now turn to the proof of Theorem 9.1.3, showing $U_2(n) = 2n - 3$. For this problem, the random permutation obtained after t swaps can be described by a weighted directed graph – i.e., a matrix (with zeros on the diagonal). It turns out that a similar proof works if we replace the invariant $f_2(t)$ by the rank of this matrix.

Proof of Theorem 9.1.3. The upper bound follows from (9.1), so it is enough to prove the lower bound $U_2(n) \geq 2n - 3$. Assume that $T_1, \ldots, T_\ell$ is a $(2, n)$-shuffle. For each $1 \leq t \leq \ell$, let us write $\sigma_t = T_{\ell-t+1} T_{\ell-t+2} \ldots T_\ell$ for the random permutation obtained from the first t lazy swaps, and let σ_0 be the identity. So we know that $(\sigma_\ell(1), \sigma_\ell(2))$ is uniformly distributed in S_n. For each t, we form a matrix $M^{(t)}$ with entries

$$M^{(t)}_{i,j} = \mathbb{P}(\sigma_t(1) = i, \sigma_t(2) = j).$$

Note that the matrix $M^{(t)}$ carries all the relevant information coming from the first t lazy transpositions. Moreover, we have

$$M^{(0)} = \begin{pmatrix} 0 & 1 & 0 & \ldots & 0 \\ 0 & 0 & 0 & \ldots & 0 \\ 0 & 0 & 0 & \ldots & 0 \\ & & \ldots & & \\ 0 & 0 & 0 & \ldots & 0 \end{pmatrix},$$

and

$$M^{(\ell)} = \begin{pmatrix} 0 & x & x & \ldots & x \\ x & 0 & x & \ldots & x \\ x & x & 0 & \ldots & x \\ & & \ldots & & \\ x & x & x & \ldots & 0 \end{pmatrix},$$

where $x = \frac{1}{n(n-1)}$.

For each t, let $S(t) \subseteq [n]$ be the set of all $i \in [n]$ such that the ith row or the ith column of $M^{(t)}$ contains a non-zero element. In other words, $S(t)$ consists of elements in $[n]$ that can be reached by 1 or 2 after t lazy swaps. Let us also write

$$f(t) = |S(t)| + \mathrm{rank}(M^{(t)}).$$

It is easy to see that $f(0) = 3$ and $f(\ell) = 2n$. So the result follows immediately from the following claim.

Claim. For each $0 \le t \le \ell - 1$, we have

$$f(t+1) \le f(t) + 1.$$

Proof. Fix some t with $0 \le t \le \ell - 1$, and let us write $M = M^{(t)}$, $M' = M^{(t+1)}$. Let $T_{\ell-t}$ have parameters (a, b, p) (where $a \neq b$). Observe that we have

$$M'_{i,j} = \begin{cases} M_{i,j} & \text{if } i \notin \{a,b\} \text{ and } j \notin \{a,b\}, \\ (1-p)M_{i,j} + pM_{\bar{i},j} & \text{if } i \in \{a,b\},\ j \notin \{a,b\} \text{ and } \{a,b\} = \{i,\bar{i}\}, \\ (1-p)M_{i,j} + pM_{i,\bar{j}} & \text{if } i \notin \{a,b\},\ j \in \{a,b\} \text{ and } \{a,b\} = \{j,\bar{j}\}, \\ 0 & \text{if } i = j \in \{a,b\}, \\ (1-p)M_{i,j} + pM_{j,i} & \text{if } \{i,j\} = \{a,b\}. \end{cases}$$

Let P be the following $n \times n$ matrix:

$$P_{i,j} = \begin{cases} 1 & \text{if } i = j \notin \{a,b\}, \\ 0 & \text{if } i \neq j \text{ and } |\{i,j\} \cap \{a,b\}| \le 1, \\ 1-p & \text{if } i = j \in \{a,b\}, \\ p & \text{if } \{i,j\} = \{a,b\}. \end{cases}$$

For example, if $a = 1$ and $b = 2$, then P is given by

$$P = \begin{pmatrix} 1-p & p & 0 & 0 & \dots & 0 \\ p & 1-p & 0 & 0 & \dots & 0 \\ 0 & 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 0 & 1 & \dots & 0 \\ 0 & 0 & 0 & 0 & \dots & 1 \end{pmatrix}.$$

Using that $M_{a,a} = M_{b,b} = 0$, we see that the matrix PMP has entries given by

$$
(PMP)_{i,j} = \begin{cases} M_{i,j} & \text{if } i \notin \{a,b\} \text{ and } j \notin \{a,b\}, \\ (1-p)M_{i,j} + pM_{\bar{i},j} & \text{if } i \in \{a,b\},\ j \notin \{a,b\} \text{ and } \{a,b\} = \{i,\bar{i}\}, \\ (1-p)M_{i,j} + pM_{i,\bar{j}} & \text{if } i \notin \{a,b\},\ j \in \{a,b\} \text{ and } \{a,b\} = \{j,\bar{j}\}, \\ p(1-p)(M_{a,b} + M_{b,a}) & \text{if } i = j \in \{a,b\}, \\ (1-p)^2 M_{i,j} + p^2 M_{j,i} & \text{if } \{i,j\} = \{a,b\}. \end{cases}
$$

So we have $M' = PMP + X$, where

$$
X_{i,j} = \begin{cases} -p(1-p)(M_{a,b} + M_{b,a}) & \text{if } i = j \in \{a,b\}, \\ p(1-p)(M_{a,b} + M_{b,a}) & \text{if } \{i,j\} = \{a,b\}, \\ 0 & \text{otherwise.} \end{cases}
$$

Observe that $\operatorname{rank}(X) \le 1$, and if $\operatorname{rank}(X) = 1$ then $M_{a,b} \neq 0$ or $M_{b,a} \neq 0$. Since $M' = PMP + X$, it follows that $\operatorname{rank}(M') \le \operatorname{rank}(M) + 1$, with equality only if $M_{a,b} \neq 0$ or $M_{b,a} \neq 0$.

On the other hand, it is easy to see that for each $i \notin \{a,b\}$, we have $i \in S(t+1)$ if and only if $i \in S(t)$. Furthermore, $a \in S(t+1)$ only if $a \in S(t)$ or $b \in S(t)$, and similarly $b \in S(t+1)$ only if $a \in S(t)$ or $b \in S(t)$. It follows that $|S(t+1)| \le |S(t)| + 1$, with equality only if exactly one of a and b belong to $S(t)$. In particular, we need $M_{a,b} = M_{b,a} = 0$ for equality.

It follows that $|S(t+1)| + \operatorname{rank}(M') \le |S(t)| + \operatorname{rank}(M) + 1$, proving the claim and hence the theorem. $\square$

9.3 Strong shuffles

In this section we prove Theorems 9.1.4 and 9.1.5 about strong (k,n)-shuffles, i.e., when we want all k-tuples to have uniformly random image. We will begin with the case $k = 1$.

9.3.1 Strong $(1,n)$-shuffles

Let us start by giving a construction of a strong $(1,n)$-shuffle of length $\frac{1}{2} n \log_2 n$ for the case when n is a power of 2. That is, we show that $\frac{1}{2} n \log_2 n$ lazy swaps suffice to make the image of every element uniform. We arrange the $n = 2^t$ elements of $[n]$ along the vertices of a t-dimensional hypercube $\{0,1\}^t$. Then our lazy swaps come in t phases, each of length 2^{t-1}, with the ith phase consisting of all possible swaps in direction i. That is, if e_i denotes the t-dimensional unit vector with ith coordinate 1, then in the ith phase we perform all the lazy swaps $(v, v + e_i, 1/2)$ (with $v \in \{0,1\}^t$ satisfying $v_i = 0$), in an arbitrary order. It is easy to check that after these $2^{t-1} t$ lazy transpositions, every point can end up everywhere else with probability exactly $1/2^t$, giving the upper bound

$$U_1^{\mathrm{all}}(2^t) \leq 2^{t-1}t.$$

This is in fact the exact value of $U_1^{\mathrm{all}}(2^t)$. To prove this, we establish the corresponding lower bound.

Lemma 9.3.1. *For any positive integer n, we have $U_1^{\mathrm{all}}(n) \geq \frac{1}{2}n\log_2 n$.*

The proof uses a rather unusual invariant: we consider the 'heaviest transversal' in the matrix $A_{i,j} = \mathbb{P}(i \text{ maps to } j)$.

Proof. Assume that $T_1, \ldots, T_\ell$ form a strong $(1,n)$-shuffle. For each $1 \leq i \leq \ell$, let us write $\sigma_i = T_{\ell-i+1}T_{\ell-i+2}\ldots T_\ell$ for the random permutation obtained from the first i lazy swaps, and let σ_0 be the identity.

For each $0 \leq i \leq \ell$ and each $\alpha \in S_n$, let

$$g(i,\alpha) = \prod_{x=1}^{n} \mathbb{P}(\sigma_i(x) = \alpha(x)),$$

and let

$$g(i) = \max_{\alpha \in S_n} g(i,\alpha).$$

Observe that $g(0) = 1$ and $g(\ell) = \frac{1}{n^n}$. So the lower bound

$$U_1^{\mathrm{all}}(n) \geq \frac{1}{2}n\log_2 n$$

follows immediately from the following claim.

Claim. For all $0 \leq i \leq \ell - 1$, we have

$$g(i+1) \geq \frac{1}{4}g(i).$$

Proof. Let $\alpha \in S_n$ be such that $g(i) = g(i,\alpha)$, and let $T_{\ell-i}$ have parameters (a,b,p). Let

$\beta = (a,b)\alpha$, $\sigma = \sigma_i$ and $\alpha^{-1}(a) = c$, $\alpha^{-1}(b) = d$. Then

$$\begin{aligned}
g(i+1,\alpha) &= \prod_{x=1}^{n} \mathbb{P}(\sigma_{i+1}(x) = \alpha(x)) \\
&= \left[\prod_{x\neq c,d} \mathbb{P}(\sigma(x) = \alpha(x))\right] \\
&\quad \times [(1-p)\mathbb{P}(\sigma(c) = a) + p\mathbb{P}(\sigma(c) = b)][(1-p)\mathbb{P}(\sigma(d) = b) + p\mathbb{P}(\sigma(d) = a)] \\
&\geq \left[\prod_{x\neq c,d} \mathbb{P}(\sigma(x) = \alpha(x))\right] [(1-p)\mathbb{P}(\sigma(c) = a)][(1-p)\mathbb{P}(\sigma(d) = b)] \\
&= (1-p)^2 g(i,\alpha),
\end{aligned}$$

and similarly

$$\begin{aligned}
g(i+1,\beta) &= \prod_{x=1}^{n} \mathbb{P}(\sigma_{i+1}(x) = \beta(x)) \\
&= \left[\prod_{x\neq c,d} \mathbb{P}(\sigma(x) = \alpha(x))\right] \\
&\quad \times [p\mathbb{P}(\sigma(c) = a) + (1-p)\mathbb{P}(\sigma(c) = b)][p\mathbb{P}(\sigma(d) = b) + (1-p)\mathbb{P}(\sigma(d) = a)] \\
&\geq \left[\prod_{x\neq c,d} \mathbb{P}(\sigma(x) = \alpha(x))\right] [p\mathbb{P}(\sigma(c) = a)][p\mathbb{P}(\sigma(d) = b)] \\
&= p^2 g(i,\alpha).
\end{aligned}$$

But either $p^2 \geq 1/4$ or $(1-p)^2 \geq 1/4$, giving $\max\{g(i+1,\alpha), g(i+1,\beta)\} \geq \frac{1}{4}g(i,\alpha) = \frac{1}{4}g(i)$, proving the claim and hence the lemma. $\square$

Curiously, giving an upper bound in the case when n is not a power of 2 is considerably more difficult that the construction in the power of 2 case. We will need the following lemma.

Lemma 9.3.2. *For any positive integers n and r, we have*

$$U_1^{\text{all}}(n+r) \leq U_1^{\text{all}}(n) + U_1^{\text{all}}(r) + n + r - 1.$$

Proof. We can take $U_1^{\text{all}}(n)$ lazy transpositions permuting $\{1,\ldots,n\}$ only such that their product σ satisfies $\mathbb{P}(\sigma(i) = j) = 1/n$ for all $i, j \in [n]$, and similarly, we can take $U_1^{\text{all}}(r)$ lazy transpositions permuting $n+1,\ldots,n+r$ only such that their product ρ satisfies $\mathbb{P}(\rho(i) = j) = 1/r$ for all $i, j \in [n+1, n+r]$. We claim that there exist lazy transpositions $T_1,\ldots,T_\ell$ with $\ell \leq n+r-1$ such that $\tau = T_\ell T_{\ell-1}\ldots T_1\rho\sigma$ satisfies $\mathbb{P}(\tau(i) = j) = 1/(n+r)$ for all $i, j \in [n+r]$.

We recursively construct the lazy transpositions $T_1, \ldots, T_\ell$ such that for all $0 \le t \le \ell$, writing f_t for the random permutation $T_t T_{t-1} \ldots T_1$, there is some $j \in [n+1]$ and $j' \in \{n+1, \ldots, n+r+1\}$ with the following properties.

- If $1 \le i < j$ or $n+1 \le i < j'$, then $\mathbb{P}(f_t^{-1}(i) \in [n]) = \frac{n}{n+r}$.
- If $j < i \le n$ or $j' < i \le n+r$, then $T_1, \ldots, T_t$ all fix i.
- Either $j = n+1$ and $j' = n+r+1$; or $j \neq n+1$, $j' \neq n+r+1$, $\mathbb{P}(f_t^{-1}(j) \in [n]) > \frac{n}{n+r}$ and $\mathbb{P}(f_t^{-1}(j') \in [n]) < \frac{n}{n+r}$.
- We have $j + (j' - n) \ge t + 2$.

When $t = 0$, these are satisfied for $j = 1$, $j' = n+1$. Now assume that $T_1, \ldots, T_t$ are already constructed and satisfy the conditions above. If we have $j = n+1$ and $j' = n+r+1$, the process terminates (and we set $\ell = t$). Otherwise we construct T_{t+1} as follows. Let us write $q = \mathbb{P}(f_t^{-1}(j) \in [n])$ and $q' = \mathbb{P}(f_t^{-1}(j') \in [n])$, so we know $q > n/(n+r) > q'$.

If $q + q' > 2n/(n+r)$, let $p = \frac{n/(n+r) - q'}{q - q'}$. Note that $0 < p < 1$, $(1-p)q' + pq = n/(n+r)$, and $(1-p)q + pq' = q + q' - ((1-p)q' + pq) > n/(n+r)$. So if we define T_{t+1} to be the lazy transposition with parameters (j, j', p), then the conditions above will be satisfied when (t, j, j') is replaced by $(t+1, j, j'+1)$. (Note that we cannot have $j' + 1 = n+r+1$, otherwise $n = \sum_{i \in [n+r]} \mathbb{P}(f_{t+1}^{-1}(i) \in [n]) > n \cdot n/(n+r) + r \cdot n/(n+r) = n$, giving a contradiction.)

Similarly, if $q + q' < 2n/(n+r)$, then let $p = \frac{q - n/(n+r)}{q - q'}$, so we have $0 < p < 1$, $(1-p)q + pq' = n/(n+r)$ and $(1-p)q' + pq < n/(n+r)$. So if T_{t+1} has parameters (j, j', p), then the conditions are satisfied when (t, j, j') is replaced by $(t+1, j+1, j')$.

Finally, if $q + q' = 2n/(n+r)$, then let T_{t+1} have parameters $(j, j', 1/2)$, so then the conditions are satisfied for $(t+1, j+1, j'+1)$. This finishes the recursive construction.

So let us take $T_1, \ldots, T_\ell$ as above. Note that for $t = \ell - 1$ we have $j \le n$ and $j' \le n+r$. By the last property, we get $\ell \le n+r-1$. Furthermore, by the first property (for $t = \ell$), we have $\mathbb{P}(f_\ell^{-1}(i) \in [n]) = \frac{n}{n+r}$ for all i. It follows that $\mathbb{P}(f_\ell \rho \sigma(x) = y) = 1/(n+r)$ for all $x, y \in [n+r]$, proving the claim. $\square$

Proof of Theorem 9.1.4. The lower bound follows from Lemma 9.3.1, and we have seen that we get the corresponding upper bound $U_1^{\text{all}}(n) \le \frac{1}{2} n \log_2 n$ when n is a power of 2. So we have

$$U_1^{\text{all}}(2^t) = 2^{t-1} t.$$

We now prove the upper bound in the case when n is not a power of 2. Let n have binomial expansion

$$n = \sum_{i=0}^{\lfloor \log_2 n \rfloor} 2^i \epsilon_i$$

(where $\epsilon_i \in \{0,1\}$ for all i). By Lemma 9.3.2, we have

$$U_1^{\mathrm{all}}(s+2^i) \leq U_1^{\mathrm{all}}(s) + U_1^{\mathrm{all}}(2^i) + 2^i + s - 1$$

for all s. Using this several times with $s = \sum_{j<i} 2^j \epsilon_j$, we get

$$\begin{aligned}
U_1^{\mathrm{all}}(n) &\leq \sum_{i:\epsilon_i=1} (U_1^{\mathrm{all}}(2^i) + \sum_{j\leq i} 2^j \epsilon_j - 1) \\
&\leq \sum_{i:\epsilon_i=1} (U_1^{\mathrm{all}}(2^i) + 2^{i+1}) \\
&= \sum_{i:\epsilon_i=1} (2^{i-1} i + 2^{i+1}) \\
&\leq \sum_{i:\epsilon_i=1} (2^{i-1} \log_2 n + 2^{i+1}) \\
&= \frac{1}{2} n \log_2 n + 2n,
\end{aligned}$$

as claimed. □

9.3.2 Strong $(2,n)$-shuffles

We now turn to the proof of Theorem 9.1.5. Our first step is to introduce the notion of a 'division shuffle'; this will be crucial for our construction. Given a positive integer n with n even, let us say that the lazy transpositions $T_1, \ldots, T_\ell$ form a *division* $(2,n)$*-shuffle* (of length ℓ) if whenever $i, j \in [n]$ are distinct, then

$$\begin{aligned}
\mathbb{P}(T_1 \ldots T_\ell(i) \in [n/2] \text{ and } T_1 \ldots T_\ell(j) \in [n/2]) &= \frac{(n/2)(n/2-1)}{n(n-1)} = \frac{1}{4} - \frac{1}{4(n-1)}, \\
\mathbb{P}(T_1 \ldots T_\ell(i) \notin [n/2] \text{ and } T_1 \ldots T_\ell(j) \in [n/2]) &= \frac{(n/2)^2}{n(n-1)} = \frac{1}{4} + \frac{1}{4(n-1)}, \\
\mathbb{P}(T_1 \ldots T_\ell(i) \notin [n/2] \text{ and } T_1 \ldots T_\ell(j) \notin [n/2]) &= \frac{(n/2)(n/2-1)}{n(n-1)} = \frac{1}{4} - \frac{1}{4(n-1)}.
\end{aligned}$$

In other words, whenever we pick two elements i, j in $[n]$, then the image under $T_1 \ldots T_\ell$ of the pair (i,j) is distributed between $\{1, \ldots, n/2\}$ and $\{n/2+1, \ldots, n\}$ in the same way as a random pair coming from $[n]$. Let $h_0(n)$ denote the shortest possible length of a division $(2,n)$-shuffle. It is not difficult to see that

$$U_2^{\mathrm{all}}(2n) \leq h_0(2n) + 2U_2^{\mathrm{all}}(n),$$

so it suffices to bound $h_0(n)$. We will do this by an inductive argument, but it will be convenient to use a slightly stronger property.

If a division $(2, n)$-shuffle $T_1, \dots, T_\ell$ is also a strong $(1, n)$-shuffle, i.e., for all $i, j \in [n]$ we have

$$\mathbb{P}(T_1 \dots T_\ell(i) = j) = 1/n,$$

let us say that $T_1, \dots, T_\ell$ is a *nice division* $(2, n)$*-shuffle*. Let $h(n)$ be the shortest possible length of a nice division $(2, n)$-shuffle. We will prove the following result.

Lemma 9.3.3. *If n is even, then*

$$h(2n) \leq 2h(n) + 2n.$$

Before we prove this lemma, we informally describe our construction. We divide the elements of $[2n]$ into two (equal sized) groups: 'top' and 'bottom' points. Furthermore, we further divide each of the top and bottom groups into 'left' and 'right'. We start with $h(n)$ lazy transpositions on the bottom points so that the image of any pair coming from the bottom will be divided between left and right in the same way as a random pair. We do the same thing for top vertices. Then, for some $q \in (0, 1)$, we will swap top-bottom pairs on the left with probability q, and top-bottom pairs on the right with probability $1 - q$, see Figure 9.1. It can be shown that there is some q such that we end up with a division $(2, 2n)$-shuffle. Indeed, if $q = 1/2$, then any two points which start in the same part of the top-bottom division are too likely to end up in the same group, whereas they are too likely to end up in opposite groups if $q = 0$, so by continuity we can pick an appropriate value of q. The construction finishes with n additional lazy transpositions guaranteeing that our division $(2, 2n)$-shuffle is also a strong $(1, 2n)$-shuffle.

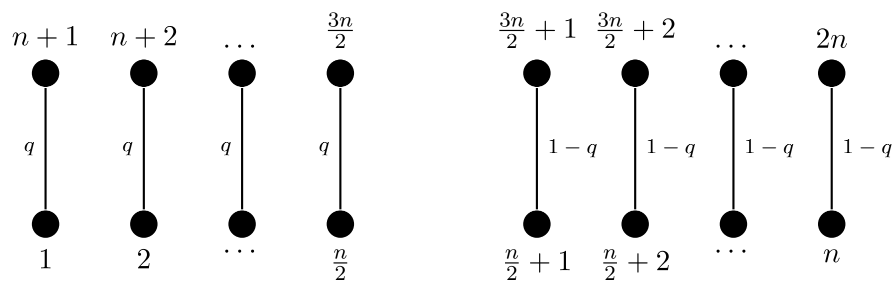

Figure 9.1: To obtain our division $(2, 2n)$-shuffle, we divide the points into top/bottom and left/right. We perform nice division $(2, n)$-shuffles at the top and at the bottom, and then take the lazy transpositions shown on this figure, for some appropriately chosen value of q.

Proof. Let $T_1, \dots, T_\ell$ be a nice division $(2, n)$-shuffle of length $\ell = h(n)$. If T_i has parameters (a, b, p) (with $a, b \in [n]$), let T'_i be the lazy transposition with parameters $(n + a, n + b, p)$. So $T'_1, \dots, T'_\ell$ form a nice division $(2, n)$-shuffle on ground set $[n + 1, 2n]$. Let us write σ for the random permutation $T_1 \dots T_\ell$, and similarly let $\sigma' = T'_1 \dots T'_\ell$.

Let q be some number in $(0,1)$ (specified later). For each $1 \leq i \leq n/2$, let S_i be the lazy transposition with parameters $(i, n+i, q)$, and for each $n/2+1 \leq i \leq n$, let S_i have parameters $(i, n+i, 1-q)$. Let $\rho = S_1 \dots S_n$.

Observe that if i, j are distinct elements of $[2n]$ such that either $i, j \in [n]$ or $i, j \in [n+1, 2n]$, then

$$\begin{aligned}
\mathbb{P}(\rho\sigma'\sigma(i) \in [n] \text{ and } \rho\sigma'\sigma(j) \in [n]) &= \left(\frac{1}{4} - \frac{1}{4(n-1)}\right)(q^2 + (1-q)^2) + \left(\frac{1}{4} + \frac{1}{4(n-1)}\right) 2q(1-q) \\
&= \frac{1}{4} - \frac{1}{4(n-1)}(1 - 4q(1-q)), \\
\mathbb{P}(\rho\sigma'\sigma(i) \notin [n] \text{ and } \rho\sigma'\sigma(j) \in [n]) &= \left(\frac{1}{4} - \frac{1}{4(n-1)}\right) 2q(1-q) + \left(\frac{1}{4} + \frac{1}{4(n-1)}\right)(q^2 + (1-q)^2) \\
&= \frac{1}{4} + \frac{1}{4(n-1)}(1 - 4q(1-q)), \\
\mathbb{P}(\rho\sigma'\sigma(i) \notin [n] \text{ and } \rho\sigma'\sigma(j) \notin [n]) &= \frac{1}{4} - \frac{1}{4(n-1)}(1 - 4q(1-q)).
\end{aligned}$$

Similarly, if $i \in [n]$ and $j \in [n+1, 2n]$ or vice versa, then

$$\begin{aligned}
\mathbb{P}(\rho\sigma'\sigma(i) \in [n] \text{ and } \rho\sigma'\sigma(j) \in [n]) &= \frac{1}{2}q\left(\frac{\frac{n}{2}-1}{n}(1-q) + \frac{1}{2}q\right) + \frac{1}{2}(1-q)\left(\frac{\frac{n}{2}-1}{n}q + \frac{1}{2}(1-q)\right) \\
&= \frac{1}{4} - \frac{1}{n}q(1-q), \\
\mathbb{P}(\rho\sigma'\sigma(i) \notin [n] \text{ and } \rho\sigma'\sigma(j) \in [n]) &= \frac{1}{2}q\left(\frac{1}{n} + \frac{n/2-1}{n}q + \frac{1}{2}(1-q)\right) \\
&\quad + \frac{1}{2}(1-q)\left(\frac{1}{n} + \frac{n/2-1}{n}(1-q) + \frac{1}{2}q\right) \\
&= \frac{1}{4} + \frac{1}{n}q(1-q), \\
\mathbb{P}(\rho\sigma'\sigma(i) \notin [n] \text{ and } \rho\sigma'\sigma(j) \notin [n]) &= \frac{1}{4} - \frac{1}{n}q(1-q).
\end{aligned}$$

Observe that $\frac{n}{4(2n-1)} \in (0, 1/4)$, so we can pick $q \in (0, 1/2)$ such that

$$q(1-q) = \frac{n}{4(2n-1)}.$$

It follows from the equations above that for this particular choice of q, we have, for all $i, j \in [2n]$

(distinct),

$$\begin{aligned}
\mathbb{P}(\rho\sigma'\sigma(i) \in [n] \text{ and } \rho\sigma'\sigma(j) \in [n]) &= \frac{1}{4} - \frac{1}{4(2n-1)},\\
\mathbb{P}(\rho\sigma'\sigma(i) \notin [n] \text{ and } \rho\sigma'\sigma(j) \in [n]) &= \frac{1}{4} + \frac{1}{4(2n-1)},\\
\mathbb{P}(\rho\sigma'\sigma(i) \notin [n] \text{ and } \rho\sigma'\sigma(j) \notin [n]) &= \frac{1}{4} - \frac{1}{4(2n-1)}.
\end{aligned}$$

So $S_1, \ldots S_n, T'_1, \ldots, T'_\ell, T_1, \ldots, T_\ell$ form a division $(2, 2n)$-shuffle.

Let W_i be the lazy transposition $(i, n/2+i, 1/2)$ for $i \in [n/2] \cup [n+1, n+n/2]$. It is easy to check that

$$W_1, \ldots, W_{n/2}, W_{n+1}, W_{n+2}, \ldots, W_{n+n/2}, S_1, \ldots S_n, T'_1, \ldots, T'_\ell, T_1, \ldots, T_\ell$$

is both a strong $(1, 2n)$-shuffle and a division $(2, 2n)$-shuffle. The result follows. □

Proof of Theorem 9.1.5. Observe that whenever $n \geq 2$ is even, then

$$h(n+2) \leq h(n) + 2n + 1. \tag{9.2}$$

Indeed, assume that $T_1, \ldots, T_\ell$ is a nice division $(2, n)$-shuffle (fixing $n+1$ and $n+2$), and let $S_1, \ldots, S_{2n+1}$ be lazy transpositions such that $S_1 \ldots S_{2n+1}$ maps $(n+1, n+2)$ uniformly to the pairs from $[n+2]$. Then $\sigma = S_{2n+1} \ldots S_1$ satisfies $\mathbb{P}(\sigma^{-1}(n+1) = i, \sigma^{-1}(n+2) = j) = \frac{1}{(n+2)(n+1)}$ for all $i, j \in [n+2]$ distinct. Conditioning on whether or not $\sigma(i)$ and $\sigma(j)$ belong to $\{n+1, n+2\}$, we see that for $\rho = T_1 \ldots T_\ell \sigma$ we have

$$\begin{aligned}
\mathbb{P}(\rho(i) \in [n/2] \cup \{n+1\} \text{ and } \rho(j) \in [n/2] \cup \{n+1\}) &= 2\frac{n}{(n+2)(n+1)}\frac{1}{2}\\
&\quad + \frac{n(n-1)}{(n+2)(n+1)}\left(\frac{1}{4} - \frac{1}{4(n-1)}\right)\\
&= \frac{1}{4} - \frac{1}{4(n+1)},\\
\mathbb{P}(\rho(i) \notin [n/2] \cup \{n+1\} \text{ and } \rho(j) \in [n/2] \cup \{n+1\}) &= \frac{1}{4} + \frac{1}{4(n+1)},\\
\mathbb{P}(\rho(i) \notin [n/2] \cup \{n+1\} \text{ and } \rho(j) \notin [n/2] \cup \{n+1\}) &= \frac{1}{4} - \frac{1}{4(n+1)}.
\end{aligned}$$

Furthermore, $T_1, \ldots, T_\ell, S_{2n+1}, S_{2n}, \ldots, S_1$ is easily seen to be a strong $(1, n+2)$-shuffle. The bound (9.2) follows easily.

We also know from Lemma 9.3.3 that

$$h(2n) \leq 2h(n) + 2n \tag{9.3}$$

and hence

$$h(2n+2) \leq 2h(n) + 6n + 1. \tag{9.4}$$

Clearly $h(2) = 1$, so it follows from (9.3) and (9.4) that whenever n is even, we have

$$h(n) \leq 3n \log_2 n. \tag{9.5}$$

Observe that if n is even, then we have

$$U_2^{\text{all}}(n) \leq h(n) + 2U_2^{\text{all}}(n/2). \tag{9.6}$$

Indeed, assume that $T_1, \ldots, T_\ell$ form a division $(2,n)$-shuffle, let $T'_1, \ldots, T'_u$ be a strong $(2, n/2)$-shuffle (fixing $n/2+1, \ldots, 2n$), and let $T''_1, \ldots, T''_v$ be another strong $(2, n/2)$-shuffle on ground set $n/2+1, \ldots, n$ (in particular, it fixes $1, \ldots, n/2$). Then

$$T''_1, \ldots, T''_v, T'_1, \ldots, T'_u, T_1, \ldots, T_\ell$$

is easily seen to be a strong $(2,n)$-shuffle, giving (9.6).

Furthermore, for all n we have

$$U_2^{\text{all}}(n+1) \leq U_2^{\text{all}}(n) + n. \tag{9.7}$$

Indeed, if $T_1, \ldots, T_\ell$ form a strong $(2,n)$-shuffle (fixing $n+1$) and $S_1, \ldots, S_n$ satisfy $\mathbb{P}(S_1 \ldots S_n(i) = n+1) = 1/(n+1)$ for all i, then $T_1, \ldots, T_\ell, S_1, \ldots, S_n$ is easily seen to be a strong $(2, n+1)$-shuffle.

It follows from (9.5), (9.6) and (9.7) that

$$U_2^{\text{all}}(n) \leq 2U_2^{\text{all}}(n/2) + 3n \log_2 n$$

if n is even, and

$$U_2^{\text{all}}(n) \leq 2U_2^{\text{all}}((n-1)/2) + 3n \log_2 n + n$$

if n is odd. Since $U_2^{\text{all}}(1) = 0$ (and $U_2^{\text{all}}(2) = 1$), we get by induction that

$$U_2^{\text{all}}(n) \leq 4n(\log_2 n)^2$$

for all n, giving the result. $\square$

9.4 Open problems

We finish this chapter with some open problems. Despite the important recent progress by Groenland, Johnston, Radcliffe and Scott [76], there is still a large gap between the upper and lower bounds for the problem raised by Fitzsimons [64] and Angel and Holroyd [12], and it would be very interesting to close this gap.

Question 9.4.1 ([12, 64])**.** *What is the asymptotic behaviour of $U(n)$?*

One of the main topics considered in this chapter was determining the shortest possible length $U_k^{\mathrm{all}}(n)$ of strong (k,n)-shuffles. We gave essentially tight bounds in the case $k=1$, and in the next case $k=2$ we gave an upper bound of $O(n\log^2 n)$. It would be interesting to decide whether or not $U_k^{\mathrm{all}}(n)$ is 'small' for all fixed values of k.

Question 9.4.2. *Given a fixed positive integer k, do we have $U_k^{\mathrm{all}}(n)=O(n^{1+\epsilon})$ for all $\epsilon>0$?*

We believe that the answer to Question 9.4.2 should be positive. However, a negative answer would also be very interesting, as it would necessarily give a significantly improved lower bound for $U(n)$.

Another related problem is to fully close the gap between the bounds for $U_2^{\mathrm{all}}(n)$. Theorem 9.1.4 gives a lower bound $U_2^{\mathrm{all}}(n)\geq U_1^{\mathrm{all}}(n)=\Theta(n\log n)$, whereas by Theorem 9.1.5 we have $U_2^{\mathrm{all}}(n)=O(n\log^2 n)$.

Question 9.4.3. *Do we have $U_2^{\mathrm{all}}(n)=\Theta(n\log n)$?*

Finally, as mentioned in the introduction, it would be interesting to consider the reachability problem that is analogous to strong shuffles. Given some k and n, let $R_k^{\mathrm{all}}(n)$ denote the minimal number of lazy transpositions $T_1,\dots,T_\ell$ such that $T_1\dots T_\ell$ maps each k-tuple from $[n]$ to any other k-tuple with positive probability. As mentioned before, the case $k=1$ is the well-known 'gossiping dons' problem, and for general k we get a generalisation of that question. As noted in the introduction (in the special case $k=2$), we have $R_k^{\mathrm{all}}(n)\leq 2U_k(n)$ for all k,n, and therefore $R_k^{\mathrm{all}}(n)=\Theta(n)$ for all k.

Question 9.4.4. *What is the value of $R_k^{\mathrm{all}}(n)$, exactly or asymptotically?*

Even the case $k=2$ seems to be a difficult problem. Farzan Byramji (personal communication, 2023) proved an upper bound of $\lceil 5n/2\rceil-4$, giving an improvement on the $3n$ noted in the introduction, and also showed that this upper bound is optimal if we only consider star transpositions (i.e., all swaps are of the form $(1,i)$ for some i).

Chapter 10

A new upper bound for cancellative pairs

10.1 Introduction

The notion of a cancellative pair was introduced by Holzman and Körner [85]. We say that a pair $(\mathcal{A}, \mathcal{B})$ of families of subsets of an n-element set S is *cancellative* if

$$\begin{gathered}\text{whenever } A, A' \in \mathcal{A} \text{ and } B \in \mathcal{B} \text{ satisfy } A \cup B = A' \cup B \text{ then } A = A' \\ \text{and whenever } A \in \mathcal{A} \text{ and } B, B' \in \mathcal{B} \text{ satisfy } A \cup B = A \cup B' \text{ then } B = B';\end{gathered} \tag{10.1}$$

or, equivalently,

$$\begin{gathered}\text{whenever } A, A' \in \mathcal{A} \text{ and } B \in \mathcal{B} \text{ satisfy } A \setminus B = A' \setminus B \text{ then } A = A' \\ \text{and whenever } A \in \mathcal{A} \text{ and } B, B' \in \mathcal{B} \text{ satisfy } B \setminus A = B' \setminus A \text{ then } B = B'.\end{gathered} \tag{10.2}$$

We will usually take $S = [n] = \{1, \ldots, n\}$ and will call a cancellative pair with $\mathcal{A} = \mathcal{B}$ a *symmetric cancellative pair*. Note that the assumption that $(\mathcal{A}, \mathcal{A})$ is a symmetric cancellative pair is slightly stronger than the assumption that $\mathcal{A}$ is a *cancellative family* [67], meaning no three distinct sets $A, B, C \in \mathcal{A}$ satisfy $A \cup B = A \cup C$. We mention that the concept of cancellative pairs corresponds to the information theoretic concept of uniquely decodable code pairs for the binary multiplying channel without feedback (see, e.g., Tolhuizen [136]).

In the case when n is a multiple of 3, we can obtain an example of a symmetric cancellative pair the following way. Partition S into $n/3$ classes of size 3, and take $\mathcal{A}$ (and $\mathcal{B}$) to be the collection of subsets of S containing exactly one element from each class. It is not hard to verify that we get a cancellative pair. Here we have $|\mathcal{A}||\mathcal{B}| = 3^{2n/3}$, where $3^{2/3} \approx 2.08$. In the symmetric case, Erdős and Katona [105] conjectured this to be the maximal size for cancellative families.

A counterexample was found by Shearer [129]. Tolhuizen [136] gave a beautiful construction to show that we can achieve $(|\mathcal{A}||\mathcal{B}|)^{1/n} \to 9/4 = 2.25$, even by symmetric pairs. This construction is (asymptotically) optimal in the symmetric case by a result of Frankl and Füredi [67].

In the general (non-symmetric) case, the exact value of $\alpha = \sup(|\mathcal{A}||\mathcal{B}|)^{1/n}$ is not known. The best known upper bound is due to Holzman and Körner [85], who showed that $|\mathcal{A}||\mathcal{B}| < \theta^n$ where $\theta \approx 2.3264$. No lower bound better than Tolhuizen's (symmetric) 2.25 is known. Our main aim in this chapter is to improve the upper bound to 2.2682^n. Our proof requires some numerical calculations by a computer.

A related concept is that of a recovering pair. A pair $(\mathcal{A}, \mathcal{B})$ of collections of subsets of an n-element set S is called *recovering* [1, 85] if for all $A, A' \in \mathcal{A}$ and $B, B' \in \mathcal{B}$ we have

$$A \setminus B = A' \setminus B' \implies A = A' \quad \text{and} \quad B \setminus A = B' \setminus A' \implies B = B'. \tag{10.3}$$

So every recovering pair is also cancellative (cf. (10.2)). A conjecture of Simonyi [1] states that $|\mathcal{A}||\mathcal{B}| \leq 2^n$ for every recovering pair. (The value of 2^n may be obtained by taking $\mathcal{A} = \mathcal{P}(S_1)$, $\mathcal{B} = \mathcal{P}(S \setminus S_1)$ for any $S_1 \subseteq S$. This conjecture is a special case of the more general 'sandglass conjecture', due to Ahlswede and Simonyi [1].) In fact, the main motivation for introducing cancellative pairs in [85] was this conjecture of Simonyi. Our upper bound of 2.2682^n was an improvement on the previously best known bounds of about 2.28^n for recovering pairs (Etkin and Ordentlich [60], using the terminology of information theory, and Soltész [132]). After the results in this chapter were published, the upper bound for recovering pairs was further improved to 2.2663^n by Nair and Yazdanpanah [118].

10.2 Proof of the upper bound

Let $h(x) = -x\log_2 x - (1-x)\log_2(1-x)$ be the binary entropy function (with the convention $0\log_2 0 = 0$). Define $\mathcal{A}_i = \{A \in \mathcal{A} \mid i \notin A\}$ and $p_i = |\mathcal{A}_i|/|\mathcal{A}|$; q_i is defined similarly for $\mathcal{B}$. We quote the following result of Holzman and Körner [85]. (We will ignore the case when $\mathcal{A}$ or $\mathcal{B}$ is empty.)

Proposition 10.2.1 (Holzman and Körner [85]). *For a cancellative pair $(\mathcal{A}, \mathcal{B})$, we have*

$$\log_2[|\mathcal{A}||\mathcal{B}|] \leq \sum_{i=1}^{n} f(p_i, q_i) \tag{10.4}$$

where $f(p,q) = ph(q) + qh(p)$.

The result above can be established by considering the entropies of each of the random variables of the form $\xi^B = A \setminus B$, where $B \in \mathcal{B}$ is fixed and $A \in \mathcal{A}$ is chosen uniformly at random (and doing the same with $\mathcal{A}$, $\mathcal{B}$ interchanged). Holzman and Körner [85] used (10.4) and induction to establish their upper bound of $|\mathcal{A}||\mathcal{B}| < \theta^n$ ($\theta \approx 2.3264$).

However, this argument can be improved. We call a cancellative pair *k-uniform* if $|A| = |B| = k$ for all $A \in \mathcal{A}$, $B \in \mathcal{B}$. As we will see, bounding $|\mathcal{A}||\mathcal{B}|$ for k-uniform families enables us to give bounds for general (non-uniform) pairs. For n/k small it is easy to give efficient bounds, and for n/k large we will use that the growth speed of the maximum of $|\mathcal{A}||\mathcal{B}|$ (with k fixed, n increasing) can be bounded.

If $(\mathcal{A}, \mathcal{B})$ and $(\mathcal{A}', \mathcal{B}')$ are cancellative pairs over disjoint ground sets S and S', define their product $(\mathcal{A}'', \mathcal{B}'')$ by

$$\mathcal{A}'' = \{A \cup A' \mid A \in \mathcal{A}, A' \in \mathcal{A}'\}$$

$$\mathcal{B}'' = \{B \cup B' \mid B \in \mathcal{B}, B' \in \mathcal{B}'\}$$

giving a cancellative pair over $S \cup S'$ with $|\mathcal{A}''||\mathcal{B}''| = |\mathcal{A}||\mathcal{B}||\mathcal{A}'||\mathcal{B}'|$.
(Note that the cancellative pair in the Introduction is just the product of cancellative pairs of the form $n = 3$, $\mathcal{A} = \mathcal{B} = \{\{1\}, \{2\}, \{3\}\}$.) Let $c(n)$ be the maximum of $|\mathcal{A}||\mathcal{B}|$ for a cancellative pair over an n-element set, and let $c_k(n)$ be the maximum considering only k-uniform pairs. Similarly to Soltész [132], we prove the following lemma.

Lemma 10.2.2. *Let M be a fixed positive integer, and suppose that $\beta > 0$ is such that $c_k(n) \leq \beta^n$ for all k divisible by M and for all $n \geq k$. Then $c(n) \leq \beta^n$ for all n.*

Proof. Suppose the conditions above are satisfied but $|\mathcal{A}||\mathcal{B}| = \omega^n$ for some $\omega > \beta$. Take the product of $(\mathcal{A}, \mathcal{B})$ with (a copy of) $(\mathcal{B}, \mathcal{A})$ to get a cancellative pair $\left(\mathcal{A}_{(1)}, \mathcal{B}_{(1)}\right)$ over some set S with $\left|\mathcal{A}_{(1)}\right| = \left|\mathcal{B}_{(1)}\right| = \omega^{|S|/2}$ and $\mathcal{A}_{(1)}$ and $\mathcal{B}_{(1)}$ containing the same number of sets of size t for any t. Also, we can take the product of $\left(\mathcal{A}_{(1)}, \mathcal{B}_{(1)}\right)$ with (copies of) itself several times to get a pair with similar properties, so we may assume that $|S|$ is large enough so that $\omega^{|S|}/(|S|+1)^2 > \beta^{|S|}$. Take $k_0 \in \{0, 1, \ldots, |S|\}$ such that $\mathcal{A}_{(1)}, \mathcal{B}_{(1)}$ each contain at least $\omega^{|S|/2}/(|S|+1)$ sets of size k_0, let $\left(\mathcal{A}_{(2)}, \mathcal{B}_{(2)}\right)$ contain only these k_0-sets. So $\left|\mathcal{A}_{(2)}\right| \left|\mathcal{B}_{(2)}\right| > \beta^{|S|}$ and $\left(\mathcal{A}_{(2)}, \mathcal{B}_{(2)}\right)$ is k_0-uniform cancellative. Take the product of $\left(\mathcal{A}_{(2)}, \mathcal{B}_{(2)}\right)$ with itself several times to obtain $\left(\mathcal{A}_{(2)}^M, \mathcal{B}_{(2)}^M\right)$, an (Mk_0)-uniform cancellative family contradicting our assumptions. $\square$

We also need a simple observation.

Lemma 10.2.3. *If k and $n \geq k$ are positive integers, then $c_k(n) \leq 2^{2(n-k)}$. In particular, $c_k(n) \leq 2^n$ for $n \leq 2k$.*

Proof. Given $A \in \mathcal{A}$, all $B \in \mathcal{B}$ have to differ on the complement of A, hence $|\mathcal{B}| \leq 2^{n-k}$. Similarly $|\mathcal{A}| \leq 2^{n-k}$. $\square$

We note that we have equality for $k \leq n \leq 2k$ (i.e., $c_k(n) = 2^{2(n-k)}$), even in the symmetric case [67]. Also, we could deduce Lemma 10.2.3 from (10.4), observing that $\sum p_i = \sum q_i = n - k$.

In order to state our key proposition, we need a definition. For $\gamma, x \geq 2$, consider the following optimisation problem:

$$\begin{aligned} \text{maximize} \quad & \frac{1}{n}\sum_{i=1}^{n} f(p_i, q_i) \\ \text{subject to} \quad & p_i q_i \leq 1/\gamma \quad \text{for } i = 1, \ldots, n \\ & \sum_{i=1}^{n} p_i = \sum_{i=1}^{n} q_i \geq n(1 - 1/x) \\ & 0 \leq p_i, q_i \leq 1 \quad \text{for } i = 1, \ldots, n \\ & n \in \mathbb{N} \end{aligned} \tag{10.5}$$

(Note that the positive integer n is not fixed.) We write $\varphi(\gamma, x)$ for the solution (i.e., the supremum) of (10.5).

Proposition 10.2.4. *Suppose k is a positive integer, $2 \leq \lambda$ such that λk is an integer, and $2 \leq r_1 \leq \gamma$. Suppose that $c_k(\lambda k) \leq r_1^{\lambda k}$ and*

$$r_1 \geq 2^{\varphi(\gamma,\lambda)}. \tag{10.6}$$

Then, for $\lambda k \leq n$,

$$c_k(n) \leq r_1^{\lambda k} \gamma^{n - \lambda k}.$$

In particular, if $\mu > \lambda$, μk is an integer and $r_2 = r_1^{\lambda/\mu}\gamma^{1-\lambda/\mu}$, then $c_k(n) \leq r_2^n$ for $\lambda k \leq n \leq \mu k$.

Proof. Notice that $\gamma \geq r_2 \geq r_1$. We know the given inequality holds for $n = \lambda k$. Suppose it is false for some n, $\lambda k + 1 \leq n$, n minimal.
Then $c_k(n)/c_k(n-1) > \gamma$. So we must have $p_i q_i < 1/\gamma$ (or else $|\mathcal{A}_i||\mathcal{B}_i| > c_k(n-1)$ and $(\mathcal{A}_i, \mathcal{B}_i)$ is cancellative over $S \setminus \{i\}$).
We also have $\sum p_i = \sum q_i = n - k = n(1 - k/n) \geq n(1 - 1/\lambda)$. Hence $\sum f(p_i, q_i) \leq n\varphi(\gamma, \lambda)$ (by the definition of φ). So then, by (10.4), we get

$$|\mathcal{A}||\mathcal{B}| \leq 2^{n\varphi(\gamma,\lambda)} \leq r_1^n \leq r_1^{\lambda k}\gamma^{n-\lambda k},$$

contradiction.
For $\lambda k \leq n \leq \mu k$, we have $c_k(n)^{1/n} \leq (r_1/\gamma)^{\lambda k/n}\gamma \leq (r_1/\gamma)^{\lambda/\mu}\gamma = r_2$. □

Proposition 10.2.4 enables us to implement the following method. Let $2 = \lambda_0 < \lambda_1 < \cdots < \lambda_N$, and let $\rho_0 = 2$. Using a computer program, we find some $\rho_1 \geq \rho_0$, then $\rho_2 \geq \rho_1$, and so on, finally ρ_N, such that the conditions of Proposition 10.2.4 hold for $\lambda = \lambda_i$, $\mu = \lambda_{i+1}$, $r_1 = \rho_i$, $r_2 = \rho_{i+1}$ and the corresponding value of γ ($i = 0, 1, \ldots, N-1$). So then $c_k(n) \leq \rho_N^n$ for $n/k \leq \lambda_N$. (Note that the values ρ_i, λ_i do not depend on k.)

To be able to apply this method, we make the following observations.

1. If λ_i is rational for all i, then we are allowed to assume that $\lambda_i k$ is an integer (since we may assume M divides k for any fixed M positive integer).
2. We do not need to consider $n/k > 3.6$. Indeed, for $n/k > 3.6$ we have $p_i + q_i > 2(1 - 1/3.6) = 13/9$ for some i, so then $p_i q_i > 1 \cdot 4/9 = 1/2.25$. Hence $c_k(n) < 2.25 c_k(n-1)$, as $(\mathcal{A}_i, \mathcal{B}_i)$ is cancellative.
3. We need to find an upper bound on $\varphi(\gamma, x)$. Details on how this is done are given in the Appendix at the end of this chapter, however, we note the following simple result (see Lemma 10.4.4). *Let $\gamma \geq 2.25$, $x \geq 2$ and let (p_0, q_0) satisfy $p_0 + q_0 = 2(1 - 1/x)$ and $p_0 q_0 = 1/\gamma$. If $0 \leq p_0, q_0 \leq 1$, $p_0 \neq q_0$, then $\varphi(\gamma, x) = f(p_0, q_0)$.*

Now we are ready to prove our result using the method described above. Choose, for example, $N = 100000$ and $\lambda_i = 2 + i(3.6 - 2)/N$. Then find appropriate values of $\rho_1, \ldots, \rho_N$ using a computer program. Details about our implementation are given in the Appendix at the end of this chapter. Our program gives $\rho_N = 2.268166\ldots$, whence $c_k(n) \leq 2.2682^n$ for all n (and k a multiple of an appropriate M). By Lemma 10.2.2, we get our main result.

Theorem 10.2.5. *For a cancellative pair $(\mathcal{A}, \mathcal{B})$ over an n-element set, we have $|\mathcal{A}||\mathcal{B}| \leq 2.2682^n$.* □

10.3 Remarks

Recovering pairs Since any recovering pair is also cancellative, the result above immediately gives the following corollary.

Corollary 10.3.1. *For a recovering pair $(\mathcal{A}, \mathcal{B})$ over an n-element set, we have $|\mathcal{A}||\mathcal{B}| \leq 2.2682^n$.* □

We remark that a bound stronger than 2^{2k} for k-uniform recovering pairs over a $2k$-element set would give a stronger bound on the maximal value of $|\mathcal{A}||\mathcal{B}|$ using the argument above (we could choose ρ_0 to be smaller). Note that the product of recovering families is recovering [132], so our arguments would still be valid.

Uniform constructions We now discuss how our upper bound on $c_k(n)$ is related to the best known k-uniform constructions as n/k varies. Tolhuizen [136] gave a family of symmetric k-uniform pairs for all values of k and n having $|\mathcal{A}| \geq \nu\binom{n}{k}2^{-k}$, where ν is a constant. It follows that for $n/k = x > 2$, we have

$$c_k(n)^{1/n} \geq 2^{2(h(1/x) - 1/x) + o(1)}.$$

This construction is known to be asymptotically optimal in the symmetric k-uniform case [67, 136]. (As pointed out after Lemma 10.2.3, the exact value of $c_k(n)$ is known for $n/k \leq 2$.)

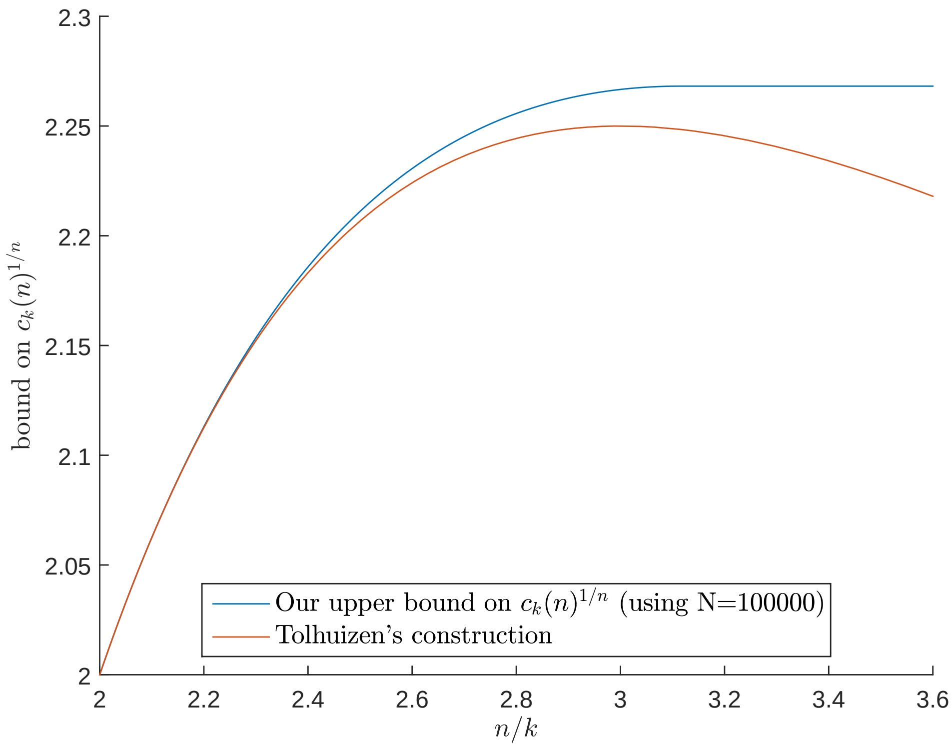

Figure 10.1: Graphical representation of the lower and upper bounds for uniform pairs.

Figure 10.1 shows the upper bound we obtain by the argument above for $c_k(n)^{1/n}$, together with the lower bound from Tolhuizen's construction (n/k fixed, n large). We note that, with a slight modification of Proposition 10.2.4, our upper bound could be decreased for n/k large (instead of becoming constant at the maximum value). However, this would not improve our constant of 2.2682, and it requires more care to find bounds for the optimization problem (10.5) when γ is small.

The symmetric case In the case $\mathcal{A} = \mathcal{B}$, an argument similar to the one considered above gives the best possible bound of 2.25^n. In fact, our argument is equivalent to that of Frankl and Füredi [67]. For convenience, we consider $G_k(n)$, the largest possible size of $\mathcal{A}$ if $(\mathcal{A}, \mathcal{A})$ is k-uniform cancellative. (So then $c_k(n) \geq G_k(n)^2$.) In this case, we have $p_i = q_i$ for each i. If $G_k(n)/G_k(n-1) = \gamma$, then $p_i \leq 1/\gamma$ for all i. But $\sum p_i = n-k$, hence $\gamma \leq \frac{n}{n-k}$. As $G_k(2k) \leq 2^k$, induction gives (for $n \geq 2k$)

$$G_k(n) \leq 2^k \binom{n}{k} \Big/ \binom{2k}{k}$$

This is exactly the formula obtained by Frankl and Füredi [67]. This is not surprising: their argument is essentially the same, but instead of removing elements one-by-one (i.e., inducting from

$n-1$ to n), they consider a random set of size $2k$. (It is not hard to deduce the bound $(3/2)^{2n}$ for symmetric pairs from here, noticing that subexponential factors can be ignored by a product argument. The asymptotic optimality of Tolhuizen's construction for k-uniform symmetric cancellative pairs ($n \to \infty$, $n/k \to x > 2$) also follows [136].)

The choice of N Increasing N over 100000 does not seem to change the first 5 digits after the decimal point in our upper bound $2.268166\ldots$, e.g. $N = 5\cdot 10^6$ gives about 2.268164. We mention that using $N = 5$ already improves the previously best known upper bound for cancellative pairs (it gives about 2.3235^n).

10.4 Appendix

This appendix contains two main parts. In the first part, we give bounds for $\varphi(\gamma, x)$. In the second part, we briefly describe how we implement our argument using a computer program.

Bounding the optimisation problem

Lemma 10.4.1. *Suppose $\gamma \geq 2.25$ and $\kappa \geq 0$. Then the maximizer (p, q) of $L_\kappa(p, q) = f(p, q) + \kappa(p+q)$ in the range $0 \leq p, q \leq 1$, $pq \leq 1/\gamma$ satisfies $pq = 1/\gamma$.*

Proof. Consider the maximizer. We may assume $p \leq q$. We show that if $pq < 1/\gamma$ then $\partial L_\kappa/\partial p > 0$. We have

$$\partial L_\kappa/\partial p = h(q) + qh'(p) + \kappa \geq h(q) + qh'(p).$$

If $p < 1/2$ then this is positive. If $p \geq 1/2$, then

$$\partial L_\kappa/\partial p \geq h\left(\frac{1}{2.25p}\right) + \frac{h'(p)}{2.25p}$$

which is positive on $[1/2, 2/3]$. □

Lemma 10.4.2. *Suppose $\kappa \geq 0$, $\gamma \geq 2.25$, $x \geq 2$ and assume that for $0 \leq p, q \leq 1$, $pq = 1/\gamma$ the maximum of $L_\kappa(p, q) = f(p, q) + \kappa(p + q)$ is $\psi(\gamma, \kappa)$. Then $\varphi(\gamma, x) \leq \psi(\gamma, \kappa) - 2\kappa(1 - 1/x)$.*

Proof. If $(p_i)_{i=1}^n$, $(q_i)_{i=1}^n$ satisfy the constraints of (10.5), then

$$\frac{1}{n}\sum_{i=1}^{n} f(p_i, q_i) \leq \frac{1}{n}\sum_{i=1}^{n} \left(f(p_i, q_i) + \kappa(p_i + q_i)\right) - \frac{1}{n}\kappa \cdot 2n(1 - 1/x).$$

Using Lemma 10.4.1 and our assumptions above, the result follows. □

Lemma 10.4.3. *Suppose $\kappa \geq 0$, $q = q(p) = 1/(\gamma p)$, and (p_0, q_0) satisfy $p_0 q_0 = 1/\gamma$, $0 \leq p_0, q_0 \leq 1$ and*

$$\kappa = \frac{p_0 q_0}{\log 2} \frac{g(p_0) - g(q_0)}{q_0 - p_0}$$

where $g(x) = \frac{\log(1-x)}{x}$. Then $L_\kappa(p, q(p))$ is maximal at (p_0, q_0).

Proof. We may assume $q > p$. As $dq/dp = -q/p$, we have (see [85] for more details)

$$\frac{d}{dp}\left[f(p, q(p)) + \kappa(p + q(p))\right] = q\left[\frac{1}{p}\log_2(1-p) - \frac{1}{q}\log_2(1-q)\right] + \kappa(1 - q/p).$$

This has the same sign as

$$\frac{pq}{\log 2}\frac{g(p) - g(q)}{q - p} - \kappa$$

where $g(x) = \frac{\log(1-x)}{x}$. As pq is constant, it suffices to show that in the range $\frac{1}{\gamma} \leq p < \frac{1}{\sqrt{\gamma}}$, the function

$$\sigma(p) = \frac{g(p) - g(q(p))}{q(p) - p}$$

is strictly decreasing. We have

$$\sigma'(p) = \frac{(g'(p) - g'(q)(-q/p))(q-p) - (g(p) - g(q))(-q/p - 1)}{(q-p)^2}.$$

Since $g'(x) = -\frac{1}{x(1-x)} - g(x)/x$, we obtain

$$\begin{aligned} p(q-p)^2\sigma'(p) &= (q-p)(pg'(p) + qg'(q)) + (p+q)(g(p) - g(q)) \\ &= (q-p)\left(-\frac{1}{1-p} - g(p) - \frac{1}{1-q} - g(q)\right) + (p+q)(g(p) - g(q)) = \\ &= -(q-p)\left(\frac{1}{1-p} + \frac{1}{1-q}\right) + 2pg(p) - 2qg(q). \end{aligned}$$

Using the substitutions $1 - p = x$, $1 - q = y$, $a = x/y > 1$, we get

$$\begin{aligned} p(q-p)^2\sigma'(p) &= -(x-y)\left(\frac{1}{x} + \frac{1}{y}\right) + 2(\log x - \log y) \\ &= -a + \frac{1}{a} + 2\log a. \end{aligned}$$

But this is negative for $a > 1$, since it is 0 at $a = 1$ and its derivative is

$$-1 - \frac{1}{a^2} + \frac{2}{a} = -\frac{(1-a)^2}{a^2}.$$

So σ is strictly decreasing. $\square$

Lemma 10.4.4. *Let $\gamma \geq 2.25$, $x \geq 2$ and let (p_0, q_0) satisfy $p_0 + q_0 = 2(1 - 1/x)$ and $p_0 q_0 = 1/\gamma$. If $0 \leq p_0, q_0 \leq 1$, $p_0 \neq q_0$, then*

$$\varphi(\gamma, x) = f(p_0, q_0).$$

Proof. Choose

$$\kappa = \frac{p_0 q_0}{\log 2} \frac{g(p_0) - g(q_0)}{q_0 - p_0}$$

(this is positive, since $g'(x) < 0$, see [85].) By Lemma 10.4.3, $\psi(\gamma, \kappa) = L_\kappa(p_0, q_0)$. By Lemma 10.4.2, $\varphi(\gamma, x) \leq f(p_0, q_0)$. Equality can be achieved by choosing $n = 2$, $p_1 = q_2 = p_0$, $p_2 = q_1 = q_0$. $\square$

Lemma 10.4.5. *Let $\gamma \geq 2.25, x \geq 2$, and assume that $\gamma \leq 1/(1 - 1/x)^2$. Then*

$$\varphi(\gamma, x) = f(1/\sqrt{\gamma}, 1/\sqrt{\gamma}).$$

Proof. It was proved in [85] that $f(p, q) \leq f(\sqrt{pq}, \sqrt{pq})$, and furthermore, that $f(t, t)$ is increasing on $[0, t_0]$, where $t_0 \approx 0.7$. If $\gamma \geq 2.25$ and $p_i q_i \leq 1/\gamma$, then $\sqrt{p_i q_i} \leq 1/\sqrt{\gamma} \leq 2/3 < t_0$, and hence $f(p_i, q_i) \leq f(\sqrt{p_i q_i}, \sqrt{p_i q_i}) \leq f(1/\sqrt{\gamma}, 1/\sqrt{\gamma})$. This gives the required upper bound for $\varphi(\gamma, x)$. For the lower bound, we can choose $n = 1$ and $p_1 = q_1 = 1/\sqrt{\gamma}$. $\square$

Implementation using a computer

The algorithm finding our final bound proceeds as follows. We choose λ_i ($i = 0, \ldots, N$) to be equally spaced between 2 and 3.6 and set $\rho_0 = 2$. In the i^{th} step, we find the smallest possible value of $\gamma = \gamma_i$ (with $\gamma_i \geq 2.25, \gamma_i \geq \rho_{i-1}$) such that Proposition 10.2.4 applies, i.e., $\rho_{i-1} \geq 2^{\varphi(\gamma_i, \lambda_{i-1})}$. Note that we can calculate $\varphi(\gamma, x)$ (for 'most' values of γ, x) using Lemma 10.4.4 and Lemma 10.4.5 (the former dealing with the cases where $\gamma > 1/(1 - 1/x)^2$, and the latter with the cases $\gamma \leq 1/(1 - 1/x)^2$. In practice, it is sufficient to use only Lemma 10.4.4 except in the first few steps.) We then set $\rho_i = \rho_{i-1}^{\lambda_{i-1}/\lambda_i} \gamma_i^{1 - \lambda_{i-1}/\lambda_i}$. Initially we have $\lambda_0 = 2$, $\rho_0 = 2$, $\gamma_0 = 4$. As i and λ_i increase, we have that ρ_i increases and γ_i decreases. The process stops when γ_i and ρ_i become equal. This happens at $\lambda_i \approx 3.12$ with $\rho_i \approx \gamma_i \approx 2.2682$.

Chapter 11

Antichains in the continuous cube

11.1 Introduction

We can partially order the continuous cube $[0,1]^n$ by writing $\mathbf{x} \leq \mathbf{y}$ if $x_i \leq y_i$ for each $i = 1, \ldots, n$. A subset $A \subseteq [0,1]^n$ is called an *antichain* if it contains no two distinct elements $\mathbf{x}, \mathbf{y}$ with $\mathbf{x} \leq \mathbf{y}$. It is natural to ask for continuous generalisations of Sperner's theorem (which describes the largest possible size of an antichain in the discrete cube $\{0,1\}^n$, see e.g. [25] for background). It is well-known that any antichain in $[0,1]^n$ must have Lebesgue measure zero. However, as shown by Engel, Mitsis, Pelekis and Reiher [51], we can make interesting statements about the sizes of the antichains if we consider their Hausdorff measure. The definition and some properties of the Hausdorff measure are recalled, following [51], in the Appendix at the end of this chapter, to make the chapter self-contained.

Engel, Mitsis, Pelekis and Reiher [51] proved the following result.

Theorem 11.1.1 (Engel, Mitsis, Pelekis and Reiher [51])**.** *Every antichain in $[0,1]^n$ has Hausdorff dimension at most $n-1$ and $(n-1)$-dimensional Hausdorff measure at most n. Moreover, these bounds cannot be improved (that is, there are antichains with $(n-1)$-dimensional Hausdorff measure arbitrarily close to n).*

They also conjectured that there are antichains in $[0,1]^n$ of $(n-1)$-dimensional Hausdorff measure exactly n. In this chapter, we construct such an antichain for each n.

Theorem 11.1.2. *For every n there is an antichain in $[0,1]^n$ of $(n-1)$-dimensional Hausdorff measure n.*

11.2 The construction

Proof of Theorem 11.1.2. Let f be a strictly increasing singular function from $[0,1]$ onto $[0,1]$, that is, a strictly increasing continuous surjective function $[0,1] \to [0,1]$ for which there is a set $S \subseteq (0,1)$ of measure 1 such that f is differentiable with derivative zero at all points of S. (See, e.g., [127] for the construction of such a function.) As noted by Engel, Mitsis, Pelekis and Reiher [51], the graph of $1-f$ has Hausdorff measure 2 (see [65]), giving an example of an antichain with the required properties when $n=2$. We will generalise this construction for any $n \geq 2$. (The case $n=1$ is trivial.)

First, we construct a continuous function $p:(0,1)^{n-1} \to (0,1)$ with the following properties.

1. The function p is differentiable at every point in $(0,1)^{n-1}$ except maybe when two of the coordinates are equal.

2. Whenever $n-2$ coordinates are fixed, then the resulting function $(0,1) \to (0,1)$ is strictly increasing and surjective. In particular, if $\mathbf{x} < \mathbf{y}$ for some $\mathbf{x}, \mathbf{y} \in (0,1)^{n-1}$, then $p(\mathbf{x}) < p(\mathbf{y})$. (We write $\mathbf{x} < \mathbf{y}$ if $\mathbf{x} \leq \mathbf{y}$ and $\mathbf{x} \neq \mathbf{y}$.)

If $n=2$ then $p(x)=x$ works. For $n \geq 3$, we define p as follows. Given $(x_1,\dots,x_{n-1}) \in (0,1)^{n-1}$, let σ be a permutation of $\{1,\dots,n-1\}$ such that $x_{\sigma(1)} \leq \cdots \leq x_{\sigma(n-1)}$, and set

$$p(x_1,\dots,x_{n-1}) = \frac{\prod_{i=1}^{n-2} x_{\sigma(i)}}{1 - x_{\sigma(n-1)} + \prod_{i=1}^{n-2} x_{\sigma(i)}}.$$

It is clear that this satisfies the conditions above.

Define $F:(0,1)^{n-1} \to (0,1)$ by

$$F(x_1,\dots,x_{n-1}) = 1 - p(f(x_1), f(x_2), \dots, f(x_{n-1})),$$

and write A for the graph of F. Note that, since f is strictly increasing, Property 2 above shows that if $\mathbf{x} < \mathbf{y}$ then $F(\mathbf{x}) > F(\mathbf{y})$ and so A is an antichain. We show that A has $(n-1)$-dimensional Hausdorff measure at least n. (The measure is then exactly n by Theorem 11.1.1.)

Write $\pi_i : \mathbb{R}^n \to \mathbb{R}^{n-1}$ for the projection $\pi(x_1,\dots,x_n) = (x_1,\dots,x_{i-1},x_{i+1},\dots,x_n)$. We find n disjoint subsets $A_1,\dots,A_n$ of A such that $\pi_i(A_i)$ has measure 1 for each i. This is sufficient, since the projections π_i cannot increase the Hausdorff measure, so each A_i must then have measure at least 1.

Recall that there is a subset $S \subseteq (0,1)$ of measure 1 such that f is differentiable with derivative zero at each point of S. We set

$$B_n = S^{n-1}$$

and, for $i = 1,\dots,n-1$,

$$B_i = \{\mathbf{x} \in (0,1)^{n-1} : x_j \in S \text{ for } j \neq i \text{ and } x_i \notin S\}.$$

Let A_i be the piece of the graph of F corresponding to domain B_i. Note that the A_i are disjoint. It is clear that $\pi_n(A_n) = B_n$ has measure 1. To show that $\pi_i(A_i)$ has measure 1 for $i = 1, \dots, n-1$, it suffices to consider the case $i = n-1$.

Let $\mathbf{x} = (x_1, \dots, x_{n-2}) \in S^{n-2}$ be a point with all coordinates distinct, we show that the set $\{x_n \in (0,1) : (x_1, \dots, x_{n-2}, x_n) \in \pi_{n-1}(A_{n-1})\}$ has measure 1. (Since the set of such points $\mathbf{x} \in S^{n-2}$ has measure 1, this then gives the result.) Write $T = \{x_1, \dots, x_{n-2}\}$. Consider the function $g_{\mathbf{x}} : (0,1) \to (0,1)$ given by $g_{\mathbf{x}}(t) = F(x_1, \dots, x_{n-2}, t)$. Note that $g_{\mathbf{x}}$ is strictly decreasing and surjective by the properties of p and f. Moreover, whenever $t \in S \setminus T$, then the chain rule gives that $g_{\mathbf{x}}$ is differentiable at t with derivative zero. It follows that $g_{\mathbf{x}}(S \setminus T)$ has measure 0 (for example, using [125, Corollary 21.5]). Hence, using that T is finite and $g_{\mathbf{x}}$ is surjective onto $(0,1)$, we get that $g_{\mathbf{x}}((0,1) \setminus S)$ has measure 1. But this says exactly that $\{x_n \in (0,1) : (x_1, \dots, x_{n-2}, x_n) \in \pi_{n-1}(A_{n-1})\}$ has measure 1, giving the result. $\square$

We remark that when $n = 3$, the function $p : (0,1)^2 \to (0,1)$ has a nice geometric interpretation: $p(\mathbf{x})$ is given by the (first) coordinate of the projection of $\mathbf{x}$ onto the diagonal $\{(t,t) : t \in (0,1)\}$ from the point $(1,0)$ if $\mathbf{x}$ is below the diagonal and from $(0,1)$ if it is above.

11.3 Appendix: The Hausdorff measure

Following [51], we recall the definition and some properties of the Hausdorff measure.

For a non-empty subset U of $\mathbb{R}^n$, let $\operatorname{diam} U$ denote its diameter. For any real number $s \geq 0$, write α_s for the volume of the s-dimensional sphere of radius 1/2. For $s \geq 0$, $\delta > 0$ and $A \subseteq \mathbb{R}^n$, write

$$\mathcal{H}^s_\delta(A) = \alpha_s \inf \left\{ \sum_{i=1}^{\infty} \operatorname{diam}(U_i)^s : A \subseteq \bigcup_{i=1}^{\infty} U_i \text{ and } \operatorname{diam}(U_i) \leq \delta \text{ for each } i \right\}.$$

Note that as δ decreases, the value of $\mathcal{H}^s_\delta(A)$ increases, so we may set

$$\mathcal{H}^s(A) = \lim_{\delta \to 0^+} \mathcal{H}^s_\delta(A).$$

It can be shown that $\mathcal{H}^s$ restricts to a measure, called the *s-dimensional Hausdorff measure*, on a σ-algebra containing all the Borel measurable sets. Note that the scaling by α_s guarantees that the Hausdorff measure and the Lebesgue measure agree when $s = n$. Furthermore, it can be shown that

$$\dim_H(A) = \inf\{s \geq 0 : \mathcal{H}^s(A) = 0\}$$

has the property that if $s > \dim_H(A)$ then $\mathcal{H}^s(A) = 0$ and if $s < \dim_H(A)$ then $\mathcal{H}^s(A) = \infty$, so the only 'interesting' value of $\mathcal{H}^s(A)$ occurs when $s = \dim_H(A)$. The value $\dim_H(A)$ is called the *Hausdorff dimension* of A. See [51] for references on the topic.

Chapter 12

Projections of antichains

12.1 Introduction

In this chapter we study another problem related to the result of Engel, Mitsis, Pelekis, and Reiher [51] mentioned in the previous chapter. A subset of $\mathbb{Z}^n$ is called a *weak antichain* if it contains no elements x and y such that for all i we have $x_i < y_i$. Let us denote by π_i the projection along the i^{th} coordinate, that is, $\pi_i : \mathbb{Z}^n \to \mathbb{Z}^{n-1}$ is given by $(x_1, \dots, x_n) \mapsto (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n)$. Engel, Mitsis, Pelekis and Reiher [51] proved the following projection inequality for weak antichains (which they used to prove an analogous result about weak antichains in the continuous cube $[0, 1]^n$, as well as the result about the Hausdorff measure of antichains mentioned in the previous chapter).

Theorem 12.1.1 (Engel, Mitsis, Pelekis and Reiher [51]). *For every finite weak antichain A in $\mathbb{Z}^n$, we have*

$$|A| \leq \sum_{i=1}^{n} |\pi_i(A)|.$$

The same authors asked the following question.

Question 12.1.2. *What is the smallest possible value $g(n, m)$ of $\operatorname{gap}(A) = \sum_{i=1}^{n} |\pi_i(A)| - |A|$ as A varies over weak antichains in $\mathbb{Z}^n$ of size m?*

Note that the question is uninteresting for (strong) antichains A in $\mathbb{Z}^n$, as we trivially have $|\pi_i(A)| = |A|$ for all i in this case. Furthermore, a weak antichain in $\{0, 1\}^n$ is just a subset of $\{0, 1\}^n$ not containing both the zero vector and the vector with all entries equal to 1. So classical results about set systems (such as Sperner's theorem, see, e.g., [25]) are not particularly relevant here.

In this chapter we answer Question 12.1.2. To state the result, we need some definitions. Let $\mathbb{Z}_{\geq 0}$ denote the set of non-negative integers, and let X_n be the subset of $\mathbb{Z}_{\geq 0}^n$ consisting of elements that have at least one coordinate which is zero. Note that any subset of X_n is a weak antichain. For given $x, y \in X_n$, let $T = \{i : x_i \neq y_i\}$, let $x' = (x_i)_{i \in T}$, $y' = (y_i)_{i \in T}$. Write $x < y$ if $\max x' < \max y'$ or ($\max x' = \max y'$ and $\max\{i : x'_i = \max x'\} < \max\{i : y'_i = \max y'\}$). Then $<$ defines a total order on X_n. We will call this the *balanced order* on X_n.

For example, the first few elements of the balanced order on X_2 are

$$(0,0),(1,0),(0,1),(2,0),(0,2),(3,0),(0,3),(4,0),(0,4),$$

and the first few elements of the balanced order on X_3 are

$$\begin{gathered}(0,0,0),(1,0,0),(0,1,0),(1,1,0),(0,0,1),(1,0,1),(0,1,1),(2,0,0),(2,1,0),(2,0,1),\\(0,2,0),(1,2,0),(0,2,1),(2,2,0),(0,0,2),(1,0,2),(0,1,2),(2,0,2),(0,2,2),(3,0,0).\end{gathered}$$

Theorem 12.1.3. *For every $n \geq 2$ and $m \geq 0$, the initial segment of size m of the balanced order on X_n minimises the gap among weak antichains in $\mathbb{Z}^n$ of size m. In particular, for every positive integer N, the set*

$$A_N = \{x \in \mathbb{Z}^n_{\geq 0} : 0 \leq x_i \leq N-1 \text{ for all } i, \text{ and } x_j = 0 \text{ for some } j\}$$

minimises the gap among weak antichains of size $|A_N| = N^n - (N-1)^n$.

In terms of asymptotic lower bounds on the gap, this gives the following result.

Theorem 12.1.4. *For every $n \geq 2$ and $m \geq 1$, we have*

$$g(n,m) \geq c_n m^{1-1/(n-1)},$$

where $c_n = \frac{1}{2}(n-1)n^{1/(n-1)}$. Moreover, for every $n \geq 2$, we have

$$g(n,m) \sim c_n m^{1-1/(n-1)} \text{ as } m \to \infty.$$

Our proofs have the following structure. Starting with any weak antichain, we modify it into a subset of X_n. This modification will be made step-by-step, and at some points our set will not be a weak antichain. However, it will always satisfy a certain weaker property, which we will call 'layer-decomposability'. Studying subsets of X_n is much simpler than studying general weak antichains, and we will finish the proof of Theorem 12.1.3 using induction on n and codimension-1 compressions. For our proof to work we will need to show that initial segments of the balanced order are extremal for another property as well. Instead of deducing the asymptotic result from Theorem 12.1.3, we will prove it directly and before Theorem 12.1.3, because its proof is simpler and motivates some of the steps in the proof of Theorem 12.1.3.

12.2 Compressing to a down-set in X_n

Recall that we denote $X_n = \{x \in \mathbb{Z}^n_{\geq 0} : x_i = 0 \text{ for some } i\}$. In this section our aim is to prove the following lemma.

Lemma 12.2.1. *If A is a finite weak antichain in $\mathbb{Z}^n$, then there is a weak antichain $A' \subseteq X_n$ of the same size having $|\pi_i(A')| \leq |\pi_i(A)|$ for each i which is a down-set, i.e., if $x, y \in \mathbb{Z}^n_{\geq 0}$, $x_i \leq y_i$ for all i and $y \in A'$ then $x \in A'$.*

We start by recalling the proof of Engel, Mitsis, Pelekis and Reiher [51] that $\text{gap}(A) \geq 0$ for every finite weak antichain. For any finite set $A \subseteq \mathbb{Z}^n$, define the i^{th} bottom layer $B_i(A)$ to be the set of elements with minimal i^{th} coordinate, i.e.,

$$B_i(A) = \{x \in A : \text{whenever } y \in A \text{ with } y_j = x_j \text{ for all } j \neq i \text{ then } y_i \geq x_i\}.$$

Furthermore, define $A_1, \ldots, A_n$ inductively by setting ($A_1 = B_1(A)$ and)

$$A_i = B_i\left(A \setminus (A_1 \cup \cdots \cup A_{i-1})\right).$$

Observe that for a weak antichain we have $A = A_1 \cup \cdots \cup A_n$. Indeed, if $x \in A \setminus (A_1 \cup \cdots \cup A_n)$ then we may inductively find $x^{(i)} \in A_{n-i}$ (for all $0 \leq i \leq n-1$) such that $x^{(i)}_j < x_j$ for all $j \geq n-i$ and $x^{(i)}_j = x_j$ for all $j < n-i$. Then $x^{(n-1)}$ has all coordinates smaller than x, giving a contradiction.

We will call a finite set A with $A = A_1 \cup \cdots \cup A_n$ *layer-decomposable*. Note that π_i restricted to A_i is injective, hence $\sum_{i=1}^n |\pi_i(A)| \geq \sum_{i=1}^n |A_i| = |A|$ for all layer-decomposable sets (and in particular for all weak antichains).

Now assume $A \subseteq \mathbb{Z}^n_{\geq 0}$. Define the *$i$-compression* $C_i(A)$ of A by replacing each $x \in B_i(A)$ by $(x_1, \ldots, x_{i-1}, 0, x_{i+1}, \ldots, x_n)$. Note that $|C_i(A)| = |A|$.

Lemma 12.2.2. *Let $A \subseteq \mathbb{Z}^n_{\geq 0}$ be any finite set. For every $i \neq j$, $\pi_j(C_i(A)) \subseteq C_i(\pi_j(A))$. In particular, $|\pi_j(C_i(A))| \leq |\pi_j(A)|$.*

(When considering $C_i(\pi_j(A))$, we mean compressing along the coordinate labelled by i, not along the i^{th} remaining coordinate.)

Proof. Suppose $(x_1, \ldots, x_{j-1}, x_{j+1}, \ldots, x_n) \in \pi_j(C_i(A))$ so that there is an $x \in C_i(A)$ with k^{th} coordinate x_k for all k.

- If $x_i = 0$ then there is a $y \in B_i(A)$ with $x_k = y_k$ for all $k \neq i$. This implies that we have $(y_1, \ldots, y_{j-1}, y_{j+1}, \ldots, y_n) \in \pi_j(A)$. But this vector and $(x_1, \ldots, x_{j-1}, x_{j+1}, \ldots, x_n)$ agree in all entries except maybe the one labelled by i. Therefore, (since $x_i = 0$) we have $(x_1, \ldots, x_{j-1}, x_{j+1}, \ldots, x_n) \in C_i(\pi_j(A))$.

- If $x_i \neq 0$ then $x \in A \setminus B_i(A)$ and there is a $y \in B_i(A)$ with $y_k = x_k$ for all $k \neq i$ and $y_i < x_i$. But then $\pi_j(y)$ and $\pi_j(x)$ agree in all coordinates except the i^{th} one, which shows $\pi_j(x) \notin B_i(\pi_j(A))$ and hence $\pi_j(x) \in C_i(\pi_j(A))$, as claimed.

$\square$

Say that A is i-compressed if $C_i(A) = A$, i.e., $B_i(A) = \{x \in A : x_i = 0\}$.

Lemma 12.2.3. *Suppose that $A \subseteq \mathbb{Z}^n_{\geq 0}$ is finite, layer-decomposable (i.e., $A = A_1 \cup \cdots \cup A_n$), and k-compressed for all $k < i$. Then $A' = C_i(A)$ satisfies the following.*

(i) A' is k-compressed for all $k \leq i$.

(ii) A' is layer-decomposable.

Proof. Let $j < i$. By Lemma 12.2.2, $|\pi_j(A')| \leq |\pi_j(A)|$. But, since $B_j(A) = \{x \in A : x_j = 0\}$, $C_i(B_j(A))$ is a subset of A' having j^{th} coordinate constant zero and j^{th} projection of size $|\pi_j(A)|$. It follows that $B_j(A') = C_i(B_j(A)) = \{x \in A' : x_j = 0\}$, giving (i).

We now show (ii). Since A is k-compressed for all $k < i$, induction on k gives

$$A_k = \{x \in A : x_k = 0 \text{ but } x_\ell \neq 0 \text{ for all } \ell < k\} \quad \text{for all } k < i. \tag{12.1}$$

Indeed, if this holds for all k' with $k' < k$, then $\bigcup_{\ell=1}^{k-1} A_\ell = \{x \in A : x_\ell = 0 \text{ for some } \ell < k\}$, so A_k contains the right hand side of (12.1), and every element of A_k has $x_\ell \neq 0$ for all $\ell < k$. Furthermore, if there is some $x \in A$ with $x_k > 0$ and $x_\ell \neq 0$ for all $\ell < k$, then there is some $y \in A$ with $y_k = 0$ and $y_j = x_j$ for all $j \neq i$ (as A is k-compressed). Then $y \in A_k$, so certainly $x \notin A_k$, giving the claim.

Similarly,

$$A'_k = \{x \in A' : x_k = 0 \text{ but } x_\ell \neq 0 \text{ for all } \ell < k\} \quad \text{for all } k \leq i. \tag{12.2}$$

But then we have

$$\begin{aligned} C_i(A \setminus (A_1 \cup \cdots \cup A_{i-1})) &= C_i(\{x \in A : x_k \neq 0 \text{ for all } k < i\}) \\ &= \{x \in C_i(A) : x_k \neq 0 \text{ for all } k < i\} \\ &= A' \setminus (A'_1 \cup \cdots \cup A'_{i-1}). \end{aligned}$$

(The first equality is immediate from (12.1). The second equality follows from the fact that C_i acts independently on each set consisting of points having a fixed value of $x_1, \ldots, x_{i-1}$. The last equality is immediate from (12.2).)

It follows that $\{x \in C_i(A \setminus (A_1 \cup \cdots \cup A_{i-1})) : x_i \neq 0\} = \{x \in A' \setminus (A'_1 \cup \cdots \cup A'_{i-1}) : x_i \neq 0\}$. But the left hand side is $A \setminus (A_1 \cup \cdots \cup A_i)$ and the right hand side is $A' \setminus (A'_1 \cup \cdots \cup A'_i)$ by (12.2).

Thus $A\setminus(A_1\cup\cdots\cup A_i) = A'\setminus(A'_1\cup\cdots\cup A'_i)$. Hence $A_j = A'_j$ for all $j > i$. Using $A = A_1\cup\cdots\cup A_n$, we have $A \setminus (A_1 \cup \cdots \cup A_i) = A_{i+1} \cup \cdots \cup A_n$ and so $A' \setminus (A'_1 \cup \cdots \cup A'_i) = A'_{i+1} \cup \cdots \cup A'_n$, giving (ii). □

Lemma 12.2.4. *If $A \subseteq \mathbb{Z}^n_{\geq 0}$ is a finite weak antichain, then $A' = C_n(C_{n-1}(\dots(C_1(A))\dots))$ satisfies*

(i) $|\pi_i(A')| \leq |\pi_i(A)|$ *for each i.*

(ii) A' *is k-compressed for all k.*

(iii) $A' = A'_1 \cup \cdots \cup A'_n$.

(iv) $A'_k = \{x \in A' : x_k = 0 \text{ but } x_\ell \neq 0 \text{ for all } \ell < k\}$ *for all k.*

(v) $A' \subseteq X_n = \{x \in \mathbb{Z}^n_{\geq 0} : x_i = 0 \text{ for some } i\}$.

Proof. The claims (i), (ii), (iii) are immediate from Lemma 12.2.2 and Lemma 12.2.3. Claim (iv) follows from (ii) exactly as in the proof of Lemma 12.2.3. Then (v) follows from (iii) and (iv). □

Note that even though some intermediate steps $C_i(C_{i-1}(\dots(C_1(A))\dots))$ need not give weak antichains, we see that after the n^{th} compression we end up with a set which is necessarily a weak antichain.

For a set $A \subseteq \mathbb{Z}^n_{\geq 0}$, define the *complete i-compression*

$$\begin{aligned} C_i^{\text{compl}}(A) = \{&(x_1,\dots,x_{i-1},a,x_{i+1},\dots,x_n) : \\ &a \in \mathbb{Z}_{\geq 0} \text{ and there are at least } a+1 \text{ elements } y \text{ of } A \text{ having for all } j \neq i \ y_j = x_j\}. \end{aligned}$$

Note that $|C_i^{\text{compl}}(A)| = |A|$.

Lemma 12.2.5. *If $A \subseteq X_n$ then for any j we have $|\pi_j(C_i^{\text{compl}}(A))| \leq |\pi_j(A)|$.*

Proof. The proof is essentially the same as for Lemma 12.2.2. Indeed, let $j \neq i$ and suppose that $(x_1,\dots,x_{j-1},x_{j+1},\dots,x_n) \in \pi_j(C_i^{\text{compl}}(A))$. So there is an $x \in C_i^{\text{compl}}(A)$ with k^{th} coordinate x_k for all k, and hence there are $y^{(0)},\dots,y^{(x_i)} \in A$ such that $y_k^{(a)} = x_k$ for all $k \neq i$ and all $0 \leq a \leq x_i$, and $y_i^{(0)} < y_i^{(1)} < \cdots < y_i^{(x_i)}$. But then $(x_1,\dots,x_{j-1},x_{j+1},\dots,x_n) \in C_i^{\text{compl}}(\pi_j(A))$. It follows that $\pi_j(C_i^{\text{compl}}(A)) \subseteq C_i^{\text{compl}}(\pi_j(A))$, giving the result.

[Alternatively, we can deduce Lemma 12.2.5 from Lemma 12.2.2 by applying C_i to A then $A \setminus B_i(A)$ then $A \setminus (B_i(A) \cup B_i(A \setminus B_i(A)))$ and so on.] □

Proof of Lemma 12.2.1. We may assume that $A \subseteq \mathbb{Z}^n_{\geq 0}$. By Lemma 12.2.4, we may also assume $A \subseteq X_n$. Keep applying complete compressions while it changes our set. These do not increase any projection by Lemma 12.2.5, and keeps our set a subset of X_n. Note that if $A' \neq C_i^{\text{compl}}(A')$ then $\sum_{x \in C_i^{\text{compl}}(A')} \sum_j x_j < \sum_{x \in A'} \sum_j x_j$, so the process must terminate. So the set A' we end up with must have $C_i^{\text{compl}}(A') = A'$ for all i, so it must be a down-set. $\square$

12.3 The asymptotic result

We now show how Lemma 12.2.1 can be used to prove the asymptotic version of our theorem. The proof of the exact version (Theorem 12.1.3) in the next section will be independent of this section, but the proof below motivates some of the steps in the proof of Theorem 12.1.3. Recall that $g(n, m)$ denotes the smallest possible value of $\text{gap}(A) = \sum_{i=1}^n |\pi_i(A)| - |A|$ as A varies over weak antichains of size m in $\mathbb{Z}^n$, and our aim is to prove the following result.

Theorem 12.1.4. *For every $n \geq 2$ and $m \geq 1$, we have*

$$g(n, m) \geq c_n m^{1-1/(n-1)},$$

where $c_n = \frac{1}{2}(n-1)n^{1/(n-1)}$. Moreover, for every $n \geq 2$, we have

$$g(n, m) \sim c_n m^{1-1/(n-1)} \text{ as } m \to \infty.$$

Proof. By Lemma 12.2.1, it suffices to consider sets $A \subseteq X_n$ which are down-sets. We prove the result by induction on n. The case $n = 2$ is trivial, since the gap is exactly 1 for any down-set in X_2. Now assume $n \geq 3$ and the result holds for $n-1$.

Define, for every $a \in \mathbb{Z}_{\geq 0}$,

$$L_a = \{(x_1, x_2, \ldots, x_{n-1}) \in \mathbb{Z}_{\geq 0}^{n-1} : (x_1, x_2, \ldots, x_{n-1}, a) \in A \text{ and } x_i = 0 \text{ for some } i < n\}.$$

Let $K = \pi_n(A) \setminus L_0$. Note that A can be written as a disjoint union of $K \times \{0\}$ and the $L_a \times \{a\}$. Also, $L_0 \supseteq L_1 \supseteq L_2 \supseteq \ldots$, and each L_a is a subset of X_{n-1} (and in particular is a weak antichain). Note furthermore that $|\pi_i(A)| = \sum_{a \geq 0} |\pi_i(L_a)|$ for all $i < n$, and $|\pi_n(A)| = |K| + |L_0|$. It follows that

$$\begin{aligned}
\sum_{i=1}^{n} |\pi_i(A)| - |A| &= \sum_{i=1}^{n-1} |\pi_i(L_0)| + \sum_{a \geq 1} \left(\sum_{i=1}^{n-1} |\pi_i(L_a)| - |L_a| \right) \\
&\geq |L_0| + \sum_{a \geq 0} g(n-1, |L_a|) \\
&\geq |L_0| + \sum_{a \geq 0, L_a \neq \emptyset} c_{n-1} |L_a|^{1-1/(n-2)}.
\end{aligned}$$

Write $|L_0| = x$. Since $|L_a| \le x$ for each a, we have $|L_a|^{1-1/(n-2)} \ge |L_a| x^{-1/(n-2)}$. It follows that

$$\sum_{i=1}^{n} |\pi_i(A)| - |A| \ge x + c_{n-1} \left(\sum_{a \ge 0} |L_a| \right) x^{-1/(n-2)}.$$

Note that $\sum_{a \ge 0} |L_a| = m - |K|$. By the (discrete) Loomis–Whitney inequality [112] (see [27] for a generalisation), and the inequality between the arithmetic and geometric mean,

$$|K|^{n-2} \le \prod_{i=1}^{n-1} |\pi_i(K)| \le \left(\frac{\sum_{i=1}^{n-1} |\pi_i(K)|}{n-1} \right)^{n-1}.$$

But $\sum_{i=1}^{n-1} |\pi_i(K)| \le |L_0|$ since we may assign to $(x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_{n-1}) \in \pi_i(K)$ the value $(x_1, \ldots, x_{i-1}, 0, x_{i+1}, \ldots, x_{n-1}) \in L_0$, giving an injective function from the disjoint union of the projections to L_0. It follows that

$$|K|^{n-2} \le \left(\frac{x}{n-1} \right)^{n-1}$$

and so

$$\begin{aligned} \sum_{i=1}^{n} |\pi_i(A)| - |A| &\ge x + c_{n-1} \left(m - \frac{1}{(n-1)^{1+1/(n-2)}} x^{1+1/(n-2)} \right) x^{-1/(n-2)} \\ &= \left(1 - \frac{c_{n-1}}{(n-1)^{1+1/(n-2)}} \right) x + c_{n-1} m x^{-1/(n-2)}. \end{aligned}$$

Differentiation shows that this is minimised at

$$x = \left(\frac{c_{n-1} m}{(n-2) \left(1 - \frac{c_{n-1}}{(n-1)^{1+1/(n-2)}} \right)} \right)^{1-1/(n-1)},$$

giving

$$\sum_{i=1}^{n} |\pi_i(A)| - |A| \ge \left(1 - \frac{c_{n-1}}{(n-1)^{1+1/(n-2)}} \right)^{\frac{1}{n-1}} (n-1)(n-2)^{1/(n-1)-1} (c_{n-1} m)^{1-1/(n-1)}.$$

But $c_{n-1} = \frac{1}{2}(n-2)(n-1)^{1/(n-2)}$, so

$$1 - \frac{c_{n-1}}{(n-1)^{1+1/(n-2)}} = \frac{n}{2(n-1)}$$

and so

$$\left(1-\frac{c_{n-1}}{(n-1)^{1+1/(n-2)}}\right)^{1/(n-1)}(n-1)(n-2)^{1/(n-1)-1}c_{n-1}^{1-1/(n-1)}=\frac{1}{2}(n-1)n^{1/(n-1)}=c_n,$$

giving $g(n,m)\geq c_n m^{1-1/(n-1)}$, as claimed.

It remains to show that for any fixed n we have $g(n,m)\leq(1+o(1))c_n m^{1-1/(n-1)}$. Let

$$A_N=\{(x_1,\dots,x_n)\in\mathbb{Z}_{\geq 0}^n:\ 0\leq x_i\leq N-1 \text{ for all } i, \text{ and there is a } j \text{ such that } x_j=0\}.$$

Note that A_N has $|\pi_i(A_N)|=N^{n-1}$ for each i, so

$$\sum_{i=1}^n|\pi_i(A_N)|=nN^{n-1}.$$

Moreover, it has size

$$m_N=|A_N|=N^n-(N-1)^n=nN^{n-1}-\binom{n}{2}N^{n-2}+O(N^{n-3}).$$

Now pick N such that $m_N\leq m<m_{N+1}$, and consider the weak antichain given as follows. Let B be an arbitrary subset of $\{0\}\times[N,N+\lfloor(m_{N+1}-m_N)^{1/(n-1)}\rfloor]^{n-1}$ of size $m-m_N$. Note that B has gap at most

$$(n-1)(m_{N+1}-m_N+1)^{(n-2)/(n-1)}=O(N^{(n-2)^2/(n-1)})$$

Put $A=A_N\cup B$. So A has size m and gap equal to the sum of gaps of A_N and B, so A has gap at most

$$\binom{n}{2}N^{n-2}+O(N^{n-3})+O(N^{(n-2)^2/(n-1)})=\binom{n}{2}N^{n-2}(1+o(1)).$$

But $m=nN^{n-1}(1+o(1))$, so the gap is $c_n m^{1-1/(n-1)}(1+o(1))$, as required. □

12.4 The exact result

Recall that we defined a total order (called the balanced order) on X_n as follows. Given $x,y\in X_n$, let $T=\{i:x_i\neq y_i\}$, let $x'=(x_i)_{i\in T}$, $y'=(y_i)_{i\in T}$. Write $x<y$ if $\max x'<\max y'$ or ($\max x'=\max y'$ and $\max\{i:x'_i=\max x'\}<\max\{i:y'_i=\max y'\}$). To see that this really is a total order, we need to show that if $x<y$ and $y<z$, then $x<z$. Set $M_x=\max x$ and $i_x=\max\{i:x_i=M_x\}$, and define M_y,M_z,i_y,i_z similarly. If $M_x<M_y$ or $M_y<M_z$, then $M_x<M_z$ and so $x<z$. If $M_x=M_y=M_z$ and either $i_x<i_y$ or $i_y<i_z$, then $i_x<i_z$, so $x<z$ again follows. Finally, if $M_x=M_y=M_z$ and $i_x=i_y=i_z$ then $x<z$ follows from induction on

n.

Recall that the result we are trying to prove is the following.

Theorem 12.1.3. *For every $n \geq 2$ and $m \geq 0$, the initial segment of size m of the balanced order on X_n minimises the gap among weak antichains in $\mathbb{Z}^n$ of size m. In particular, for every positive integer N, the set*

$$A_N = \{x \in \mathbb{Z}^n_{\geq 0} : 0 \leq x_i \leq N-1 \text{ for all } i, \text{ and } x_j = 0 \text{ for some } j\}$$

minimises the gap among weak antichains of size $|A_N| = N^n - (N-1)^n$.

If $A \subseteq X_n$, we define the balanced-i-compression $\mathcal{C}_i^<(A)$ as follows. For each a, write

$$L^i_a(A) = \{(x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n) \in X_{n-1} : (x_1, \ldots, x_{i-1}, a, x_{i+1}, \ldots, x_n) \in A\}.$$

Also write

$$K^i(A) = \{(x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n) \in \mathbb{Z}^n_{>0} : (x_1, \ldots, x_{i-1}, 0, x_{i+1}, \ldots, x_n) \in A\}.$$

(Here $\mathbb{Z}_{>0}$ denotes the set of positive integers.) Note that for each $a > 0$, $L^i_a(A)$ corresponds to all points of A with i^{th} coordinate equal to a, but for $a = 0$ such points are partitioned into $L^i_0(A)$ and $K^i(A)$ according to whether or not they have another zero coordinate. We define $A' = \mathcal{C}_i^<(A)$ to be the set for which $L^i_a(A')$, $K^i(A')$ are given as follows. Let $L^i_a(A')$ be the initial segment of the balanced order on X_{n-1} of size $|L^i_a(A)|$ for each a, and let $K^i(A')$ be the first $|K^i(A)|$ elements of the ordering $\prec$ on $\mathbb{Z}^{n-1}_{>0}$ given by $(x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n) \prec (y_1, \ldots, y_{i-1}, y_{i+1}, \ldots, y_n)$ if and only if $(x_1, \ldots, x_{i-1}, 0, x_{i+1}, \ldots, x_n) < (y_1, \ldots, y_{i-1}, 0, y_{i+1}, \ldots, y_n)$ (in the balanced order) on X_n. (Note that this is independent of the choice of i, and in fact the relation $\prec$ is given by the same rules as the balanced order.) Observe that $|A'| = |A|$.

It is not immediately clear that $\mathcal{C}_i^<(A)$ is a down-set when A is a down-set. For this we will have to establish another extremal property of initial segments. For $A \subseteq X_n$, write

$$S(A) = \{(x_1, \ldots, x_n) \in \mathbb{Z}^n_{>0} : \text{ for all } j \text{ we have } (x_1, \ldots, x_{j-1}, 0, x_{j+1}, \ldots, x_n) \in A\}.$$

We will prove that initial segments maximise $|S(A)|$ and minimise the gap by induction on the dimension n. The following lemma will be essential in the induction step.

Lemma 12.4.1. *Let $n \geq 3$. Suppose that initial segments I of the balanced order maximise $|S(I)|$ among down-sets in X_{n-1} of given size. Then whenever A is a down-set in X_n and $i \in \{1, \ldots, n\}$, then $A' = \mathcal{C}_i^<(A)$ satisfies the following.*

(i) A' is a down-set.

(ii) $|S(A')| \geq |S(A)|$.

(iii) If it is also true that initial segments of the balanced order minimise the gap among subsets of X_{n-1} of given size, then $\mathrm{gap}(A') \leq \mathrm{gap}(A)$.

Proof. (i) It is clear that $L_0^i(A') \supseteq L_1^i(A') \supseteq \ldots$, and that the $L_a^i(A')$ and $K^i(A')$ are down-sets (in X_{n-1} and $\mathbb{Z}_{>0}^{n-1}$, respectively), since initial segments of the balanced order are down-sets. So it remains to show that $K^i(A') \subseteq S(L_0^i(A'))$. Note that we know this is true for A instead of A' since A is a down-set.

We claim that if I is an initial segment of the balanced order on X_{n-1}, then $S(I)$ is an initial segment of the ordering $\prec$ of $\mathbb{Z}_{>0}^{n-1}$ defined earlier. To see this, suppose that $x, y \in \mathbb{Z}_{>0}^{n-1}$, $y \in S(I)$ and $x \prec y$, we want to show that $x \in S(I)$. Let $T = \{j : x_j \neq y_j\}$ and $k = \min\{\ell \in T : y_\ell = \min_{j \in T} y_j\}$. Then we have the following, for each j.

- If $j \notin T$ then $(x_1, \ldots, x_{j-1}, 0, x_{j+1}, \ldots, x_{n-1}) < (y_1, \ldots, y_{j-1}, 0, y_{j+1}, \ldots, y_{n-1})$.
- If $j \in T$ then $(x_1, \ldots, x_{j-1}, 0, x_{j+1}, \ldots, x_{n-1}) \leq (y_1, \ldots, y_{k-1}, 0, y_{k+1}, \ldots, y_{n-1})$. Indeed, let us write $\bar{x}, \bar{y}$ for these vectors (respectively) and let $\bar{T} = \{\ell : \bar{x}_\ell \neq \bar{y}_\ell\}$. Note that $\bar{T} = T$ if $k \neq j$ and $\bar{T} = T \setminus \{j\}$ otherwise. It $\bar{T} = \emptyset$ then $\bar{x} = \bar{y}$, now assume $\bar{T} \neq \emptyset$. So $\max_{\ell \in \bar{T}} \bar{x}_\ell \leq \max_{\ell \in T} x_\ell \leq \max_{\ell \in T} y_\ell = \max_{\ell \in \bar{T}} \bar{y}_\ell$ and if we have equality throughout then $\max\{\ell \in \bar{T} : \bar{y}_\ell = \max_{s \in \bar{T}} \bar{y}_s\} = \max\{\ell \in T : y_\ell = \max_{s \in T} y_s\} \geq \max\{\ell \in T : x_\ell = \max_{s \in T} x_s\} \geq \max\{\ell \in \bar{T} : \bar{x}_\ell = \max_{s \in \bar{T}} \bar{x}_s\}$. These imply $\bar{x} \leq \bar{y}$.

Using that $y \in S(I)$ and that I is an initial segment, the above shows that $x \in S(I)$.

So $S(L_0^i(A'))$ and $K^i(A')$ are both initial segments. But we have $|S(L_0^i(A'))| \geq |S(L_0^i(A))| \geq |K^i(A)| = |K^i(A')|$, proving (i).

(ii) Note that for any $B \subseteq X_n$ and $a > 0$ we have

$$\begin{aligned}\{(x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n) \in \mathbb{Z}_{>0}^{n-1} : (x_1, \ldots, x_{i-1}, a, x_{i+1}, \ldots, x_n) \in S(B)\}& \\ = K^i(B) \cap S(L_a^i(B)).&\end{aligned}$$

But for each a we have $|S(L_a^i(A))| \leq |S(L_a^i(A'))|$, and $K^i(A'), S(L_a^i(A'))$ are nested (since both of them are initial segments of $\prec$). This implies that

$$\begin{aligned}|K^i(A) \cap S(L_a^i(A))| &\leq \min(|K^i(A)|, |S(L_a^i(A))|) \\ &\leq \min(|K^i(A')|, |S(L_a^i(A'))|) = |K^i(A') \cap S(L_a^i(A'))|.\end{aligned}$$

We get (ii) by summing over all values of a.

(iii) For any down-set $B \subseteq X_n$, we have $|\pi_j(B)| = \sum_{a \geq 0} |\pi_j(L_a^i(B))|$ for all $j \neq i$ and $|\pi_i(B)| = |L_0^i(B)| + |K^i(B)|$. It follows that $\mathrm{gap}(B) = |L_0^i(B)| + \sum_{a \geq 0} \mathrm{gap}(L_a^i(B))$. (Note that on the right hand side we have the gaps of $(n-1)$-dimensional sets.) But then (iii) follows trivially from the assumption that initial segments of the balanced order minimise the gap on X_{n-1}. $\square$

The following lemma will be useful when considering sets satisfying $\mathcal{C}_i^<(A) = A$ for all i.

Lemma 12.4.2. *Suppose $n \geq 3$ and $A \subseteq X_n$ is a down-set having $\mathcal{C}_i^<(A) = A$ for all i. Assume that $x < y$ with $x \notin A$ and $y \in A$. Then*

(i) x has a unique coordinate which is zero.

(ii) if $x_\ell = y_\ell$ for some ℓ, then $x_\ell = y_\ell = 0$ and y has at least one other coordinate which is zero.

Proof. Assume first that $x_\ell = y_\ell$ for some ℓ. Since $\mathcal{C}_\ell^<(A) = A$, it must be the case that $x_\ell = y_\ell = 0$ and exactly one of x, y have a zero coordinate not at the ℓ^{th} position. It follows that if we write $i = \max\{j : y_j = \max y\}$ then $y_i \neq x_i$. Using $y > x$, we get that

$$y_i \geq x_j \text{ for all } j, \text{ and} \tag{12.3}$$

$$y_i > x_j \text{ for all } j \geq i. \tag{12.4}$$

Pick some $k \neq i, \ell$. Then the vector y' obtained by replacing the k^{th} coordinate of y by 0 is in A (since A is a down-set), and we have $y' > x$ (by (12.3) and (12.4)). By the same argument as above, we deduce from $x_\ell = y'_\ell$ and $\mathcal{C}_\ell^<(A) = A$ that $x_\ell = y'_\ell = 0$, and – since $y'_k = 0$ – that it must be the case that x has no zero coordinates other than the ℓ^{th} one. Hence $x_\ell = y_\ell = 0$, $x_s \neq 0$ for all $s \neq \ell$, and there is an $s \neq \ell$ such that $y_s = 0$. This proves the lemma in this case.

Now assume that $x_\ell \neq y_\ell$ for all ℓ. Writing $i = \max\{j : y_j = \max y\}$ again, (12.3) and (12.4) still hold. We only need to show that x has at most one coordinate which is zero. Assume that $x_k = x_\ell = 0$ with $k \neq \ell$, we may assume that $\ell \neq i$ (otherwise swap k and ℓ). Let y' be obtained from y by replacing the ℓ^{th} coordinate by 0. Then $y' \in A$ (since A is a down-set) and $y' > x$ (by (12.3) and (12.4)). But also $y'_\ell = x_\ell$, so by the first case (applied to x and y') we know that x has exactly one zero coordinate, giving a contradiction. $\square$

Lemma 12.4.3. *For every $n \geq 2$, initial segments I of the balanced order maximise $|S(I)|$ among down-sets in X_n of given size.*

Proof. We prove the statement by induction on n. If $n = 2$, then any down-set in X_n of size m is of the form $B_N = \{(i, 0) : i \in \mathbb{Z}_{\geq 0}, i \leq N\} \cup \{(0, i) : i \in \mathbb{Z}_{\geq 0}, i \leq m - 1 - N\}$ for some $0 \leq N \leq m - 1$ integer. We have $S(B_N) = \{(i, j) \in \mathbb{Z}^2_{>0} : 1 \leq i \leq N, 1 \leq j \leq m - 1 - N\}$, so $|S(B_N)| = N(m - 1 - N)$. Over the integers, this attains a maximum at $N = \lceil (m-1)/2 \rceil$, which corresponds to the initial segment of the balanced ordering.

Now assume that $n \geq 3$ and the result holds for smaller values of n. Let A be any subset of X_n, we show the initial segment I of same size has $|S(I)| \geq |S(A)|$. Taking a down-set A'

in X_n minimising $\sum_{x \in A'}$ (position of x in the balanced order) among sets with $|A'| = |A|$ and $|S(A')| \geq |S(A)|$, we may assume that $\mathcal{C}_i^{<}(A) = A$ for each i (by Lemma 12.4.1). Suppose that there are $x, y \in X_n$ with $x < y$, $y \in A$ and $x \notin A$.

Take y to be maximal (in the balanced order). Let $i = \max\{j : y_j = \max y\}$. If there is an $x \notin A$ with $x < y$ and the unique zero coordinate not being at the i^{th} position, pick the minimal of these (in the balanced order). Otherwise pick $x \notin A$ which is minimal. Consider $A' = A \setminus \{y\} \cup \{x\}$. Note that A' is again a down-set.

We show that $|S(A')| \geq |S(A)|$. (This would give a contradiction.) If y has more than one zero coordinates, then $S(A) \setminus S(A') = \emptyset$, so the claim is clear. Otherwise y has a unique zero coordinate y_t, and we must have $x_\ell \neq y_\ell$ for all ℓ by Lemma 12.4.2. In particular, $y_i \neq x_i$. Thus $y_i > x_\ell$ for all $\ell \geq i$ and $y_i \geq \max x$. Observe that

$$\begin{aligned} S(A) \setminus S(A') = \{&(y_1, \ldots, y_{t-1}, a, y_{t+1}, \ldots, y_n) : \\ &a \in \mathbb{Z}_{>0} \text{ and replacing any coordinate by 0 we get an element of } A\}. \end{aligned}$$

Recall that there is a unique s such that $x_s = 0$. We claim that $S(A) \setminus S(A')$ is empty unless $s = i$. Indeed, suppose $s \neq i$ and $S(A) \setminus S(A')$ has an element z corresponding to $a \geq 1$. Let z' be obtained from z by setting the s^{th} coordinate to be zero. Then $z' \in A$, $z' > x$ (as $z'_i = y_i$ so $z'_i > x_\ell$ for all $\ell \geq i$ and $z' \geq \max x$), $x_s = z'_s = 0$ and there is a unique coordinate at which z' is zero. This contradicts Lemma 12.4.2.

So we may assume $s = i$. Note that if $a \geq x_t$ and the corresponding vector appears in the set above, then A has an element z with $z_i = y_i$ and $z_t = x_t \neq 0$ (using that $n \geq 3$ and that A is a down-set. Note that $x_t \neq 0$ since $x_\ell \neq y_\ell$ for all ℓ.) But then $z > x$, so this contradicts Lemma 12.4.2. It follows that $|S(A) \setminus S(A')| \leq x_t - 1 \leq y_i - 1$.

Furthermore, since $i = s$,

$$\begin{aligned} S(A') \setminus S(A) = \{&(x_1, \ldots, x_{i-1}, a, x_{i+1}, \ldots, x_n) : \\ &a \in \mathbb{Z}_{>0} \text{ and replacing any coordinate by 0 we get an element of } A'\}. \end{aligned}$$

Also, by our choice of x, any $z \notin A$ with $z < y$ has ($z_i = 0$ and) $z_\ell \neq 0$ for all $\ell \neq i$. But this easily shows that for all $1 \leq a \leq y_i - 1$, the corresponding vector lies in $S(A') \setminus S(A)$. So $|S(A') \setminus S(A)| \geq y_i - 1 \geq |S(A) \setminus S(A')|$.

So we get a contradiction, finishing the proof. □

Lemma 12.4.4. *For every $n \geq 2$, initial segments I of the balanced order minimise* $\text{gap}(I)$ *among down-sets in X_n of given size.*

Proof. Again we prove this by induction on n. The case $n = 2$ is trivial, since any down-set in X_2 has gap 1.

Now assume that $n \geq 3$ and the result holds for smaller values of n. Let A be any subset of X_n, we show that the initial segment of same size has a gap which is not greater. Taking a down-set A' in X_n minimising $\sum_{x \in A'}$ (position of x in the balanced order) among sets with $|A'| = |A|$ and $\text{gap}(A') \leq \text{gap}(A)$, we may assume that $\mathcal{C}_i^{<}(A) = A$ for each i (by Lemma 12.4.1 and Lemma 12.4.3). Suppose that there are $x, y \in X_n$ with $x < y$, $y \in A$ and $x \notin A$. Take y to be maximal and x to be minimal (in the balanced order). Let $A' = A \setminus \{y\} \cup \{x\}$. Note that A' is a down-set.

By Lemma 12.4.2, there is a unique s such that $x_s = 0$. Then $\pi_j(A') \setminus \pi_j(A) = \emptyset$ if $j \neq s$ and $|\pi_s(A') \setminus \pi_s(A)| = 1$. On the other hand, if t is such that $y_t = 0$ then $|\pi_t(A) \setminus \pi_t(A')| = 1$. It follows that $\text{gap}(A') \leq \text{gap}(A)$, giving a contradiction. □

Proof of Theorem 12.1.3. The result follows immediately from Lemmas 12.4.4 and 12.2.1. □

Chapter 13

Orientation covering numbers

13.1 Introduction

Given a non-empty graph G and k orientations $\vec{G}_1, \ldots, \vec{G}_k$ of G, we say that $\vec{G}_1, \ldots, \vec{G}_k$ is an *orientation covering* of G if whenever $x, y, z \in V(G)$ with $xy, xz \in E(G)$ then there is an orientation in which both xy and xz are oriented away from x (i.e., there is some i such that $(x, y), (x, z) \in E(\vec{G}_i)$). The *orientation covering number* $\sigma(G)$ of G is the smallest positive integer k such that there is a list of k orientations forming an orientation covering of G. Orientation coverings were introduced by Esperet, Gimbel and King [58], who showed that there is a natural connection between orientation coverings and the minimal number $\mathrm{eq}(H)$ of equivalence subgraphs (i.e., subgraphs that are disjoint unions of cliques) needed to cover a line graph H. In particular, they showed that if $L(G)$ is the line graph of G then $\mathrm{eq}(L(G)) \leq \sigma(G) \leq 3\,\mathrm{eq}(L(G))$, and if G is triangle-free then $\mathrm{eq}(L(G)) = \sigma(G)$.

Esperet, Gimbel and King [58] showed that $\sigma(G) \leq \sigma\left(K_{\chi(G)}\right)$ for any graph G, where χ denotes the chromatic number. They asked whether we always have $\sigma(G) = \sigma\left(K_{\chi(G)}\right)$. In this chapter we answer this question in the positive.

Theorem 13.1.1. *For any non-empty graph G, we have* $\sigma(G) = \sigma\left(K_{\chi(G)}\right)$.

The value of $\sigma(K_n)$ has been investigated by Esperet, Gimbel and King [58], who determined its order of magnitude and the exact values for small values of n. An observation of Gyárfás (see [58]) shows that we have $\chi(DS_n) \leq \sigma(K_n) \leq \chi(DS_n) + 2$, where DS_n is the double-shift graph on n vertices. (The graph DS_n is defined as follows. Its vertices are the 3-element subsets of $\{1, \ldots, n\}$, with the vertices $\{i, j, k\}$ and $\{j, k, \ell\}$ joined by an edge whenever $i < j < k < \ell$.) Using the results of Füredi, Hajnal, Rödl and Trotter [70] on the chromatic number of DS_n, this gives $\sigma(K_n) = \log\log n + \frac{1}{2}\log\log\log n + O(1)$. (All logarithms in this chapter are base 2.) In this chapter we will also determine the value of $\sigma(K_n)$ exactly in terms of a certain sequence of positive integers sometimes called the Hoşten–Morris numbers. As a corollary, we get the following improved estimate.

Theorem 13.1.2. *We have* $\sigma(K_n) = \lceil \log\log n + \frac{1}{2}\log\log\log n + \frac{1}{2}(\log \pi + 1) + o(1) \rceil$ *as* $n \to \infty$.

Given a positive integer k, let $[k]$ denote $\{1, \ldots, k\}$, as usual, and let $\mathcal{P}([k])$ denote the set of subsets of $[k]$. Given a family $\mathcal{A} \subseteq \mathcal{P}([k])$ of subsets of $[k]$, we say that $\mathcal{A}$ is *intersecting* if whenever $S, T \in \mathcal{A}$ then $S \cap T \neq \emptyset$. We say that $\mathcal{A}$ is *maximal intersecting* if $\mathcal{A}$ is intersecting and whenever $\mathcal{B} \supseteq \mathcal{A}$ and $\mathcal{B}$ is intersecting then $\mathcal{B} = \mathcal{A}$. (Equivalently, if $\mathcal{A}$ is intersecting and $|\mathcal{A}| = 2^{k-1}$.) The following characterisation of $\sigma(G)$ is the key to our results.

Theorem 13.1.3. *For any non-empty graph* G*,* $\sigma(G)$ *is the smallest positive integer* k *such that there are at least* $\chi(G)$ *maximal intersecting families over* $[k]$*.*

Clearly, Theorem 13.1.3 implies Theorem 13.1.1. Let $\lambda(k)$ denote the number of maximal intersecting families over $[k]$. The numbers $\lambda(k)$ are sometimes called Hoşten–Morris numbers, after a paper of Hoşten and Morris [84] in which they showed that the order dimension of K_n is the smallest positive integer k with $\lambda(k) \geq n$. (Given a partial order P, its order dimension is the smallest integer d such that there exist total orders $L_1, \ldots, L_d$ with the property that $x < y$ in P if and only if $x < y$ in each L_i. The complete graph K_n can be viewed as a partially ordered set on points $V(K_n) \cup E(K_n)$ with ordering given by inclusion, i.e., $x < e$ whenever e is an edge with x as an endpoint.) Using Theorem III in [48], it is easy to see that an equivalent formulation of the result of Hoşten and Morris is that the minimal number of transitive orientations forming an orientation covering of K_n is the smallest positive integer k with $\lambda(k) \geq n$. Note that by Theorem 13.1.3 this number is the same as the orientation covering number of K_n.

Although no exact or asymptotic formula is known for $\lambda(k)$, it was shown by Brouwer, Mills, Mills and Verbeek [34] that

$$\log \lambda(k) \sim \frac{2^k}{\sqrt{2\pi k}}. \tag{13.1}$$

Furthermore, the exact values of $\lambda(k)$ are known [34] for k up to 9, with $\lambda(9) \approx 4 \times 10^{20}$.

Theorem 13.1.2 follows from Theorem 13.1.3 and (13.1). Indeed, taking logarithms in (13.1) shows that $\sigma(K_n)$ is the smallest positive integer k with $\log\log n \leq k - \frac{1}{2}(\log \pi + 1) - \frac{1}{2}\log k + o(1)$, which gives $\sigma(K_n) = \lceil \log\log n + \frac{1}{2}\log\log\log n + \frac{1}{2}(\log \pi + 1) + o(1) \rceil$.

13.2 Proof of Theorem 13.1.3

The proof is based on the following observation.

Lemma 13.2.1. *For any non-empty graph* G*,* $\sigma(G)$ *is the smallest positive integer* k *with the property that there is a collection* $(\mathcal{A}_v)_{v \in V(G)}$ *of subsets of* $\mathcal{P}([k])$ *(i.e.,* $\mathcal{A}_v \subseteq \mathcal{P}([k])$ *for all* v*) such that the following two conditions hold.*

1. *If $uv \in E(G)$, then there exists $S \in \mathcal{A}_u$ and $T \in \mathcal{A}_v$ such that $S \cap T = \emptyset$.*

2. *For all $v \in V(G)$ and $S, T \in \mathcal{A}_v$, we have $S \cap T \neq \emptyset$. (I.e., $\mathcal{A}_v$ is intersecting.)*

Proof. First assume that $\sigma(G) = k$ and $\vec{G}_1, \ldots, \vec{G}_k$ form an orientation covering of G. For each ordered pair (x, y) of vertices such that xy is an edge of G, let $S_{(x,y)} = \{i \in [k] : (x, y) \in E(\vec{G}_i)\}$, i.e., the set of indices i such that in $\vec{G}_i$ the edge xy is oriented from x to y. Let $\mathcal{A}_v = \{S_{(v,w)} : vw \in E(G)\}$. Clearly $S_{(v,w)} \cap S_{(w,v)} = \emptyset$, so Condition 1 holds. Also, we have $S_{(v,w)} \cap S_{(v,w')} \neq \emptyset$ whenever $vw, vw' \in E(G)$, since by assumption there is an i such that $(v, w), (v, w') \in E(\vec{G}_i)$. So Condition 2 holds as well.

Conversely, suppose that we have such a collection $(\mathcal{A}_v)_{v \in V(G)}$ with $\mathcal{A}_v \subseteq \mathcal{P}([k])$ for all v. For each $uv \in E(G)$, pick $S_{(u,v)} \in \mathcal{A}_u$ and $S_{(v,u)} \in \mathcal{A}_v$ such that $S_{(u,v)} \cap S_{(v,u)} = \emptyset$. Define the orientations $\vec{G}_1, \ldots, \vec{G}_k$ of G by orienting the edge uv from u to v in $\vec{G}_i$ if $i \in S_{(u,v)}$, and from v to u if $i \in S_{(v,u)}$, and arbitrarily otherwise. Since we picked $S_{(u,v)}$ and $S_{(v,u)}$ to be disjoint for each uv, this gives a valid orientation for each i. Furthermore, whenever $uv, uw \in E(G)$, then $S_{(u,v)} \cap S_{(u,w)} \neq \emptyset$ (by Condition 2), so we get an orientation covering. This gives $\sigma(G) \leq k$, as claimed. $\square$

Proof of Theorem 13.1.3. We first show the lower bound for $\sigma(G)$. Let G be any non-empty graph, and let $(\mathcal{A}_v)_{v \in V(G)}$ be as in Lemma 13.2.1 for $k = \sigma(G)$. For each $v \in V(G)$, let $\mathcal{B}_v$ be a maximal intersecting family with $\mathcal{B}_v \supseteq \mathcal{A}_v$. Note that the families $(\mathcal{B}_v)_{v \in V(G)}$ still satisfy both conditions in Lemma 13.2.1. Indeed, Condition 1 holds, since given $uv \in E(G)$, we can find $S \in \mathcal{A}_u$ and $T \in \mathcal{A}_v$ with $S \cap T = \emptyset$, and then the same sets S and T satisfy $S \in \mathcal{B}_u, T \in \mathcal{B}_v$ and $S \cap T = \emptyset$. Condition 2 holds because we picked each $\mathcal{B}_v$ to be intersecting.

Furthermore, we can colour the vertex v with 'colour' $\mathcal{B}_v$ (i.e., vertices v and w receive the same colour if and only if $\mathcal{B}_v = \mathcal{B}_w$). Then adjacent vertices receive distinct colours (since each $\mathcal{B}_v$ is intersecting but $\mathcal{B}_v \cup \mathcal{B}_w$ is not whenever $vw \in E(G)$), and each colour comes from the set of maximal intersecting families. It follows that the number of maximal intersecting families over $[k]$ is at least $\chi(G)$.

Conversely, assume that k is a positive integer such that there are at least $\chi(G)$ distinct maximal intersecting families $\mathcal{B}_1, \ldots, \mathcal{B}_{\chi(G)}$ over $[k]$. Let $c : V(G) \to [\chi(G)]$ be a proper vertex-colouring of G, and set $\mathcal{A}_v = \mathcal{B}_{c(v)}$ for each v. Certainly each $\mathcal{A}_v$ is intersecting. Furthermore, by maximality, no $\mathcal{A}_v \cup \mathcal{A}_w$ can be intersecting when $c(v) \neq c(w)$, and hence $\mathcal{A}_v \cup \mathcal{A}_w$ is not intersecting when $vw \in E(G)$. It follows that $(\mathcal{A}_v)_{v \in V(G)}$ satisfies both conditions in Lemma 13.2.1 and so $\sigma(G) \leq k$. $\square$

Example. Consider the orientation covering of K_4 shown on Figure 13.1. Following the proof of Lemma 13.2.1, this corresponds to the families $\mathcal{A}_a = \{\{1,3\}, \{1,2\}, \{1\}\}$, $\mathcal{A}_b = \{\{2\}, \{1,2\}\}$, $\mathcal{A}_c = \{\{3\}\}, \mathcal{A}_d = \{\{2,3\}, \{1,3\}, \{1,2\}\}$. As in the proof of Theorem 13.1.3, we can extend these to get the maximal intersecting families $\mathcal{B}_a = \{S : S \subseteq [3], 1 \in S\}$, $\mathcal{B}_b = \{S : S \subseteq [3], 2 \in S\}$, $\mathcal{B}_c = \{S : S \subseteq [3], 3 \in S\}$ and $\mathcal{B}_d = \{\{1,2\},\{1,3\},\{2,3\},\{1,2,3\}\}$. We can also follow the proofs of Theorem 13.1.3 and Lemma 13.2.1 and use these maximal intersecting families to recover an orientation covering of K_4. Note that $\mathcal{B}_a, \mathcal{B}_b, \mathcal{B}_c, \mathcal{B}_d$ is a complete list of maximal intersecting families over [3], so we need more that 3 orientations for K_m when $m > 4$ (by Theorem 13.1.3).

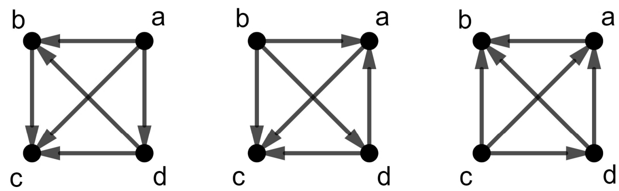

Figure 13.1: An orientation covering of K_4.

Bibliography

[1] R. Ahlswede and G. Simonyi. On the optimal structure of recovering set pairs in lattices: the sandglass conjecture. *Discrete Mathematics*, 128(1-3):389–394, 1994.

[2] N. Alon. An extremal problem for sets with applications to graph theory. *Journal of Combinatorial Theory, Series A*, 40(1):82–89, 1985.

[3] N. Alon and I. Ben-Eliezer. Local rainbow colorings. *Journal of Combinatorics*, 2(2): 293–304, 2011.

[4] N. Alon and A. Shapira. A characterization of easily testable induced subgraphs. *Combinatorics, Probability and Computing*, 15(6):791–805, 2006.

[5] N. Alon and A. Shapira. On an extremal hypergraph problem of Brown, Erdős and Sós. *Combinatorica*, 26(6):627–645, 2006.

[6] N. Alon and C. Shikhelman. Many T copies in H-free graphs. *Journal of Combinatorial Theory, Series B*, 121:146–172, 2016.

[7] N. Alon, H. Lefmann, and V. Rödl. On an anti-Ramsey type result. In *Sets, Graphs and Numbers*, volume 60 of *Colloquia Mathematica Societatis János Bolyai*, pages 9–22. North-Holland, Amsterdam, 1991.

[8] N. Alon, A. Moitra, and B. Sudakov. Nearly complete graphs decomposable into large induced matchings and their applications. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 1079–1090, 2012.

[9] H. J. van Alphen. Uitbreiding van een stelling von Besicovitch. *Mathematica (Zutphen) B*, 10:144–157, 1942.

[10] D. Amir. Chebyshev centers and uniform convexity. *Pacific Journal of Mathematics*, 77 (1):1–6, 1978.

[11] D. Amir. Best simultaneous approximation (Chebyshev centers). In *Parametric Optimization and Approximation*, pages 19–35. Springer, 1984.

[12] O. Angel and A. E. Holroyd. Perfect shuffling by lazy swaps. *Electronic Communications in Probability*, 23:1–11, 2018.

[13] J. Balogh, M. Delcourt, E. Heath, and L. Li. Generalized rainbow Turán numbers of odd cycles. *Discrete Mathematics*, 345(2):112663, 2022.

[14] J. Balogh, C. Chen, K. Hendrey, B. Lund, H. Luo, C. Tompkins, and T. Tran. Maximal 3-wise intersecting families. *Combinatorica*, to appear.

[15] J. Balogh, G. Kronenberg, A. Pokrovskiy, and T. Szabó. The maximum length of K_r-bootstrap percolation. *Proceedings of the American Mathematical Society*, to appear.

[16] F. A. Behrend. On sets of integers which contain no three terms in arithmetical progression. *Proceedings of the National Academy of Sciences of the United States of America*, 32(12): 331–332, 1946.

[17] F. Benevides and M. Przykucki. On slowly percolating sets of minimal size in bootstrap percolation. *Electronic Journal of Combinatorics*, 20(2):P46, 2013.

[18] F. Benevides and M. Przykucki. Maximum percolation time in two-dimensional bootstrap percolation. *SIAM Journal on Discrete Mathematics*, 29(1):224–251, 2015.

[19] A. Besicovitch. Sur deux questions de l'integrabilité. *Journal of the Physical mathematical society of the State Perm University*, 2, 1920.

[20] A. Besicovitch. On Kakeya's problem and a similar one. *Mathematische Zeitschrift*, 27(1): 312–320, 1928.

[21] A. Blokhuis. More on maximal intersecting families of finite sets. *Journal of Combinatorial theory, Series A*, 44(2):299–303, 1987.

[22] B. Bollobás. On generalized graphs. *Acta Mathematica Academiae Scientiarum Hungaricae*, 16(3-4):447–452, 1965.

[23] B. Bollobás. Weakly k-saturated graphs. In *Beiträge zur Graphentheorie (Kolloquium, Manebach, 1967)*, pages 25–31. Teubner, Leipzig, 1968.

[24] B. Bollobás. *Extremal graph theory.* Academic Press, 1978.

[25] B. Bollobás. *Combinatorics: set systems, hypergraphs, families of vectors, and combinatorial probability.* Cambridge University Press, 1986.

[26] B. Bollobás and E. Győri. Pentagons vs. triangles. *Discrete Mathematics*, 308(19):4332–4336, 2008.

[27] B. Bollobás and A. Thomason. Projections of bodies and hereditary properties of hypergraphs. *Bulletin of the London Mathematical Society*, 27(5):417–424, 1995.

[28] B. Bollobás, M. Przykucki, O. Riordan, and J. Sahasrabudhe. On the maximum running time in graph bootstrap percolation. *Electronic Journal of Combinatorics*, 24(2):P2.16, 2017.

[29] B. Bollobás and P. Erdős. On a Ramsey–Turán type problem. *Journal of Combinatorial Theory, Series B*, 21(2):166–168, 1976.

[30] J. A. Bondy and M. Simonovits. Cycles of even length in graphs. *Journal of Combinatorial Theory, Series B*, 16(2):97–105, 1974.

[31] E. Boros, Z. Füredi, and J. Kahn. Maximal intersecting families and affine regular polygons in $PG(2, q)$. *Journal of Combinatorial Theory, Series A*, 52(1):1–9, 1989.

[32] R. C. Bose and S. Chowla. Theorems in the additive theory of numbers. *Commentarii Mathematici Helvetici*, 37:141–147, 1962.

[33] J. Bourgain. On the dimension of Kakeya sets and related maximal inequalities. *Geometric and Functional Analysis*, 9(2):256–282, 1999.

[34] A. E. Brouwer, C. F. Mills, W. H. Mills, and A. Verbeek. Counting families of mutually intesecting sets. *Electronic Journal of Combinatorics*, 20(2):P8, 2013.

[35] D. Bulavka, M. Tancer, and M. Tyomkyn. Weak saturation of multipartite hypergraphs. *Combinatorica*, to appear.

[36] R. T. Bumby. A problem with telephones. *SIAM Journal on Algebraic Discrete Methods*, 2(1):13–18, 1981.

[37] X. Cheng and Z. Xu. Local rainbow colorings for various graphs. *arXiv preprint arXiv:2207.07532*, 2022.

[38] D. Conlon and J. Fox. Graph removal lemmas. In *Surveys in Combinatorics 2013*, pages 1–50. Cambridge University Press, 2013.

[39] D. Conlon, J. Fox, C. Lee, and B. Sudakov. The Erdős–Gyárfás problem on generalized Ramsey numbers. *Proceedings of the London Mathematical Society*, 110(1):1–18, 2015.

[40] D. Conlon, J. Fox, C. Lee, and B. Sudakov. On the grid Ramsey problem and related questions. *International Mathematics Research Notices*, 2015(17):8052–8084, 2015.

[41] D. Conlon, J. Fox, and B. Sudakov. Recent developments in graph Ramsey theory. In *Surveys in Combinatorics 2015*, pages 49–118. Cambridge University Press, 2015.

[42] M. Csörnyei, K. Héra, and M. Laczkovich. Closed sets with the Kakeya property. *Mathematika*, 63(1):184–195, 2017.

[43] F. Cunningham Jr. The Kakeya problem for simply connected and for star-shaped sets. *The American Mathematical Monthly*, 78(2):114–129, 1971.

[44] A. Czumaj. Random permutations using switching networks. In *Proceedings of the forty-seventh annual ACM symposium on Theory of Computing*, pages 703–712, 2015.

[45] S. Das, C. Lee, and B. Sudakov. Rainbow Turán problem for even cycles. *European Journal of Combinatorics*, 34(5):905–915, 2013.

[46] R. O. Davies. Some remarks on the Kakeya problem. *Mathematical Proceedings of the Cambridge Philosophical Society*, 69(3):417–421, 1971.

[47] S. J. Dow, D. A. Drake, Z. Füredi, and J. A. Larson. A lower bound for the cardinality of a maximal family of mutually intersecting sets of equal size. In *Proceedings of the sixteenth Southeastern international conference on combinatorics, graph theory and computing*, volume 48, pages 47–48, 1985.

[48] B. Dushnik. Concerning a certain set of arrangements. *Proceedings of the American Mathematical Society*, 1(6):788–796, 1950.

[49] Z. Dvir. On the size of Kakeya sets in finite fields. *Journal of the American Mathematical Society*, 22(4):1093–1097, 2009.

[50] D. Eichhorn and D. Mubayi. Edge-coloring cliques with many colors on subcliques. *Combinatorica*, 20(3):441–444, 2000.

[51] K. Engel, T. Mitsis, C. Pelekis, and C. Reiher. Projection inequalities for antichains. *Israel Journal of Mathematics*, 238(1):61–90, 2020.

[52] P. Erdős and D. J. Kleitman. Extremal problems among subsets of a set. *Discrete Mathematics*, 8(3):281–294, 1974.

[53] P. Erdős. Problems and results on finite and infinite graphs. In *Recent advances in graph theory (Proc. Second Czechoslovak Sympos., Prague, 1974)*, pages 183–192. Academia, Prague, 1975.

[54] P. Erdős and A. Gyárfás. A variant of the classical Ramsey problem. *Combinatorica*, 17: 459–467, 1997.

[55] P. Erdős and L. Lovász. Problems and results on 3-chromatic hypergraphs and some related questions. In *Infinite and Finite Sets*, volume 10 of *Colloquia Mathematica Societatis János Bolyai*, pages 609–627. North-Holland, Amsterdam, 1975.

[56] P. Erdős and P. Turán. On a problem of Sidon in additive number theory, and on some related problems. *Journal of the London Mathematical Society*, 16(4):212–215, 1941.

[57] B. Ergemlidze, E. Győri, and A. Methuku. On the Rainbow Turán number of paths. *Electronic Journal of Combinatorics*, 26(1):P1.17, 2019.

[58] L. Esperet, J. Gimbel, and A. King. Covering line graphs with equivalence relations. *Discrete Applied Mathematics*, 158(17):1902–1907, 2010.

[59] A. Espuny Díaz, B. Janzer, G. Kronenberg, and J. Lada. Long running times for hypergraph bootstrap percolation. *European Journal of Combinatorics*, 115:103783, 2024.

[60] R. H. Etkin and E. Ordentlich. Analysis of deterministic binary interference channels via a general outer bound. *IEEE Transactions on Information Theory*, 57(5):2597–2604, 2011.

[61] D. Fabian, P. Morris, and T. Szabó. Slow graph bootstrap percolation I: Cycles. *arXiv preprint arXiv:2308.00498*, 2023.

[62] K. J. Falconer. Continuity properties of k-plane integrals and Besicovitch sets. *Mathematical Proceedings of the Cambridge Philosophical Society*, 87(2):221–226, 1980.

[63] C. Fefferman. The multiplier problem for the ball. *Annals of Mathematics*, 94(2):330–336, 1971.

[64] J. Fitzsimons. What is the most efficient way to generate a random permutation from probabilistic pairwise swaps? Theoretical Computer Science Stack Exchange, 2011. `https://cstheory.stackexchange.com/q/5321`.

[65] J. Foran. The length of the graph of a one to one function from [0, 1] to [0, 1]. *Real Analysis Exchange*, 25(2):809–816, 1999.

[66] P. Frankl. An extremal problem for two families of sets. *European Journal of Combinatorics*, 3(2):125–127, 1982.

[67] P. Frankl and Z. Füredi. Union-free hypergraphs and probability theory. *European Journal of Combinatorics*, 5(2):127–131, 1984.

[68] Z. Füredi and T. Jiang. Hypergraph Turán numbers of linear cycles. *Journal of Combinatorial Theory, Series A*, 123(1):252–270, 2014.

[69] Z. Füredi and M. Simonovits. The history of degenerate (bipartite) extremal graph problems. In *Erdős Centennial*, pages 169–264. Springer, 2013.

[70] Z. Füredi, P. Hajnal, V. Rödl, and W. T. Trotter. Interval orders and shift graphs. In *Sets, Graphs and Numbers*, volume 60 of *Colloquia Mathematica Societatis János Bolyai*, pages 297–313. North-Holland, Amsterdam, 1991.

[71] D. Gerbner, E. Győri, A. Methuku, and M. Vizer. Generalized Turán problems for even cycles. *Journal of Combinatorial Theory, Series B*, 145:169–213, 2020.

[72] D. Gerbner, A. Methuku, T. Mészáros, and C. Palmer. Generalized rainbow Turán numbers. *Electronic Journal of Combinatorics*, 29(2):P2.44, 2022.

[73] L. Gishboliner and A. Shapira. A generalized Turán problem and its applications. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 760–772, 2018.

[74] W. T. Gowers and B. Janzer. Generalizations of the Ruzsa–Szemerédi and rainbow Turán problems for cliques. *Combinatorics, Probability and Computing*, 30:591–608, 2021.

[75] R. L. Graham, B. L. Rothschild, and J. H. Spencer. *Ramsey theory*. Wiley, 2nd edition, 1990.

[76] C. Groenland, T. Johnston, J. Radcliffe, and A. Scott. Perfect shuffling with fewer lazy transpositions. *arXiv preprint arXiv:2208.06629*, 2022.

[77] C. Groenland, T. Johnston, J. Radcliffe, and A. Scott. Short reachability networks. *arXiv preprint arXiv:2208.06630*, 2022.

[78] K. Gunderson, S. Koch, and M. Przykucki. The time of graph bootstrap percolation. *Random Structures & Algorithms*, 51(1):143–168, 2017.

[79] E. Győri. Triangle-free hypergraphs. *Combinatorics, Probability and Computing*, 15(1-2): 185–191, 2006.

[80] E. Győri and N. Lemons. Hypergraphs with no cycle of a given length. *Combinatorics, Probability and Computing*, 21(1-2):193–201, 2012.

[81] E. Győri and H. Li. The maximum number of triangles in C_{2k+1}-free graphs. *Combinatorics, Probability and Computing*, 21(1-2):187–191, 2012.

[82] I. Hartarsky and L. Lichev. The maximal running time of hypergraph bootstrap percolation. *arXiv preprint arXiv:2208.13489*, 2022.

[83] K. Héra and M. Laczkovich. The Kakeya problem for circular arcs. *Acta Mathematica Hungarica*, 150(2):479–511, 2016.

[84] S. Hoşten and W. D. Morris, Jr. The order dimension of the complete graph. *Discrete Mathematics*, 201(1-3):133–139, 1999.

[85] R. Holzman and J. Körner. Cancellative pairs of families of sets. *European Journal of Combinatorics*, 16(3):263–266, 1995.

[86] H. Huang and J. Ma. On tight cycles in hypergraphs. *SIAM Journal on Discrete Mathematics*, 33(1):230–237, 2019.

[87] B. Janzer. A new upper bound for cancellative pairs. *Electronic Journal of Combinatorics*, 25(2):P2.13, 2018.

[88] B. Janzer. A note on antichains in the continuous cube. *Mathematika*, 66:514–516, 2020.

[89] B. Janzer. Projections of antichains. *Electronic Journal of Combinatorics*, 27(1):P1.54, 2020.

[90] B. Janzer. A note on the orientation covering number. *Discrete Applied Mathematics*, 304: 349–351, 2021.

[91] B. Janzer. Large hypergraphs without tight cycles. *Combinatorial Theory*, 1:Paper No. 12, 2021.

[92] B. Janzer. Rotation inside convex Kakeya sets. *arXiv preprint arXiv:2209.09728*, 2022.

[93] B. Janzer. A note on saturation for k-wise intersecting families. *Combinatorial Theory*, 2(2):Paper No. 11, 2022.

[94] B. Janzer. The generalized rainbow Turán problem for cycles. *SIAM Journal on Discrete Mathematics*, 36:436–448, 2022.

[95] B. Janzer and O. Janzer. On locally rainbow colourings. *arXiv preprint arXiv:2304.12260*, 2023.

[96] B. Janzer, J. R. Johnson, and I. Leader. Partial shuffles by lazy swaps. *arXiv preprint arXiv:2210.13286*, 2022.

[97] O. Janzer. Rainbow Turán number of even cycles, repeated patterns and blow-ups of cycles. *Israel Journal of Mathematics*, 253:813–840, 2023.

[98] E. Järvenpää, M. Järvenpää, T. Keleti, and A. Máthé. Continuously parametrized Besicovitch sets in $\mathbf{R}^n$. *Annales AcademiæScientiarum Fennicæ. Mathematica*, 36(2):411–421, 2011.

[99] T. Jiang and J. Ma. Cycles of given lengths in hypergraphs. *Journal of Combinatorial Theory, Series B*, 133:54–77, 2018.

[100] D. Johnston, C. Palmer, and A. Sarkar. Rainbow Turán problems for paths and forests of stars. *Electronic Journal of Combinatorics*, 24(1):P1.34, 2017.

[101] G. Kalai. Weakly saturated graphs are rigid. *Annals of Discrete Mathematics*, 20:189–190, 1984.

[102] G. Kalai. Hyperconnectivity of graphs. *Graphs and Combinatorics*, 1(1):65–79, 1985.

[103] M. Karchmer. On proving lower bounds for circuit size. In *Proceedings of the Eigth Annual Structure in Complexity Theory Conference*, pages 112–118. IEEE, 1993.

[104] M. Karchmer and A. Wigderson. On span programs. In *Proceedings of the Eigth Annual Structure in Complexity Theory Conference*, pages 102–111. IEEE, 1993.

[105] G. O. H. Katona. Extremal problems for hypergraphs. In *Combinatorics*, volume 56 of *Mathematics Centre Tracts*, pages 13–42. Amsterdam, 1974.

[106] N. Katz and T. Tao. New bounds for Kakeya problems. *Journal d'Analyse Mathématique*, 87(1):231–263, 2002.

[107] P. Keevash, D. Mubayi, B. Sudakov, and J. Verstraëte. Rainbow Turán problems. *Combinatorics, Probability and Computing*, 16(1):109–126, 2007.

[108] L. Kolasa and T. Wolff. On some variants of the Kakeya problem. *Pacific Journal of Mathematics*, 190(1):111–154, 1999.

[109] A. Kostochka, D. Mubayi, and J. Verstraëte. Turán problems and shadows I: Paths and cycles. *Journal of Combinatorial Theory, Series A*, 129:57–79, 2015.

[110] G. Kronenberg, T. Martins, and N. Morrison. Weak saturation numbers of complete bipartite graphs in the clique. *Journal of Combinatorial Theory, Series A*, 178:105357, 2021.

[111] S. Letzter. Hypergraphs with no tight cycles. *Proceedings of the American Mathematical Society*, 151(02):455–462, 2023.

[112] L. H. Loomis and H. Whitney. An inequality related to the isoperimetric inequality. *Bulletin of the American Mathematical Society*, 55(10):961–962, 1949.

[113] L. Lovász. Flats in matroids and geometric graphs. In *Combinatorial Surveys (Proc. 6th British Combinatorial Conference)*, pages 45–86, 1977.

[114] K. Matzke. The saturation time of graph bootstrap percolation. *arXiv preprint arXiv:1510.06156*, 2015.

[115] G. Moshkovitz and A. Shapira. Exact bounds for some hypergraph saturation problems. *Journal of Combinatorial Theory, Series B*, 111:242–248, 2015.

[116] D. Mubayi. Edge-coloring cliques with three colors on all 4-cliques. *Combinatorica*, 18(2): 293–296, 1998.

[117] D. Mubayi, O. Pikhurko, and B. Sudakov. Hypergraph Turán problem: Some open questions. Manuscript (AIM workshop problem list), 2011.

[118] C. Nair and M. Yazdanpanah. On the AND-OR interference channel and the sandglass conjecture. In *2020 IEEE International Symposium on Information Theory (ISIT)*, pages 1540–1545. IEEE, 2020.

[119] J. A. Noel and A. Ranganathan. On the running time of hypergraph bootstrap percolation. *Electronic Journal of Combinatorics*, 30(2):P2.46, 2023.

[120] J. Pál. Ein minimumproblem für ovale. *Mathematische Annalen*, 83(3):311–319, 1921.

[121] O. Pikhurko. Uniform families and count matroids. *Graphs and Combinatorics*, 17(4): 729–740, 2001.

[122] O. Pikhurko. Weakly saturated hypergraphs and exterior algebra. *Combinatorics, Probability and Computing*, 10(5):435–451, 2001.

[123] M. Przykucki. Maximal percolation time in hypercubes under 2-bootstrap percolation. *Electronic Journal of Combinatorics*, 19(2):P41, 2012.

[124] A. A. Razborov. On the method of approximations. In *Proceedings of the twenty-first annual ACM symposium on Theory of computing*, pages 167–176, 1989.

[125] A. C. M. van Rooij and W. H. Schikhof. *A Second Course on Real Functions*. Cambridge University Press, 1982.

[126] I. Z. Ruzsa and E. Szemerédi. Triple systems with no six points carrying three triangles. In *Combinatorics, Vol. II*, volume 18 of *Colloquia Mathematica Societatis János Bolyai*, pages 939–945. North-Holland, Amsterdam-New York, 1978.

[127] R. Salem. On some singular monotonic functions which are strictly increasing. *Transactions of the American Mathematical Society*, 53(3):427–439, 1943.

[128] A. Shapira and M. Tyomkyn. Weakly saturated hypergraphs and a conjecture of Tuza. *Proceedings of the American Mathematical Society*, 151(07):2795–2805, 2023.

[129] J. B. Shearer. A new construction for cancellative families of sets. *Electronic Journal of Combinatorics*, 3(1):R15, 1996.

[130] S. Shelah. Primitive recursive bounds for van der Waerden numbers. *Journal of the American Mathematical Society*, 1(3):683–697, 1988.

[131] M. Sipser. A topological view of some problems in complexity theory. In *Theory of Algorithms*, volume 44 of *Colloquia Mathematica Societatis János Bolyai*, pages 387–391. North-Holland, Amsterdam, 1985.

[132] D. Soltész. New bounds on Simonyi's conjecture. *European Journal of Combinatorics*, 70: 251–267, 2018.

[133] V. Sós, P. Erdős, and W. Brown. On the existence of triangulated spheres in 3-graphs, and related problems. *Periodica Mathematica Hungarica*, 3(3-4):221–228, 1973.

[134] J. Spencer. Minimal scrambling sets of simple orders. *Acta Mathematica Academiae Scientiarum Hungaricae*, 22:349–353, 1971.

[135] B. Sudakov and I. Tomon. The extremal number of tight cycles. *International Mathematics Research Notices*, 2022(13):9663–9684, 2022.

[136] L. M. Tolhuizen. New rate pairs in the zero-error capacity region of the binary multiplying channel without feedback. *IEEE Transactions on Information Theory*, 46(3):1043–1046, 2000.

[137] J. Verstraëte. Extremal problems for cycles in graphs. In *Recent trends in combinatorics*, pages 83–116. Springer, 2016.

[138] A. Wigderson. The fusion method for lower bounds in circuit complexity. In *Combinatorics, Paul Erdos is Eighty*, volume 1, pages 453–468. János Bolyai Mathematical Society, 1993.