Extremal Combinatorics, Graph Limits and Computational Complexity

Jonathan A. Noel

Christ Church

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2016

Acknowledgements

First and foremost, I would like to thank my supervisor Alex Scott for his constant support and encouragement. I have learned a great number of important things from Alex over the past three years which will, without a doubt, have a positive impact on my future career and life.

At every stage of my academic life, I have benefitted greatly from the generosity of people who had little to gain by helping me. Deserving of special mention are (in somewhat chronological order) Robb Fry, Rick Brewster, Sean McGuinness, Bruce Reed, Doug West, Xuding Zhu, Dan Král', Tomáš Kaiser, Zdeněk Dvořák, Benny Sudakov and Rob Morris. While it is unlikely that I will ever be able repay any of these individuals, I hope to be able to 'pay it forward' someday.

It is a very exciting time to be a combinatorialist at Oxford, and I have been incredibly fortunate to be a part of it. I would like to thank everyone in the combinatorics group for their contribution to maintaining such a stimulating research environment.

Graduate studies can feel, at times, like a solitary experience. As I come to the end of it, however, I am realising that I have never been alone – not for one minute. For that, I would like to thank all of my friends for their support. I would especially like to acknowledge Tash, without whom the past few years would have been excessively dull. When I look back on the best and most memorable moments from my time here, you will be in all of them. Thank you for being there for me.

Lastly, I would like to thank my family. Moving (even further) away from home has been a huge sacrifice for me, just as I know that watching me go, and supporting me through it, has been a huge sacrifice for them. In the end, I think that we will all agree that it was worth it.

Abstract

This thesis is primarily focused on problems in extremal combinatorics, although we will also consider some questions of analytic and algorithmic nature.

The *d-dimensional hypercube* is the graph with vertex set $\{0,1\}^d$ where two vertices are adjacent if they differ in exactly one coordinate. In Chapter 2 we obtain an upper bound on the 'saturation number' of Q_m in Q_d. Specifically, we show that for $m \geq 2$ fixed and d large there exists a subgraph G of Q_d of bounded average degree such that G does not contain a copy of Q_m but, for every G' such that $G \subsetneq G' \subseteq Q_d$, the graph G' contains a copy of Q_m. This result answers a question of Johnson and Pinto and is best possible up to a factor of $O(m)$.

In Chapter 3, we show that there exists $\varepsilon > 0$ such that for all k and for n sufficiently large there is a collection of at most $2^{(1-\varepsilon)k}$ subsets of $[n]$ which does not contain a chain of length $k+1$ under inclusion and is maximal subject to this property. This disproves a conjecture of Gerbner, Keszegh, Lemons, Palmer, Pálvölgyi and Patkós. We also prove that there exists a constant $c \in (0,1)$ such that the smallest such collection is of cardinality $2^{(1+o(1))ck}$ for all k.

In Chapter 4, we obtain an exact expression for the 'weak saturation number' of Q_m in Q_d. That is, we determine the minimum number of edges in a spanning subgraph G of Q_d such that the edges of $E(Q_d)\setminus E(G)$ can be added to G, one edge at a time, such that each new edge completes a copy of Q_m. This answers another question of Johnson and Pinto. We also obtain a more general result for the weak saturation of 'axis aligned' copies of a multidimensional grid in a larger grid.

In the r*-neighbour bootstrap process*, one begins with a set A_0 of 'infected' vertices in a graph G and, at each step, a 'healthy' vertex becomes infected if it has at least r infected neighbours. If every vertex of G is eventually infected, then we say that A_0 *percolates*. In Chapter 5, we apply ideas from weak saturation to prove that, for fixed $r \geq 2$, every percolating set in Q_d has cardinality at least $\frac{1+o(1)}{r}\binom{d}{r-1}$. This confirms a conjecture of Balogh and Bollobás and is asymptotically best possible. In addition, we determine the minimum cardinality exactly in the case $r = 3$ (the minimum cardinality in the case $r = 2$ was already known).

In Chapter 6, we provide a framework for proving lower bounds on the number of comparable pairs in a subset S of a partially ordered set (poset) of prescribed size. We apply this framework to obtain an explicit bound of this type for the poset $\mathcal{V}(q, n)$ consisting of all subspaces of $\mathbb{F}_q^n$ ordered by inclusion which is best possible when S is not too large.

In Chapter 7, we apply the result from Chapter 6 along with the recently developed 'container method,' to obtain an upper bound on the number of antichains in $\mathcal{V}(q, n)$ and a bound on the size of the largest antichain in a p-random subset of $\mathcal{V}(q, n)$ which holds with high probability for p in a certain range.

In Chapter 8, we construct a 'finitely forcible graphon' W for which there exists a sequence $(\varepsilon_i)_{i=1}^{\infty}$ tending to zero such that, for all $i \geq 1$, every weak ε_i-regular partition of W has at least $\exp\left(\varepsilon_i^{-2}/2^{5\log^* \varepsilon_i^{-2}}\right)$ parts. This result shows that the structure of a finitely forcible graphon can be much more complex than was anticipated in a paper of Lovász and Szegedy.

For positive integers p, q with $p/q \geq 2$, a *circular* (p, q)*-colouring* of a graph G is a mapping $V(G) \to \mathbb{Z}_p$ such that any two adjacent vertices are mapped to elements of $\mathbb{Z}_p$ at distance at least q from one another. The reconfiguration problem for circular colourings asks, given two (p, q)-colourings f and g of G, is it possible to transform f into g by recolouring one vertex at a time so that every intermediate mapping is a (p, q)-colouring? In Chapter 9, we show that this question can be answered in polynomial time for $2 \leq p/q < 4$ and is PSPACE-complete for $p/q \geq 4$.

Statement of Originality

This thesis contains no material which has been accepted in whole, or in part, for any other degree or diploma. Chapters 2, 3 and 4 are based on two joint papers with Natasha Morrison and Alex Scott [170, 171]. Chapter 5 is based on a joint paper with Natasha Morrison [169]. Chapters 6 and 7 are based on a joint paper with Alex Scott and Benny Sudakov [176]. Chapter 8 is based on a joint paper with Jacob Cooper, Tomáš Kaiser and Daniel Král' [74]. Chapter 9 is based on a joint paper with Richard Brewster, Sean McGuinness and Benjamin Moore [55].

Contents

List of Figures

Chapter 1

Introduction

The central focus of extremal combinatorics is on understanding the characteristics (size, structure, etc) of combinatorial objects which maximise or minimise a particular parameter subject to certain constraints. A standard question in extremal combinatorics is the following.

For a fixed graph H, what is the maximum number of edges in a graph on n vertices which does not contain a copy of H? (1.1)

This maximum number, known as the *extremal number* for H, is denoted by $\mathrm{ex}(n, H)$. One of the cornerstones of extremal combinatorics is Turán's Theorem [215] which provides an exact expression for $\mathrm{ex}(n, K_r)$ and says that the unique graph attaining the maximum is the complete multipartite graph in which every part is of size $\lfloor n/(r-1) \rfloor$ or $\lceil n/(r-1) \rceil$. The special case $r = 3$ of Turán's Theorem was first proved by Mantel [164] in 1907 and is among the oldest theorems in extremal combinatorics.

In this thesis, we study several topics in extremal combinatorics which are vaguely related to the *Turán-type* problems discussed above: namely, *saturation*, *weak saturation*, *bootstrap percolation* and *supersaturation*. We also prove a theorem on *finitely forcible* graph limits which is analytic in nature, but has its roots in extremal graph theory. In addition, we study an algorithmic *reconfiguration* problem for graph homomorphisms which, admittedly, has little to do with extremal combinatorics, but is interesting in its own right. We provide a brief introduction to these topics below; additional background to each topic can be found at the beginning of the relevant chapter.

1.1 Saturated Subgraphs of Minimum Size

Given graphs F and H, a spanning subgraph G of F is said to be (F,H)-*saturated* if G does not contain a copy of H but, for every edge $e \in E(F) \setminus E(G)$, the graph $G+e$ contains a copy of H. It is easy to see that every H-free graph with $\mathrm{ex}(n,H)$ edges is (K_n,H)-saturated and, therefore, question (1.1) can be rephrased as asking for the maximum number of edges in a (K_n,H)-saturated graph. We are interested in a similar question, but with the word 'maximum' replaced by 'minimum.'

For fixed graphs F and H, what is the minimum number of edges in an (F,H)-saturated graph? (1.2)

This minimum number is known as the *saturation number* of H in F and is denoted by $\mathrm{sat}(F,H)$.

In Chapter 2, we prove an upper bound on $\mathrm{sat}(Q_d,Q_m)$ where Q_d is the d-*dimensional hypercube*; i.e. the graph with vertex set $\{0,1\}^d$ where two vertices are adjacent if they differ on exactly one coordinate. Our proof is constructive; it exploits the existence of an efficient dominating set in hypercubes known as a Hamming code [118] and a construction of Conder [70] of a 3-colouring of the edges of Q_d with no monochromatic cycle of length 4 or 6. The following result answers a question of Johnson and Pinto [133] and is best possible up to a factor of $O(m)$; see (1.9).

Theorem 1.3 (Morrison, Noel and Scott [170])**.** *For fixed $m \geq 2$ and $d \to \infty$, we have*

$$\mathrm{sat}(Q_d,Q_m) \leq (1+o(1))72m^2 2^d.$$

1.2 A Saturation Problem for Set Systems

Define $[n] := \{1,\dots,n\}$ and let $\mathcal{P}(n)$ be the poset of subsets of $[n]$ ordered by inclusion. An *antichain* in $\mathcal{P}(n)$ (or, in any poset) is a set of pairwise incomparable elements, and a k-*chain* is a set of k pairwise comparable elements. Another cornerstone of extremal combinatorics is Sperner's Theorem [210] which says that the size of the largest antichain in $\mathcal{P}(n)$ is $\binom{n}{\lfloor n/2\rfloor}$. Erdős [83] generalised Sperner's Theorem by proving that the cardinality of a collection $\mathcal{F} \subseteq \mathcal{P}(n)$ containing no $(k+1)$-chain is at most the sum of the k largest binomial coefficients $\binom{n}{i}$.

In Chapter 3, we study a 'saturation analogue' of Erdős' result, first considered by Gerbner, Keszegh, Lemons, Palmer, Pálvölgyi and Patkós [104]. Say that a collection $\mathcal{F} \subseteq \mathcal{P}(n)$ is $(k+1)$-*chain saturated* if $\mathcal{F}$ does not contain a $(k+1)$-chain but $\mathcal{F} \cup \{S\}$ contains a $(k+1)$-chain for every set $S \in \mathcal{P}(n) \setminus \mathcal{F}$. We let chain-sat$(n,k)$ denote the size of the smallest $(k+1)$-chain saturated collection in $\mathcal{P}(n)$. In [104] it is proved that chain-sat(n,k) = chain-sat(m,k) if n and m are sufficiently large with respect to k, and so we can define

$$\text{chain-sat}(k) := \lim_{n\to\infty} \text{chain-sat}(n,k).$$

Gerbner et al. [104] conjectured that chain-sat$(k) = 2^{k-1}$ for all $k \geq 1$ and proved the upper bound. The conjecture is trivially true for $k \in \{1,2\}$ and was shown to be true for $k = 3$ in [104] and for $k \in \{4,5\}$ in [171].

In Chapter 3, we disprove the conjecture of [104] for all $k \geq 6$ in a strong sense. First, we show that there is a 7-chain saturated collection in $\mathcal{P}(n)$ of cardinality $30 < 2^5$ for all n sufficiently large. Note that, since chain-sat$(k+1) \leq 2\,$chain-sat(k) for all $k \geq 1$ (see [104] for a proof), this is already enough to disprove the conjecture for all $k \geq 6$. We go a step further by obtaining the following exponential improvement on the upper bound chain-sat$(k) \leq 2^{k-1}$ proved in [104].

Theorem 1.4 (Morrison, Noel and Scott [171])**.** *There exists $\varepsilon > 0$ such that, for all k,* chain-sat$(k) \leq 2^{(1-\varepsilon)k}$.

The proof of the above theorem involves proving that chain-sat(k) satisfies a certain submultiplicativity bound. The following result is a simple by-product of this bound (via Fekete's Lemma).

Theorem 1.5 (Morrison, Noel and Scott [171])**.** *There exists a positive constant $c \in (0,1)$ such that* chain-sat$(k) = 2^{(1+o(1))ck}$.

1.3 Weakly Saturated Subgraphs of Minimum Size

We also consider the following 'relaxation' of question (1.2).

> *For fixed graphs F and H, what is the minimum number of edges in a spanning subgraph G of F such that the edges of $E(F) \setminus E(G)$ can be added to G, one edge at a time, in such a way that each edge completes a copy of H when it is added?* (1.6)

Such a graph G is said to be *weakly (F, H)-saturated*. The *weak saturation number* of F in H, denoted by $\text{wsat}(F, H)$, is the minimum number of edges in a weakly (F, H)-saturated graph.

Trivially, every (F, H)-saturated graph is also weakly (F, H)-saturated and, therefore, for all graphs F and H,

$$\text{sat}(F, H) \geq \text{wsat}(F, H). \tag{1.7}$$

On the other hand, for $n > r \geq 3$, there exist weakly (K_n, K_r)-saturated graphs with $\text{sat}(K_n, K_r)$ edges which are not (K_n, K_r)-saturated. For example, every n-vertex tree is weakly (K_n, K_3)-saturated, but the only (K_n, K_3)-saturated graph with $n - 1$ edges is the star $K_{1,n-1}$.

In Chapter 4, we determine the weak saturation number of Q_m in Q_d. This answers another question of Johnson and Pinto [133]. In the proof, we apply a linear algebraic lemma of Balogh, Bollobás, Morris and Riordan [21], which we shall state and prove in Section 4.1.

Theorem 1.8 (Morrison, Noel and Scott [170])**.** *For $d \geq m \geq 1$, we have*

$$\text{wsat}\left(Q_d, Q_m\right) = (m-1)2^d - \sum_{j=0}^{m-2} (m-1-j)\binom{d}{j}.$$

Combining Theorem 1.8 with (1.7), we obtain the following asymptotic lower bound on the saturation number for Q_m in Q_d. For fixed $m \geq 2$ and $d \to \infty$,

$$\text{sat}(Q_d, Q_m) \geq (m - 1 - o(1))2^d. \tag{1.9}$$

This improves on a lower bound of $(m + 1 - o(1))2^{d-1}$ proved by Johnson and Pinto [133]. Theorem 1.8 is actually a special case of a more general theorem on the weak saturation number of 'axis-aligned' copies of a multi-dimensional grid in a larger grid. This result is stated and proved in Chapter 4.

1.4 Bootstrap Percolation in the Hypercube

Given a positive integer r and a (finite or infinite) graph G, the *r-neighbour bootstrap process* begins with an initial set of 'infected' vertices in G and, at each step of the process, a 'healthy' vertex becomes infected if it has at least r infected neighbours.

More formally, if A_0 is the initial set of infected vertices, then, for $j \geq 1$, the set of vertices that are infected after the jth step of the process is defined by

$$A_j := A_{j-1} \cup \{v \in V(G) : |N_G(v) \cap A_{j-1}| \geq r\}$$

where $N_G(v)$ denotes the set of all vertices of G which are adjacent to v. We say that A_0 *percolates* if $\bigcup_{j=0}^{\infty} A_j = V(G)$. The minimum cardinality of a set which percolates in G under the r-neighbour bootstrap process is denoted by $m(G, r)$.

Motivated by potential applications in a probabilistic setting, Balogh and Bollobás [15] conjectured that $m(Q_d, r) = \frac{1+o(1)}{r}\binom{d}{r-1}$ for fixed r and $d \to \infty$ (see also [19, 21]). As was demonstrated in [19], the upper bound is not difficult to prove; see Construction 5.3. In Chapter 5, we prove a matching asymptotic lower bound. In the following result note that the expression in the lower bound is dominated by the $j = 1$ term in the sum; i.e., by $\frac{1}{r}\binom{d-2}{r-1}$.

Theorem 1.10 (Morrison and Noel [169]). *For $d \geq r \geq 2$, we have*

$$m(Q_d, r) \geq 2^{r-1} + \sum_{j=1}^{r-1} \binom{d-j-1}{r-j} \frac{j2^{j-1}}{r}$$

where, by convention, $\binom{a}{b} = 0$ when $a < b$.

The key to our proof is a connection between bootstrap percolation and weak saturation which is illustrated by the following bound. For any graph G and $r \geq 1$,

$$m(G, r) \geq \frac{\text{wsat}(G, K_{1,r+1})}{r}. \tag{1.11}$$

We prove (1.11) in Section 5.1. Given (1.11) we see that, in order to prove the conjecture of Balogh and Bollobás [15], it suffices to show that $\text{wsat}(Q_d, K_{1,r+1}) \geq (1+o(1))\binom{d}{r-1}$. We apply the linear algebraic lemma of Balogh, Bollobás, Morris and Riordan [21] once again to obtain an exact expression for this quantity. Note that the following result implies Theorem 1.10 via (1.11).

Theorem 1.12 (Morrison and Noel [169]). *For $d \geq r \geq 0$, we have*

$$\text{wsat}\left(Q_d, K_{1,r+1}\right) = r2^{r-1} + \sum_{j=1}^{r-1} \binom{d-j-1}{r-j} j2^{j-1}.$$

We also generalise Theorem 1.12 to general multidimensional grids; a precise statement can be found in Section 5.3. In addition, we provide a recursive construction which improves on the best known upper bound on $m(Q_d, r)$. With a bit more work, we use this construction to show that Theorem 1.10 is tight in the case $r = 3$.

Theorem 1.13 (Morrison and Noel [169]). *For $d \geq 3$, we have*

$$m(Q_d, 3) = \left\lceil \frac{d(d+3)}{6} \right\rceil + 1.$$

1.5 Supersaturation in Posets

A natural extension of Turán's problem (question (1.1)) is to ask, given a graph H and integers n and m such that $m > \mathrm{ex}(n, H)$, what is the minimum number of copies of H in a graph with n vertices and m edges? This is known as the *supersaturation* problem for graphs, and is very well studied. In Chapter 6, we investigate supersaturation problems for posets which take the following form.

Given a poset P and an integer m greater than the size of the largest antichain in P, what is the minimum number of comparable pairs in a subset S of P of cardinality m? (1.14)

An early result in this direction is a theorem of Kleitman [141] which says that the number of comparable pairs in a subset of $\mathcal{P}(n)$ of prescribed size is minimised by a collection consisting of sets whose cardinalities are as close to $n/2$ as possible.[1] This result strengthens Sperner's Theorem [210] and settled a conjecture of Erdős and Katona (see [141] or [76]). Quantitatively speaking, a weak form of Kleitman's result says that if the cardinality of $S \subseteq \mathcal{P}(n)$ exceeds the sum of the k largest binomial coefficients $\binom{n}{i}$ by t, then S must contain at least $t\binom{\lceil (n+k)/2 \rceil}{k} = \Omega(tn^k)$ comparable pairs.

In Chapter 6, we provide a general framework for proving supersaturation results in posets based on counting comparable pairs relative to a random chain. We then apply this framework to obtain (a) an alternative proof of an approximate version of the aforementioned theorem of Kleitman [141], and (b) a lower bound on the number of comparable pairs in a subset S of the poset $\mathcal{V}(q, n)$ consisting of all subspaces of $\mathbb{F}_q^n$ ordered by inclusion. In order to state the latter result, we require a few definitions.

[1] Das, Gan and Sudakov [76] further refined Kleitman's argument and characterised the tight examples.

Definition 1.15. Given a subset S of a poset P, let $\mathrm{comp}(S)$ denote the number of comparable pairs in S.

Definition 1.16. Given a prime power q and integers n and k with $0 \leq k \leq n$, define $\begin{bmatrix} n \\ k \end{bmatrix}_q$ to be the number of subspaces of $\mathbb{F}_q^n$ of dimension k.

Remark 1.17. Elementary counting shows that

$$\begin{bmatrix} n \\ k \end{bmatrix}_q := \prod_{i=0}^{k-1} \frac{1-q^{n-i}}{1-q^{i+1}}.$$

One may also observe that, for fixed n and q, the sequence $\begin{bmatrix} n \\ 0 \end{bmatrix}_q, \dots, \begin{bmatrix} n \\ n \end{bmatrix}_q$ is unimodal and satisfies $\begin{bmatrix} n \\ k \end{bmatrix}_q = \begin{bmatrix} n \\ n-k \end{bmatrix}_q$ for all $0 \leq k \leq n$.

Given an integer m, a natural way of constructing a subset S of $\mathcal{V}(q,n)$ of cardinality m with a small number of comparable pairs is to take m subspaces of dimension as close to $n/2$ as possible (analogous to the construction in $\mathcal{P}(n)$ mentioned above). If m is of the form $\begin{bmatrix} n \\ \lceil \frac{n}{2} \rceil \end{bmatrix}_q + t$ for $0 \leq t \leq \begin{bmatrix} n \\ \lfloor (n-1)/2 \rfloor \end{bmatrix}_q$, then such a collection will contain precisely $t \begin{bmatrix} \lceil (n+1)/2 \rceil \\ 1 \end{bmatrix}_q$ comparable pairs. In Chapter 6, we prove following theorem which, in particular, implies that this construction is best possible for all m in this range. Our method of proof involves counting comparisons relative to a random chain and is likely to be applicable to other posets as well.

Theorem 1.18 (Noel, Scott and Sudakov [176])**.** *Let q be a prime power and let k be a fixed positive integer. Then there exists a constant $n_0(k)$ such that if $n \geq n_0(k)$ and $S \subseteq \mathcal{V}(q,n)$ has cardinality at least*

$$\sum_{r=0}^{k-1} \begin{bmatrix} n \\ \lceil \frac{n-k+1+2r}{2} \rceil \end{bmatrix}_q + t,$$

then

$$\mathrm{comp}(S) \geq t \begin{bmatrix} \lceil (n+k)/2 \rceil \\ k \end{bmatrix}_q.$$

1.6 Containers for Antichains

The motivation for combinatorial supersaturation has increased in recent years due to its applications involving the container method. The container method was developed by Saxton and Thomason [205] and, independently, by Balogh, Morris and

Samotij [26][2] and has lead to several substantial results on counting combinatorial objects with forbidden substructures and extremal problems for sparse random structures.

In Chapter 7, we apply Theorem 1.18 and the container method to prove an upper bound on the number of antichains in $\mathcal{V}(q,n)$ and an upper bound on the largest antichain in a 'not too sparse' random subset of $\mathcal{V}(q,n)$ which holds with high probability.[3]

Theorem 1.19 (Noel, Scott and Sudakov [176])**.** *Let q be a fixed prime power. Then the number of antichains in $\mathcal{V}(q,n)$ is at most $2^{\left(1+O\left(\sqrt{n}q^{-n/4}\right)\right)\left[{n \atop \lfloor n/2 \rfloor}\right]_q}$.*

Definition 1.20. Given a set X, a *p-random subset* of X is a subset of X obtained by including each element of X with probability p independently of one another.

Theorem 1.21 (Noel, Scott and Sudakov [176])**.** *Let q be a fixed prime power and $\varepsilon > 0$. There exists a positive constant $c(\varepsilon,q)$ such that if $p \geq c(\varepsilon,q)/q^{n/2}$, then with high probability every antichain in a p-random subset of $\mathcal{V}(q,n)$ has cardinality at most $(1+\varepsilon)p\left[{n \atop \lfloor n/2 \rfloor}\right]_q$.*

Theorem 1.21 is best possible up to the choice of the constant $c(\varepsilon,q)$; this is made explicit in Section 7.1.

1.7 Finitely Forcible Graph Limits

One of the most important results in extremal combinatorics is the Erdős–Simonovits Stability Theorem [86, 208]. In rough terms, what this theorem says is that if G is a graph which does not contain a complete graph on r vertices and has 'close' to $\left(1-\frac{1}{r-1}\right)\binom{n}{2}$ edges, then, in some specific sense, the structure of G is 'close' to that of a complete multipartite graph in which every part has size $\lfloor n/(r-1)\rfloor$ or $\lceil n/(r-1)\rceil$. A short and elegant proof of this result can be found in a recent paper of Füredi [98].

In other words, the Erdős–Simonovits Stability Theorem says that the property that G is close to a balanced complete $(r-1)$-partite graph can be 'forced' by specifying that G contains no clique on r vertices and that the edge density of G is close

[2] Some precursors to the container method can be found in the work of Kleitman and Winston [144], Sapozhenko [199, 201, 202], Green [112] and Saxton and Thomason [204].

[3] Throughout the thesis, the term 'with high probability' is taken to mean 'with probability tending to 1 as $n \to \infty$.'

to $\left(1 - \frac{1}{r-1}\right)$. A typical *finite forcibility* problem focuses on determining approximate 'global' structures[4] (e.g. being close to a balanced $(r-1)$-partite graph) that can be forced in a graph by specifying the asymptotic densities of a finite number of subgraphs (e.g. K_r and K_2).

The notion of finite forcibility is defined more formally in Chapter 8 in the language of graph limits. For the purposes of this introductory discussion, we stick with an informal description. In what follows, a *graphon* is an analytic object which represents the limit of a convergent sequence of dense graphs and a graphon is said to be *finitely forcible* if it is the unique (in a certain sense) graphon which attains some finite list of subgraph densities. For example, a sequence of balanced complete bipartite graphs of increasing order is a convergent sequence of graphs, and the corresponding graphon is finitely forcible as it is the unique graphon with edge density equal to $1/2$ and triangle density equal to 0.

In [159], Lovász and Szegedy made several conjectures regarding the properties of finitely forcible graphons. At the time, every graphon that was known to be finitely forcible had a relatively simple structure, which lead them to believe that this was the case in general. In particular, one of their conjectures (interpreted in a particular way) would have implied that every finitely forcible graphon has a 'weak ε-regular partition' in which the number of parts is bounded above by a polynomial in ε^{-1} (see Section 8.2 for background on weak ε-regularity). This conjecture was disproved by Glebov, Klimošová and Král' [107], but the problem of determining the minimum number of parts in a weak ε-regular partition of a finitely forcible graphon remained. In Chapter 8, we use a construction of Conlon and Fox [72] and techniques developed in [107, 108] (which apply the notion of 'flag algebras' due to Razborov [187]) to construct a finitely forcible graphon for which there exists a sequence $(\varepsilon_i)_{i=1}^{\infty}$ tending to zero such that, for each $i \geq 1$, the number of parts in any weak ε_i-regular partition of this graphon is at least exponential in $\varepsilon_i^{-2}/2^{5\log^* \varepsilon_i^{-2}}$ (see Theorem 8.4). This result nearly matches a known upper bound for general graphons (finitely forcible or not) and is, in a certain sense, best possible.

[4] We remark that we are using the word 'structure' somewhat loosely here. For example, a graph G with $\sim \frac{1}{2}\binom{n}{2}$ edges distributed sufficiently uniformly (i.e. a 'quasirandom graph') will be thought of as having 'structure' close to that of a typical instance of the Erdős–Rényi random graph $G(n, 1/2)$.

1.8 Reconfiguring Graph Colourings

The final topic of this thesis is on the computational complexity of 'reconfiguration problems' for graph colourings and homomorphisms. Generally speaking, a reconfiguration problem has the following form: given two solutions to a fixed combinatorial problem (e.g. two cliques of order k in a graph or two satisfying assignments of a specific 3-SAT instance) is it possible to transform one of these solutions into the other by applying a sequence of allowed modifications such that every intermediate object is also a solution to the problem?

As a specific example, for a fixed integer k and a graph G, one may ask the following: given two (proper) k-colourings f and g of G, is it possible to transform f into g by changing the colour of one vertex at a time such that every intermediate mapping is a k-colouring? In the affirmative we say that f *reconfigures* to g. Trivially, this problem is solvable in polynomial time for $k \leq 2$. Rather surprisingly, Cereceda, van den Heuvel and Johnson [58] proved that it is also solvable in polynomial time for $k = 3$ despite the fact that determining if a graph admits a 3-colouring is NP-complete. On the other hand, Bonsma and Cereceda [48] proved that, when $k \geq 4$, the problem is PSPACE-complete. Combining these two results yields the following dichotomy theorem.

Theorem 1.22 (Cereceda, van den Heuvel and Johnson [58]; Bonsma and Cereceda [48])**.** *The reconfiguration problem for k-colourings is solvable in polynomial time for $k \leq 3$ and is PSPACE-complete for $k \geq 4$.*

In Chapter 9, we study the complexity of the reconfiguration problem for circular colourings. Given a graph G and positive integers p and q with $p/q \geq 2$, a *(circular) (p, q)-colouring* of G is a mapping $f : V(G) \to \{0, \ldots, p-1\}$ such that, for every edge $uv \in E(G)$,

$$q \leq |f(u) - f(v)| \leq p - q. \tag{1.23}$$

Clearly, a $(p, 1)$-colouring is nothing more than a p-colouring and therefore (p, q)-colourings generalise classical graph colourings. Circular colourings were introduced by Vince [219], and have been studied extensively; see the survey of Zhu [226]. Analogous to that of classical graph colourings, the reconfiguration problem for circular colourings asks, given (p,q)-colourings f and g of G, whether it is possible to reconfigure f into g by recolouring one vertex at a time while maintaining (1.23) for every $uv \in E(G)$ throughout. We determine the complexity of this problem for all fixed values of p and q.

Theorem 1.24 (Brewster, McGuinness, Moore and Noel [55]). *Let p, q be fixed positive integers with $p/q \geq 2$. Then the (p,q)-*Recolouring *problem is solvable in polynomial time for $2 \leq p/q < 4$ and is PSPACE-complete for $p/q \geq 4$.*

For the 'polynomial cases' $2 \leq p/q < 4$, we provide an explicit polynomial time algorithm. It is clear the the reconfiguration problem for (p,q)-colourings is in PSPACE. To show that it is PSPACE-complete for $p/q \geq 4$, we reduce the reconfiguration problem for $\lfloor p/q \rfloor$-colourings to the reconfiguration problem for (p,q)-colourings (in polynomial time) and apply Theorem 1.22.

Chapter 2

Minimum Saturation in the Hypercube

We open this chapter with a (non-exhaustive) survey of saturation problems in extremal combinatorics; see Faudree, Faudree and Schmitt [93] for a more extensive survey. The earliest result in this area is due to Zykov [227] and Erdős, Hajnal and Moon [89] who independently proved that

$$\operatorname{sat}(K_n, K_r) = \binom{n}{2} - \binom{n-r+2}{2} = (r-2)(n-r+2) + \binom{r-2}{2} \tag{2.1}$$

for all $n \geq r \geq 2$. In addition they proved that the only graph attaining the minimum is the graph obtained from K_n by deleting the edges of a clique on $n-r+2$ vertices.

Motivated by an extension of (2.1) to hypergraphs, Bollobás [35] proved the following result which has come to be known as the 'Two Families Theorem' and is of fundamental importance throughout extremal combinatorics.

Two Families Theorem (Bollobás [35])**.** *Let $A_1, \ldots, A_m$ and $B_1, \ldots, B_m$ be two families of finite sets such that*

(a) $A_i \cap B_i = \emptyset$ for $1 \leq i \leq m$, and

(b) $A_i \cap B_j \neq \emptyset$ for $1 \leq i, j \leq m$ with $i \neq j$.

Then $\sum_{i=1}^{m} \binom{|A_i|+|B_i|}{|A_i|}^{-1} \leq 1$.

The special case in which $|A_1| = \cdots = |A_m|$ and $|B_1| = \cdots = |B_m|$ was rediscovered and posed as a conjecture by Ehrenfeucht and Mycielski (see [139]), and proofs were published by Katona [139], Jaeger and Payan [128] and Tarján [211]. We remark that the Two Families Theorem implies the LYM Inequality [160,166,225] in extremal set theory which, itself, is a strengthening of Sperner's Theorem [210].

Let us apply the Two Families Theorem to obtain an analogue of (2.1) for hypergraphs, as was done in [35]. Given two k-uniform hypergraphs F and H, a subhypergraph G of F is said to be (F, H)-*saturated* if it contains no copy of H, but $G + e$ contains a copy of H for every hyperedge $e \in E(F) \setminus E(G)$. We denote the complete k-uniform hypergraph on n vertices by $K_n^{(k)}$.

Now, let G be a $\left(K_n^{(k)}, K_r^{(k)}\right)$-saturated hypergraph with $\binom{n}{k} - m$ hyperedges and let $B_1, \ldots, B_m$ be the hyperedges which are absent from G. For $1 \leq i \leq m$, let V_i be the vertex set of a copy of $K_r^{(k)}$ contained in $G + B_i$. If we define $A_i := V\left(K_n^{(k)}\right) \setminus V_i$ for $1 \leq i \leq m$, then we see that $A_1, \ldots, A_m$ and $B_1, \ldots, B_m$ satisfy the conditions of the Two Families Theorem. Thus, since $|A_i| = n - r$ and $|B_i| = k$ for all i, we get $m\binom{n+k-r}{k}^{-1} \leq 1$ or, in other words, $m \leq \binom{n+k-r}{k}$. This proves

$$\operatorname{sat}\left(K_n^{(k)}, K_r^{(k)}\right) \geq \binom{n}{k} - \binom{n+k-r}{k}.$$

The matching upper bound is easy to prove: simply let G be a k-uniform hypergraph obtained from $K_n^{(k)}$ by deleting the hyperedges of a complete k-uniform hypergraph on $n+k-r$ vertices. Bollobás [35] also proved that, up to isomorphism, this hypergraph is the only tight example. In Chapter 4, we will discuss so-called 'skew versions' of the Two Families Theorem which can be applied, in a similar manner, to obtain lower bounds for weak saturation problems in graphs and hypergraphs.

In their original paper on saturation, Erdős, Hajnal and Moon [89] also considered the problem of determining $\operatorname{sat}(K_{n,n}, K_{r,r})$ for $n \geq r$. As they pointed out, one can obtain an upper bound

$$\operatorname{sat}(K_{n,n}, K_{r,r}) \leq n^2 - (n-r+1)^2 = (2r-2)n - (r-1)^2 \tag{2.2}$$

by considering the graph obtained from $K_{n,n}$ by deleting the edges of a copy of $K_{n-r+1,n-r+1}$. They conjectured that this upper bound is tight. This conjecture was proved independently by Bollobás [36] and Wessel [220, 221]. However, the related problem of determining $\operatorname{sat}(K_{n,n}, K_{r,s})$ for $r \neq s$ is surprisingly less well-understood. In [172], Moshkovitz and Shapira proposed the following conjecture and provided a construction attaining the upper bound.

Conjecture 2.3 (Moshkovitz and Shapira [172]). *Let $1 \le r \le s$ be integers. Then there exists $n_0 = n_0(r,s)$ such that if $n \ge n_0(r,s)$, then*

$$\mathrm{sat}\left(K_{n,n}, K_{r,s}\right) = (r+s-2)n - \left\lfloor \left(\frac{r+s-2}{2}\right)^2 \right\rfloor.$$

Of course, the aforementioned result of Bollobás [36] and Wessel [220,221] implies the conjecture when $r = s$. Gan, Korándi and Sudakov [103] proved the conjecture in the case $r = 2$ and $s = 3$ and obtained a lower bound of $\mathrm{sat}\left(K_{n,n}, K_{r,s}\right) \ge (r + s-2)n - (r+s-2)^2$ for $n \ge s \ge r \ge 1$. However, the conjecture remains open in general. We remark that an 'ordered version' of this problem was conjectured in [89] and solved in [36,220,221]; see [103,172] for discussion regarding the ordered problem and how it differs from Conjecture 2.3.

With regards to the asymptotics of $\mathrm{sat}(K_n, H)$ for general graphs H, Kászonyi and Tuza [138] provided an explicit upper bound which, in particular, implies that $\mathrm{sat}(K_n, H) = O(n)$ for any fixed graph H. Pikhurko [182] proved an analogous result for hypergraphs which confirmed a conjecture of Tuza [216]. Another tantalising conjecture of Tuza [216], which remains open, is as follows.

Conjecture 2.4 (Tuza [216]). $\lim_{n\to\infty} \frac{\mathrm{sat}(K_n,H)}{n}$ *exists for every fixed graph H.*

Evidence towards a positive answer to this conjecture can be found in [214], while evidence towards a negative answer can be found in [181,184].

There are also several papers on determining the minimum number of edges in a (K_n, H)-saturated graph G with additional constraints on its maximum or minimum degree [5,41,77,80,90,102,117,184]. Recently, Day [77] used the LYM Inequality and an ingenious greedy argument to prove that, for every integer d, there exists $f(d)$ such that every (K_n, K_r)-saturated graph with minimum degree at least d contains at least $dn - f(d)$ edges. This improves a bound of Pikhurko [184] and answers a question of Bollobás [41].

In a recent paper, Korándi and Sudakov [146] considered the problem of determining $\mathrm{sat}(G(n,p), K_r)$ where $G(n,p)$ is the Erdős–Rényi random graph with density p (see [42,131] for background on random graphs). In particular, it is shown that, for fixed $r \ge 3$ and fixed $p \in (0,1)$,

$$\mathrm{sat}(G(n,p), K_r) = (1+o(1))n \log_{1/(1-p)}(n)$$

with high probability. Rather surprisingly, this implies that, for fixed p, the asymptotics of $\mathrm{sat}(G(n,p),K_r)$ are independent of the value of r. As was noted in [146], some analogous problems for non-complete graphs H, hypergraphs, and non-constant densities p remain open.

As was mentioned in Section 1.1, our goal in this chapter is to prove an upper bound on $\mathrm{sat}(Q_d,Q_m)$. This problem was first considered by Choi and Guan [65] in the case $m=2$ who constructed a (Q_d,Q_2)-saturated graph with at most $\left(\frac{1}{4}+o(1)\right)d2^{d-1}$ edges. Santolupo (see [93]) conjectured that this construction is best possible. This was disproved (in a strong sense) by Johnson and Pinto [133].

Theorem 2.5 (Johnson and Pinto [133]). *For every fixed $m\geq 2$, there exists $0<\varepsilon_m<1$ such that*

$$\mathrm{sat}(Q_d,Q_m)=O\left(d^{1-\varepsilon_m}2^{d-1}\right)=O\left(|E(Q_d)|/d^{\varepsilon_m}\right).$$

In the case $m=2$, Johnson and Pinto [133] obtained a stronger bound; namely,

$$\mathrm{sat}(Q_d,Q_2)\leq 10\cdot 2^d. \tag{2.6}$$

That is, for every d, they proved that there exists a (Q_d,Q_2)-saturated graph of bounded average degree. Motivated by this result, they asked the following: for which fixed values of m is it true that $\mathrm{sat}(Q_d,Q_m)=O\left(2^d\right)$? The result of this chapter (Theorem 1.3) shows that this is the case for every fixed $m\geq 2$.

The rest of the chapter is structured as follows. In Sections 2.1 and 2.2, we review some fundamental properties of hypercubes which will be used in the proof of Theorem 1.3. Then, in Section 2.3, we prove Theorem 1.3 by giving an explicit construction of a (Q_d,Q_m)-saturated graph with bounded average degree. We conclude the chapter in Section 2.4 by mentioning some open problems.

2.1 Preliminaries

Let e_j denote the jth standard basis vector in $\mathbb{F}_2^d$; i.e., the vector in which the jth coordinate is equal to one and every other coordinate is zero. Given a vertex v of Q_d, let $|v|$ denote the sum of its coordinates.

A basic fact about hypercubes is that, for $d\geq m$, every copy Q of Q_m in Q_d is induced by a set of vertices of the form $\left\{v+\sum_{j\in J'}e_j : J'\subseteq J\right\}$ where v is any vertex of Q and $J \subseteq [d]$ has cardinality m (here, addition is in $\mathbb{F}_2^d$). We say that the coordinates of J are *variable* under Q and that the coordinates in $[d] \setminus J$ are *fixed* under Q.

We will need a standard result from Coding Theory, due to Hamming [118] (for another reference, see [161,162]). This result was also used by Johnson and Pinto [133] (and we use some ideas from [133] in our proof).

Theorem 2.7 (Hamming [118]). *There is a stable set $C \subseteq V(Q_{2^t-1})$ such that every vertex of $V(Q_{2^t-1}) \setminus C$ has a unique neighbour in C.*

A set C as in Theorem 2.7 is often referred to as a *Hamming code*. We remark that, since every vertex of Q_{2^t-1} has exactly $2^t - 1$ neighbours,

$$|C| = \frac{2^{2^t-1}}{(2^t-1)+1} = 2^{2^t-t-1}. \tag{2.8}$$

Another ingredient in the proof of Theorem 1.3 is the following result of Conder [70].

Theorem 2.9 (Conder [70]). *For every $s \geq 1$, there is a 3-colouring of the edges of Q_s in which there is no monochromatic cycle of length 4 or 6.*

2.2 Definitions and Proof Outline

Our construction will use Theorems 2.7 and 2.9 to build a spanning subgraph G of Q_d which contains no copy of Q_m with the property that any (Q_d, Q_m)-saturated graph G' containing G must satisfy the bound in Theorem 1.3. In particular, G itself will not be (Q_d, Q_m)-saturated (in fact, it will even contain some isolated vertices).

Throughout the proof, let $m \geq 2$ be fixed and let d be an integer which we may choose to be sufficiently large. Let t and s be the unique integers such that

$$d = 6m(2^t - 1) + s \text{ and } 0 \leq s < 6m2^t. \tag{2.10}$$

We view the set $[d]$ as a union of m intervals of length $6(2^t-1)$ indexed by the elements of $[m]$ followed by one interval of length s. Given $v \in V(Q_d)$ and $1 \leq i \leq m$, let $v(i)$ denote the vertex of $V(Q_{6(2^t-1)})$ obtained by restricting v to the coordinates of the ith interval, and let $v(m+1)$ be the vertex of $V(Q_s)$ obtained by restricting v to its last

s coordinates. Each of the first m intervals of $[d]$ is further divided into 6 subintervals of length (2^t-1) indexed by the pairs (r,γ) where $r\in\{0,1\}$ and $\gamma\in\{0,1,2\}$. For $(i,r,\gamma)\in[m]\times\{0,1\}\times\{0,1,2\}$, let $v(i,r,\gamma)$ be the vertex of $V(Q_{2^t-1})$ obtained by restricting $v(i)$ to the subinterval corresponding to the pair (r,γ).

Let C be a subset of $V\left(Q_{2^t-1}\right)$ as in Theorem 2.7. We will treat the vertices $v\in V\left(Q_d\right)$ differently depending on which of the triples $(i,r,\gamma)\in[m]\times\{0,1\}\times\{0,1,2\}$ satisfy $v(i,r,\gamma)\in C$. For starters, let X denote the set of all vertices v for which there exists some $i\in[m]$ and $(r,\gamma)\neq(r',\gamma')$ such that $v(i,r,\gamma)$ and $v(i,r',\gamma')$ are both contained in C. The vertices of X will be isolated in G and will not play a large role in the construction. Now, divide the vertices of $V\left(Q_d\right)\setminus X$ into sets $A_0,\dots,A_m$ where A_j is defined to be the set of all vertices of $V\left(Q_d\right)\setminus X$ for which there are exactly j triples (i,r,γ) such that $v(i,r,\gamma)\in C$. We remark that, by (2.8) and (2.10),

$$\left|V\left(Q_d\right)\setminus A_0\right|=O\left(2^d/d\right). \tag{2.11}$$

In other words, A_0 contains an *overwhelming* majority of the vertices of Q_d.

The key properties of the graph G that we construct are the following.

(a) it does not contain Q_m as a subgraph,

(b) it does not contain any edge uv of Q_d where $u,v\in A_0$ and

(c) it has the property that, for any $uv\in E(Q_d)$ with $u,v\in A_0$, the graph $G+uv$ contains a copy of Q_m.

Given such a graph G, let G' be *any* graph obtained from G by adding a maximal set of edges from $E(Q_d)\setminus E(G)$ to G without creating a copy of Q_m. Then G' is (Q_d,Q_m)-saturated and, by (b) and (c), every edge of G' has at least one endpoint in $V\left(Q_d\right)\setminus A_0$. Every vertex of Q_d has degree d and so, by (2.11), the total number of edges incident to $V\left(Q_d\right)\setminus A_0$ is at most $O\left(2^d\right)$. Thus, we obtain $\operatorname{sat}(Q_d,Q_m)\le\left|E\left(G'\right)\right|=O\left(2^d\right)$, as desired.

The difficulty of the proof comes from ensuring that (a) and (c) hold *simultaneously*. Let us describe, in rough terms, how we achieve this. In what follows, let $\phi_1:E\left(Q_{6(2^t-1)}\right)\to\{0,1,2\}$ and $\phi_2:E\left(Q_s\right)\to\{0,1,2\}$ be colourings as in Theorem 2.9, and let $\phi:E\left(Q_{6(2^t-1)}\right)\sqcup E\left(Q_s\right)\to\{0,1,2\}$ be the mapping which is equal to ϕ_1 on $E\left(Q_{6(2^t-1)}\right)$ and equal to ϕ_2 on $E\left(Q_s\right)$.

First, we let G contain every edge of Q_d joining a vertex of A_j to a vertex of A_{j+1} for $0 \leq j \leq m-2$. It follows that every vertex of A_0 is contained in at least $m6^{m-1}$ copies of Q_{m-1}. To see this, let $v \in A_0$, let I be a subset of $[m]$ of cardinality $m-1$ and, for every $i \in I$, let $(r_i, \gamma_i) \in \{0,1\} \times \{0,1,2\}$. Since $v \in A_0$, we have that $v(i, r_i, \gamma_i) \notin C$ and so $v(i, r_i, \gamma_i)$ has a unique neighbour $x_i \in C$. For each subset S of I, let v_S denote the vertex of Q_d such that $v_S(i, r_i, \gamma_i) = x_i$ for $i \in S$ and v_S is equal to v on all other coordinates. Then it is easily observed that

- $v_\emptyset = v$,
- $v_S \in A_{|S|}$ for every $S \subseteq I$, and
- v_S is adjacent to $v_{S \cup \{j\}}$ for every $S \subsetneq I$ and $j \in I \setminus S$.

Thus, the vertices of $\{v_S : S \subseteq I\}$ form a copy of Q_{m-1} in G containing v.

Now, in order to ensure that (c) holds, we also add edges within the sets A_j for $j = 1, \ldots, m-1$ so that, for every pair $u, v \in A_0$ which are adjacent in Q_d, there exists a choice of the set $I \subseteq [m]$ and the pairs (r_i, γ_i) so that v_S is adjacent to u_S for every non-empty set $S \subseteq I$ (here, u_S is defined analogously to v_S). If k is the unique element of $[m+1]$ such that $u(k) \neq v(k)$, then the set I is chosen so that $k \notin I$. The choice of the pairs (r_i, γ_i) depends on a 'parity condition' as well as the colour of the edge $u(k)v(k)$ under ϕ. The reason that we require these conditions is to 'break' any potential copies of Q_m which could occur in G as a result of adding edges within the sets A_j for $1 \leq j \leq m-1$.

2.3 The Construction

We now give the details of our construction. Let G be a spanning subgraph of Q_d in which the edges of G are defined by the following two steps.

Step 1. *Let G contain every edge of Q_d which joins a vertex of A_j to a vertex of A_{j+1} for $0 \leq j \leq m-2$.*

Step 2. *Suppose that uv is an edge of Q_d such that $u, v \in A_j$ for some $1 \leq j \leq m-1$. Let k be the unique element of $[m+1]$ such that $u(k) \neq v(k)$. We let G contain the edge uv if the following two conditions are satisfied:*

(i) either $k = m+1$ or there does not exist $(r', \gamma') \in \{0,1\} \times \{0,1,2\}$ such that $u(k, r', \gamma') = v(k, r', \gamma') \in C$, and

(ii) for every $(i, r, \gamma) \in [m] \times \{0,1\} \times \{0,1,2\}$ such that $i \neq k$ and $v(i, r, \gamma) \in C$ we have $\gamma = \phi(u(k)v(k))$ and $|v| + |v(k)| + |v(i)| + j - 1 \equiv r \mod 2$.

Note that $|v| + |v(k)| + |v(i)| + j - 1 \equiv |u| + |u(k)| + |u(i)| + j - 1 \mod 2$ since the only coordinate on which u and v differ is in the kth interval. Therefore, condition (ii) is invariant under swapping u and v.

Before moving on we make a few simple observations, each of which can be verified by looking carefully at Steps 1 and 2.

Observation 2.12. $A_m \cup X$ is a set of isolated vertices in G.

Observation 2.13. A_0 is a stable set of G.

Observation 2.14. If uv is an edge of G such that $u(i, r, \gamma) = v(i, r, \gamma) \in C$, then $u(i, r', \gamma') = v(i, r', \gamma')$ for every pair $(r', \gamma') \in \{0,1\} \times \{0,1,2\}$.

Observation 2.15. If uv is an edge of G such that $u(m+1) \neq v(m+1)$, then $u, v \in A_j$ for some $1 \leq j \leq m-1$.

The next observation follows from the fact that C is a stable set in Q_{2^t-1}.

Observation 2.16. If uv is an edge of G such that $u(i, r, \gamma) \neq v(i, r, \gamma)$, then at most one of $u(i, r, \gamma)$ or $v(i, r, \gamma)$ is in C.

To complete the proof of Theorem 1.3, it suffices to establish the following two claims. We state these claims now and show that they imply Theorem 1.3 before proving the claims themselves.

Claim 2.17. *For every edge uv of Q_d such that $u, v \in A_0$, the graph $G + uv$ contains a copy of Q_m.*

Claim 2.18. *G does not contain Q_m as a subgraph.*

Proof of Theorem 1.3. If G is not (Q_d, Q_m)-saturated, then by Claim 2.18 we can extend G to a (Q_d, Q_m)-saturated graph G' by adding a maximal set of edges from $E(Q_d) \setminus E(G)$ to G which do not create a copy of Q_m. By Claim 2.17 none of these additional edges are between vertices in A_0 and so, by Observation 2.13, A_0 is a stable set in G'. Thus, every edge of G' has at least one endpoint in $(\bigcup_{k=1}^{m} A_k) \cup X$. The total number of edges incident to vertices of this set is at most

$$d\left|\left(\bigcup_{k=1}^{m} A_k\right) \cup X\right| = d|A_1| + d\left|\left(\bigcup_{k=2}^{m} A_k\right) \cup X\right|$$

Note that $|(\bigcup_{k=2}^{m} A_k) \cup X| = O\left(|C|^2 2^{d-2(2^t-1)}\right)$, which is $O\left(2^d/d^2\right)$ by (2.8) and (2.10). Therefore the second term of the above expression is $o\left(2^d\right)$ and it suffices to bound $d|A_1|$. We have

$$\begin{aligned} d|A_1| = d\left(6m|C|2^{d-(2^t-1)}\right) &= d\left(6m2^{2^t-t-1}2^{d-(2^t-1)}\right) = 6md2^{d-t} \\ &= 6m\left(6m(2^t-1)+s\right)2^{d-t} < 72m^2 2^d \end{aligned}$$

by (2.10). The result follows. □

Remark 2.19. Note that if d is of the form $6m(2^t - 1)$ for some t, then we obtain a better bound: $\mathrm{sat}(Q_d, Q_m) \leq (1 + o(1))36m^2 2^d$.

Thus, it suffices to prove Claims 2.17 and 2.18. We begin by presenting a short proof of Claim 2.17, which was foreshadowed in the previous section.

Proof of Claim 2.17. Let uv be an edge of Q_d where $u, v \in A_0$ and let k be the unique element of $[m+1]$ for which $u(k) \neq v(k)$. Define $\gamma := \phi(u(k)v(k))$ and let $I \subseteq [m] \setminus \{k\}$ be a set of size $m - 1$. For $i \in I$, pick $r_i \in \{0, 1\}$ so that

$$|v| + |v(k)| + |v(i)| \equiv r_i \mod 2.$$

For each $i \in I$, let c_i be the unique neighbour of $v(i, r_i, \gamma)$ in Q_{2^t-1} contained in C. Given $I' \subseteq I$, we let $v_{I'}$ and $u_{I'}$ be the vertices of Q_d such that $u_{I'}(i, r_i, \gamma) = v_{I'}(i, r_i, \gamma) = c_i$ for all $i \in I'$ and, on all other coordinates, $u_{I'}$ and $v_{I'}$ agree with u and v, respectively. In particular, $v_\emptyset = v$ and $u_\emptyset = u$. If $I' \neq I$, then for any $i' \in I \setminus I'$ we see that $v_{I'}$ is adjacent to $v_{I' \cup \{i'\}}$ and $u_{I'}$ is adjacent to $u_{I' \cup \{i'\}}$ in G via edges added in Step 1. Also, for $I' \subseteq I$, we have $v_{I'}, u_{I'} \in A_{|I'|}$ and, if $I' \neq \emptyset$, then for each $i \in I'$,

$$|v_{I'}| + |v_{I'}(k)| + |v_{I'}(i)| \equiv r_i + |I'| - 1 \mod 2.$$

Thus, for $I' \neq \emptyset$, $v_{I'}$ is adjacent to $u_{I'}$ via an edge of G added in Step 2. This implies that $\{v_{I'} : I' \subseteq I\} \cup \{u_{I'} : I' \subseteq I\}$ induces a copy of Q_m in $G + uv$, which completes the proof. $\square$

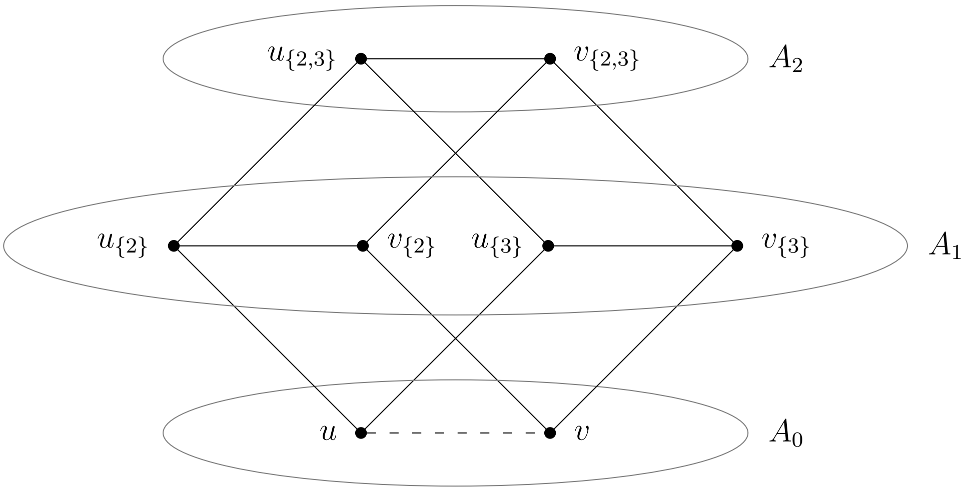

Figure 2.1: An illustration of the proof of Claim 2.17 in the case that $m = 3$ and $u(1) \neq v(1)$.

Proof of Claim 2.18. Suppose, to the contrary, that G contains a copy Q of Q_m. Let $J \subseteq [d]$ be the set of coordinates which are variable under Q. For each $i \in [m+1]$, let $J(i)$ be the set of coordinates of J which are in the interval of $[d]$ corresponding to i. Moreover, given $(i, r, \gamma) \in [m] \times \{0,1\} \times \{0,1,2\}$, let $J(i, r, \gamma)$ be the set of coordinates of J which are in the interval of $[d]$ corresponding to (i, r, γ).

Subclaim 2.20. *For each $i \in [m]$ there is a pair $(r_i, \gamma_i) \in \{0,1\} \times \{0,1,2\}$ such that $J(i) = J(i, r_i, \gamma_i)$.*

Proof. Suppose to the contrary that there exists $(r, \gamma) \neq (r', \gamma')$ such that both $J(i, r, \gamma)$ and $J(i, r', \gamma')$ are non-empty.

First, suppose that there exists a vertex $v \in V(Q)$ with $v(i, r, \gamma) \in C$. Let u be a neighbour of v in Q obtained by changing a coordinate in $J(i, r', \gamma')$. By

Observation 2.14, the edge uv is not contained in G, which is a contradiction. Thus, $v(i,r,\gamma), v(i,r',\gamma') \notin C$ for all $v \in V(Q)$.

Now, let Q' be a copy of Q_2 in Q obtained by starting with an arbitrary vertex of Q and varying one coordinate of $J(i,r,\gamma)$ and one coordinate of $J(i,r',\gamma')$. By the result of the previous paragraph, we see that $V(Q') \subseteq A_j$ for some j and so every edge of Q' was added in Step 2. By taking the edges of Q' and restricting them to the interval of $[d]$ corresponding to i, we see that all of the resulting edges must receive the same colour under ϕ. This contradicts the fact that there is no copy of Q_2 which is monochromatic under ϕ and completes the proof of the subclaim. □

Subclaim 2.21. *Suppose that $|J(i)| \geq 2$ for some $i \in [m]$ and let Q' be a copy of Q_2 in Q obtained by starting with an arbitrary vertex of Q and varying two coordinates of $J(i)$. Then there is a unique vertex $v \in V(Q')$ with $v(i,r_i,\gamma_i) \in C$.*

Proof. By Subclaim 2.20 we have $J(i) = J(i,r_i,\gamma_i)$. By Observation 2.16, if uv is an edge of Q', then we cannot have $u(i,r_i,\gamma_i), v(i,r_i,\gamma_i) \in C$. Thus, if there are two vertices u,v of Q' for which $u(i,r_i,\gamma_i), v(i,r_i,\gamma_i) \in C$, then they must be non-adjacent. So, in this case, the vertices of Q' alternate between A_j and A_{j+1} for some j and every edge of Q' was added in Step 1. However, this implies that for $w \in V(Q')$ such that $w(i,r_i,\gamma_i) \notin C$, the vertex $w(i,r_i,\gamma_i)$ of Q_{2^t-1} must have two distinct neighbours in C, which is a contradiction.

Now we assume that every vertex v of Q' satisfies $v(i,r_i,\gamma_i) \notin C$. However, in this case, we see that every edge of Q' was added in Step 2. As before, we obtain a copy of Q_2 in $Q_{6(2^t-1)}$ which is monochromatic under ϕ, a contradiction. □

Subclaim 2.22. *$|J(i)| \leq 2$ for every $i \in [m]$.*

Proof. Suppose not. Let Q' be a copy of Q_3 in Q obtained by starting with an arbitrary vertex and varying three coordinates of $J(i)$. If we vary any pair of these coordinates, leaving the third one fixed, we obtain a copy of Q_2 in Q' which must obey Subclaim 2.21. This implies that there are precisely two vertices $x,y \in V(Q')$ such that $x(i,r_i,\gamma_i), y(i,r_i,\gamma_i) \in C$ and they are at distance 3 in Q'. However, now we get that the vertices of $V(Q') \setminus \{x,y\}$ induce a copy F of C_6 in Q, where every edge of F was added in Step 2. Taking the edges of F and restricting them to the interval of $[d]$ corresponding to i, we see that all such edges must receive the same colour under ϕ.

This contradicts the fact that there is no copy of C_6 which is monochromatic under ϕ and completes the proof of the subclaim. $\square$

Subclaim 2.23. $|J(m+1)| \leq 1$.

Proof. If not, then let Q' be a copy of Q_2 in Q obtained by starting at an arbitrary vertex and varying two coordinates of $J(m+1)$. Then, restricting the edges of Q' to the last s coordinates, we obtain a copy of Q_2 in Q_s which is monochromatic under ϕ, a contradiction. $\square$

Subclaim 2.24. *Suppose that uv and vw are distinct edges of Q where $u, v, w \in A_j$ for some j. Then there is some $k \in [m]$ such that $u(k) \neq v(k)$ and $v(k) \neq w(k)$.*

Proof. We show that there cannot exist distinct integers $k, k' \in [m+1]$ such that $u(k) \neq v(k)$ and $v(k') \neq w(k')$. This is sufficient to prove the subclaim since, by Subclaim 2.23, we cannot have $u(m+1) \neq v(m+1)$ and $v(m+1) \neq w(m+1)$.

Suppose that such integers k, k' exist, and let x be the unique vertex of Q_d distinct from v which is joined to both u and w in Q_d. Note that both of the edges xu and wx are present in Q. Also, $x \in A_j$ and so all of the edges uv, vw, wx, xu were added to G in Step 2. This implies that $j \geq 1$ and that there is a triple (i, r, γ) with $i \notin \{k, k'\}$ such that

$$u(i, r, \gamma) = v(i, r, \gamma) = w(i, r, \gamma) = x(i, r, \gamma) \in C.$$

However, one can easily check that the parity of $|v| + |v(k)| + |v(i)|$ is different from the parity of $|x| + |x(k)| + |x(i)|$ modulo 2 and so only one can be equivalent to r. By definition of Step 2, only one of the edges uv and xw can exist in G. This contradiction completes the proof. $\square$

Subclaim 2.25. *If $|J(m+1)| = 1$, then $|J(i)| \leq 1$ for all $i \in [m]$.*

Proof. If not, let Q' be a copy of Q_3 in Q obtained by starting at an arbitrary vertex and varying one coordinate of $J(m+1)$ and two of $J(i)$. Then, by Subclaim 2.21, there is an edge $uv \in E(Q')$ such that $u(i) \neq v(i)$ and $u, v \in A_j$ for some j (in fact, there are many such edges). Now, let w be the neighbour of v in Q' obtained by changing the coordinate of $J(m+1)$. By Observation 2.15, we have that $w \in A_j$ as well. However, this contradicts Subclaim 2.24. $\square$

Subclaim 2.26. *There is at most one $i \in [m]$ for which $|J(i)| = 2$.*

Proof. If not, let $i, i' \in [m]$ such that $|J(i)| = |J(i')| = 2$ and let Q' be a copy of Q_4 obtained by starting at an arbitrary vertex and varying two coordinates of $J(i)$ and two coordinates of $J(i')$. By Subclaim 2.21, we see that there must exist edges $uv, vw \in E(Q')$ such that $u(i) \neq v(i)$, $v(i') \neq w(i')$ and $u, v, w \in A_j$ for some j. This contradicts Subclaim 2.24 and completes the proof. $\square$

Now, let us complete the proof of the claim. Throughout, for each $i \in [m]$, we let (r_i, γ_i) be a pair such that $J(i) = J(i, r_i, \gamma_i)$, which exists by Subclaim 2.20. In what follows, we let j^* be the minimum integer j such that $V(Q) \cap A_j \neq \emptyset$ and let $v^* \in V(Q) \cap A_{j^*}$. Note that for every $i \in [m]$ such that $J(i) \neq \emptyset$ we must have $v^*(i, r_i, \gamma_i) \notin C$. If not, then starting with v^* and changing a coordinate of $J(i)$ yields a vertex of $V(Q) \cap A_{j^*-1}$ by Observations 2.14 and 2.16, contradicting our choice of j^*. By Subclaims 2.25 and 2.26, there is at most one $i \in [m]$ such that $J(i) = \emptyset$. This implies that

$$j^* \in \{0, 1\}, \tag{2.27}$$

where j^* must equal zero if no such i exists. We divide the proof into cases.

Case 1: $|J(i)| = 1$ for every $i \in [m]$.

In this case, we must have $j^* = 0$ and so $v^* \in A_0$. However, if we start with v and change the coordinate of $J(i)$ for any i, then we obtain a vertex u with $u(i, r_i, \gamma_i) \in C$ by Observation 2.13. Thus, changing all m of these coordinates yields a vertex w such that $w(i, r_i, \gamma_i) \in C$ for all $i \in [m]$; that is, $w \in A_m$. This is a contradiction since, by Observation 2.12, w is an isolated vertex and therefore cannot belong to Q.

Case 2: $|J(m + 1)| = 1$.

Let u be the neighbour of v^* in Q obtained by changing the coordinate in $J(m+1)$. We must have $u, v^* \in A_1$ by (2.27) and Observation 2.15. This implies that there is some $i \in [m]$ such that $J(i) = \emptyset$ and $v^*(i, r_i, \gamma_i) \in C$. Also, by Subclaim 2.25 and the Pigeonhole Principle, we must have $|J(i')| = 1$ for every $i' \in [m] \setminus \{i\}$.

Now, for each $i' \in [m] \setminus \{i\}$, let $w_{i'}$ be the neighbour of v^* in Q obtained by changing the coordinate in $J(i')$. By Subclaim 2.24 and the fact that $u, v^* \in A_1$, we cannot have $w_{i'} \in A_1$. So, by our choice of j^*, we must have $w_{i'}(i', r_{i'}, \gamma_{i'}) \in C$. As in the proof of Case 1, if we start with v^* and change the coordinate of $J(i')$ for every $i' \in [m] \setminus \{i\}$, we obtain a vertex of A_m contained in Q, contradicting Observation 2.12. This completes the proof in this case.

Case 3: $|J(i)| = 2$ for some $i \in [m]$.

By Subclaim 2.21, there is a neighbour u of v^* in Q obtained by changing a coordinate in $J(i)$ such that $u \in A_{j^*}$. This immediately implies that $j^* = 1$ by (2.27) and Observation 2.13.

Now, for each $i' \in [m]$ for which $|J(i')| = 1$, let $w_{i'}$ be the neighbour of v^* in Q obtained by changing the coordinate in $J(i')$. By Subclaim 2.24 and the fact that $u, v^* \in A_1$, we must have $w_{i'}(i', r_{i'}, \gamma_{i'}) \in C$. Thus, if we let x be the vertex obtained from v^* by changing the coordinate of $J(i')$ for every such i', then $x \in A_{m-1}$.

Let Q' be the copy of Q_2 in Q obtained by starting at x and varying the two coordinates of $J(i)$. Then, by Subclaim 2.21, there must be some vertex y of $V(Q')$ such that $y(i, r_i, \gamma_i) \in C$. However, this implies that $y \in A_m$, contradicting Observation 2.12. This completes the proof of Claim 2.18 and of Theorem 1.3. $\square$

2.4 Open Problems

We conclude this chapter by stating several open problems. Recall that we obtained a better upper bound on $\mathrm{sat}(Q_d, Q_m)$ when d is of the form $6m(2^t - 1)$ than for general values of d (see Remark 2.19). It is not clear whether this is simply an artefact of our proof, or part of a more general phenomenon. We ask the following (c.f. Conjecture 2.4).

Question 2.28 (Morrison, Noel and Scott [170]). *For fixed $m \geq 2$, does $\frac{\mathrm{sat}(Q_d, Q_m)}{2^d}$ converge as $d \to \infty$?*

As it stands, the best known lower bound on $\mathrm{sat}(Q_d, Q_m)$ is that given by (1.9). We doubt that this bound is tight. In particular, we ask the following.

Question 2.29 (Morrison, Noel and Scott [170]). *For fixed $m \geq 2$, let $c_m :=$ $\liminf_{d\to\infty} \mathrm{sat}(Q_d, Q_m)/2^d$. Does $\frac{c_m}{m} \to \infty$?*

For $d \geq 1$ and $a_1, \ldots, a_d \geq 2$, the *d-dimensional* $a_1 \times \cdots \times a_d$ *grid* is the graph with vertex set $\prod_{i=1}^{d}[a_i]$ where two vertices are adjacent if they differ by one in exactly one coordinate. As a slight abuse of notation, we denote this graph by $\prod_{i=1}^{d}[a_i]$. In the symmetric case $a_1 = \cdots = a_d = n$, we write $[n]^d$. Note that, in particular, $[2]^d$ is isomorphic to Q_d.

For $n \geq r \geq 2$ and $d \geq m \geq 1$, we say that a copy of $[r]^m$ in $[n]^d$ is *axis-aligned* if it is induced by a set of vertices of the form $I_1 \times \cdots \times I_d$ where exactly m of the sets I_i are intervals of length r in $[n]$ and the rest are singletons. We remark that if $r = 2$ or $d = m$, then every copy of $[r]^m$ in $[n]^d$ is axis-aligned. Let $\mathrm{sat}^*\left([n]^d, [r]^m\right)$ denote the minimum number of edges in a spanning subgraph G of $[n]^d$ containing no axis-aligned copy of $[r]^m$ such that, for every edge $e \in E\left([n]^d\right) \setminus E(G)$, the graph $G + e$ contains an axis-aligned copy of $[r]^m$. It would be interesting to obtain bounds on $\mathrm{sat}\left([n]^d, [r]^m\right)$ and $\mathrm{sat}^*([n]^d, [r]^m)$. In particular, we ask the following questions.

Question 2.30. *Fix $m \geq 2$ and $n \geq r \geq 2$, and let $d \to \infty$. Is it true that*

$$\mathrm{sat}\left([n]^d, [r]^m\right) = O\left(n^d\right)?$$

Question 2.31 (Morrison, Noel and Scott [170])**.** *Fix $m \geq 2$ and $n \geq r \geq 2$, and let $d \to \infty$. Is it true that*

$$\mathrm{sat}^*\left([n]^d, [r]^m\right) = O\left(n^d\right)?$$

It would also be interesting to study the cases where d is fixed and n is large, or where n is a function of d. More generally, one could consider similar problems for the asymmetric grid $\prod_{i=1}^{d}[a_i]$.

Let C_ℓ denote the cycle of length ℓ. The problem of determining $\mathrm{sat}\left(K_n, C_\ell\right)$ has been extensively studied [33, 39, 63, 64, 99, 110, 178, 217]. Also, in [93], the problem of determining $\mathrm{sat}\left(K_{n,m}, C_{2\ell}\right)$ is proposed. We also find it natural to ask about the saturation number of cycles in the hypercube.

Problem 2.32 (Morrison, Noel and Scott [170])**.** *For $2 \leq \ell \leq 2^{d-1}$, determine $\mathrm{sat}(Q_d, C_{2\ell})$.*

A similar question is also asked for cycles in $[n]^d$ in [170]. We remark that the related problem of determining $\mathrm{ex}\left(Q_d, C_{2\ell}\right)$ was proposed by Erdős [88] and is very well studied; see, e.g., [9, 22, 54, 66, 70, 71, 100, 101]. Here, $\mathrm{ex}\,(Q_d, C_{2\ell})$ refers to the maximum number of edges in a $C_{2\ell}$-free subgraph of Q_d.

We close by stating a saturation problem in K_n which, to our knowledge, has not been investigated apart from the case $d = 2$ [178, 217].

Problem 2.33. *For $d \geq 2$ and $n \geq 2^d$, determine* $\mathrm{sat}(K_n, Q_d)$.

The related problem of determining $\mathrm{ex}(n, Q_3)$ was proposed by Turán [215] and is still open; see [97].

Chapter 3

Chain Saturation in Set Systems

As was mentioned in Section 1.2, Gerbner et al. [104] conjectured that chain-sat$(k) = 2^{k-1}$ for all $k \geq 1$ and the conjecture is known to hold for $1 \leq k \leq 5$ [104, 171]. Gerbner et al. [104] proved that $2^{k/2-1} \leq$ chain-sat$(k) \leq 2^{k-1}$ for all k, where the upper bound is implied by the following simple construction.

Construction 3.1 (Gerbner et al. [104])**.** For $n \geq k-1$, define

$$\mathcal{G} := \mathcal{P}(k-2) \cup \{S \cup ([n] \setminus [k-2]) : S \in \mathcal{P}(k-2)\}.$$

It is easily verified that $\mathcal{G}$ is $(k+1)$-chain saturated and that $|\mathcal{G}| = 2^{k-1}$.

In this chapter we prove Theorems 1.4 and 1.5. These results disprove the conjecture of [104] and show that there exists a constant $c \in (0,1)$ such that chain-sat$(k) = 2^{(1+o(1))ck}$. The determination of this constant is, however, an open problem (Problem 3.19 below).

In Section 3.1, we prove some preliminary results which will be used throughout the chapter. In particular, we provide conditions under which a $(k+1)$-chain saturated collection $\mathcal{F}$ can be decomposed into, or constructed from, a sequence of k disjoint saturated antichains. The main ingredient in the proofs of Theorems 1.4 and 1.5 is a submultiplicativity bound on chain-sat(k) which follows from a 'combining lemma' for chain saturated collections proved in Section 3.2. In Section 3.3, we use this lemma and an explicit construction of a 7-chain saturated collection to prove Theorem 1.4. We use the submultiplicativity bound to prove Theorem 1.5 in Section 3.4.

3.1 Preliminaries

Given a collection $\mathcal{F} \subseteq \mathcal{P}(n)$, we say that a set $A \subseteq [n]$ is an *atom* for $\mathcal{F}$ if A is maximal with respect to the property that

$$\text{for every set } S \in \mathcal{F}, \text{ we have } S \cap A \in \{\emptyset, A\}. \tag{3.2}$$

We say that an atom A with $|A| \geq 2$ is *homogeneous* for $\mathcal{F}$. The following lemma of [104] says that a $(k+1)$-chain saturated collection can have at most one homogeneous set. We include a proof for completeness.

Lemma 3.3 (Gerbner et al. [104]). *If $\mathcal{F} \subseteq \mathcal{P}(n)$ is a $(k+1)$-chain saturated collection and H_1 and H_2 are homogeneous for $\mathcal{F}$, then $H_1 = H_2$.*

Proof. Suppose to the contrary that H_1 and H_2 are homogeneous for $\mathcal{F}$ and that $H_1 \neq H_2$. Then, since each of H_1 and H_2 are maximal with respect to (3.2), we have that $H_1 \cup H_2$ is not homogeneous for $\mathcal{F}$. Therefore, there is a set $S \in \mathcal{F}$ which contains some, but not all, of $H_1 \cup H_2$. Without loss of generality, we have $S \cap H_1 = H_1$ and $S \cap H_2 = \emptyset$. Now, pick $x \in H_1$ and $y \in H_2$ arbitrarily and define

$$T := (S \setminus \{x\}) \cup \{y\}.$$

Clearly T cannot be in $\mathcal{F}$ since $T \cap H_1 = H_1 \setminus \{x\}$ and H_1 is homogeneous for $\mathcal{F}$. Since $\mathcal{F}$ is $(k+1)$-chain saturated, there must exist sets $A_1, \dots, A_k \in \mathcal{F}$ and $t \in \{0, \dots, k\}$ such that

$$A_1 \subsetneq \cdots \subsetneq A_t \subsetneq T \subsetneq A_{t+1} \subsetneq \cdots \subsetneq A_k.$$

Since H_1 and H_2 are homogeneous for $\mathcal{F}$, and neither H_1 nor H_2 is contained in T, we get that $A_t \subsetneq T \setminus (H_1 \cup H_2) \subseteq S$. Similarly, $A_{t+1} \supsetneq S$. However, this implies that $\{A_1, \dots, A_k\} \cup \{S\}$ is a $(k+1)$-chain in $\mathcal{F}$, a contradiction. $\square$

By Lemma 3.3, if $\mathcal{F}$ is a $(k+1)$-chain saturated collection for which there exists a homogeneous set, then the homogeneous set must be unique. Throughout the rest of the chapter, it will be useful to distinguish the elements of $\mathcal{F}$ which contain the homogeneous set from those that do not.

Definition 3.4. Let $\mathcal{F} \subseteq \mathcal{P}(n)$ be $(k+1)$-chain saturated with homogeneous set H. We say that a set $S \in \mathcal{F}$ is *large* if $H \subseteq S$ or *small* if $S \cap H = \emptyset$. Let $\mathcal{F}^{\text{large}}$ and $\mathcal{F}^{\text{small}}$ denote the collection of large and small sets of $\mathcal{F}$, respectively. Thus, $\mathcal{F} = \mathcal{F}^{\text{small}} \sqcup \mathcal{F}^{\text{large}}$.

Lemma 3.5. *Let $\mathcal{A} \subseteq \mathcal{P}(n)$ be a saturated antichain with homogeneous set H. Then every set $S \in \mathcal{P}(n) \setminus \mathcal{A}$ either contains a set in $\mathcal{A}^{\text{small}}$ or is contained in a set of $\mathcal{A}^{\text{large}}$.*

Proof. Suppose, to the contrary, that $S \in \mathcal{P}(n) \setminus \mathcal{A}$ does not contain a set of $\mathcal{A}^{\text{small}}$ and is not contained in a set of $\mathcal{A}^{\text{large}}$. Since $\mathcal{A}$ is saturated, we get that either

(a) there exists $A \in \mathcal{A}^{\text{large}}$ such that $A \subsetneq S$, or

(b) there exists $B \in \mathcal{A}^{\text{small}}$ such that $S \subsetneq B$.

Suppose that (a) holds. Let $x \in H$ and $y \in S \setminus A$ and define $T := (A \setminus \{x\}) \cup \{y\}$. Since H is homogeneous for $\mathcal{A}$ and $T \cap H = H \setminus \{x\}$, we must have $T \notin \mathcal{A}$. Also, since H is homogeneous for $\mathcal{A}$, any set $T' \in \mathcal{A}$ containing T would have to contain $T \cup \{x\} \supsetneq A$. Therefore, since $\mathcal{A}$ is an antichain, no such set T' can exist. Thus, there is a set $T'' \in \mathcal{A}$ such that $T'' \subsetneq T \subseteq S$. Since H is homogeneous for $\mathcal{A}$ and $T \cap H \neq H$, we get that $T'' \in \mathcal{A}^{\text{small}}$, contradicting our assumption on S.

Note that we are also done in the case that (b) holds by considering the saturated antichain $\{X \setminus A : A \in \mathcal{A}\}$ and applying the argument of the previous paragraph. $\square$

There is a natural way to partition a $(k+1)$-chain free collection $\mathcal{F} \subseteq \mathcal{P}(n)$ into a sequence of k pairwise disjoint antichains. Specifically, for $0 \leq i \leq k-1$, let $\mathcal{A}_i$ be the collection of all minimal elements of $\mathcal{F} \setminus \left(\bigcup_{j<i} \mathcal{A}_j\right)$ under inclusion. We say that $(\mathcal{A}_i)_{i=0}^{k-1}$ is the *canonical decomposition* of $\mathcal{F}$ into antichains.

Our next goal is to provide conditions under which a sequence of k pairwise disjoint saturated antichains can be united to obtain a $(k+1)$-chain saturated collection. Later we will prove a partial converse: if $\mathcal{F} \subseteq \mathcal{P}(n)$ is a $(k+1)$-chain saturated collection with a homogeneous set, then every antichain of the canonical decomposition of $\mathcal{F}$ is saturated. We also provide an example which shows that this is not necessarily the case if we remove the condition that $\mathcal{F}$ has a homogeneous set.

Definition 3.6. We say that a sequence $(\mathcal{D}_i)_{i=0}^{t}$ of subsets of $\mathcal{P}(n)$ is *layered* if, for $1 \leq i \leq t$, every $D \in \mathcal{D}_i$ strictly contains some $D' \in \mathcal{D}_{i-1}$ as a subset.

Note that the canonical decomposition of any collection is layered.

Lemma 3.7. *If $(\mathcal{A}_i)_{i=0}^{t}$ is a layered sequence of pairwise disjoint saturated antichains, then every $A \in \mathcal{A}_i$ is strictly contained in some $B \in \mathcal{A}_{i+1}$*

Proof. Let $A \in \mathcal{A}_i$. Since $\mathcal{A}_{i+1}$ is a saturated antichain disjoint from $\mathcal{A}_i$, there exists some $B \in \mathcal{A}_{i+1}$ such that either $B \subsetneq A$ or $A \subsetneq B$. In the latter case we are done, so suppose $B \subsetneq A$. Since $(\mathcal{A}_i)_{i=0}^{t}$ is layered, there exists some $A' \in \mathcal{A}_i$ such that $A' \subsetneq B$. Hence we have $A' \subsetneq B \subsetneq A$, contradicting the fact that $\mathcal{A}_i$ is an antichain and completing the proof. □

Lemma 3.8. *If $(\mathcal{A}_i)_{i=0}^{k-1}$ is a layered sequence of pairwise disjoint saturated antichains in $\mathcal{P}(n)$, then $\mathcal{F} := \bigcup_{i=0}^{k-1} \mathcal{A}_i$ is $(k+1)$-chain saturated.*

Proof. Clearly, $\mathcal{F}$ cannot contain a $(k+1)$-chain since $\mathcal{A}_0, \dots, \mathcal{A}_{k-1}$ are antichains. Let $S \in \mathcal{P}(n) \setminus \mathcal{F}$ be arbitrary and define $t := \max\{i : S \supsetneq A \text{ for some } A \in \mathcal{A}_i\}$ (if no such i exists, set $t := -1$). If $t \geq 0$, then S strictly contains some set $A_t \in \mathcal{A}_t$. As $(\mathcal{A}_i)_{i=0}^{k-1}$ is layered, for $0 \leq i \leq t-1$, there exist sets $A_i \in \mathcal{A}_i$ such that

$$A_0 \subsetneq \cdots \subsetneq A_t \subsetneq S.$$

Now, if $t \geq k-2$, then since $\mathcal{A}_{t+1}$ is a saturated antichain and S does not contain a set of $\mathcal{A}_{t+1}$, there must exist $A_{t+1} \in \mathcal{A}_{t+1}$ such that $S \subsetneq A_{t+1}$. By Lemma 3.7, we see that for $t+2 \leq i \leq k-1$ there exists $A_i \in \mathcal{A}_i$ such that

$$S \subsetneq A_{t+1} \subsetneq \cdots \subsetneq A_{k-1}.$$

Thus $\{A_0, \dots, A_{k-1}\} \cup \{S\}$ is a $(k+1)$-chain, as desired. □

In Lemma 3.8, we require the sequence $(\mathcal{A}_i)_{i=0}^{k-1}$ of saturated antichains to be layered. As it turns out, if each antichain $\mathcal{A}_i$ has a homogeneous set, then $(\mathcal{A}_i)_{i=0}^{k-1}$ is layered if and only if $\left(\mathcal{A}_i^{\text{small}}\right)_{i=0}^{k-1}$ is layered. This will be useful in the proof of Theorem 1.4.

Lemma 3.9. *Let $(\mathcal{A}_i)_{i=0}^{k-1}$ be a sequence of pairwise disjoint saturated antichains in $\mathcal{P}(n)$, each of which has a homogeneous set. Then $(\mathcal{A}_i)_{i=0}^{k-1}$ is layered if and only if $\left(\mathcal{A}_i^{\text{small}}\right)_{i=0}^{k-1}$ is layered.*

Proof. Suppose that $(\mathcal{A}_i)_{i=0}^{k-1}$ is layered and, for some $i \geq 0$, let $A \in \mathcal{A}_{i+1}^{\text{small}}$ be arbitrary. We show that A contains a set of $\mathcal{A}_i^{\text{small}}$. Otherwise, since $(\mathcal{A}_i)_{i=0}^{k-1}$ is layered, we get that there is some $B \in \mathcal{A}_i^{\text{large}}$ such that $B \subsetneq A$. Therefore, since $\mathcal{A}_i$ is an antichain, A cannot be contained in an element of $\mathcal{A}_i^{\text{large}}$. By Lemma 3.5 and the fact that $\mathcal{A}_i$ and $\mathcal{A}_{i+1}$ are disjoint, we get that A contains a set of $\mathcal{A}_i^{\text{small}}$, as desired.

Now, suppose that $\left(\mathcal{A}_i^{\text{small}}\right)_{i=0}^{k-1}$ is layered. Given $i \geq 0$ and $S \in \mathcal{A}_{i+1}^{\text{large}}$, we show that S contains a set of $\mathcal{A}_i$, which will complete the proof. If not, then since $\mathcal{A}_i$ is saturated and disjoint from $\mathcal{A}_{i+1}$, there must exist $T \in \mathcal{A}_i$ such that $S \subsetneq T$. Since $\mathcal{A}_{i+1}$ is an antichain, S cannot be strictly contained in a set of $\mathcal{A}_{i+1}^{\text{large}}$, and so neither can T. Therefore, by Lemma 3.5, there is a set $A \in \mathcal{A}_{i+1}^{\text{small}}$ contained in T. However, since $\left(\mathcal{A}_i^{\text{small}}\right)_{i=0}^{k-1}$ is layered, there exists $A' \in \mathcal{A}_i^{\text{small}}$ such that $A' \subsetneq A$. But then, $A' \subsetneq T$, which contradicts the assumption that $\mathcal{A}_i$ is an antichain. The result follows. $\square$

It is natural to wonder whether a converse to Lemma 3.8 is true. That is: *if $\mathcal{F}$ is $(k+1)$-chain saturated, does it follow that the canonical decomposition of $\mathcal{F}$ is a layered sequence of pairwise disjoint saturated antichains?* The following example shows that this is not always the case. For illustrative purposes, it will sometimes be useful to consider set systems on a general ground set X, as opposed to $[n]$. For a general set X, let $\mathcal{P}(X)$ denote the collection of all subsets of X.

Example 3.10. Let $X := \{x_1, x_2, x_3\}$, $Y := \{y_1, y_2, y_3\}$ and $Z := X \cup Y$. We define

$$\mathcal{B}_0 := \{\{x_i, x_j\} : i \neq j\} \cup \{\{x_i, y_i\} : i \in \{1, 2, 3\}\} \cup \{\{x_k, y_i, y_j\} : i, j, k \text{ distinct }\} \cup \{Y\},$$

$$\mathcal{B}_1 := \{X, \{x_1, x_2, y_1\}, \{x_1, x_3, y_3\}, \{x_2, x_3, y_2\}, \{x_1, y_1, y_3\}, \{x_2, y_1, y_2\}, \{x_3, y_2, y_3\},$$
$$\{x_1, x_2, y_2, y_3\}, \{x_1, x_3, y_1, y_2\}, \{x_2, x_3, y_1, y_3\}\}.$$

One can easily check by hand that $\mathcal{B}_1$ and $\mathcal{B}_2$ are antichains and that every element in $\mathcal{B}_2$ contains an element of $\mathcal{B}_1$; that is, $(\mathcal{B}_i)_{i=0}^{1}$ is a layered sequence of disjoint antichains. In fact, $(\mathcal{B}_i)_{i=0}^{1}$ is the canonical decomposition of the collection $\mathcal{F} := \mathcal{B}_0 \cup \mathcal{B}_1$. Clearly $\mathcal{B}_1$ is not a saturated antichain as $Y \notin \mathcal{B}_1$ and Y is not comparable to any element of $\mathcal{B}_1$. We claim that $\mathcal{F}$ is 3-chain saturated.

Consider any $S \in \mathcal{P}(Z) \setminus \mathcal{F}$. We will show that $\mathcal{F} \cup \{S\}$ contains a 3-chain. It is easy to check that every element of $\mathcal{B}_0 \setminus \{Y\}$ is contained in a set of $\mathcal{B}_1$. Hence if S is contained in some set $B \in \mathcal{B}_0 \setminus \{Y\}$, then $\mathcal{F} \cup \{S\}$ contains a 3-chain. In particular, this completes the proof when $|S| \in \{0, 1, 2\}$. Similarly, since $(\mathcal{B}_i)_{i=0}^{1}$ is layered, if S contains some set $B \in \mathcal{B}_1$, then $\mathcal{F} \cup \{S\}$ contains a 3-chain. Therefore, we are done if $|S| \in \{4, 5, 6\}$.

It remains to consider the case $|S| = 3$. Since $X, Y \in \mathcal{F}$, we must have $|S \cap Y| = 2$, or $|S \cap X| = 2$. If $|S \cap Y| = 2$, we have $S \in \{\{x_1, y_1, y_2\}, \{x_2, y_2, y_3\}, \{x_3, y_1, y_3\}\}$. This implies that S is contained in a set $B \in \mathcal{B}_1$ and contains a set $B' \in \mathcal{B}_0 \cap \mathcal{P}(n)$. If $|S \cap X| = 2$, then S contains some set $\{x_i, x_j\} \in \mathcal{B}_0$. Also, it is easily verified that S is contained in a set of $\mathcal{B}_1$. Thus, $\mathcal{F}$ is a 3-chain saturated collection whose canonical decomposition contains an antichain which is not saturated, namely, $\mathcal{B}_1$.

However, for a $(k+1)$-chain saturated collection with a homogeneous set, the converse to Lemma 3.8 does hold; we can partition $\mathcal{F}$ into a layered sequence of k pairwise disjoint saturated antichains.

Lemma 3.11. *Let $\mathcal{F} \in \mathcal{P}(n)$ be a $(k+1)$-chain saturated collection with homogeneous set H and canonical decomposition $(\mathcal{A}_i)_{i=0}^{k-1}$. Then $\mathcal{A}_i$ is saturated for all i.*

Proof. Fix i and let $S \in \mathcal{P}(n) \setminus \mathcal{A}_i$. Let $x \in H$ and define

$$T := (S \setminus H) \cup \{x\}.$$

Then $T \notin \mathcal{F}$ since $T \cap H = \{x\}$ and H is homogeneous for $\mathcal{F}$. Therefore, there exists $\{A_0, \dots, A_{k-1}\} \subseteq \mathcal{F}$ and $t \in \{0, \dots, k\}$ such that

$$A_0 \subsetneq \cdots \subsetneq A_{t-1} \subsetneq T \subsetneq A_t \subsetneq \cdots \subsetneq A_{k-1}.$$

By definition of the canonical decomposition, we must have $A_j \in \mathcal{A}_j$ for all j. Also, since H is homogeneous for $\mathcal{F}$ and $T \cap H \notin \{\emptyset, H\}$, we must have $A_{t-1} \subseteq T \setminus H \subseteq S$ and $A_t \supseteq T \cup H \supseteq S$. Therefore,

$$A_0 \subsetneq \cdots \subsetneq A_{t-1} \subseteq S \subseteq A_t \subsetneq \cdots \subsetneq A_{k-1}.$$

Since $S \neq A_i$, we must have either $A_i \subsetneq S$ or $S \subsetneq A_i$ depending on whether or not $i < t$. Therefore, $\mathcal{A}_i$ is saturated for all i. $\square$

3.2 Combining Chain Saturated Collections

Our first goal in this section is to prove that, under certain conditions, a (k_1+1)-chain saturated collection $\mathcal{F}_1 \subseteq \mathcal{P}(X_1)$ and a (k_2+1)-chain saturated collection $\mathcal{F}_2 \subseteq \mathcal{P}(X_2)$ can be combined to yield a (k_1+k_2-1)-chain saturated collection in $\mathcal{P}(X_1 \cup X_2)$. This will be used to prove Theorem 1.4, as well as Theorem 1.5.

Lemma 3.12. *Let X_1 and X_2 be disjoint sets. For $i \in \{1,2\}$, let $\mathcal{F}_i \subseteq \mathcal{P}(X_i)$ be a (k_i+1)-chain saturated collection containing $\{\emptyset, X_i\}$ and let $H_i \subsetneq X_i$ be homogeneous for $\mathcal{F}_i$. If $\mathcal{G} \subseteq \mathcal{P}(X_1 \cup X_2)$ is defined by*

$$\mathcal{G} := \left\{A \cup B : A \in \mathcal{F}_1^{\text{small}}, B \in \mathcal{F}_2^{\text{small}}\right\} \cup \left\{S \cup T : S \in \mathcal{F}_1^{\text{large}}, T \in \mathcal{F}_2^{\text{large}}\right\},$$

then $\mathcal{G}$ is a (k_1+k_2-1)-chain saturated collection which contains $\{\emptyset, X_1 \cup X_2\}$ and $H_1 \cup H_2$ is homogeneous for $\mathcal{G}$.

Proof. It is clear that $\mathcal{G}$ contains $\{\emptyset, X_1 \cup X_2\}$ and that $H_1 \cup H_2$ is homogeneous for $\mathcal{G}$. We show that $\mathcal{G}$ is a (k_1+k_2-1)-chain saturated collection.

First, let us show that $\mathcal{G}$ does not contain a chain of length k_1+k_2-1. Suppose that $\{A_1, \ldots, A_r\}$ is an r-chain in $\mathcal{G}$. We can assume that $A_1 = \emptyset$ and $A_r = X_1 \cup X_2$. Define

$$I_1 := \{i : A_i \cap X_1 \subsetneq A_{i+1} \cap X_1\}, \text{ and}$$

$$I_2 := \{i : A_i \cap X_2 \subsetneq A_{i+1} \cap X_2\}.$$

Clearly, $I_1 \cup I_2 = \{1, \ldots, r-1\}$. Also, for $i \in \{1,2\}$, since $\mathcal{F}_i$ has no (k_i+1)-chain, we must have $|I_i| \leq k_i - 1$. Let t be the maximum index such that $A_t \cap X_1 \in \mathcal{F}_1^{\text{small}}$. Note that t exists and is less than r since $A_1 = \emptyset$ and $A_r = X_1 \cup X_2$. By construction of $\mathcal{G}$, $A_t \cap X_2$ is a small set for $\mathcal{F}_2$ and, for $i \in \{1,2\}$, $A_{t+1} \cap X_i$ is a large set for $\mathcal{F}_i$. This implies that $t \in I_1 \cap I_2$ and so

$$r - 1 = |I_1 \cup I_2| = |I_1| + |I_2| - |I_1 \cap I_2| \leq k_1 + k_2 - 3$$

as required.

Now, let $S \in \mathcal{P}(X_1 \cup X_2) \setminus \mathcal{G}$. We show that $\mathcal{G} \cup \{S\}$ contains a (k_1+k_2-1)-chain. Fix $x_1 \in H_1$ and $x_2 \in H_2$ and define

$$T := (S \setminus (H_1 \cup H_2)) \cup \{x_1, x_2\}.$$

For $i \in \{1, 2\}$, let $T_i := T \cap X_i$. Then $T_i \notin \mathcal{F}_i$ since $T_i \cap H_i = \{x_i\}$. Therefore, there exists $A^i_1, \ldots, A^i_{k_i} \in \mathcal{F}_i$ and $t_i \in \{1, \ldots, k_i - 1\}$ such that

$$\emptyset = A^i_1 \subsetneq \cdots \subsetneq A^i_{t_i} \subsetneq T_i \subsetneq A^i_{t_i+1} \subsetneq \cdots \subsetneq A^i_{k_i} = X_i$$

Note that $A^i_j \in \mathcal{F}^{\text{small}}_i$ for $j \le t_i$ and $A^i_j \in \mathcal{F}^{\text{large}}_i$ for $j \ge t_i + 1$. This implies that $A^1_{t_1} \cup A^2_{t_2} \subsetneq S$ and $A^1_{t_1+1} \cup A^2_{t_2+1} \supsetneq S$. Therefore,

$$A^1_1 \cup A^2_1 \subsetneq A^1_1 \cup A^2_2 \subsetneq \cdots \subsetneq A^1_1 \cup A^2_{t_2} \subsetneq A^1_2 \cup A^2_{t_2} \subsetneq \cdots \subsetneq A^1_{t_1} \cup A^2_{t_2} \subsetneq S$$

$$\subsetneq A^1_{t_1+1} \cup A^2_{t_2+1} \subsetneq A^1_{t_1+1} \cup A^2_{t_2+2} \subsetneq \cdots \subsetneq A^1_{t_1+1} \cup A^2_{k_2} \subsetneq A^1_{t_1+2} \cup A^2_{k_2} \subsetneq \cdots \subsetneq A^2_{k_1} \cup A^2_{k_2}$$

and so $\mathcal{G} \cup \{S\}$ contains a $(k_1 + k_2 - 1)$-chain. The result follows. □

Remark 3.13. If $\mathcal{F}_1$, $\mathcal{F}_2$ and $\mathcal{G}$ are as in Lemma 3.12, then

$$|\mathcal{G}| = \left|\mathcal{F}^{\text{small}}_1\right| \left|\mathcal{F}^{\text{small}}_2\right| + \left|\mathcal{F}^{\text{large}}_1\right| \left|\mathcal{F}^{\text{large}}_2\right|.$$

3.3 Proof of Theorem 1.4

We apply Lemma 3.12 to prove Theorem 1.4. The first part of the proof of Theorem 1.4 is to exhibit an infinite family of 7-chain saturated collections with cardinality $30 < 2^5$.

Proposition 3.14. *There is a 7-chain saturated collection $\mathcal{F} \subseteq \mathcal{P}(n)$ with a homogeneous set such that $\left|\mathcal{F}^{\text{small}}\right| = \left|\mathcal{F}^{\text{large}}\right| = 15$.*

Proof. Let $n \ge 8$. Let x_1, x_2, y_1, y_2, w and z be distinct elements of $[n]$ and define $H := [n] \setminus \{x_1, x_2, y_1, y_2, w, z\}$. We apply Lemma 3.8 to construct a 7-chain saturated collection $\mathcal{F} \subseteq \mathcal{P}(n)$ of order 30. Naturally, we define $\mathcal{A}_0 = \{\emptyset\}$ and $\mathcal{A}_5 := \{[n]\}$. Also, define

$$\mathcal{A}_1 := \{\{x_1\}, \{x_2\}, \{y_1\}, \{w\}, H \cup \{y_2, z\}\}, \text{ and}$$

$$\mathcal{A}_4 := \{[n] \setminus A : A \in \mathcal{A}_1\}.$$

It is easily observed that $\mathcal{A}_1$ and $\mathcal{A}_4$ are saturated antichains. We define $\mathcal{A}_2$ and $\mathcal{A}_3$ by first specifying their small sets. Define

$$\mathcal{A}^{\text{small}}_2 := \{\{x_i, y_j\} : 1 \le i, j \le 2\} \cup \{\{w, z\}\}, \text{ and}$$

$$\mathcal{A}_3^{\text{small}} := \{\{x_1, y_1, w\}, \{x_1, y_1, z\}, \{x_2, y_2, w\}, \{x_2, y_2, z\}\}.$$

Given any collection $\mathcal{B} \subseteq \mathcal{P}(n)$, a set $S \subseteq [n]$ is said to be *stable* for $\mathcal{B}$ if S does not contain an element of $\mathcal{B}$. For $i = 2, 3$, define $\mathcal{A}_i^{\text{large}}$ to be the collection consisting of all maximal stable sets for $\mathcal{A}_i^{\text{small}}$ and let $\mathcal{A}_i := \mathcal{A}_i^{\text{small}} \cup \mathcal{A}_i^{\text{large}}$. Note that every element of $\mathcal{A}_i^{\text{large}}$ contains H. It is clear that $\mathcal{A}_i$ is an antichain for $i = 2, 3$. Moreover, $\mathcal{A}_i$ is saturated since every set $A \in \mathcal{P}(n)$ either contains an element of $\mathcal{A}_i^{\text{small}}$ or is contained in an element of $\mathcal{A}_i^{\text{large}}$.

One can easily verify by hand that $\left(\mathcal{A}_i^{\text{small}}\right)_{i=0}^5$ is layered. Therefore, by Lemma 3.9, $(\mathcal{A}_i)_{i=0}^5$ is a layered sequence of pairwise disjoint saturated antichains. By Lemma 3.8, $\mathcal{F} := \bigcup_{i=0}^5 \mathcal{A}_i$ is a 7-chain saturated collection. Also,

$$\left|\mathcal{F}^{\text{small}}\right| = \sum_{i=0}^{5} \left|\mathcal{A}_i^{\text{small}}\right| = (1 + 5 + 9 + 0) = 15, \text{ and}$$

$$\left|\mathcal{F}^{\text{large}}\right| = \sum_{i=0}^{5} \left|\mathcal{A}_i^{\text{large}}\right| = (0 + 9 + 5 + 1) = 15,$$

as desired. □

We remark that the construction in Proposition 3.14 bears some similarity to one which was used in [104] to prove that $\text{chain-sat}(k, k) \leq \frac{15}{16} 2^{k-1}$ for every $k \geq 6$.

For the proof of Theorem 1.4 we will use the fact that

$$\text{chain-sat}(k) \leq 2\, \text{chain-sat}(k-1). \tag{3.15}$$

This was proved in [104] using the fact that if $\mathcal{F} \subseteq \mathcal{P}(n)$ is k-chain saturated, then $\mathcal{F} \cup \{A \cup \{n+1\} : A \in \mathcal{F}\}$ is a $(k+1)$-chain saturated collection in $\mathcal{P}(n+1)$.

Proof of Theorem 1.4. First, we prove that the result holds when k is of the form $4j+2$ for some $j \geq 1$. In this case, we repeatedly apply Lemma 3.12 and Proposition 3.14 to obtain a $(k+1)$-chain saturated collection $\mathcal{F}$ on an arbitrarily large ground set X such that

$$\left|\mathcal{F}^{\text{small}}\right| + \left|\mathcal{F}^{\text{large}}\right| = 15^j + 15^j = 2 \cdot 15^j.$$

Therefore, if $k = 4j + 2$, then $\text{chain-sat}(k) \leq 2 \cdot 15^j$.

For k of the form $4j + 2 + s$ for $j \geq 1$ and $1 \leq s \leq 3$, apply (3.15) to obtain $\text{chain-sat}(k) \leq 2^s\, \text{chain-sat}(4j+2) \leq 2^{s+1} \cdot 15^j$. Thus, we are done by setting ε slightly smaller than $\left(1 - \frac{\log_2(15)}{4}\right)$. □

3.4 Asymptotic Behaviour of chain-sat(k)

To prove Theorem 1.5, we require the following facts, which are proved in [104]. We include proofs for completeness.

Lemma 3.16 (Gerbner et al. [104]). *For $n \geq k \geq 1$, there is a $(k+1)$-chain saturated collection $\mathcal{F} \subseteq \mathcal{P}(n)$ such that $|\mathcal{F}| = \text{chain-sat}(n,k)$ and $\{\emptyset, [n]\} \subseteq \mathcal{F}$.*

Proof. Let $\mathcal{F} \subseteq \mathcal{P}(n)$ be any $(k+1)$-chain saturated collection such that $|\mathcal{F}| = \text{chain-sat}(n,k)$. We let $(\mathcal{A}_i)_{i=0}^{k-1}$ denote the canonical decomposition of $\mathcal{F}$ and define

$$\mathcal{F}' := (\mathcal{F} \setminus (\mathcal{A}_0 \cup \mathcal{A}_{k-1})) \cup \{\emptyset, [n]\}.$$

It is clear that $\mathcal{F}'$ is $(k+1)$-chain saturated and $|\mathcal{F}'| \leq |\mathcal{F}| = \text{chain-sat}(n,k)$, which proves the result. $\square$

Claim 3.17 (Gerbner et al. [104]). *Given a collection $\mathcal{F} \subseteq \mathcal{P}(n)$, if $n > 2^{|\mathcal{F}|}$, then there is a homogeneous set for $\mathcal{F}$.*

Proof. We observe that every atom A for $\mathcal{F}$ corresponds to a subcollection $\mathcal{F}_A := \{S \in \mathcal{F} : A \subseteq S\}$ of $\mathcal{F}$ such that $\mathcal{F}_A \neq \mathcal{F}_{A'}$ whenever $A \neq A'$. This implies that the number of atoms for $\mathcal{F}$ is at most $2^{|\mathcal{F}|}$. Therefore, since $n > 2^{|\mathcal{F}|}$, there must be a non-singleton atom, i.e. a homogeneous set, for $\mathcal{F}$. $\square$

Proof of Theorem 1.5. We show that, for all k, ℓ,

$$\text{chain-sat}(k+\ell) \leq 4\,\text{chain-sat}(k)\,\text{chain-sat}(\ell). \tag{3.18}$$

Letting $f(k) := 4\,\text{chain-sat}(k)$, we see that (3.18) implies that $f(k+\ell) \leq f(k)f(\ell)$ for every k, ℓ. It follows by Fekete's Lemma that $f(k)^{1/k}$ converges, and so $\text{chain-sat}(k)^{1/k}$ converges as well.

For $n > 2^{2^{k+\ell-2}}$, let X and Y be disjoint sets of size n and let $\mathcal{F}_k \subseteq \mathcal{P}(X)$ and $\mathcal{F}_\ell \subseteq \mathcal{P}(Y)$ be $(k+1)$-chain saturated and $(\ell+1)$-chain saturated collections of cardinalities chain-sat(k) and chain-sat(ℓ), respectively. By Claim 3.17, we can assume that $\mathcal{F}_k$ and $\mathcal{F}_\ell$ have homogeneous sets and, by Lemma 3.16, we can assume that $\{\emptyset, X\} \subseteq \mathcal{F}_k$ and $\{\emptyset, Y\} \subseteq \mathcal{F}_\ell$. We apply Lemma 3.12 and Remark 3.13 to obtain a $(k + \ell - 1)$-chain saturated collection $\mathcal{G} \subseteq \mathcal{P}(X \cup Y)$ of order at most $|\mathcal{F}_k||\mathcal{F}_\ell| = \text{chain-sat}(k)\,\text{chain-sat}(\ell)$. Therefore, by (3.15), we have

$$\text{chain-sat}(k+\ell) \le 4\,\text{chain-sat}(k+\ell-2) \le 4|\mathcal{G}| \le 4\,\text{chain-sat}(k)\,\text{chain-sat}(\ell)$$

as required. $\square$

A result of Gerbner et al. [104] implies that the constant c in Theorem 1.5 is at least $1/2$, and Theorem 1.2 implies that $c \le \frac{\log_2(15)}{4}$. We wonder about the precise value of c.

Problem 3.19 (Morrison, Noel and Scott [171])**.** *Determine the value of c such that* $\text{chain-sat}(k) = 2^{(1+o(1))ck}$.

Chapter 4

Weak Saturation in the Hypercube and the Grid

The study of weak saturation was initiated by Bollobás [37], who conjectured in [39] that $\mathrm{wsat}(K_n, K_r) = \mathrm{sat}(K_n, K_r)$ for all $n \geq r \geq 1$. A proof of this conjecture can be derived from the following result of Frankl [95], which is usually referred to as the 'Skew Two Families Theorem.' For us, the most important difference between the Skew Two Families Theorem and Bollobás' Two Families Theorem is that the condition $i \neq j$ has been replaced with $i < j$ in (b).

Skew Two Families Theorem (Frankl [95])**.** *Let $A_1, \dots, A_m$ and $B_1, \dots, B_m$ be two families of sets such that $|A_i| \leq k$ and $|B_i| \leq \ell$ for $1 \leq i \leq m$ and the following hold:*

(a) $A_i \cap B_i = \emptyset$ for $1 \leq i \leq m$, and

(b) $A_i \cap B_j \neq \emptyset$ for $1 \leq i < j \leq m$.

Then $m \leq \binom{k+\ell}{k}$.

Remark 4.1. Of course, another difference between the Skew Two Families Theorem and Bollobás' Two Families is that the conclusion $\sum_{i=1}^{m} \binom{|A_i|+|B_i|}{|A_i|}^{-1} \leq 1$ has been replaced with the weaker statement $m \leq \binom{k+\ell}{k}$. In fact, this difference is essential, as the bound $\sum_{i=1}^{m} \binom{|A_i|+|B_i|}{|A_i|}^{-1} \leq 1$ does not necessarily hold under the hypotheses of the Skew Two Families Theorem. For example, suppose that $A_1, \dots, A_{2^n}$ consists of all subsets of $[n]$ ordered by decreasing cardinality and set $B_i := [n] \setminus A_i$ for $1 \leq i \leq 2^n$. Then clearly $A_1, \dots, A_{2^n}$ and $B_1, \dots, B_{2^n}$ satisfy (a) and (b), but $\sum_{i=1}^{2^n} \binom{|A_i|+|B_i|}{|A_i|}^{-1} = n$.

Interestingly, the idea behind Frankl's proof originated in a paper of Lovász [152] which, on its surface, had very little to do with weak saturation. In that paper, Lovász used multilinear techniques (exterior algebra) to generalise Bollobás' Two Families Theorem to a statement regarding flats of rank r in a matroid representable over a commutative field. These methods have since been applied by Alon [4], Kalai [136] and Pikhurko [183] to obtain various weak saturation results for graphs and hypergraphs.

In the previous chapter, we used Bollobás' Two Families Theorem to obtain lower bounds on saturation numbers of complete hypergraphs. In much the same way, one can use the Skew Two Families Theorem to obtain the same lower bounds for the analogous weak saturation numbers; in particular, as mentioned above, the Skew Two Families Theorem implies the conjecture of Bollobás [39] that $\text{wsat}(K_n, K_r) = \text{sat}(K_n, K_r)$ for all $n \geq r \geq 1$.

In [4], Alon obtained the following 'multipartite' generalisation of the Skew Two Families Theorem. As mentioned above, his proof also makes use of multilinear techniques.

Partite Skew Two Families Theorem (Alon [4])**.** *Let $X_1, \ldots, X_k$ be disjoint sets and let $a_1, \ldots, a_k$ and $b_1, \ldots, b_k$ be non-negative integers. Suppose that $A_1, \ldots, A_m$ and $B_1, \ldots, B_m$ are two families of subsets of $X_1 \cup \cdots \cup X_k$ such that*

(a) $|A_i \cap X_j| \leq a_j$ and $|B_i \cap X_j| \leq b_j$ for $1 \leq i \leq m$ and $1 \leq j \leq k$,

(b) $A_i \cap B_i = \emptyset$ for $1 \leq i \leq m$, and

(c) $A_i \cap B_j \neq \emptyset$ for $1 \leq i < j \leq m$.

Then $m \leq \prod_{i=1}^{k} \binom{a_i+b_i}{a_i}$.

It is eay to see that the Skew Two Families Theorem is equivalent to the case $k = 1$ of the Partite Skew Two Families Theorem. As a simple application of the Partite Skew Two Families Theorem, let us argue that $\text{wsat}\left(K_{n,n}, K_{r,r}\right) = n^2 - (n-r+1)^2$, as was done in [4].

First, note that an upper bound of $n^2 - (n-r+1)^2$ follows immediately from (2.2) via (1.7). Now, let G be a weakly $(K_{n,n}, K_{r,r})$-saturated graph with $n^2 - m$ edges and label the edges of $E\left(K_{n,n}\right) \setminus E(G)$ by $B_1, \ldots, B_m$ so that, for each $i \geq 1$, B_i is contained in a copy of $K_{r,r}$ in $G + \{B_1, \ldots, B_i\}$. Let V_i be the vertex set of that copy of $K_{r,r}$ and define $A_i := V(K_{n,n}) \setminus V_i$. Applying the Partite Two Families Theorem to $A_1, \ldots, A_m$ and $B_1, \ldots, B_m$ with $k = 2$, $a_1 = a_2 = n - r$, $b_1 = b_2 = 1$ and X_1 and X_2 being the two parts of $K_{n,n}$, we get $m \leq (n - r + 1)^2$. This proves that $\text{wsat}(K_{n,n}, K_{r,r}) = n^2 - (n - r + 1)^2$. This argument extends naturally to give a weak saturation result for k-partite k-uniform hypergraphs in which each hyperedge contains exactly one vertex from each part; see [4].

One may notice that the argument of the previous paragraph does not adapt easily to give a tight lower bound on $\text{wsat}(K_{n,n}, K_{r,s})$ for $1 \leq r < s \leq n$. The issue is that a copy of $K_{r,s}$ can appear in $K_{n,n}$ in two different ways: with r vertices on the left and s on the right, or vice versa. This issue was resolved by Moshkovitz and Shapira [172], who proved a version of the Partite Skew Two Families Theorem which allows for a permutation of the partite sets $X_1, \ldots, X_k$. Their proof involves a clever, and purely combinatorial, reduction to the Partite Skew Two Families Theorem. We will not state their result here, but mention that it can be applied to prove that

$$\text{wsat}(K_{n,n}, K_{r,s}) \geq n^2 - (n - r + 1)^2 + (s - r)^2$$

for $1 \leq r \leq s \leq n$. This is easily seen to be tight by considering the graph obtained from $K_{n,n}$ by deleting the edges of a copy of $K_{n-r+1,n-r+1}$ and then adding the edges of a copy of $K_{s-r,s-r}$. We remark that the special case $r = 1$ can also be derived from a result of Balogh, Bollobás, Morris and Riordan [21]. The result of Moshkovitz and Shapira [172] can also be applied to obtain tight lower bounds on a wide range of 'asymmetric' weak saturation problems in k-partite k-uniform hypergraphs.

There has been some recent interest in weak saturation problems in random settings. Korándi and Sudakov [146] proved that, for fixed $r \geq 3$ and $p \in (0, 1)$,

$$\text{wsat}(G(n, p), K_r) = \binom{n}{2} - \binom{n - r + 2}{2}$$

with high probability. This is in stark contrast to their result, mentioned in Chapter 2, that $\text{sat}(G(n, p), K_r) = \Theta(n \log n)$.

Another probabilistic problem, first studied by Balogh, Bollobás and Morris [20], is the following: *given a graph H and n large, what is the infimum over all $p \in (0, 1)$ such that the probability that $G(n, p)$ is weakly (K_n, H)-saturated is at least $1/2$?* This infimal density is known as the *critical density* and is denoted $p_c(n, H)$. It is easily observed that an n-vertex graph G is weakly (K_n, K_3)-saturated if and only if it is connected. Therefore, $p_c(K_n, K_3) = \log n/n + \Theta(1/n)$ by a classical result of Erdős and Rényi [91]. For more general cliques, Balogh, Bollobás and Morris [20] proved the following bounds.

$$p_c(n, K_4) = \Theta\left(\sqrt{\frac{1}{n \log n}}\right) \text{ and}$$

$$\Omega\left(n^{-\frac{r-2}{\binom{r}{2}-2}} (\log n)^{-1}\right) \leq p_c(n, K_r) \leq n^{-\frac{r-2}{\binom{r}{2}-2}} \log n \text{ for fixed } r \geq 5.$$

In addition, they obtain bounds on $p_c(n, H)$ for graphs H which are 'balanced,' in the sense that $\frac{|E(H)|-2}{|V(H)|-2} \geq \frac{|E(H')|-1}{|V(H')|-2}$ for every proper subgraph H' of H with at least three vertices.

There are also a number of recent results regarding the amount of 'time' that a weakly (K_n, K_r)-saturated graph requires for all of its edges to be added. To make this precise, it is useful to redefine the notion of a weakly (F, H)-saturated graph in terms of an infection process, known as the *graph bootstrap process*. In the (F, H)-bootstrap process, one begins with an initial set $E_0 \subseteq E(F)$ of 'infected edges' and, at each step, a 'healthy' edge of F becomes infected if it completes an infected copy of H. We say that E_0 *percolates* if every edge of F is eventually infected. It is clear that E_0 percolates if and only if the subgraph of F induced by E_0 is weakly (F, H)-saturated. The problem of determining the *slowest* percolating set in the (K_n, K_r)-bootstrap process (i.e. that which requires the maximum number of steps) was investigated independently by Bollobás, Przykucki, Riordan and Sahasrabudhe [45] and Matzke [165]. It is proved that, for $r = 3, 4$, the maximum time for the (K_n, K_r)-bootstrap process is $\lceil \log(n-1) \rceil$ and $n - 3$, respectively. The following intriguing conjecture was posed in [45].

Conjecture 4.2 (Bollobás, Przykucki, Riordan and Sahasrabudhe [45])**.** *For fixed $r \geq 5$, the maximum time for the (K_n, K_r)-bootstrap process is $o(n^2)$.*

Gunderson, Koch and Przykucki [114] consider a 'percolation time' problem for weak saturation in a random setting. Specifically, they bound the critical density for $E(G(n,p))$ to percolate by a specified time under the (K_n, K_r)-bootstrap process.

Our goal in this chapter is to determine the weak saturation number of axis-aligned grids[1] in a larger grid. Theorem 1.8 will follow as a special case since every copy of Q_m

[1] See Section 2.4 for a definition of an 'axis-aligned' copy of $[r]^m$ in $[n]^d$.

in Q_d is axis-aligned. In what follows, the notation $\mathrm{wsat}^*\left([n]^d, [r]^m\right)$ indicates that we are considering weak saturation with respect to axis-aligned copies of $[r]^m$ in $[n]^d$. We note that the special case $\mathrm{wsat}(Q_d, Q_2) = 2^d - 1$ was proved earlier by Johnson and Pinto [133] by exhibiting the existence of a weakly (Q_d, Q_2)-saturated spanning tree of Q_d and noticing that every (Q_d, Q_2)-saturated graph must be connected.

Theorem 4.3 (Morrison, Noel and Scott [170]). *For $n \geq r \geq 2$ and $d \geq m \geq 1$,*

$$\mathrm{wsat}^*\left([n]^d, [r]^m\right) = \sum_{j=0}^{d}\sum_{i=0}^{d-j}(m-1+i)\binom{d}{j}\binom{d-j}{i}(n-r+1)^j(r-2)^i$$

$$-\sum_{j=0}^{m-2}\sum_{i=0}^{d-j}(m-1-j)\binom{d}{j}\binom{d-j}{i}(n-r+1)^j(r-2)^i,$$

where, by convention, $0^0 = 1$.

In the next section, we state and prove a linear algebraic lemma of Balogh, Bollobás, Morris and Riordan [21] which is our main tool in this chapter and the next. In Section 4.2, we apply this lemma to prove Theorem 4.3. We pose several open problems in Section 4.3.

4.1 The Linear Algebraic Lemma

Given a graph F and a set $\mathcal{H}$ of subgraphs of F, let $\mathrm{wsat}\,(F, \mathcal{H})$ be the minimum number of edges in a spanning subgraph G of F such that the edges of $E(F) \setminus E(G)$ can be added to G, one edge at a time, in such a way that each added edge increases the number of graphs of $\mathcal{H}$ contained in G; such a graph is said to be *weakly* $(F, \mathcal{H})$-*saturated*. In this language, $\mathrm{wsat}^*\left([n]^d, [r]^m\right) = \mathrm{wsat}\left([n]^d, \mathcal{H}\right)$ where $\mathcal{H}$ consists of all axis-aligned copies of $[r]^m$ in $[n]^d$.

Generally speaking, to prove an upper bound on a weak saturation number, one only needs to construct a *single* example of a weakly saturated graph of small size. It would seem that proving lower bounds is more difficult, as one must show that no graph with 'too few' edges can be weakly (F, H)-saturated. A major advantage of the following lemma is that it allows us to prove lower bounds in a constructive manner as well. In addition, as we will see in our applications of this lemma, the upper bound construction often guides the lower bound construction. This lemma (in a slightly different form) was stated and proved by Balogh, Bollobás, Morris and Riordan [21], but the idea behind it was already present in the work of Kalai [136] and Pikhurko [183]; we include a proof for completeness.

Lemma 4.4 (Balogh, Bollobás, Morris and Riordan [21]). *Let F be a graph, let $\mathcal{H}$ be a collection of subgraphs of F, and let W be a vector space. Suppose that there exists a set $\{f_e : e \in E(F)\} \subseteq W$ such that for every $H \in \mathcal{H}$ there are non-zero scalars $\{c_e : e \in E(H)\}$ such that $\sum_{e \in E(H)} c_e f_e = 0$. Then*

$$\operatorname{wsat}(F, \mathcal{H}) \geq \dim\left(\operatorname{span}\{f_e : e \in E(F)\}\right).$$

Proof. Let G be a weakly $(F, \mathcal{H})$-saturated graph. Define $G_0 := G$ and label the edges of $E(F) \setminus E(G)$ by $e_1, \ldots, e_k$ such that for $1 \leq i \leq k$ there is a subgraph $H_i \in \mathcal{H}$ of $G_i := G_{i-1} + e_i$ such that H_i contains e_i. Thus, by hypothesis, f_{e_i} can be written as a linear combination of the vectors in $\{f_e : e \in E(H_i) \setminus \{e_i\}\}$. This implies that

$$\operatorname{span}\{f_e : e \in E(G_0)\} = \operatorname{span}\{f_e : e \in E(G_1)\} = \cdots = \operatorname{span}\{f_e : e \in E(G_k)\}.$$

Since $G = G_0$ and $F = G_k$, we must have $|E(G)| \geq \dim\left(\operatorname{span}\{f_e : e \in E(F)\}\right)$. Since G was an arbitrary weakly $(F, \mathcal{H})$-saturated graph, this completes the proof. □

4.2 Proof of Theorem 4.3

Given a vertex v of $[n]^d$, say that a coordinate of v is *large* if it is at least r and *small* otherwise. We let $L(v)$ denote the number of large coordinates of v and let $|v|$ denote the sum of the coordinates of v. Given $i \in [d]$, a *line* $\mathcal{L}$ in direction i is a path of length $n-1$ in $[n]^d$ in which any two vertices of $\mathcal{L}$ differ on the ith coordinate.

Proof of Theorem 4.3. Let G be a spanning subgraph of $[n]^d$ in which, for each vertex $v \in [n]^d$,

- we add every edge uv of $[n]^d$ such that $|u| = |v| - 1$ and u and v differ on a small coordinate of v, and
- we add $\min\{L(v), m-1\}$ edges uv of $[n]^d$ such that $|u| = |v| - 1$ and u and v differ on a large coordinate of v.

Some elementary counting gives us

$$E(G) = \sum_{j=0}^{d}\sum_{i=0}^{d-j}(m-1+i)\binom{d}{j}\binom{d-j}{i}(n-r+1)^j(r-2)^i$$

$$-\sum_{j=0}^{m-2}\sum_{i=0}^{d-j}(m-1-j)\binom{d}{j}\binom{d-j}{i}(n-r+1)^j(r-2)^i.$$

So, to prove the upper bound, we need only show that G is weakly saturated with respect to axis-aligned copies of $[r]^m$. In order of increasing $|v|$, we add all edges of $E([n]^d) \setminus E(G)$ from v to vertices u with $|u| = |v| - 1$ one edge at a time. By construction, for every added edge uv with $|u| = |v| - 1$, the coordinate on which u and v differ is large in v. For any vertex x, if $L(x) \leq m - 1$, then all edges xy with $|x| = |y| - 1$ are already present in G. Also, if v is a vertex with $L(v) \geq m$, then G contains $m - 1$ edges wv with $|w| = |v| - 1$ such that the coordinate on which w and v differ is large in v. Putting this together, an easy inductive argument shows that, if we add edges in this order, then each edge added from v to a vertex u with $|u| = |v| - 1$ creates a new axis-aligned copy of $[r]^m$ in which v is the 'top' vertex.

For the lower bound, we apply Lemma 4.4 where $\mathcal{H}$ consists of all axis-aligned copies of $[r]^m$ in $[n]^d$. So, it suffices to show that there exists a vector space W and a set $\left\{f_e : e \in E\left([n]^d\right)\right\} \subseteq W$ which satisfy the hypotheses of Lemma 4.4 such that

$$\dim\left(\operatorname{span}\left\{f_e : e \in E\left([n]^d\right)\right\}\right) \geq |E(G)|. \tag{4.5}$$

The space W that we choose is the direct sum of $[n]^d$ copies of $\mathbb{R}^{m-1}$, one for each vertex of $[n]^d$, and dn^{d-1} copies of $\mathbb{R}^{r-2}$, one for each line in $[n]^d$. That is,

$$W := \left(\bigoplus_x \mathbb{R}^{m-1}\right) \oplus \left(\bigoplus_{\mathcal{L}} \mathbb{R}^{r-2}\right)$$

where x ranges over all vertices of $[n]^d$ and $\mathcal{L}$ ranges over all lines in $[n]^d$. Given a vector w of W, a vertex x of $[n]^d$ and a line $\mathcal{L}$ of $[n]^d$, let $\pi_x(w)$ denote the projection of w onto the copy of $\mathbb{R}^{m-1}$ corresponding to x and let $\pi_{\mathcal{L}}(w)$ denote the projection of w onto the copy of $\mathbb{R}^{r-2}$ corresponding to $\mathcal{L}$. Note that w is determined by its projections.

Let $Z := \{z_1, \ldots, z_d\}$ be a collection of d vectors of $\mathbb{R}^{m-1}$ in general position. Also, let $Y := \{y_1, \ldots, y_{n-1}\}$ be a set of $n - 1$ vectors of $\mathbb{R}^{r-2}$ such that $y_1, \ldots, y_{r-2}$

are linearly independent and, for $r-1 \leq t \leq n-1$,

$$y_t := -\sum_{j=t-r+2}^{t-1} y_j. \tag{4.6}$$

Thus any consecutive $r-2$ vectors $y_i, \ldots, y_{i+r-2}$ are linearly independent and any consecutive $r-1$ vectors $y_j, \ldots, y_{j+r-1}$ sum to zero. Suppose that $e = uv$ is an edge of $[n]^d$ such that u and v differ on coordinate $i \in [d]$ and let t be the minimum of the ith coordinates of u and v. Further, let $\mathcal{L}$ be the unique line of $[n]^d$ containing e. We define f_e to be the vector of W such that

- $\pi_u(f_e) = \pi_v(f_e) = z_i$ and $\pi_x(f_e) = 0$ for every $x \in V(Q_d) \setminus \{u, v\}$, and
- $\pi_{\mathcal{L}}(f_e) = y_t$ and $\pi_{\mathcal{L}'}(f_e) = 0$ for every line $\mathcal{L}' \neq \mathcal{L}$.

In order to apply Lemma 4.4, we need to show that for every axis-aligned copy H of $[r]^m$ in $[n]^d$ there are non-zero scalars $\{c_e : e \in E(H)\}$ such that $\sum_{e\in E(H)} c_e f_e = 0$. Let H be an axis-aligned copy of $[r]^m$ in $[n]^d$ and let $I \subseteq [d]$ be the set of m coordinates which vary under H. Let $\{c_i : i \in I\}$ be a set of scalars such that $\sum_{i\in I} c_i z_i = 0$. Note that $c_i \neq 0$ for all $i \in I$ since the vectors of Z are in general position.

For each vertex v of H, let $M(v) \subseteq [d]$ be the set of indices j such that both $v - e_j$ and $v + e_j$ are contained in H. Define the *lines* of H to be the paths of length $r-1$ in H obtained by taking the intersection of a line of $[n]^d$ with $V(H)$. Note that $|M(v)|$ is precisely the number of lines of H in which v has degree two. For each line $\mathcal{L}$ of H, define $m(\mathcal{L}) := |M(v)|$ where v is an endpoint of $\mathcal{L}$ (clearly the two endpoints give the same value). For each $i \in I$, let E_i be the set of all edges of H for which i is the variable coordinate. For each edge $e \in E_i$ contained in a line $\mathcal{L}$ of H, define $d_e := 2^{m(\mathcal{L})} c_i$. We claim that $\sum_{e\in E(H)} d_e f_e = 0$.

First, let $\mathcal{L}$ be a line of H in direction i. Then, for some $t \geq r-1$, we have

$$\pi_{\mathcal{L}}\left(\sum_{e\in E(H)} d_e f_e\right) = \sum_{e\in E(H)} d_e \pi_{\mathcal{L}}(f_e) = \sum_{e\in E(\mathcal{L})} 2^{m(\mathcal{L})} c_i \pi_{\mathcal{L}}(f_e) = 2^{m(\mathcal{L})} c_i \sum_{j=t-r+2}^{t} y_j.$$

However, this sum is equal to zero by (4.6).

Now, fix a vertex v of H and for each $i \in I$ let $\mathcal{L}_i$ be the line of H in direction i containing v. We have

$$\pi_v\left(\sum_{e\in E(H)} d_e f_e\right) = \sum_{e\in E(H)} d_e \pi_v(f_e)$$

$$= \sum_{i\in M(v)} \sum_{e\in E_i} 2^{m(\mathcal{L}_i)} c_i \pi_v(f_e) + \sum_{i\in I\setminus M(v)} \sum_{e\in E_i} 2^{m(\mathcal{L}_i)} c_i \pi_v(f_e)$$

which is equal to

$$\sum_{i\in M(v)} 2^{m(\mathcal{L}_i)+1} c_i \pi_v(f_e) + \sum_{i\in I\setminus M(v)} 2^{m(\mathcal{L}_i)} c_i \pi_v(f_e)$$

by definition of $M(v)$ and the fact that $\pi_v(f_e) = 0$ if e is not incident to v. We observe that, for any $i \in M(v)$, we have $m(\mathcal{L}_i) = |M(v)| - 1$ and for any $i \in I \setminus M(v)$ we have $m(\mathcal{L}_i) = |M(v)|$. Therefore, the above sum is equal to $2^{|M(v)|} \sum_{i\in I} c_i z_i$ which is zero by our choice of $\{c_i : i \in I\}$. Combining this with the result of the previous paragraph, we get $\sum_{e\in E(H)} d_e f_e = 0$, as desired.

To complete the proof, it suffices to prove (4.5). To do so, we let G be a graph as in the proof of the upper bound and show that the vectors of $\{f_e : e \in E(G)\}$ are linearly independent. Let $\{c_e : e \in E(G)\}$ be any set of scalars, not all of which are zero, and let F be the spanning subgraph of G containing all edges $e \in E(G)$ such that $c_e \neq 0$. Let v be a vertex of non-zero degree in F such that $|v|$ is maximum.

First, consider the case that v has degree at most $m - 1$ in F. Let J denote the set of $d_F(v)$ coordinates such that, for each $j \in J$, the edge from v to $v - e_j$ is present in F. Then $\sum_{e\in E(G)} c_e \pi_v(f_e)$ is a linear combination of the vectors in $\{z_j : j \in J\}$ in which not all of the coefficients are zero. Since $|J| \leq m - 1$ and the vectors of Z are in general position, this sum is non-zero and therefore $\sum_{e\in E(G)} c_e f_e \neq 0$.

Now, suppose that v has degree at least m in F. By construction of G, this implies that there is a coordinate j which is small for v such that the edge e from v to $v - e_j$ is present in F. Let $\mathcal{L}$ be the unique line of $[n]^d$ containing e. Then, by maximality of $|v|$, every edge of $\mathcal{L}$ contained in F joins two vertices for which j is a small coordinate. This implies that $\sum_{e\in E(F)} c_e \pi_{\mathcal{L}}(f_e)$ is a linear combination of the vectors in $\{y_i : 1 \leq i \leq r-1\}$ in which not all of the coefficients are zero. Thus, since the vectors $y_1, \ldots, y_{r-1}$ are linearly independent, this sum is non-zero and we obtain $\sum_{e\in E(G)} c_e f_e \neq 0$, which completes the proof. □

Remark 4.7. It is clear that, for most values of $n \geq r \geq 2$ and $d \geq m \geq 1$, there is some flexibility in the construction of the graph G in Theorem 4.3. Therefore, the tight examples are usually not unique. For an explicit example in two dimensions, see Figure 4.1.

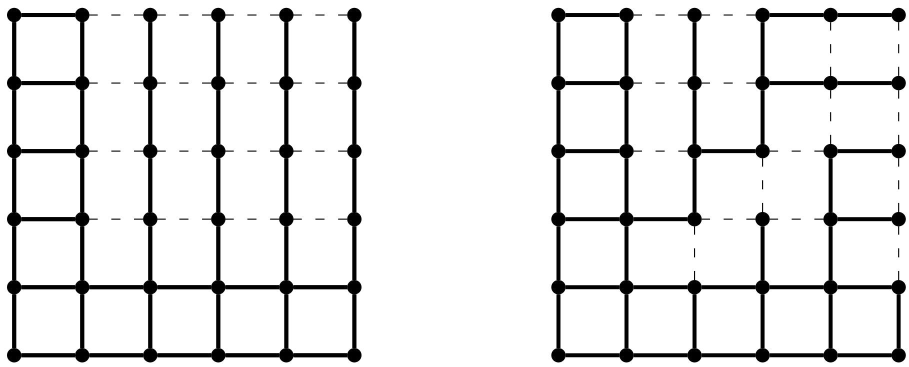

Figure 4.1: A pair of non-isomorphic weakly $([6]^2, [3]^2)$-saturated graphs, each with $\mathrm{wsat}\left([6]^2, [3]^2\right)$ edges.

4.3 Open Problems

We conclude the chapter by stating some open problems. Theorem 4.3 provides an exact expression for the weak saturation number of axis aligned copies of $[r]^m$ in $[n]^d$. However, for most values of r, m, n, d, our upper bound constructions contain several 'bent' copies of $[r]^m$. We ask about the weak saturation number of general copies of $[r]^m$ in $[n]^d$.

Problem 4.8 (Morrison, Noel and Scott [170]). *For general* $1 \le m \le d$ *and* $2 \le r \le n$*, determine* $\mathrm{wsat}\left([n]^d, [r]^m\right)$.

One could also consider the generalisations of Theorem 4.3 and Problem 4.8 to asymmetric grids.

Another problem is to determine the weak saturation number Q_d in K_n for $d \ge 2^d$; this can be viewed as a weak saturation version of Problem 2.33. Note that the case $d = 2$ was solved in [53] but, to our knowledge, it is open for every fixed $d \ge 3$. For fixed d and n sufficiently large with respect to d, one can obtain a weakly (K_n, Q_d)-saturated graph by taking a weakly $(K_{n-2^{d-1}}, Q_d)$-saturated graph G with the minimum number of edges and adding a set S of 2^{d-1} vertices such that

- $|S| = 2^{d-1}$,
- the vertices of S induce a copy of Q_{d-1}, and
- the edges between $V(G)$ and S form a matching of size $2^{d-1} - 1$.

This construction implies that $\mathrm{wsat}(K_n, Q_d) \leq \left(\frac{(d+1)2^{d-2}-1}{2^{d-1}}\right) n + O(1)$. We conjecture that this is tight.

Conjecture 4.9. *For fixed $d \geq 3$ and $n \to \infty$,*

$$\mathrm{wsat}(K_n, Q_d) = \left(\frac{(d+1)2^{d-2}-1}{2^{d-1}}\right) n + O(1).$$

Chapter 5

Bootstrap Percolation in the Hypercube and the Grid

Bootstrap percolation is a simple, but important, example of a cellular automaton. Roughly speaking, *cellular automata* are dynamic processes on graphs in which each vertex of G is initially assigned to one of a finite number of 'states' and, after each time step, the states of the vertices are simultaneously updated based on *homogeneous* and *local* update rules; that is, the update rules are the same for every vertex and depend only on the states of vertices at a bounded distance from it. The notion of a cellular automaton was formally introduced by von Neumann [174] but grew out of his discussions with Ulam during the 1940's and 50's (see [218] or [194]). Cellular automata serve as mathematical models for a wide range of natural and sociological processes including, for example, traffic flow [163], the growth of brain tumors [137] and the spread of opinions through a social network [12].

The r-neighbour bootstrap process was introduced by Chalupa, Leath and Reich [62] as a mathematical simplification of existing dynamic models of ferromagnetism, but has also found applications in the study other physical phenomena such as crack formation and hydrogen mixtures (see [1]). Advancements in bootstrap percolation have been highly influential in the analysis of more complex cellular automata; for examples of this, we refer the reader to [13, 14, 47, 61, 68, 94, 168].

The central problem in bootstrap percolation is to estimate the critical density at which a randomly generated set becomes likely to percolate. To be more precise, for $p \in [0, 1]$, suppose that A_0^p is a subset of $V(G)$ obtained by including each vertex randomly with probability p independently of all other vertices. The *critical probability* for the r-neighbour bootstrap process on G is defined to be

$$p_c(G,r) := \inf\{p : \mathbb{P}(A_0^p \text{ percolates}) \geq 1/2\}.$$

While bootstrap percolation has now been studied for many different graph classes and random graph distributions (see, e.g., [6, 7, 28, 30, 34, 44, 57, 115, 129, 132, 203]), most classical problems focus on induced subgraphs of the integer lattice $\mathbb{Z}^d$. Two elements (vertices) of $\mathbb{Z}^d$ are adjacent if they are at Euclidean distance one. One of the first rigorous results in the study of the bootstrap percolation model is due to Schonmann [207] who proved that, for all d and r, the critical probability for the r-neighbour bootstrap process in $\mathbb{Z}^d$ attains one of the 'trivial' values. Specifically,

$$p_c\left(\mathbb{Z}^d, r\right) = \begin{cases} 0 & \text{if } r \leq d, \\ 1 & \text{if } r \geq d+1. \end{cases}$$

An earlier result of van Enter [82] settled the case $d = 2$ and $r = 3$.

Aizenmann and Lebowitz [3] were the first to consider the behaviour of the bootstrap percolation model on finite sublattices of $\mathbb{Z}^d$. Their main result is that

$$p_c\left([n]^d, 2\right) = \Theta\left(\frac{1}{\log(n)^{d-1}}\right)$$

for fixed d and $n \to \infty$ where, as usual, $[n] := \{1, \ldots, n\}$. An analogous result was proved for $r = d = 3$ by Cerf and Cirillo [59] and then generalised to all fixed $2 \leq r \leq d$ by Cerf and Manzo [60].

Shortly after this, Holroyd [123] obtained a remarkably sharp estimate in the case $d = r = 2$; namely,

$$p_c\left([n]^2, 2\right) = \frac{(1+o(1))\pi^2}{18\log(n)},$$

where log denotes the natural (base e) logarithm. Moreover, he showed that the probability that $A_0^p \subseteq [n]^2$ percolates with respect to the 2-neighbour bootstrap process changes abruptly from $o(1)$ to $1 - o(1)$ when p is in a small window around $p_c\left([n]^2, 2\right)$; this phenomenon is known as a *sharp threshold*. More precise asymptotics for $p_c\left([n]^2, 2\right)$ were later obtained by Gravner, Holroyd and Morris [111]. In [17], Balogh, Bollobás and Morris obtained a sharp threshold for the case $r = d = 3$ and laid some of the groundwork for further generalisations. This rewarding line of research culminated in a paper of Balogh, Bollobás, Duminil-Copin and Morris [16] in which a sharp threshold was determined for all fixed values of d and $2 \leq r \leq d$.

Comparably, far less is known about the critical probability when d tends to infinity. A first step in this direction was made by Balogh and Bollobás [15], who proved that

$$p_c(Q_d, 2) = \Theta\left(d^{-2}2^{-2\sqrt{d}}\right)$$

as $d \to \infty$. This result was substantially improved by Balogh, Bollobás and Morris [19], who obtained a sharp threshold for the 2-neighbour bootstrap process in $[n]^d$ whenever $d \gg \log(n) \geq 1$. In addition, Balogh, Bollobás and Morris [18] proved sharp asymptotic bounds on $p_c(Q_d, \lceil d/2 \rceil)$ and $p_c([n]^d, d)$ for n large and $d \geq (\log\log(n))^2 \log\log\log(n)$. Beyond this, very little has been established for high dimensions. For example, for fixed $r \geq 3$, the logarithm of $p_c(Q_d, r)$ is not even known to within a constant factor; that is, the following conjecture of [19] is open.

Conjecture 5.1 (Balogh, Bollobás and Morris [19])**.** *For fixed $r \geq 3$,*

$$p_c(Q_d, r) = \exp\left(-\Theta\left(d^{1/2^{r-1}}\right)\right).$$

When bounding the probability that a randomly generated set percolates, it is often useful to understand the structure and behaviour of percolating sets in $[n]^d$ with certain extremal characteristics. Of particular importance is the minimum cardinality of a percolating set, i.e., $m(G, r)$. A rewarding 'coffee time problem' [43, Problem 34] is to prove that $m([n]^2, 2) = n$ for all $n \geq 1$. This exercise turns out to have important extensions. In particular, the proof of Balogh, Bollobás and Morris [19] of a tight bound on $p_c([n]^d, 2)$ for large d relies on the following generalisation to rectangular grids proved in [15]:

$$m\left(\prod_{i=1}^{d}[a_i], 2\right) = \left\lceil \frac{\sum_{i=1}^{d}(a_i - 1)}{2} \right\rceil + 1. \tag{5.2}$$

As was mentioned in [21], a 'stumbling block' for obtaining good estimates for the critical probability when $d \to \infty$ is has been a lack of a tight asymptotic lower bound on $m(Q_d, r)$ for general r. Our goal in this chapter is to prove Theorem 1.10, which provides such a bound.

We mentioned in Section 1.4 that an upper bound of $m(Q_d, r) \leq \frac{1+o(1)}{r}\binom{d}{r-1}$ is not difficult to prove. Let us verify this now. In what follows, *level* k of Q_d is defined to be the set of all vertices whose coordinates sum to k.

Construction 5.3 (Balogh, Bollobás and Morris [19]). Let d be large with respect to r and let A_0 consist of all vertices on level $r-2$ of Q_d and a set S of $\frac{1+o(1)}{r}\binom{d}{r-1}$ vertices on level r such that every vertex on level $r-1$ has at least one neighbour in S. Note that the existence of such a set S, known as an 'approximate Steiner system,' is guaranteed by a seminal result of Rödl [192]. It is easily observed that $|A_0| = \frac{1+o(1)}{r}\binom{d}{r-1}$ and that A_0 percolates.

Remark 5.4. By an important result of Keevash [140], for certain values of d and r, the approximate Steiner system in Construction 5.3 can be replaced by an exact Steiner system. For these values of r and d, we obtain

$$m(Q_d, r) \leq \frac{1}{r}\binom{d}{r-1} + \binom{d}{r-2} = \frac{d^{r-1}}{r!} + \frac{d^{r-2}(r+2)}{2r(r-2)!} + O\left(d^{r-3}\right). \tag{5.5}$$

On the other hand, Theorem 1.10 implies a lower bound of

$$m(Q_d, r) \geq \frac{d^{r-1}}{r!} + \frac{d^{r-2}(6-r)}{2r(r-2)!} + O\left(d^{r-3}\right) \tag{5.6}$$

which differs from (5.5) by an additive term of order $\Theta\left(d^{r-2}\right)$.

As was mentioned in Section 1.4, the key step of the proof is to establish a relationship between bootstrap percolation and weak saturation; specifically, (1.11). We provide a proof of this in the next section. Once this is proved, we will focus our attention on proving lower bounds on the weak saturation number of stars in grids. We focus on the hypercube in Section 5.2, proving Theorem 1.12. Our most general result provides an exact expression for the weak saturation number of $K_{1,r+1}$ in $\prod_{i=1}^d [a_i]$. This result is stated and proved in Section 5.3. In addition, we provide a recursive upper bound on $m(Q_d, r)$ for general $d \geq r$ in Section 5.4 which improves on the bound in (5.5). In the case $r = 3$, we combine this general bound with some ad hoc constructions in small dimensions to prove Theorem 1.13.

5.1 The Key Inequality

We open this section by proving the following lemma, which implies (1.11) and improves it slightly for graphs containing vertices of degree less than r (including, for example, the graph $\prod_{i=1}^d [a_i]$ when $d < r$).

Lemma 5.7. *Let G be a graph and let F be a spanning subgraph of G such that the set*

$$A_0 := \{v \in V(G) : d_F(v) \geq \min\{r, d_G(v)\}\}$$

percolates with respect to the r-neighbour bootstrap process on G. Then F is weakly $(G, K_{1,r+1})$-saturated.

Proof. By hypothesis, we can label the vertices of G by $v_1, \ldots, v_n$ in such a way that

- $\{v_1, \ldots, v_{|A_0|}\} = A_0$, and
- for $|A_0| + 1 \leq i \leq n$, the vertex v_i has at least r neighbours in $\{v_1, \ldots, v_{i-1}\}$.

Let us show that F is weakly $(G, K_{1,r+1})$-saturated. We begin by adding to F every edge of $E(G) \setminus E(F)$ which is incident to a vertex in A_0 (one edge at a time in an arbitrary order). For every vertex $v \in A_0$, we have that either

- there are at least r edges of F incident to v, or
- every edge of G incident with v is already present in F.

Therefore, every edge of $E(G) \setminus E(F)$ incident to a vertex in A_0 completes a copy of $K_{1,r+1}$ when it is added.

Now, for each $i = |A_0| + 1, \ldots, n$ in turn, we add every edge incident to v_i which has not already been added (one edge at a time in an arbitrary order). Since v_i has at least r neighbours in $\{v_1, \ldots, v_{i-1}\}$ and every edge incident to a vertex in $\{v_1, \ldots, v_{i-1}\}$ is already present, we get that every such edge completes a copy of $K_{1,r+1}$ when it is added. The result follows. □

For completeness, we will now deduce (1.11) from Lemma 5.7.

Proof of (1.11). Let A_0 be a set of cardinality $m(G, r)$ which percolates with respect to the r-neighbour bootstrap process on G and let F be a spanning subgraph of G such that $d_F(v) \geq \min\{d_G(v), r\}$ for each $v \in A_0$. Note that this can be achieved by adding at most r edges per vertex of A_0 and so we can assume that $|E(F)| \leq r|A_0| = r \cdot m(G, r)$. By Lemma 5.7, F is weakly $(G, K_{1,r+1})$-saturated and so

$$\mathrm{wsat}(G, K_{1,r+1}) \leq |E(F)| \leq r \cdot m(G, r)$$

as required. □

5.2 The Hypercube Case

Our goal in this section is to prove Theorem 1.12. First, we require some definitions.

Definition 5.8. Given $k \geq 1$, an index $i \in [k]$ and $x \in \mathbb{R}^k$, let x_i denote the ith coordinate of x. The *support* of x is defined by $\operatorname{supp}(x) := \{i \in [k] : x_i \neq 0\}$.

Definition 5.9. The *direction* of an edge $e = uv \in E(Q_d)$ is the unique index $i \in [d]$ such that $u_i \neq v_i$. Given a vertex $v \in V(Q_d)$, we define $e(v, i)$ to be the unique edge in direction i that is incident to v.

Note that each edge of Q_d receives two labels (one for each of its endpoints). Our approach will make use of the following simple linear algebraic fact.

Lemma 5.10. *Let $k \geq \ell \geq 0$ be fixed. Then there exists a subspace X of $\mathbb{R}^k$ of dimension $k - \ell$ such that $|\operatorname{supp}(x)| \geq \ell + 1$ for every $x \in X \setminus \{0\}$.*

Proof. Define X to be the span of a set $\{v_1, \ldots, v_{k-\ell}\}$ of unit vectors of $\mathbb{R}^k$ chosen independently and uniformly at random with respect to the standard measure on the unit sphere S^{k-1}. Given a fixed subspace W of $\mathbb{R}^k$ of dimension at most ℓ and $1 \leq i \leq k - \ell$, the space

$$\operatorname{span}(W \cup \{v_1, \ldots, v_{i-1}\})$$

has dimension less than k. Thus, the unit sphere of this space has measure zero in S^{k-1} and so, with probability one, $v_i \notin \operatorname{span}(W \cup \{v_1, \ldots, v_{i-1}\})$. It follows that $\dim(X) = k - \ell$ and $X \cap W = \{0\}$ almost surely. In particular, if we let $T \subseteq [k]$ be a fixed set of cardinality ℓ and define

$$W_T := \left\{x \in \mathbb{R}^k : \operatorname{supp}(x) \subseteq T\right\},$$

then $X \cap W_T = \{0\}$ almost surely. Since there are only finitely many sets $T \subseteq [k]$ of cardinality ℓ, we can assume that X is chosen so that $X \cap W_T = \{0\}$ for every such set. This completes the proof. $\square$

In the appendix, we provide an explicit (ie. non-probabilistic) example of a vector space X satisfying Lemma 5.10. The following lemma highlights an important property of the space X.

Lemma 5.11. *Let $k \geq \ell \geq 0$ and let X be a subspace of $\mathbb{R}^k$ of dimension $k - \ell$ such that $|\operatorname{supp}(x)| \geq \ell + 1$ for every $x \in X \backslash \{0\}$. For every set $T \subseteq [k]$ of cardinality $\ell + 1$, there exists $x \in X$ with $\operatorname{supp}(x) = T$.*

Proof. Let $T \subseteq [k]$ with $|T| = \ell + 1$. Clearly, the space $\{x \in \mathbb{R}^k : \operatorname{supp}(x) \subseteq T\}$ has dimension $\ell + 1$. Therefore, since $\dim(X) = k - \ell$, there must be a non-zero vector $x \in X$ with $\operatorname{supp}(x) \subseteq T$. However, this inclusion must be equality since $|\operatorname{supp}(x)| \geq \ell + 1$. $\square$

We are now in position to prove Theorem 1.12. For notational convenience, we write

$$w := r2^{r-1} + \sum_{j=1}^{r-1} \binom{d-j-1}{r-j} j2^{j-1}.$$

Also, using Lemma 5.10, let X be a subspace of $\mathbb{R}^d$ of dimension $d - r$ such that $|\operatorname{supp}(x)| \geq r+1$ for every $x \in X \setminus \{0\}$. We deduce Theorem 1.12 from the following lemma, after which we will prove the lemma itself.

Lemma 5.12. *There is a spanning subgraph F of Q_d and a set $\{f_e : e \in E(Q_d)\}$ of vectors of $\mathbb{R}^w$ such that*

(Q1) F is weakly $(Q_d, K_{1,r+1})$-saturated and $|E(F)| = w$,

(Q2) $\sum_{i=1}^d x_i f_{e(v,i)} = 0$ for every $v \in V(Q_d)$ and $x \in X$, and

(Q3) $\operatorname{span}\{f_e : e \in E(Q_d)\} = \mathbb{R}^w$.

Proof of Theorem 1.12. Clearly, the existence of a graph F satisfying (Q1) implies the upper bound $\operatorname{wsat}(Q_d, K_{1,r+1}) \leq w$. We show that the lower bound follows from (Q2), (Q3) and Lemma 4.4. Note that the edge sets of copies of $K_{1,r+1}$ in Q_d are precisely the sets of the form $\{e(v,i) : i \in T\}$ where v is a fixed vertex of Q_d and T is a subset of $[d]$ of cardinality $r+1$. By Lemma 5.11 we know that there exists some $x \in X$ with $\operatorname{supp}(x) = T$. By (Q2) we have

$$\sum_{i=1}^d x_i f_{e(v,i)} = \sum_{i \in T} x_i f_{e(v,i)} = 0.$$

Therefore, by Lemma 4.4,

$$\operatorname{wsat}(Q_d, K_{1,r+1}) \geq \dim(\operatorname{span}\{f_e : e \in E(Q_d)\})$$

which equals w by (Q3). The result follows. $\square$

Proof of Lemma 5.12. We proceed by induction on d. We begin by settling some easy boundary cases before explaining the inductive step.

Case 1: $r = 0$.

In this case, $K_{1,r+1} \simeq K_2$. Also, $w = 0$ and $X = \mathbb{R}^d$. We let F be a spanning subgraph of Q_d with no edges and set $f_e := 0$ for every $e \in Q_d$. It is trivial to check that (Q1), (Q2) and (Q3) are satisfied.

Case 2: $d = r \geq 1$.

In this case, $w = d2^{d-1} = |E(Q_d)|$ and $X = \{0\}$. We define $F := Q_d$ and let $\{f_e : e \in E(Q_d)\}$ be a basis for $\mathbb{R}^w$. Clearly (Q1), (Q2) and (Q3) are satisfied.

Case 3: $d > r \geq 1$.

We begin by constructing F in such a way that (Q1) is satisfied. For $i \in \{0, 1\}$, let Q_{d-1}^i denote the subgraph of Q_d induced by $\{0,1\}^{d-1} \times \{i\}$. Note that both Q_{d-1}^0 and Q_{d-1}^1 are isomorphic to Q_{d-1}. Let F be a spanning subgraph of Q_d such that

- the subgraph F_0 of F induced by $V\left(Q_{d-1}^0\right)$ is a weakly $(Q_{d-1}, K_{1,r+1})$-saturated graph of minimum size,
- the subgraph F_1 of F induced by $V\left(Q_{d-1}^1\right)$ is a weakly $(Q_{d-1}, K_{1,r})$-saturated graph of minimum size, and
- F contains no edge in direction d.

Figure 5.1 contains a specific instance of this construction. Now, let us define $w_0 := \mathrm{wsat}(Q_{d-1}, K_{1,r+1})$ and $w_1 := \mathrm{wsat}(Q_{d-1}, K_{1,r})$. By construction, we have $|E(F)| = w_0 + w_1$ which is equal to w by the inductive hypothesis. Let us verify that F is weakly $(Q_d, K_{1,r+1})$-saturated. To see this, we add the edges of $E(Q_d) \setminus E(F)$ to F in three stages. By construction, we can begin by adding all edges of Q_{d-1}^0 which are not present in F_0 in such a way that each edge completes a copy of $K_{1,r+1}$ in Q_{d-1}^0 when it is added. In the second stage, we add all edges of Q_d in direction d one by one in any order. Since every vertex of Q_d has degree $d \geq r + 1$ and every edge of Q_{d-1}^0 has already been added, we get that every edge added in this stage completes a copy of $K_{1,r+1}$ in Q_d. Finally, we add the edges of Q_{d-1}^1 which are not present in

F_1 in such a way that each added edge completes a copy of $K_{1,r}$ in Q_{d-1}^1. Since the edges in direction d have already been added, we see that every such edge completes a copy of $K_{1,r+1}$ in Q_d. Therefore, (Q1) holds.

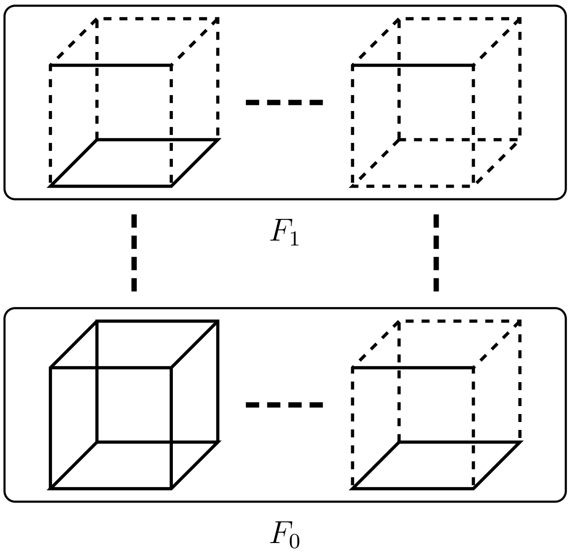

Figure 5.1: A weakly $(Q_5, K_{1,4})$-saturated graph F constructed inductively from a weakly $(Q_4, K_{1,4})$-saturated graph F_0 and a weakly $(Q_4, K_{1,3})$-saturated graph F_1, each of which is also constructed inductively.

Thus, all that remains is to construct $\{f_e : e \in E(Q_d)\}$ in such a way that (Q2) and (Q3) are satisfied. Let $\pi : X \to \mathbb{R}^{d-1}$ be the standard projection defined by $\pi : (x_1, \ldots, x_d) \mapsto (x_1, \ldots, x_{d-1})$. Let $z \in X$ be an arbitrary vector such that $d \in \operatorname{supp}(z)$ (such a vector exists by Lemma 5.11) and let $T_z : X \to X$ be the linear map defined by

$$T_z(x) := x - \frac{x_d}{z_d} z$$

for $x \in X$. Define

$$X_0 := \pi(T_z(X)) \text{ and}$$

$$X_1 := \pi(X).$$

Clearly, $\ker(T_z) = \operatorname{span}\{z\}$ and, since every $x \in X \setminus \{0\}$ has $|\operatorname{supp}(x)| \geq r+1 \geq 2$, we have $\ker(\pi) = \{0\}$. This implies that X_0 has dimension $d-r-1$ and that X_1 has dimension $d-r$. Also, by construction, we have that $|\operatorname{supp}(x)| \geq r+1$ for every non-zero $x \in X_0$ and $|\operatorname{supp}(x)| \geq r$ for every non-zero $x \in X_1$.

Therefore, by the inductive hypothesis, there exists $\{f_e^0 : e \in E(Q_{d-1}^0)\}$ in $\mathbb{R}^{w_0}$ and $\{f_e^1 : e \in E(Q_{d-1}^1)\}$ in $\mathbb{R}^{w_1}$ such that

(Q2.0) $\sum_{i=1}^{d-1} x_i f^0_{e(v,i)} = 0$ for every $v \in V(Q^0_{d-1})$ and $x \in X_0$,

(Q2.1) $\sum_{i=1}^{d-1} x_i f^1_{e(v,i)} = 0$ for every $v \in V(Q^1_{d-1})$ and $x \in X_1$,

(Q3.0) span $\{f_e^0 : e \in E(Q_{d-1}^0)\} = \mathbb{R}^{w_0}$, and

(Q3.1) span $\{f_e^1 : e \in E(Q_{d-1}^1)\} = \mathbb{R}^{w_1}$.

We will define the vectors $\{f_e : e \in E(Q_d)\} \subseteq \mathbb{R}^{w_0} \oplus \mathbb{R}^{w_1} \simeq \mathbb{R}^w$ satisfying (Q2) and (Q3) in three stages. First, if $e \in E(Q^0_{d-1})$, then we set

$$f_e := f_e^0 \oplus 0.$$

Next, for each edge of the form $e = e(v,d)$ for $v \in V(Q^0_{d-1})$, we let

$$f_e := -\frac{1}{z_d}\sum_{i=1}^{d-1} z_i f_{e(v,i)} \tag{5.13}$$

(recall the definition of z above). Finally, if $e = uv \in E(Q^1_{d-1})$, then we let $e' = u'v'$ where u' and v' are the unique neighbours of u and v in $V(Q^0_{d-1})$ and define

$$f_e := f^0_{e'} \oplus f^1_e.$$

It is easily observed that $\dim(\operatorname{span}\{f_e : e \in E(Q_d)\}) = w_0 + w_1 = w$ by (Q3.0), (Q3.1) and the construction of f_e given above. Therefore, (Q3) holds.

Finally, we prove that (Q2) is satisfied. First, let $v \in V(Q^0_{d-1})$ and let $x \in X$ be arbitrary. Define $x^\dagger := T_z(x)$ and note that $d \notin \operatorname{supp}(x^\dagger)$. We have

$$\sum_{i=1}^{d} x_i f_{e(v,i)} = \sum_{i=1}^{d-1} x_i^\dagger f_{e(v,i)} + \frac{x_d}{z_d}\sum_{i=1}^{d} z_i f_{e(v,i)}$$

by definition of T_z. Both of the sums on the right side are zero by (Q2.0) and (5.13). Now, suppose that $v \in V(Q^1_{d-1})$ and let v' be the unique neighbour of v in $V(Q^0_{d-1})$. Given $x \in X$, we have

$$\sum_{i=1}^{d} x_i f_{e(v,i)} = \sum_{i=1}^{d} x_i f_{e(v',i)} + \sum_{i=1}^{d-1} x_i \left(0 \oplus f^1_{e(v,i)}\right)$$

which is zero by (Q2.1) and the fact that $\sum_{i=1}^{d} x_i f_{e(v',i)} = 0$, which was proved above (as $v' \in V(Q^0_{d-1})$). Therefore, (Q2) holds. This completes the proof of the lemma. □

5.3 General Grids

Our objective in this section is to determine the weak saturation number of $K_{1,r+1}$ in $\prod_{i=1}^{d}[a_i]$ in full generality. We express this weak saturation number in terms of the following recurrence relation.

Definition 5.14. Let d and r be integers such that $0 \le r \le 2d$ and let $a_1, \ldots, a_d \ge 2$. Define $w_r(a_1, \ldots, a_d)$ to be

- 0, if $r = 0$;
- $\displaystyle\sum_{j=1}^{d}(a_j - 1)\prod_{i \neq j} a_i$, if $r = 2d$;
- $d2^{d-1}$, if $a_1 = \cdots = a_d = 2$ and $d + 1 \le r \le 2d - 1$;
- $\displaystyle r2^{r-1} + \sum_{j=1}^{r-2}\binom{d-j-1}{r-j} j2^{j-1}$, if $a_1 = \cdots = a_d = 2$ and $1 \le r \le d$; and
- $\displaystyle w_r(a_1, \ldots, a_{i-1}, a_i - 1, a_{i+1}, \ldots a_d) + w_{r-1}(a_1, \ldots, a_{i-1}, a_{i+1}, \ldots a_d)$
$\displaystyle + \sum_{\substack{S \subseteq [d]\setminus\{i\} \\ |S| \ge 2d-r}} 2^{|S|}\prod_{j \notin S}(a_j - 2)$, if $1 \le r \le 2d - 1$ and $a_i \ge 3$.

We prove the following.

Theorem 5.15. *For $0 \le r \le 2d$ and $a_1, \ldots, a_d \ge 2$, we have*

$$\mathrm{wsat}\left(\prod_{i=1}^{d}[a_i], K_{1,r+1}\right) = w_r(a_1, \ldots, a_d).$$

Let $a_1, \ldots, a_d \ge 2$ and define $G := \prod_{i=1}^{d}[a_i]$. In proving of Theorem 5.15, we employ an inductive approach similar to the one used in the proof of Theorem 1.12. The main difference is that a vertex v of G may be incident to either one or two edges in direction $i \in [d]$ depending on whether or not $v_i \in \{1, a_i\}$. With this in mind, we define a labelling of the edges of G.

Definition 5.16. Say that an edge $e = uv \in E(G)$ in direction $i \in [d]$ is *odd* if $\min\{u_i, v_i\}$ is odd and *even* otherwise. We label e by $e(v, 2i-1)$ if e is odd and $e(v, 2i)$ if e is even.

Note that each edge of G receives two labels, one for each of its endpoints.

Definition 5.17. For $v \in V(G)$, define $I_v^G := \{j \in [2d] : e(v,j) \in E(G)\}$.

We are now in position to prove Theorem 5.15. Using Lemma 5.10, we let X be a subspace of $\mathbb{R}^{2d}$ of dimension $2d-r$ such that $|\operatorname{supp}(x)| \geq r+1$ for every $x \in X \setminus \{0\}$. Define $w := w_r(a_1, \ldots, a_d)$. As with the proof of Theorem 1.12, we state a lemma from which we deduce Theorem 5.15, and then we prove the lemma.

Lemma 5.18. *There is a spanning subgraph F of G and a set $\{f_e : e \in E(G)\}$ of vectors of $\mathbb{R}^w$ such that*

(G1) F is weakly $(G, K_{1,r+1})$-saturated and $|E(F)| = w$,

(G2) $\sum_{i=1}^{2d} x_i f_{e(v,i)} = 0$ for every $v \in V(G)$ and $x \in X$ such that $\operatorname{supp}(x) \subseteq I_v^G$, and

(G3) $\operatorname{span}\{f_e : e \in E(G)\} = \mathbb{R}^w$.

Proof of Theorem 5.15. First observe that the existence of a graph F satisfying (G1) implies $\operatorname{wsat}(G, K_{1,r+1}) \leq w$. To obtain a matching lower bound, we apply Lemma 4.4 as we did in the hypercube case. The edge sets of copies of $K_{1,r+1}$ in G are the sets of the form $\{e(v,i) : i \in T\}$, where $v \in V(G)$ and T is a subset of I_v^G of cardinality $r+1$. By applying Lemma 5.11 together with (G2), we see that the conditions of Lemma 4.4 are satisfied. Thus by (G3), $\operatorname{wsat}(G, K_{1,r+1}) \geq w$. □

Proof of Lemma 5.18. We proceed by induction on $|V(G)|$. We begin with the boundary cases.

Case 1: $r = 0$.

In this case, $K_{1,r+1} \simeq K_2$. Also, $w = 0$ and $X = \mathbb{R}^{2d}$. We let F be a spanning subgraph of G with no edges and set $f_e := 0$ for every $e \in Q_d$. Properties (G1), (G2) and (G3) are satisfied trivially.

Case 2: $r = 2d \geq 2$.

In this case, $w = |E(G)|$ and $X = \{0\}$. We define $F := G$ and let $\{f_e : e \in E(G)\}$ be a basis for $\mathbb{R}^w$. Clearly (G1), (G2) and (G3) are satisfied.

Case 3: $a_1 = \ldots = a_d = 2$ and $1 \leq r \leq 2d-1$.

In this case, $G = [2]^d$ is isomorphic to Q_d and every edge of G is odd. First, suppose that $d+1 \leq r \leq 2d-1$. Then we have $w = |E(G)|$ and we define $F := G$ and let $\{f_e : e \in E(G)\}$ be a basis for $\mathbb{R}^w$.

On the other hand, if $1 \leq r \leq d$, then we let X' be the subspace of X consisting of all vectors x of X such that every element of $\mathrm{supp}(x)$ is odd. It is not hard to show that X' has dimension $d-r$ and that every vector $x \in X'$ has $|\mathrm{supp}(x)| \geq r+1$. Thus, we are done by Lemma 5.12.

Case 4: $a_i \geq 3$ for some $i \in [d]$ and $1 \leq r \leq 2d-1$.

Without loss of generality, assume that $a_d \geq 3$. Define

$$G_1 := \prod_{i=1}^{d-1}[a_i] \times [a_d - 1], \text{ and}$$

$$G_2 := G \setminus G_1.$$

Observe that every vertex of G_2 has a unique neighbour in $V(G_1)$. The edges with one endpoint in G_1 and the other in G_2 will play a particular role in the proof. We define

$$\tau := \begin{cases} 2d-1 & a_d - 1 \text{ is odd} \\ 2d & a_d - 1 \text{ is even,} \end{cases}$$

and we write $\bar{\tau}$ for the unique element of $\{2d-1, 2d\} \setminus \{\tau\}$. Observe that for $v \in V(G_2)$, we have that $\bar{\tau} \notin I_v^G$, and that $I_v^{G_2} = I_v^G \setminus \{\tau\}$. On the other hand, if $v \in V(G_1)$, then

$$I_v^{G_1} = \begin{cases} I_v^G \setminus \{\tau\} & \text{if } v_d = a_d - 1, \\ I_v^G & \text{otherwise.} \end{cases}$$

Define

$$Y := \{v \in V(G_1) : v_d = a_d - 1 \text{ and } d_{G_1}(v) < r\}.$$

It is not hard to see that

$$|Y| = \sum_{\substack{S \subseteq [d-1] \\ |S| \geq 2d-r}} 2^{|S|} \prod_{j \notin S} (a_j - 2).$$

For brevity we write $y := |Y|$ and

$$w_1 := \mathrm{wsat}(G_1, K_{1,r+1}),$$

$$w_2 := \mathrm{wsat}(G_2, K_{1,r}).$$

We construct a graph F satisfying (G1). Define F to be a spanning subgraph of G such that

- the subgraph F_1 of F induced by $V(G_1)$ is a weakly $(G_1, K_{1,r+1})$-saturated graph of minimum size,

- the subgraph F_2 of F induced by $V(G_2)$ is a weakly $(G_2, K_{1,r})$-saturated graph of minimum size, and

- an edge e from $V(G_1)$ to $V(G_2)$ is contained in F if and only if e is of the form $e(v,\tau)$ for $v \in Y$.

Applying the inductive hypothesis and Definition 5.14, we see that $|E(F)| = w_1 + w_2 + y = w$, as required. To see that F is weakly $(G, K_{1,r+1})$-saturated, we add the edges of $E(G)\backslash E(F)$ to F in three stages. First, by definition of F_1, we can add the edges that are not present in $E(F_1)$ in such a way that every added edge completes a copy of $K_{1,r+1}$ in G_1. Next, we can add the edges of the form $e(v,\tau)$, where $v \notin Y$ and $v_d = a_d - 1$, in any order. By definition of Y, we see that every such v has at least r neighbours in G_1. As every edge in $E(G_1)$ has already been added, the addition of $e(v,\tau)$ completes a copy of $K_{1,r+1}$ in G. Finally, we add the edges of G_2 that are not present in F_2 in such a way that each added edge completes a copy of $K_{1,r}$ in G_2. Every such edge completes a copy of $K_{1,r+1}$ in G since every vertex in G_2 has a neighbour in G_1 and every edge between G_1 and G_2 is already present. Thus, (G1) holds.

It remains to find a collection $\{f_e : e \in E(G)\}$ satisfying (G2) and (G3). Let $\pi : X \to \mathbb{R}^{2d-2}$ be the projection defined by $\pi : (x_1, \dots, x_{2d}) \mapsto (x_1, \dots, x_{2d-2})$. Let z be a fixed vector of X such that $\bar{\tau} \in \operatorname{supp}(z)$ and define $T_z : X \to X$ by

$$T_z(x) := x - \frac{x_{\bar{\tau}}}{z_{\bar{\tau}}}.$$

Define $X_1 := X$ and $X_2 := \pi\left(T_z(X)\right)$. Since $\ker\left(T_z\right) = \operatorname{span}\{z\}$ and $\ker(\pi) = \{0\}$ we see that X_2 has dimension $2d - r - 1 = 2(d-1) - (r-1)$. Also, by construction, we have $|\operatorname{supp}(x)| \geq r$ for every non-zero $x \in X_2$. By applying the inductive hypothesis to both G_1 and G_2, we can find collections $\{f_e^1 : e \in E(G_1)\}$ in $\mathbb{R}^{w_1}$ and $\{f_e^2 : e \in E(G_2)\}$ in $\mathbb{R}^{w_2}$ such that

(G2.1) $\sum_{i=1}^{2d} x_i f^1_{e(v,i)} = 0$ for every $v \in V(G_1)$ and $x \in X_1$ with $\operatorname{supp}(x) \subseteq I_v^{G_1}$,

(G2.2) $\sum_{i=1}^{2d-2} x_i f^2_{e(v,i)} = 0$ for every $v \in V(G_2)$ and $x \in X_2$ with $\operatorname{supp}(x) \subseteq I_v^{G_2}$, (G3.1) $\operatorname{span}\{f_e^1 : e \in E(G_1)\} = \mathbb{R}^{w_1}$, and

(G3.2) $\operatorname{span}\{f_e^2 : e \in E(G_2)\} = \mathbb{R}^{w_2}$.

Using this, we will now construct a collection $\{f_e : e \in E(G)\} \subseteq \mathbb{R}^{w_1} \oplus \mathbb{R}^{w_2} \oplus \mathbb{R}^{y} \simeq \mathbb{R}^{w}$ in four steps. First, for $e \in E(G_1)$, we define

$$f_e := f_e^1 \oplus 0 \oplus 0.$$

Let $\{f_y^3 : y \in Y\}$ be a basis of $\mathbb{R}^y$. Next, we consider edges $e = uv$, where $v \in V(G_1)$, and $u \in V(G_2)$. If v is in Y, then we let

$$f_e := 0 \oplus 0 \oplus f_v^3.$$

If v is not in Y, then let $z^v \in X$ be a vector such that $\operatorname{supp}(z^v) \subseteq I_v^G$ and $\tau \in \operatorname{supp}(z^v)$, which exists by Lemma 5.10. Define

$$f_e := -\frac{1}{z_\tau^v} \sum_{i \in [2d] \setminus \{\tau\}} z_i^v f_{e(v,i)}. \tag{5.19}$$

Finally if $e = uv \in E(G_2)$, then let $e' = u'v'$ where $u'v'$ are the unique neighbours of u and v in $V(G_1)$ and define

$$f_e := f_{e'}^1 \oplus f_e^2 \oplus 0.$$

It is clear from (G3.1), (G3.2) and the construction of f_e, that the dimension of $\operatorname{span}\{f_e : e \in E(G)\}$ is $w_1 + w_2 + y = w$. Thus (G3) is satisfied.

It remains to show that (G2) holds. Firstly, suppose $v \in V(G_1)$ and let $x \in X$ be such that $\operatorname{supp}(x) \subseteq I_v^G$. If $v_d < a_d - 1$, then $\sum_{i=1}^{2d} x_i f_{e(v,i)} = 0$ by (G2.1). If $v_d = a_d - 1$ and $v \in Y$, then, by definition of Y, we have $|I_v^G| \leq r$ and so it must be the case that $x = 0$ and we are done. Now suppose that $v \notin Y$ and that $v_d = a_d - 1$. Define

$$x^\dagger := x - \frac{x_\tau}{z_\tau^v} z^v.$$

We have,

$$\sum_{i=1}^{2d} x_i f_{e(v,i)} = \sum_{i=1}^{2d} x_i^\dagger f_{e(v,i)} + \frac{x_\tau}{z_\tau^v} \sum_{i=1}^{2d} z_i^v f_{e(v,i)}. \tag{5.20}$$

Note that $\tau \notin \operatorname{supp}(x^\dagger)$ and thus $\operatorname{supp}(x^\dagger) \subseteq I_v^{G_1}$. Therefore, both of the sums on the right side of (5.20) are zero by (G2.1) and (5.19).

Finally, consider $v \in V(G_2)$. Let v' be the unique neighbour of v in $V(G_2)$. Given $x \in X$, with $\operatorname{supp}(x) \subseteq I_v^G$ we have

$$\sum_{i=1}^{2d} x_i f_{e(v,i)} = \sum_{i=1}^{2d} x_i f_{e(v',i)} + \sum_{i=1}^{2d-2} x_i \left(0 \oplus f^2_{e(v,i)} \oplus 0\right).$$

We have that $\sum_{i=1}^{2d} x_i f_{e(v',i)} = 0$ for $v' \in V(G_1)$, as proved above. The second sum on the right side is zero by (G2.2), which is applicable as $\bar{\tau} \notin I_v^G \supseteq \operatorname{supp}(x)$, and so $x \in T_z(X)$. This completes the proof of the lemma. $\square$

5.4 Upper Bound Constructions

In this section, we prove a recursive upper bound on $m(Q_d, r)$ for general $d \geq r \geq 1$ and then apply it to obtain an exact expression for $m(Q_d, 3)$.

Lemma 5.21. *For $d \geq r \geq 1$,*

$$m\left(Q_d, r\right) \leq m\left(Q_{d-r}, r\right) + (r-1) m\left(Q_{d-r}, r-1\right) + \sum_{j=1}^{\lceil r/2 \rceil - 1} \binom{r}{2j+1} m\left(Q_{d-r}, r-2j\right).$$

Proof. Let $d \geq r$ be fixed positive integers. For $1 \leq t \leq r$, let B_t be a subset of $V\left(Q_{d-r}\right)$ of cardinality $m\left(Q_{d-r}, t\right)$ which percolates with respect to the t-neighbour bootstrap process in Q_{d-r}.

Given $x \in V(Q_d)$, let $[x]_r$ and $[x]_{d-r}$ denote the vectors obtained by restricting x to its first r coordinates and last $d-r$ coordinates, respectively. We partition $\{0,1\}^r$ into $r+1$ sets $L_0, \ldots, L_r$ such that L_i consists of the vectors whose coordinate sum is equal to i. We construct a percolating set A_0 in Q_d. Given $x \in V(Q_d)$, we include x in A_0 if one of the following holds:

- $[x]_r \in L_1$ and either
  - $[x]_r = (1, 0, \ldots, 0)$ and $[x]_{d-r} \in B_r$.
  - $[x]_r \neq (1, 0, \ldots, 0)$ and $[x]_{d-r} \in B_{r-1}$.
- $[x]_r \in L_{2j+1}$ for some $1 \leq j \leq \lceil r/2 \rceil - 1$ and $[x]_{d-r} \in B_{r-2j}$.

It is clear that

$$|A_0| = m\left(Q_{d-r}, r\right) + (r-1)m\left(Q_{d-r}, r-1\right) + \sum_{j=1}^{\lceil r/2\rceil - 1} \binom{r}{2j+1} m\left(Q_{d-r}, r-2j\right)$$

by construction. We will be done if we can show that A_0 percolates with respect to the r-neighbour bootstrap process.

We begin by showing that every vertex x with $[x]_r \in L_0 \cup L_1$ is eventually infected. First, we can infect every vertex x such that $[x]_r = (1, 0, \ldots, 0)$, one by one in some order, by definition of B_r. Next, consider a vertex x such that $[x]_r \in L_0$ and $[x]_{d-r} \in B_{r-1}$. Then x has $r-1$ neighbours $z \in A_0$ such that $[z]_r \neq (1, 0, \ldots, 0)$, by construction, and one infected neighbour y such that $[y]_r = (1, 0, \ldots, 0)$. Thus, every such x becomes infected. Now, by definition of B_{r-1}, the remaining vertices x such that $[x]_r \in L_0$ can be infected since every such vertex has an infected neighbour y such that $[y]_r = (1, 0, \ldots, 0)$. Finally, each vertex x such that $x \neq (1, 0, \ldots, 0)$ and $[x]_r \in L_1$ becomes infected using the definition of B_{r-1} and the fact that every vertex y with $[y]_r \in L_0$ is already infected.

Now, suppose that, for some $1 \leq j \leq \lceil r/2\rceil - 1$ every vertex x such that $[x]_r \in L_0 \cup \cdots \cup L_{2j-1}$ is already infected. We show that every vertex x with $[x]_r \in L_{2j} \cup L_{2j+1}$ is eventually infected. First, consider a vertex x with $[x]_r \in L_{2j}$ and $[x]_{d-r} \in B_{r-2j}$. Such a vertex has $2j$ infected neighbours y such that $[y]_r \in L_{2j-1}$ and $r-2j$ neighbours z such that $[z]_r \in L_{2j+1} \cap A_0$. Therefore, every such x becomes infected. Now, by definition of B_{r-2j}, the remaining vertices x such that $[x]_r \in L_{2j}$ can be infected since every such vertex has $2j$ infected neighbours y such that $[y]_r \in L_{2j-1}$. Finally, each vertex x such that $[x]_r \in L_{2j+1}$ becomes infected using the definition of B_{r-2j} and the fact that every vertex y with $[y]_r \in L_{2j-1}$ is already infected.

Finally, if r is even, then we need to show that every vertex of L_r becomes infected. Every such vertex has precisely r neighbours in L_{r-1}. Thus, given that every vertex of L_{r-1} is infected, x becomes infected as well. This completes the proof. $\square$

We remark that the recursion in Lemma 5.21 gives a bound of the form $m\left(Q_d, r\right) \leq \frac{1+o(1)d^{r-1}}{r!}$ where the second order term is better than the one in (5.5). Next, we prove Theorem 1.13.

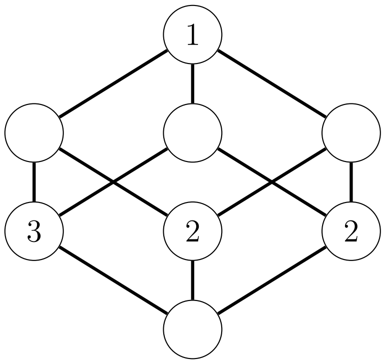

Figure 5.2: An illustration of the set A_0 constructed in the proof of Lemma 5.21 in the case $r = 3$. Each node represents a copy of Q_{d-3}. The set A_0 consists of a copy of B_i on each node labelled $i \in \{1, 2, 3\}$.

Proof of Theorem 1.13. The lower bound follows from Theorem 1.10. We prove the upper bound by induction on d. First, we settle the cases $d \in \{3, \ldots, 8\}$. For notational convenience, we associate each element v of $\{0,1\}^d$ with the subset of $[d]$ for which v is the characteristic vector. Moreover, we identify each non-empty subset of $[d]$ with the concatenation of its elements (e.g. $\{1,3,7\}$ is written 137). One can verify (by hand or by computer) that the set A_0^d, defined below, percolates with respect to the 3-neighbour bootstrap process in Q_d and that it has the cardinality $\left\lceil \frac{d(d+3)}{6} \right\rceil + 1$.

$$
\begin{aligned}
A_0^3 &:= \{1, 2, 3, 123\}, \\
A_0^4 &:= \left(A_0^3 \setminus \{3\}\right) \cup \{134, 4, 234\}, \\
A_0^5 &:= \left(A_0^4 \setminus \{134\}\right) \cup \{135, 245, 12345\}, \\
A_0^6 &:= \left(A_0^5 \setminus \{135, 245\}\right) \cup \{346, 12356, 456, 23456\}, \\
A_0^7 &:= \left(A_0^6 \setminus \{346\}\right) \cup \{13457, 24567, 12367, 1234567\}, \\
A_0^8 &:= \left(A_0^7 \setminus \{13457, 24567\}\right) \cup \{34568, 1234578, 34678, 25678, 2345678\}.
\end{aligned}
$$

Now, suppose $d \geq 9$ and that the theorem holds for smaller values of d. If d is odd, then we apply Lemma 5.21 to obtain

$$m\left(Q_d, 3\right) \leq m\left(Q_{d-3}, 3\right) + 2m\left(Q_{d-3}, 2\right) + m\left(Q_{d-3}, 1\right).$$

Clearly, $m\left(Q_{d-3}, 1\right) = 1$ and it is easy to show that $m\left(Q_{d-3}, 2\right) \leq \frac{d-3}{2} + 1$ (since $d - 3$

is even). Therefore, by the inductive hypothesis and (5.2),

$$m\left(Q_d,3\right) \le \left\lceil \frac{(d-3)d}{6} \right\rceil + 1 + 2\left(\frac{d-3}{2}+1\right) + 1 = \left\lceil \frac{d(d+3)}{6} \right\rceil + 1.$$

Now, suppose that $d \ge 10$ is even. For $t \in \{1,2,3\}$, let B_t be a subset of $V\left(Q_{d-6}\right)$ of cardinality $m\left(Q_{d-6},t\right)$ which percolates with respect to the t-neighbour bootstrap process on Q_{d-6} and let A_0^6 be as above. Given a vector $x \in V(Q_d)$, let $[x]_6$ be the restriction of x to its first six coordinates and $[x]_{d-6}$ be the restriction of x to its last $d-6$ coordinates. We define a subset A_0 of $V(Q_d)$. We include a vertex $x \in V(Q_d)$ in A_0 if $[x]_6 \in A_0^6$ and one of the following holds:

- $[x]_6 = (0,0,1,1,0,1)$ and $[x]_{d-6} \in B_3$.
- $[x]_6 \ne (0,0,1,1,0,1)$ and we have $x_5 = 1$ and $[x]_{d-6} \in B_2$.
- $x_5 = x_6 = 0$ and $[x]_{d-6} \in B_1$.

The fact that A_0 percolates follows from arguments similar to those given in the proof of Lemma 5.21; we omit the details. By construction,

$$|A_0| = m\left(Q_{d-6},3\right) + 4m\left(Q_{d-6},2\right) + 5m\left(Q_{d-6},1\right)$$

which equals

$$\left\lceil \frac{(d-6)(d-3)}{6} \right\rceil + 1 + 4\left(\frac{d-6}{2}+1\right) + 5 = \left\lceil \frac{d(d+3)}{6} \right\rceil + 1$$

by (5.2) and the inductive hypothesis. The result follows. □

5.5 Open Problems

In this chapter, we have determined the main asymptotics of $m\left(Q_d,r\right)$ for fixed r and d tending to infinity and obtained a sharper result for $r=3$. We wonder whether sharper asymptotics are possible for general r.

Question 5.22. *For fixed $r \ge 4$ and $d \to \infty$, does*

$$\frac{m\left(Q_d,r\right) - \frac{d^{r-1}}{r!}}{d^{r-2}}$$

converge? If so, what is the limit?

As Theorem 1.13 illustrates, it may be possible to obtain an exact expression for $m(Q_d, r)$ for some small fixed values of r. The first open case is the following.

Problem 5.23. *Determine $m(Q_d, 4)$ for all $d \geq 4$.*

Using a computer, we have determined that $m(Q_5, 4) = 14$, which is greater than the lower bound of 13 implied by Theorem 1.10. Thus, Theorem 1.10 is not tight for general d and r. However, we wonder whether it could be tight when r is fixed and d is sufficiently large.

Question 5.24. *For fixed $r \geq 4$, is it true that*

$$m(Q_d, r) = 2^{r-1} + \left\lceil \sum_{j=1}^{r-1} \binom{d-j-1}{r-j} \frac{j2^{j-1}}{r} \right\rceil$$

provided that d is sufficiently large?

Another direction that one could take is to determine $\mathrm{wsat}(G, K_{1,r+1})$ for other graphs G. For example, one could consider the d-dimensional torus $\mathbb{Z}_n^d$.

Problem 5.25. *Determine $\mathrm{wsat}\left(\mathbb{Z}_n^d, K_{1,r+1}\right)$ for all n, d and r.*

5.6 Appendix: An Explicit Linear Algebraic Construction

Given integers k and ℓ with $k \geq \ell \geq 0$, we construct an explicit subspace X of $\mathbb{R}^k$ of dimension $k - \ell$ such that $|\mathrm{supp}(x)| \geq \ell + 1$ for every $x \in X \setminus \{0\}$. This can be seen as an alternative proof of Lemma 5.10.

The construction is based on a so-called *Vandermonde matrix*. For $1 \leq i \leq k - \ell$ we let $\alpha_i \in \mathbb{R}^\ell$ be the vector such that, for $1 \leq j \leq \ell$,

$$\alpha_{i,j} := i^j.$$

Now, for $1 \leq i \leq k - \ell$ let e_i be the ith standard unit basis vector of $\mathbb{R}^{k-\ell}$ and define

$$v_i := \alpha_i \oplus e_i.$$

The space X is defined to be $\mathrm{span}\{v_1, \ldots, v_{k-\ell}\}$. It is clear that $\dim(X) = k - \ell$ by construction. All that remains is to show that $|\mathrm{supp}(x)| \geq \ell+1$ for every $x \in X\setminus\{0\}$. We require a few definitions.

Definition 5.26. Given a set $T \subseteq [k]$, let $\pi_T : \mathbb{R}^k \to \mathbb{R}^{|T|}$ be the standard projection $\pi_T : (x_1, \ldots, x_k) \mapsto (x_i : i \in T)$.

Definition 5.27. For $n \geq 1$, a collection $Z \subseteq \mathbb{R}^n$ is in *general position* in $\mathbb{R}^n$ if any set of at most n vectors from Z is linearly independent.

Our proof of the following proposition follows an argument of Moshonkin.[1]

Proposition 5.28. *For any set $T \subseteq [\ell]$, the vectors $\{\pi_T(\alpha_i) : 1 \leq i \leq k - \ell\}$ are in general position in $\mathbb{R}^{|T|}$.*

Proof. We assume that $|T| \geq 1$; otherwise, the result is trivial. Let $t := |T|$. Suppose that the proposition is false and let $I \subseteq [k - \ell]$ be a set of cardinality t for which there exists $\{c_i : i \in I\}$, not all of which are zero, such that

$$\sum_{i \in I} c_i \pi_T(\alpha_i) = 0.$$

Equivalently, for each $j \in T$,

$$\sum_{i \in I} c_i i^j = 0.$$

Since the determinant of a square matrix is equal to the determinant of its transpose, there must also exist scalars $\{c'_j : j \in T\}$, not all zero, such that

$$\sum_{j \in T} c'_j i^j = 0$$

for every $i \in I$.

Let $p(x)$ denote the real polynomial $\sum_{j \in T} c'_j x^j$. Then $p(x)$ is a polynomial with between one and t non-zero terms and at least t positive real roots (namely, each $i \in I$). We show, by induction on t, that no such polynomial can exist. The base case $t = 1$ is trivial. Now, let $p(x)$ be a polynomial with $t \geq 2$ non-zero terms and at least t positive real roots. Define $q(x)$ to be the polynomial of smallest degree such that $q(x) = x^s p(x)$ for some $s \geq 0$. It is clear that $q(x)$ has at least as many positive real roots as $p(x)$. However, the derivative of $q(x)$ has at most $t - 1$ positive terms and at least $t - 1$ positive real roots, contradicting the inductive hypothesis. This completes the proof. $\square$

[1] A. G. Moshonkin, *Concerning Hall's theorem*, Mathematics in St. Petersburg, Amer. Math. Soc. Transl. Ser. 2, vol. 174, Amer. Math. Soc., Providence, RI, 1996, pp. 73–77.

Now, suppose that $x \in X \setminus \{0\}$ such that $|\operatorname{supp}(x)| \leq \ell$. Define

$$U_1 := \operatorname{supp}(x) \cap [\ell],$$

$$U_2 := \operatorname{supp}(x) \setminus [\ell] \text{ and}$$

$$T := [\ell] \setminus U_1.$$

Since $x \in X \setminus \{0\}$, we can write

$$x = \sum_{i=1}^{k-\ell} c_i v_i$$

for scalars c_i which are not all zero. However, it must be the case that $c_i = 0$ for each i such that $i + \ell \notin U_2$. Thus,

$$0 = \pi_T(x) = \sum_{i:i+\ell \in U_2} c_i \pi_T(v_i) = \sum_{i:i+\ell \in U_2} c_i \pi_T(\alpha_i). \tag{5.29}$$

However, since $|U_1| + |U_2| \leq \ell$ we have $|T| = \ell - |U_1| \geq |U_2|$. Thus, (5.29) contradicts Proposition 5.28. It follows that $|\operatorname{supp}(x)| \geq \ell + 1$ for every $x \in X \setminus \{0\}$, as required.

Chapter 6

Supersaturation in Posets

A typical supersaturation problem for graphs asks, given a fixed graph H and an integer $m > \mathrm{ex}(n, H)$, what is the minimum number of copies of H in a graph with n vertices and m edges? This question can be seen as a quantitative version of Turán's original problem (question (1.1)). Perhaps the earliest supersaturation result is a strengthening of Mantel's Theorem [164] due to Rademacher (unpublished, see [85]) which says that every graph with n vertices and more than $n^2/4$ edges must contain at least $\lfloor n/2 \rfloor$ triangles. This result was generalised by Erdős [85, 87], who obtained an exact value for the minimum number of cliques of order r in a graph with n vertices and m edges where $\mathrm{ex}(n, K_r) < m \leq \mathrm{ex}(n, K_r) + C_r n$ for some constant $C_r > 0$.

Since the work of Rademacher and Erdős, there has been a great deal of interest in determining asymptotic bounds on the minimum number of cliques of order r contained in a graph with n vertices and $m = \rho\binom{n}{2}$ edges where ρ is a fixed real number in $\left(1 - \frac{1}{r-1}, 1\right)$. Motivated by some illuminating early work in [38, 109, 154, 167] (see also [92, 116, 177]), Lovász and Simonovits [155] conjectured that this minimum number is asymptotically equal to the number of cliques of order r in a complete multipartite graph which is 'nearly balanced' in the sense that all parts are of the same size except for one (which is smaller). This important conjecture was finally proved in the case $r = 3$ by Razborov [188], $r = 4$ by Nikiforov [175] and for general $r \geq 5$ by Reiher [191]. All three of these proofs rely on analytic methods.

Returning momentarily to classical results, the following general theorem on the number of copies of a fixed graph H in an n-vertex graph with a prescribed number of edges was proved by Erdős and Simonovits [92]. Some important applications of this theorem will be discussed in Chapter 7.

Theorem 6.1 (Erdős and Simonovits [92]). *Let H be a fixed graph. For every $\varepsilon > 0$ there exists $\delta = \delta(H, \varepsilon) > 0$ such that if G is an n-vertex graph with at least $\mathrm{ex}(n, H) + \varepsilon n^2$ edges, then G contains at least $\delta n^{|V(H)|}$ copies of H.*

There are now many papers on supersaturation for other combinatorial structures such as monotone subsequences of finite integer sequences [23, 173, 197], substructures in finite set systems [75, 76, 79, 141, 180] and certain arithmetic structures in subsets of abelian groups [198].

In [141], Kleitman applied standard ideas from matching theory to prove that the number of comparable pairs in a subset of $\mathcal{P}(n)$ of prescribed size is minimised by a collection consisting of sets whose cardinalities are as close to $n/2$ as possible. Interestingly, though, these ideas cannot be easily extended to give a supersaturation theorem for longer chains. Kleitman [141] conjectured the following supersaturation version of the result of Erdős [83]: *for any $k \geq 3$, the number of k-chains in a subset of $\mathcal{P}(n)$ of prescribed size is minimised by a collection consisting of sets whose cardinalities are as close to $n/2$ as possible.* Das, Gan and Sudakov [76] and Dove, Griggs, Kang and Sereni [79] independently proved this conjecture in some restricted cases; however, the conjecture is open in general. The main idea in both of these papers is that one can sometimes obtain a lower bound on the number of k-chains in a subset S of a poset P by considering the expected size of the intersection of S with a chain in P chosen randomly according to some distribution. In some sense, this idea traces back to Lubell's proof [160] of the LYM inequality [35, 160, 166, 225] (see also [40]).

In the next section, we state a general lemma for proving lower bounds on the number of comparable pairs in a subset S of a poset P using the strategy from [76, 79]. While it is unlikely that such a lemma can be used to prove a structural result as precise as Kleitman's [141], the main advantages of this approach are that it is widely applicable and that it can often be used to obtain bounds which are good enough for applications. We illustrate this in Section 6.2 by proving an approximate version of Kleitman's result [141] and Theorem 1.18; both of these results will be applied in Chapter 7.

6.1 Comparison Counting Via Random Chains

The following lemma describes our main approach; before stating it, we require a definition.

Definition 6.2. Let P be a poset. A *comparability digraph* for P is a digraph D with $V(D) = P$ such that

- for every arc (x, y) of D, we have that x is comparable to y in P, and
- for every comparable pair $x, y \in P$, at least one of the arcs (x, y) or (y, x) is in $E(D)$.

Lemma 6.3. *Let P be a poset, let D be a comparability digraph for P, let μ be a distribution on the chains of P and let $\mathcal{C}_\mu$ be a random chain chosen according to μ. Suppose further that μ has the property that $\mathbb{P}(x \in \mathcal{C}_\mu) > 0$ for every $x \in P$. Then, for every set $S \subseteq P$, we have*

$$\operatorname{comp}(S) \geq \left(\max_{(x,y)\in E(D)} \mathbb{P}\left(y \in \mathcal{C}_\mu \mid x \in \mathcal{C}_\mu\right)\right)^{-1} \left(|S| - \left(\min_{x\in P} \mathbb{P}\left(x \in \mathcal{C}_\mu\right)\right)^{-1}\right).$$

Proof. We proceed by induction on $|S|$; the base case $|S| \leq (\min_{x\in P} \mathbb{P}(x \in \mathcal{C}_\mu))^{-1}$ is trivial. Given $w \in S$, we define $d_S^+(w)$ to be the number of arcs $(w, z) \in E(D)$ such that $z \in S$. Suppose, for contradiction, that there exists $w \in S$ such that

$$d_S^+(w) \geq \left(\max_{(x,y)\in E(D)} \mathbb{P}\left(y \in \mathcal{C}_\mu \mid x \in \mathcal{C}_\mu\right)\right)^{-1}.$$

In this case, we apply the inductive hypothesis to $S \setminus \{w\}$ to obtain at least

$$\left(\max_{(x,y)\in E(D)} \mathbb{P}\left(y \in \mathcal{C}_\mu \mid x \in \mathcal{C}_\mu\right)\right)^{-1} \left(|S \setminus \{w\}| - \left(\min_{x\in P} \mathbb{P}\left(x \in \mathcal{C}_\mu\right)\right)^{-1}\right).$$

comparable pairs in S which do not involve w. We then obtain the desired bound on $\operatorname{comp}(S)$ by considering these comparable pairs along with those of the form $(w, z) \in E(D)$ with $z \in S$. Thus, from here forward, we assume that every $w \in S$ satisfies

$$d_S^+(w) < \left(\max_{(x,y)\in E(D)} \mathbb{P}\left(y \in \mathcal{C}_\mu \mid x \in \mathcal{C}_\mu\right)\right)^{-1}. \tag{6.4}$$

The chain $\mathcal{C}_\mu$ contains precisely $|\mathcal{C}_\mu \cap S|$ elements of S and $\binom{|\mathcal{C}_\mu \cap S|}{2}$ of the comparable pairs of S (since every pair in $\mathcal{C}_\mu$ is comparable). Since $k - \binom{k}{2} \leq 1$ for all k, we get that

$$1 \geq \mathbb{E}\left(|\mathcal{C}_\mu \cap S| - \binom{|\mathcal{C}_\mu \cap S|}{2}\right) = \sum_{w \in S} \mathbb{P}\left(w \in \mathcal{C}_\mu\right) - \sum_{\substack{w,z \in S \\ (w,z) \in E(D)}} \mathbb{P}\left(w, z \in \mathcal{C}_\mu\right)$$

$$= \sum_{w \in S} \mathbb{P}\left(w \in \mathcal{C}_\mu\right) \left(1 - \sum_{\substack{z \in S \\ (w,z) \in E(D)}} \mathbb{P}\left(z \in \mathcal{C}_\mu \mid w \in \mathcal{C}_\mu\right)\right)$$

$$\geq \sum_{w \in S} \mathbb{P}\left(w \in \mathcal{C}_\mu\right) \left(1 - d_S^+(w) \left(\max_{(x,y) \in E(D)} \mathbb{P}\left(y \in \mathcal{C}_\mu \mid x \in \mathcal{C}_\mu\right)\right)\right).$$

By (6.4), all of the summands are positive, and so this expression is bounded below by

$$\sum_{w \in S} \left(\min_{x \in P} \mathbb{P}\left(x \in \mathcal{C}_\mu\right)\right) \left(1 - d_S^+(w) \left(\max_{(x,y) \in E(D)} \mathbb{P}\left(y \in \mathcal{C}_\mu \mid x \in \mathcal{C}_\mu\right)\right)\right)$$

Rearranging, we obtain

$$\mathrm{comp}(S) = \sum_{w \in S} d_S^+(w)$$

$$\geq \left(\max_{(x,y) \in E(D)} \mathbb{P}\left(y \in \mathcal{C}_\mu \mid x \in \mathcal{C}_\mu\right)\right)^{-1} \left(|S| - \left(\min_{x \in P} \mathbb{P}\left(x \in \mathcal{C}_\mu\right)\right)^{-1}\right)$$

as desired. $\square$

Remark 6.5. We remark that the proof of this lemma generalises to give a bound on the number of edges induced by a set of m vertices in a general graph G (in terms of m). That is, one could orient the edges of the graph and select a random clique according to a probability distribution μ.

In order for Lemma 6.3 be effective, one requires a distribution on the chains in P which is sufficiently 'balanced' in the sense that a random chain chosen with according to this distribution is not too *unlikely* to contain any fixed element of P and, at the same time, not too *likely* to contain any fixed comparable pair of P. In both of the applications provided in this chapter, it suffices to let $\mathcal{C}_\mu$ be a maximal chain in P chosen uniformly at random. However, in other posets, the problem of finding an appropriate distribution can be much more involved. One such example is the poset of divisors of a fixed integer with n prime factors of bounded multiplicity. In [176], we use Lemma 6.3 to obtain some supersaturation results for such posets under some (fairly restrictive) conditions.

6.2 Supersaturation in $\mathcal{P}(n)$ and $\mathcal{V}(q, n)$

As a first application of Lemma 6.3, we prove the following approximate version of the result of Kleitman [141].

Theorem 6.6 (Kleitman [141])**.** *Let k be a fixed positive integer. Then there exists a constant $n_0(k)$ such that if $n \geq n_0(k)$ and $S \subseteq \mathcal{P}(n)$ is a set of cardinality at least*

$$\sum_{r=0}^{k-1} \binom{n}{\lceil \frac{n-k+1+2r}{2} \rceil} + t,$$

then

$$\mathrm{comp}(S) \geq t \binom{\lceil (n+k)/2 \rceil}{k}.$$

Proof. For $0 \leq r \leq k-1$, let P_r be the poset consisting all subsets of $[n]$ of cardinality equivalent to $\lceil \frac{n-k+1+2r}{2} \rceil$ modulo k ordered by inclusion. Note that $P_0 \cup \cdots \cup P_{k-1}$ partitions $\mathcal{P}(n)$. Define $S_r := S \cap P_r$ and

$$t_r := |S_r| - \binom{n}{\lceil \frac{n-k+1+2r}{2} \rceil}.$$

Clearly, $\sum_{r=0}^{k-1} t_r = t$. Our goal is to show that, for $0 \leq r \leq k-1$, the following holds:

$$\mathrm{comp}(S_r) \geq t_r \binom{\lceil (n+k)/2 \rceil}{k}. \tag{6.7}$$

This will imply the result since $\mathrm{comp}(S) \geq \sum_{r=0}^{k-1} \mathrm{comp}(S_r)$.

Let $r \in \{0, \dots, k-1\}$ be fixed. We prove (6.7) by induction on t_r. Note that (6.7) holds trivially if $t_r \leq 0$. Now, for $t_r \geq 1$, if S_r contains an element x which is comparable to at least $\binom{\lceil (n+k)/2 \rceil}{k}$ elements of S_r, then we apply the induction hypothesis to find at least $(t_r - 1)\binom{\lceil (n+k)/2 \rceil}{k}$ comparable pairs in $S_r \setminus \{x\}$ and we are done. Therefore, there can be no such $x \in S_r$.

Let $s \in \{0, \dots, k-1\}$ so that $\lceil \frac{n-k+1+2r}{2} \rceil \equiv s \bmod k$. Let X be the set of all elements of S_r of cardinality s. Suppose that $X \neq \emptyset$ and for each $x \in X$ let $N_{k+s}(x)$ be the elements of $P_r \setminus S_r$ of cardinality $k+s$ which are comparable to x. By the result of the previous paragraph, x is comparable to fewer than $\binom{\lceil (n+k)/2 \rceil}{k}$ elements of S_r, and so

$$|N_{k+s}(x)| > \binom{n-s}{k} - \binom{\lceil (n+k)/2 \rceil}{k} \geq \binom{n-k+1}{k} - \binom{\lceil (n+k)/2 \rceil}{k}$$

which, if $n_0(k)$ is chosen large enough with respect to k, is at least

$$\binom{n}{k-1} \geq \binom{n}{s} \geq |X|.$$

Therefore, we can greedily associate each element x of X to an element y_x of $P_r \setminus S_r$ of cardinality $k+s$ such that x is comparable to y_x and the y_x are distinct from one another. It is clear that the set $S_r' := (S_r \setminus X) \cup \{y_x : x \in X\}$ has the same cardinality as S_r and at most as many comparable pairs as S_r. So, in what follows, we can assume that $X = \emptyset$; that is, every element of S_r has cardinality at least $k+s$.

Now, define P_r^* to be the poset obtained from P_r by deleting all sets of cardinality s. Let D be a digraph with $V(D) = P_r^*$ such that $E(D)$ consists of the pairs (x, y) where x and y are comparable and the cardinality of y is at least as close to $n/2$ as the cardinality of x. Let $\mathcal{C}_\mu$ be a random chain in P_r^* obtained by taking a maximal chain in $\mathcal{P}(n)$ uniformly at random and intersecting it with P_r^*. Then any two elements of P_r^* of the same cardinality are equally likely to be contained in $\mathcal{C}_\mu$, and so

$$\mathbb{P}(x \in \mathcal{C}_\mu) \geq \binom{n}{\lceil \frac{n-k+1+2r}{2} \rceil}^{-1} \text{ for all } x \in P_r^*.$$

Also, given $(x, y) \in E(D)$, we have

$$\mathbb{P}(y \in \mathcal{C}_\mu \mid x \in \mathcal{C}_\mu) = \begin{cases} \binom{|x|}{|y|}^{-1} & \text{if } x > y, \\ \binom{n-|x|}{|y|-|x|}^{-1} & \text{otherwise.} \end{cases}$$

By definition of D and the fact that $X = \emptyset$, we get that this quantity is always at most $\binom{\lceil (n+k)/2 \rceil}{k}^{-1}$. Thus, (6.7) holds by Lemma 6.3. $\square$

We remark that the idea of splitting the levels of $\mathcal{P}(n)$ into equivalence classes modulo k is also used by Collares Neto and Morris [69]. The proof of Theorem 6.6 given above exploits the following two properties of a uniformly random maximal chain $\mathcal{C}_\mu$ in the boolean lattice.

- any two subsets of $[n]$ of the same cardinality are equally likely to be contained in $\mathcal{C}_\mu$, and

- for any two comparable pairs (x_1, y_1) and (x_2, y_2) in $\mathcal{P}(n)$ such that $|x_1| = |x_2|$ and $|y_1| = |y_2|$, we have $\mathbb{P}(x_1, y_1 \in \mathcal{C}_\mu) = \mathbb{P}(x_2, y_2 \in \mathcal{C}_\mu)$.

It is clear that one can obtain a supersaturation result for any poset satisfying these two conditions by taking a maximal chain uniformly at random and applying Lemma 6.3. We demonstrate this by outlining a proof of Theorem 1.18.

Proof of Theorem 1.18. For $0 \leq r \leq k-1$, let P_r be the poset consisting all subspaces of $\mathbb{F}_q^n$ of dimension equivalent to $\lceil \frac{n-k+1+2r}{2} \rceil$ modulo k ordered by inclusion. Analogous to the proof of Theorem 6.6, we define $S_r := S \cap P_r$ and

$$t_r := |S_r| - \begin{bmatrix} n \\ \lceil \frac{n-k+1+2r}{2} \rceil \end{bmatrix}_q .$$

Clearly, $\sum_{r=0}^{k-1} t_r = t$. We prove that, for each $r = 0, \dots, k-1$,

$$\mathrm{comp}(S_r) \geq t_r \begin{bmatrix} \lceil (n+k)/2 \rceil \\ k \end{bmatrix}_q , \tag{6.8}$$

We prove (6.8) by induction on t_r where the case $t_r \leq 0$ is trivial. By the inductive hypothesis, we can assume that every element of S_r is comparable to fewer than $\begin{bmatrix} \lceil (n+k)/2 \rceil \\ k \end{bmatrix}_q$ other elements of S_r. If S_r contains an element of dimension $s < k$, then we can shift all such elements up to subspaces of dimension $k + s$ without increasing the number of comparable pairs (similar to the proof of Theorem 6.6) since, if $n_0(k)$ is chosen large enough with respect to k,

$$\begin{bmatrix} n-k+1 \\ k \end{bmatrix}_q - \begin{bmatrix} \lceil (n+k)/2 \rceil \\ k \end{bmatrix}_q = \left(\prod_{i=0}^{k-1} \frac{1-q^{n-i}}{1-q^{i+1}} \right) - \left(\prod_{i=0}^{k-1} \frac{1-q^{\lceil (n+k)/2 \rceil - i}}{1-q^{i+1}} \right)$$

$$\geq q^{nk - \left(\sum_{i=0}^{k-1} 2i+1\right)} - q^{\lceil (n+k)/2 \rceil k} \geq q^{n(k-1)} \geq \begin{bmatrix} n \\ k-1 \end{bmatrix}_q .$$

Now, we let P_r^* be the poset obtained from P_r by deleting all elements of dimension less than k and let D be a digraph with $V(D) = P_r^*$ where $E(D)$ consists of all pairs (x, y) such that x is comparable to y and the dimension of y is at least as close to $n/2$ as the dimension of x. Let $\mathcal{C}_\mu$ be a random chain in P_r^* obtained by taking a random maximal chain in $\mathcal{V}(q, n)$ and intersecting it with P_r^*. It is not hard to show that

$$\mathbb{P}(x \in \mathcal{C}_\mu) \geq \begin{bmatrix} n \\ \lceil \frac{n-k+1+2r}{2} \rceil \end{bmatrix}_q^{-1} \text{ for all } x \in P_r^*, \text{ and}$$

$$\mathbb{P}\left(y \in \mathcal{C}_\mu \mid x \in \mathcal{C}_\mu\right) \leq \begin{bmatrix} \lceil (n+k)/2 \rceil \\ k \end{bmatrix}_q^{-1} \text{ for all } (x, y) \in E(D).$$

The result follows. $\square$

Chapter 7

Applications of the Container Method in Posets

In the past few years, there has been some considerable success in extending classical results in extremal combinatorics to sparse random settings. An example of a question of this type, first investigated by Babai, Simonovits and Spencer [8], is roughly stated as follows: *Typically speaking, what is the maximum number of edges in a subgraph of $G(n, p)$ containing no copy of a fixed graph H?* This can be thought of as a random analogue of the original question of Turán [215] (question (1.1)). For a graph H with $\chi(H) \geq 3$, one obvious way of constructing such a subgraph is to split the vertices of $G(n, p)$ arbitrarily into $\chi(H) - 1$ sets of roughly equal size and remove all of the edges within the sets. Provided that p is not too small as a function of n we get that, with high probability, the resulting graph has $p\left(1 - \frac{1}{\chi(H)-1} + o(1)\right)\binom{n}{2} = (1 + o(1))p\,\mathrm{ex}(n, H)$ edges.[1]

Haxell, Kohayakawa and Łuczak [121] (see also [145]) conjectured that this construction is asymptotically best possible provided that $p = p(n)$ is large enough so that, with high probability, the number of copies of H in $G(n, p)$ greatly exceeds the number of edges in $G(n, p)$. Note that this is, in some sense, necessary; if the number of copies of H is small with respect to the number of edges in $G(n, p)$, then it is possible to obtain a better lower bound by simply deleting one edge from each copy of H. After a great deal of interest, this conjecture was finally proved by Conlon and Gowers [73] (under the additional assumption that H is 'strictly 2-balanced') and by

[1]Here, we used the famous theorem of Erdős and Stone [84] which says that $\mathrm{ex}(n, H) = \left(1 - \frac{1}{r-1} + o(1)\right)\binom{n}{2}$ for any graph H containing an edge.

Schacht [206] using quite different methods. While an obvious advantage of Schacht's method is that it applies to all graphs H, the method of Conlon and Gowers lends itself more naturally to proving stronger stability-type results. This apparent gap between the two methods was bridged by Samotij [195], who refined the approach of Schacht [206] to prove a stability version of the conjecture in full generality.

Shortly thereafter another, very different, method for proving extremal theorems in random structures was developed independently by Saxton and Thomason [205] and Balogh, Morris and Samotij [26]. As was demonstrated in [26, 205] this method, known as the 'container method,' can be used to prove many of the results of Conlon–Gowers [73], Schacht [206] and Samotij [195]. This includes, for example, the general stability version of the Haxell–Kohayakawa–Łuczak Conjecture [121] and an even stronger conjecture of Kohayakawa, Łuczak and Rödl [145] (the so-called KŁR Conjecture). Two advantages of the container method are that it is usually relatively straightforward to apply and that it often provides 'counting' results which can be much stronger than the corresponding probabilistic statements. Supersaturation results form a key ingredient in most proofs involving the container method; for instance, all of the applications mentioned in this paragraph relied on the classical supersaturation theorem of Erdős and Simonovits [92] (Theorem 6.1). For additional applications of the container method, see [24, 25, 27, 29, 31, 32, 69] and the references therein.

Our applications of the container method in this chapter were mainly inspired by the work of Balogh, Mycroft and Treglown [27] and Collares Neto and Morris [69] who independently applied an approximate version of the supersaturation theorem of Kleitman [141] (see Theorem 6.6) to prove the following 'random version' of Sperner's Theorem [210]; this result confirms a conjecture of Osthus [179].

Theorem 7.1 (Balogh, Mycroft and Treglown [27]; Collares Neto and Morris [69])**.** *Let $\varepsilon > 0$ be fixed. Then there exists a positive constant $c(\varepsilon)$ such that if $p \geq c(\varepsilon)/n$, then with high probability every antichain in a p-random subset of $\mathcal{P}(n)$ has cardinality at most $(1+\varepsilon)p\binom{n}{\lfloor n/2 \rfloor}$.*

Our Theorem 1.21, proved using Theorem 1.18 in Section 7.3, is a natural analogue of Theorem 7.1 for the poset $\mathcal{V}(q,n)$. Using ideas from [27], Balogh, Treglown and Wagner [31] recently applied the container method to give an alternative proof of a well known theorem of Kleitman [142] which says that the number of antichains in $\mathcal{P}(n)$ is at most $2^{(1+o(1))\binom{n}{\lfloor n/2 \rfloor}}$.[2] The problem of approximating the number of antichains in $\mathcal{P}(n)$ is known as Dedekind's Problem and has a long history; see [135]. In Section 7.2, we show how a 'multi-stage' container lemma for $\mathcal{P}(n)$, first used by Balogh, Mycroft and Treglown [27], implies a stronger upper bound of

$$2^{\left(1+O\left(\sqrt{\log(n)/n}\right)\right)\binom{n}{\lceil n/2 \rceil}}. \tag{7.2}$$

We remark that much stronger bounds than (7.2) are known. In particular, Korshunov [147] obtained precise asymptotics of the number of antichains in $\mathcal{P}(n)$; see also [135, 143, 200]. We prove an upper bound on the number of antichains in $\mathcal{V}(q, n)$ (Theorem 1.19) using a similar approach. Our proofs use a container lemma for general posets which is stated and proved in Section 7.1. The proof of this lemma relies heavily on ideas from [27] (some of which are based on ideas of Kleitman and Winston [144]).

7.1 A Container-Type Lemma for General Posets

In this section, we prove a container-type lemma for general posets which generalises a result of Balogh, Mycroft and Treglown [27]. The main idea in the proof originated in the work of Kleitman and Winston [144] and has now been used many times; see, e.g., [27, 31, 196]. In essence, what this lemma says is that if P satisfies a certain supersaturation bound, then there is a collection of 'not very large' subsets of P indexed by 'very small' subsets of P with the property that every antichain in P is contained in at least one set from the collection. The elements of such a collection are referred to as *containers*. The fact that the containers are indexed by small subsets of P is typically used to argue that the total number of containers is very small compared, for example, to the total number of antichains.

Definition 7.3. Given a set X and an integer j, we let $\binom{X}{\leq j}$ denote the set of all subsets of m of cardinality at most j. Also, for $n \geq j$, define $\binom{n}{\leq j} := \sum_{r=0}^{j} \binom{n}{r}$.

[2]In addition, the authors of [31] obtained a 'two-coloured' generalisation of this result, obtained bounds on the number of '(p, q)-tilted Sperner families' (see [151]), and proved a number of results on intersecting families and error-correcting codes.

Lemma 7.4. *Let d and m be positive integers and let P be a poset such that $|P| > m$ and every subset S of P of cardinality greater than m contains at least $|S|d$ comparable pairs. Then there exists a function*

$$f : \binom{P}{\leq |P|/(2d+1)} \to \binom{P}{\leq m}$$

such that for every antichain $I \subseteq P$ there exists a subset T of I of cardinality at most $|P|/(2d+1)$ with $T \cap f(T) = \emptyset$ and $I \subseteq T \cup f(T)$.

Actually, we will prove something more general. A key idea from [27] is that it can often be advantageous to prove a 'multi-stage' container lemma, where the purpose of the early stages is to 'thin out' the poset which helps to reduce the total number of containers (see [31] for another application of this approach). We prove the following lemma of this type; Lemma 7.4 corresponds to the case $k = 1$. We remark that we will only actually apply this lemma in the cases $k = 1$ and $k = 2$.

Lemma 7.5. *For $k \geq 1$ let $d_1 > \cdots > d_k$ and $m_0 > \cdots > m_k$ be positive integers and let P be a poset such that $|P| = m_0$ and, for $1 \leq j \leq k$, every subset S of P of cardinality greater than m_j contains at least $|S|d_j$ comparable pairs. Then there exists functions $f_1, \dots, f_k$ where*

$$f_j : \binom{P}{\leq \sum_{r=1}^{j} (m_{r-1}/(2d_r+1))} \to \binom{P}{\leq m_j}$$

such that for every antichain $I \subseteq P$ there exists disjoint subsets $T_1, \dots, T_k$ of I with

(i) $|T_j| \leq m_{j-1}/(2d_j+1)$ *for* $1 \leq j \leq k$,

(ii) $T_j \subseteq f_{j-1}\left(\bigcup_{r=1}^{j-1} T_r\right)$ *for* $2 \leq j \leq k-1$,

(iii) $\left(\bigcup_{r=1}^{j} T_r\right) \cap f_j\left(\bigcup_{r=1}^{j} T_r\right) = \emptyset$ *for* $1 \leq j \leq k$, *and*

(iv) $I \subseteq \left(\bigcup_{r=1}^{k} T_r\right) \cup f_k\left(\bigcup_{r=1}^{k} T_r\right)$.

Proof. Fix an arbitrary total order $x_1, \dots, x_{|P|}$ on the elements of P (this will only be used to 'break ties' later in the proof). The proof of this lemma amounts to a simple application of the so-called Kleitman–Winston algorithm [144]. This algorithm takes, as an input, an antichain I of P and produces the sets $T_1, \dots, T_k$ as well as the sets $f_1(T_1), \dots, f_k\left(\bigcup_{r=1}^{k} T_r\right)$. After describing this algorithm, we will verify that these sets have all of the desired properties and that the functions $f_1, \ldots, f_k$ are well defined; i.e., we show that $f_j\left(\bigcup_{r=1}^{j} T_r\right)$ depends only on $\bigcup_{r=1}^{j} T_r$ and not on the input antichain I or on the sets $T_1, \ldots, T_j$ individually.

Set $P_0 := P$. Now, given $i \geq 1$ and the set P_{i-1}, the ith step of the algorithm will either terminate or it will produce a set $P_i \subsetneq P_{i-1}$ which it will pass to the next step. For each $x \in P_{i-1}$, let $N_{P_{i-1}}(x)$ be the set of elements of P_{i-1} that are comparable to x and let $d_{P_{i-1}}(x) := |N_{P_{i-1}}(x)|$. Let u_{i-1} be the element of P_{i-1} for which $d_{P_{i-1}}(u_{i-1})$ is maximum; if there is a tie, then we let u_{i-1} be the choice which comes earliest in the total order. Now, given u_{i-1}, we proceed differently depending on the value of $d_{P_{i-1}}(u_{i-1})$.

Case 1: $d_{P_{i-1}}(u_{i-1}) \geq 2d_k$.

Let j be the smallest integer in $\{1, \ldots, k\}$ such that $d_{P_{i-1}}(u_{i-1}) \geq 2d_j$. First, if $i = 1$, then since $|P| = m_0 > m_1$ we must have that P contains at least $|P|d_1$ comparable pairs. This implies that $d_{P_0}(u_0) \geq 2d_1$ and so $j = 1$. If $i \geq 2$, then we let ℓ be the smallest integer in $\{1, \ldots, k\}$ such that $d_{P_{i-2}}(u_{i-2}) \geq 2d_\ell$. If $\ell < j$, then we define

$$f_\ell\left(\bigcup_{r=1}^{\ell} T_r\right) = \cdots = f_{j-1}\left(\bigcup_{r=1}^{j-1} T_r\right) := P_{i-1}.$$

Now, if $u_{i-1} \notin I$, then we simply define $P_i := P_{i-1} \setminus \{u_{i-1}\}$ and proceed to the next step. On the other hand, if $u_{i-1} \in I$, then we add u_{i-1} to T_j, define $P_i := P_{i-1} \setminus \left(\{u_{i-1}\} \cup N_{P_{i-1}}(u_{i-1})\right)$ and proceed to the next step. It is clear that $P_i \subsetneq P_{i-1}$.

Case 2: $d_{P_{i-1}}(u_{i-1}) < 2d_k$.

In this case, we will terminate the algorithm at the end of this step. Before terminating, though, we will need to define all of the sets $f_j\left(\bigcup_{r=1}^{j} T_r\right)$ which were not defined in previous steps.

Note that $i \geq 2$ since $|P| = m_0 > m_k$ which implies that P contains at least $|P|d_k$ comparable pairs and, hence, we must have $d_{P_0}(u_0) \geq 2d_k$. So, since $i \geq 2$ and the algorithm terminates on the first step in which Case 2 is satisfied (in particular, Case 1 was satisfied at the previous step), we can define ℓ to be the smallest integer in $\{1,\ldots,k\}$ such that $d_{P_{i-2}}(u_{i-2}) \geq 2d_\ell$. We set

$$f_\ell\left(\bigcup_{r=1}^{\ell} T_r\right) = \cdots = f_k\left(\bigcup_{r=1}^{k} T_r\right) := P_{i-1}$$

and terminate the algorithm. This concludes the description of the algorithm.

A key feature of the algorithm is that the sequence $d_{P_0}(u_0), d_{P_1}(u_1), \ldots$ is non-increasing, which follows easily from the fact that $P_i \subsetneq P_{i-1}$ for all $i \geq 1$. Since this sequence is non-increasing and $d_1 > \cdots > d_k$ we get that, for each $j \in \{1,\ldots,k\}$, there is a unique step of the algorithm such that $d_{P_{i-2}}(u_{i-2}) \geq 2d_j$ and $d_{P_{i-1}}(u_{i-1}) < 2d_j$. It is precisely at this step when $f_j\left(\bigcup_{r=1}^{j} T_r\right)$ is defined and when the definition of T_j is finalised. In particular, the set $f_j\left(\bigcup_{r=1}^{j} T_r\right)$ is defined exactly once during the running of the algorithm.

It is clear by construction that (ii), (iii) and (iv) hold and that $T_j \subseteq I$ for $1 \leq j \leq k$. Also, by construction, every element of $f_j\left(\bigcup_{r=1}^{j} T_r\right)$ is comparable to fewer than $2d_j$ other elements in $f_j\left(\bigcup_{r=1}^{j} T_r\right)$. Therefore, $\mathrm{comp}\left(f_j\left(\bigcup_{r=1}^{j} T_r\right)\right) < \left|f_j\left(\bigcup_{r=1}^{j} T_r\right)\right| d_j$ which implies that $\left|f_j\left(\bigcup_{r=1}^{j} T_r\right)\right| \leq m_j$ by hypothesis. This verifies that f_j does, indeed, map into $\binom{P}{\leq m_j}$. For each element that we added to T_1 we deleted at least $2d_1+1$ elements of P. Therefore, $|T_1| \leq |P|/(2d_1+1) = m_0/(2d_1+1)$. Similarly, for $j \geq 2$, for each element that we added to T_j we deleted at least $2d_j+1$ elements of $f_{j-1}\left(\bigcup_{r=1}^{j-1} T_r\right)$ and so $|T_j| \leq m_{j-1}/(2d_j+1)$. Thus, (i) holds.

Finally, we argue that the functions $f_1,\ldots,f_k$ are well defined. Let I and I' be antichains yielding the same set $\bigcup_{r=1}^{j} T_r$ for some j with $1 \leq j \leq k$. Let $u_0, u_1, \ldots$ and $P_0, P_1, \ldots$ be the vertices and sets produced while running the algorithm with input I and let $u_0', u_1', \ldots$ and $P_0', P_1', \ldots$ be the vertices and sets produced while running the algorithm with input I'. Let i denote the minimum integer such that $d_{P_i}(u_i) < 2d_j$. We prove that $u_t = u_t'$ and $P_t = P_t'$ for all $0 \leq t \leq i-1$. If not, let t be the smallest such integer for which it fails. Clearly, $t \geq 1$. The fact that $P_{t-1} = P_{t-1}'$ implies that $u_t = u_t'$ and so we must have $P_t \neq P_t'$. However, since $u_{t-1} = u_{t-1}'$ and $P_{t-1} = P_{t-1}'$, the only way that we can have $P_t \neq P_t'$ is if $u_{t-1} \in I$ and $u_{t-1}' \notin I'$ or vice versa. By definition of i, we have $d_{P_{t-1}}(u_{t-1}) \geq 2d_j$ and so $u_{t-1} \in I$ implies $u_{t-1} \in \bigcup_{r=1}^{j} T_r$ by the description of the algorithm. However, $u_{t-1}' \notin I'$ implies that the algorithm will not add u_{t-1}' to any of the sets $T_1,\ldots,T_k$ and so $u_{t-1}' \notin \bigcup_{r=1}^{j} T_r$, contradicting the fact that $u_{t-1} = u'_{t-1}$. Therefore $f_j\left(\bigcup_{r=1}^{j} T_r\right)$ does not depend on I. This argument also proves that it depends only on the union $\bigcup_{r=1}^{j} T_r$ and not on the sets $T_1, \dots, T_j$ individually. This completes the proof. $\square$

The following lemma is a consequence of Lemma 7.5 which we will use to count antichains in the next section.

Lemma 7.6. *For $k \geq 1$ let $d_1 > \cdots > d_k$ and $m_0 > \cdots > m_k$ be positive integers and let P be a poset such that $|P| = m_0$ and, for $1 \leq j \leq k$, every subset S of P of cardinality greater than m_j contains at least $|S|d_j$ comparable pairs. Then there is a collection $\mathcal{F}$ of subsets of P such that*

(a) $|\mathcal{F}| \leq \prod_{r=1}^{k} \binom{m_{r-1}}{\leq m_{r-1}/(2d_r+1)}$,

(b) $|A| \leq m_k + \sum_{r=1}^{k} \frac{m_{r-1}}{2d_r+1}$ *for every $A \in \mathcal{F}$, and*

(c) for every antichain I of P, there exists $A \in \mathcal{F}$ such that $I \subseteq A$.

Proof. Apply Lemma 7.5 to obtain the functions $f_1, \dots, f_k$. For each antichain I, let $T_{1,I}, \dots, T_{k,I}$ be disjoint sets as in Lemma 7.5 such that $I \subseteq \left(\bigcup_{r=1}^{k} T_{r,I}\right) \cup f\left(\bigcup_{r=1}^{k} T_{r,I}\right)$. Define

$$\mathcal{F} := \left\{\left(\bigcup_{r=1}^{k} T_{r,I}\right) \cup f\left(\bigcup_{r=1}^{k} T_{r,I}\right) : I \subseteq P \text{ is an antichain}\right\}.$$

Then (b) and (c) hold by construction. Let us show that (a) holds. The number of ways of choosing $T_{1,I}$ is at most $\binom{|P|}{\leq |P|/(2d_1+1)}$. Now, for $2 \leq j \leq k$, given that $T_{1,I}, \dots, T_{j-1,I}$ have been chosen, we know that $T_{j,I}$ is a subset of $f_{j-1}\left(\bigcup_{r=1}^{j-1} T_{r,I}\right)$. Therefore, there are at most $\binom{m_{j-1}}{\leq m_{j-1}/(2d_j+1)}$ ways to choose $T_{j,I}$. This argument proves (a). $\square$

7.2 Counting Antichains

We begin by applying Lemma 7.6 and Theorem 6.6 to give a short proof of a known upper bound on the number of antichains in $\mathcal{P}(n)$.

Theorem 7.7 (See [135]). *The number of antichains in $\mathcal{P}(n)$ is at most*

$$2^{\left(1+O\left(\sqrt{\log(n)/n}\right)\right)\binom{n}{\lceil n/2\rceil}}.$$

Proof. Set $d_1 := n\sqrt{\log n}$, $d_2 := \sqrt{n \log n}$, $m_0 := 2^n$, $m_1 := \frac{2\binom{n}{\lfloor n/2 \rfloor}}{1-8d_1/n^2}$ and $m_2 := \frac{\binom{n}{\lfloor n/2 \rfloor}}{1-2d_2/n}$. If S is a subset of $\mathcal{P}(n)$ of cardinality at least m_1, then, by Theorem 6.6,

$$\operatorname{comp}(S) \geq \left(|S| - 2\binom{n}{\lceil n/2 \rceil}\right)\left(\frac{n^2}{8}\right)$$

$$= (|S| - m_1)\left(\frac{n^2}{8}\right) + \left(m_1 - 2\binom{n}{\lceil n/2 \rceil}\right)\left(\frac{n^2}{8}\right) \geq (|S| - m_1)d_1 + m_1 d_1 = |S|d_1$$

since $\left(m_1 - 2\binom{n}{\lfloor n/2 \rfloor}\right)\left(\frac{n^2}{8}\right) = m_1 d_1$. Similarly, if S has cardinality at least m_2, then $\operatorname{comp}(S) \geq |S|d_2$. Therefore, we can apply Lemma 7.6 to $\mathcal{P}(n)$ with $k = 2$ and these values of d_1, d_2, m_0, m_1 and m_2 to obtain a collection $\mathcal{F}$ of containers. We get that the number of antichains in $\mathcal{P}(n)$ is at most

$$\begin{aligned}
\binom{2^n}{\leq 2^n/d_1}\binom{m_1}{\leq m_1/d_2} 2^{m_2 + \frac{m_1}{d_2} + \frac{2^n}{d_1}} &\leq 2^{2n}\binom{2^n}{2^n/d_1}\binom{m_1}{m_1/d_2} 2^{m_2 + \frac{m_1}{d_2} + \frac{2^n}{d_1}} \\
&\leq 2^{2n}\left(e \cdot d_1\right)^{2^n/d_1}\left(e \cdot d_2\right)^{m_1/d_2} 2^{m_2 + \frac{m_1}{d_2} + \frac{2^n}{d_1}} \\
&= 2^{2n + \frac{2^n \log_2(e \cdot d_1)}{d_1} + \frac{m_1 \log_2(e \cdot d_2)}{d_2} + m_2 + \frac{m_1}{d_2} + \frac{2^n}{d_1}} \\
&\leq 2^{m_2 + O\left(\frac{2^n \log n}{d_1} + \frac{m_1 \log n}{d_2}\right)}
\end{aligned}$$

for n sufficiently large. Now, bounding the exponent, we get

$$\begin{aligned}
&m_2 + O\left(\frac{2^n \log n}{d_1} + \frac{m_1 \log n}{d_2}\right) \\
&= \frac{\binom{n}{\lfloor n/2 \rfloor}}{1 - 2d_2/n} + O\left(\frac{2^n\sqrt{\log n}}{n} + \binom{n}{\lfloor n/2 \rfloor}\sqrt{\log n/n}\right) \\
&\leq \binom{n}{\lfloor n/2 \rfloor} + \frac{4d_2}{n}\binom{n}{\lfloor n/2 \rfloor} + O\left(\binom{n}{\lfloor n/2 \rfloor}\sqrt{\log n/n}\right) \\
&= \left(1 + O\left(\sqrt{\log n/n}\right)\right)\binom{n}{\lfloor n/2 \rfloor}
\end{aligned}$$

since $\binom{n}{\lfloor n/2 \rfloor} = \Theta\left(\frac{2^n}{\sqrt{n}}\right)$ by Stirling's Approximation. The result follows. $\square$

It is possible that one could apply the full power of the supersaturation theorem of Kleitman [141] alongside Lemma 7.6 to obtain a better bound than that of Theorem 7.7; however, we will refrain from this here as this is unlikely to yield the best known results.

Theorem 1.19 can be proved by following along similar lines to the proof of Theorem 7.7 given above. In proving it, the following theorem of Wilf [222] will be useful for estimating relative orders of magnitude.

Theorem 7.8 (Wilf [222])**.** *For fixed q and $n \to \infty$, we have*

$$|\mathcal{V}(q,n)| = \Theta\left(q^{n^2/2}\right), \text{ and}$$

$$\begin{bmatrix} n \\ \lfloor n/2 \rfloor \end{bmatrix}_q = \Theta\left(q^{n^2/2}\right).$$

In particular, $|\mathcal{V}(q,n)| = \Theta\left(\begin{bmatrix} n \\ \lfloor n/2 \rfloor \end{bmatrix}_q\right)$.

We prove Theorem 1.19.

Proof of Theorem 1.19. Let $d := \sqrt{n}q^{n/4}$ and $m := \frac{\begin{bmatrix} n \\ \lfloor n/2 \rfloor \end{bmatrix}_q}{1-d\begin{bmatrix} \lceil (n+1)/2 \rceil \\ 1 \end{bmatrix}_q^{-1}}$. By Theorem 1.18, if S is a subset of $\mathcal{V}(q,n)$ of cardinality at least m, then $\mathrm{comp}(S) \geq |S|d$. Thus, by Lemma 7.6, the number of antichains in $\mathcal{V}(q,n)$ is at most

$$\binom{|\mathcal{V}(q,n)|}{\leq |\mathcal{V}(q,n)|/d} 2^{m+\frac{|\mathcal{V}(q,n)|}{d}} \leq |\mathcal{V}(q,n)| \binom{|\mathcal{V}(q,n)|}{|\mathcal{V}(q,n)|/d} 2^{m+\frac{|\mathcal{V}(q,n)|}{d}}$$

$$\leq |\mathcal{V}(q,n)| (e \cdot d)^{\frac{|\mathcal{V}(q,n)|}{d}} 2^{m+\frac{|\mathcal{V}(q,n)|}{d}} = 2^{m+\log|\mathcal{V}(q,n)|+\left(\frac{|\mathcal{V}(q,n)|}{d}\right)(\log(e\cdot d)+1)}$$

Since q is fixed, we can use Theorem 7.8 to bound the exponent as follows:

$$m + \log|\mathcal{V}(q,n)| + \left(\frac{|\mathcal{V}(q,n)|}{d}\right)(\log(e \cdot d) + 1)$$

$$\leq \frac{\begin{bmatrix} n \\ \lfloor n/2 \rfloor \end{bmatrix}_q}{1 - d\begin{bmatrix} \lceil (n+1)/2 \rceil \\ 1 \end{bmatrix}_q^{-1}} + O\left(n^2 + \frac{|\mathcal{V}(q,n)| \cdot n}{d}\right)$$

$$\leq \begin{bmatrix} n \\ \lfloor n/2 \rfloor \end{bmatrix}_q + \frac{2d\begin{bmatrix} n \\ \lfloor n/2 \rfloor \end{bmatrix}_q}{\begin{bmatrix} \lceil (n+1)/2 \rceil \\ 1 \end{bmatrix}_q} + O\left(\frac{|\mathcal{V}(q,n)| \cdot n}{d}\right)$$

This completes the proof since $|\mathcal{V}(q,n)| = \Theta\left(\begin{bmatrix} n \\ \lfloor n/2 \rfloor \end{bmatrix}_q\right)$ by Theorem 7.8 and, for fixed q, we have $\begin{bmatrix} \lceil (n+1)/2 \rceil \\ 1 \end{bmatrix}_q = \Theta\left(q^{n/2}\right)$. □

7.3 Antichains in Random Subsets

In this section, we use Lemma 7.6 to prove Theorem 1.21. We remark that our approach is very similar to that of Balogh, Mycroft and Treglown [27].

Proof of Theorem 1.21. Let $p = c(\varepsilon, q)/q^{n/2}$ for some constant $c(\varepsilon, q)$, to be chosen later, and let A_p be a subset of $\mathcal{V}(q,n)$ obtained by including each element independently with probability p. Let $m = (1+\varepsilon/2)\begin{bmatrix} n \\ \lceil n/2\rceil \end{bmatrix}_q$. By Theorem 1.18, every subset S of $\mathcal{V}(q,n)$ of cardinality at least m satisfies

$$\operatorname{comp}(S) \geq \left(|S| - \begin{bmatrix} n \\ \lceil n/2\rceil \end{bmatrix}_q\right)\begin{bmatrix} \lceil (n+1)/2\rceil \\ 1 \end{bmatrix}_q$$

$$\geq \left(1 - \frac{1}{1+\varepsilon/2}\right)\begin{bmatrix} \lceil (n+1)/2\rceil \\ 1 \end{bmatrix}_q |S| \geq \left(\frac{\varepsilon}{3}\right)\begin{bmatrix} \lceil (n+1)/2\rceil \\ 1 \end{bmatrix}_q |S|$$

for $\varepsilon < 1$. So, we set $d = \left(\frac{\varepsilon}{3}\right)\begin{bmatrix} \lceil (n+1)/2\rceil \\ 1 \end{bmatrix}_q$ and let f be the function resulting from applying Lemma 7.4 to $\mathcal{V}(q,n)$ with these values of m and d.

Fix $T \subseteq \mathcal{V}(q,n)$ such that $|T| \leq \frac{|\mathcal{V}(q,n)|}{2d+1}$. The probability that A_p contains an antichain I of size at least $(1+\varepsilon)p\begin{bmatrix} n \\ \lceil n/2\rceil \end{bmatrix}_q$ with $T \subseteq I \subseteq T \cup f(T)$ is at most

$$\mathbb{P}(T \subseteq A_p) \cdot \mathbb{P}\left(|A_p \cap f(T)| \geq p(1+\varepsilon)\begin{bmatrix} n \\ \lceil n/2\rceil \end{bmatrix}_q - |T|\right)$$

since $T \cap f(T) = \emptyset$. By Theorem 7.8, $|T| \leq \frac{|\mathcal{V}(q,n)|}{2d+1} = O\left(p\begin{bmatrix} n \\ \lceil n/2\rceil \end{bmatrix}_q\right)$. By choosing $c(\varepsilon, q)$ sufficiently large, we get that the above probability is at most

$$p^{|T|}\mathbb{P}\left(|A_p \cap f(T)| \geq p(1+3\varepsilon/4)\begin{bmatrix} n \\ \lceil n/2\rceil \end{bmatrix}_q\right).$$

By the Chernoff bound and the fact that $|f(T)| \leq (1+\varepsilon/2)m$, the second factor can be bounded above by, say, $e^{-\varepsilon^2 mp/100}$. Therefore, the probability that A_p contains an antichain of size at least $(1+\varepsilon)p\begin{bmatrix} n \\ \lceil n/2\rceil \end{bmatrix}_q$ is at most

$$\sum_{\substack{T \subseteq \mathcal{V}(q,n) \\ |T| \leq |\mathcal{V}(q,n)|/(2d+1)}} p^{|T|} e^{-\varepsilon^2 mp/100} \leq \sum_{a=0}^{|\mathcal{V}(q,n)|/(2d+1)} \binom{|\mathcal{V}(q,n)|}{a} p^a e^{-\varepsilon^2 mp/100}.$$

Given that $c(\varepsilon, q)$ is large enough, the quantity $\binom{|\mathcal{V}(q,n)|}{a}p^a$ is increasing for $0 \leq a \leq \frac{|\mathcal{V}(q,n)|}{2d+1}$. Thus, the above expression is bounded above by

$$\left(\frac{|\mathcal{V}(q,n)|}{2d+1} + 1\right)(e(2d+1)p)^{|\mathcal{V}(q,n)|/(2d+1)} e^{-\varepsilon^2 mp/100}.$$

Finally, for $c(\varepsilon, q)$ sufficiently large, this expression is $o(1)$. The result follows. $\square$

Let us argue that Theorem 1.21 is best possible up to the choice of the constant $c(\varepsilon, q)$. Let A_p be a p-random subset of $\mathcal{V}(q,n)$ where $p = c/q^{n/2}$ for some constant $c > 0$ and let n be large with respect to q and c. Given $x \in \mathcal{V}(q,n)$ of dimension $\lceil (n+1)/2 \rceil$, the expected number of subspaces of x of dimension $\lfloor n/2 \rfloor$ contained in A_p is constant. Thus, with probability bounded away from zero, a constant proportion of the elements of $\mathcal{V}(q,n) \cap A_p$ of dimension $\lceil (n+1)/2 \rceil$ have no subspaces of dimension $\lfloor n/2 \rfloor$ contained in A_p. Taking this set along with all subspaces of $\mathbb{F}_q^n$ of dimension $\lfloor n/2 \rfloor$ gives us an antichain of size at least $(1+\varepsilon)p \begin{bmatrix} n \\ \lfloor n/2 \rfloor \end{bmatrix}_q$ for some $\varepsilon = \varepsilon(q,c) > 0$. Here, we used the fact that $\begin{bmatrix} n \\ \lceil (n+1)/2 \rceil \end{bmatrix}_q = \Theta\left(\begin{bmatrix} n \\ \lfloor n/2 \rfloor \end{bmatrix}_q \right)$ for fixed q.

7.4 Open Problems

It would be interesting to improve our bound on the number of antichains in $\mathcal{V}(q,n)$ (Theorem 1.19). For example, we propose the following.

Problem 7.9 (Noel, Scott and Sudakov [176]). *Determine the precise asymptotics for the number of antichains in $\mathcal{V}(q,n)$.*

Short of this, it would be interesting, and probably quite challenging, to determine the second order term in the asymptotics of the logarithm of the number of antichains in $\mathcal{V}(q,n)$.

Chapter 8

Weak Regularity and Finitely Forcible Graph Limits

The theory of combinatorial limits has recently attracted a significant amount of attention. This line of research was sparked by limits of dense graphs [50–52, 157], which we focus on here, followed by limits of other structures, e.g. permutations [105, 124, 125, 150], sparse graphs [46, 81] and partial orders [130]. Methods related to combinatorial limits have led to substantial results in many areas of mathematics and computer science, particularly in extremal combinatorics. For example, the notion of flag algebras, which is strongly related to combinatorial limits, resulted in progress on many important problems in extremal combinatorics [9–11, 22, 113, 119, 120, 148, 149, 185–187, 189, 190]. Theory of combinatorial limits also provided a new perspective on existing concepts in other areas, e.g. property testing algorithms in computer science [106, 126, 158].

A convergent sequence of dense graphs can be represented by an analytic object called a *graphon*. Let $d(H, W)$ be the density of a graph H in a graphon W (a formal definition is given in Section 8.1). A graphon W is said to be *finitely forcible* if it is determined by finitely many subgraph densities, i.e. there exist graphs $H_1, \ldots, H_k$ and reals $d_1, \ldots, d_k$ such that W is the unique graphon with $d(H_i, W) = d_i$. Finitely forcible graphons appear in many different settings, one of which is in extremal combinatorics. It is known that if a graphon is finitely forcible, then it is the unique graphon which minimises a fixed finite linear combination subgraph densities, i.e. finitely forcible graphons are extremal points of the space of all graphons. The following conjecture [159, Conjecture 7] claims that the converse is also true.

Conjecture 8.1. *Let $H_1, \ldots, H_k$ be finite graphs and $\alpha_1, \ldots, \alpha_k$ reals. There exists a finitely forcible graphon W that minimises the sum $\sum_{i=1}^{k} \alpha_i d(H_i, W)$.*

Finitely forcible graphons are also related to quasirandomness in graphs as studied e.g. by Chung, Graham and Wilson [67], Rödl [193] and Thomason [212,213]. In the language of graph limits, results on quasirandom graphs state that every constant graphon is finitely forcible. A generalisation of this statement was proved by Lovász and Sós [156]: every step graphon (i.e. a multipartite graphon with a finite number of parts and uniform edge densities between its parts) is finitely forcible.

In [159], Lovász and Szegedy carried out a more systematic study of finitely forcible graphons. The examples of finitely forcible graphons that they constructed led to a belief that finitely forcible graphons must have a simple structure. To formalise this, they introduced the (topological) space $T(W)$ of typical vertices of a graphon W and conjectured the following [159, Conjectures 9 and 10].

Conjecture 8.2. *The space of typical vertices of every finitely forcible graphon is compact.*

Conjecture 8.3. *The space of typical vertices of every finitely forcible graphon has a finite dimension.*

Both conjectures were disproved through counterexample constructions in [107,108].

Conjecture 8.3 is a starting point of our work. Analogously to weak regularity of graphs, every graphon has a weak ε-regular partition with at most $2^{O\left(\varepsilon^{-2}\right)}$ parts. (See Section 8.2 for the necessary definitions.) If the space of typical vertices of a graphon is equipped with an appropriate metric, then its Minkowski dimension is linked to the number of parts in its weak regular partitions. In particular, if its Minkowski dimension is d, then the graphon has weak ε-regular partitions with $O\left(\varepsilon^{-d}\right)$ parts. Consequently, if Conjecture 8.3 were true, the number of parts of a weak ε-regular partitions of a finitely forcible graphon would be bounded by a polynomial of ε^{-1}. The number of parts in weak ε-regular partitions of a graphon constructed in [107] as a counterexample to Conjecture 8.3 is $2^{\Theta(\log^2 \varepsilon^{-1})}$, which is superpolynomial in ε^{-1}, but is much smaller than the general upper bound of $2^{O\left(\varepsilon^{-2}\right)}$. We construct a finitely forcible graphon almost matching the upper bound.

Theorem 8.4 (Cooper, Kaiser, Král and Noel [74]). *There exist a finitely forcible graphon W and positive reals ε_i tending to 0 such that every weak ε_i-regular partition of W has at least $2^{\Omega\left(\varepsilon_i^{-2}/2^{5\log^*\varepsilon_i^{-2}}\right)}$ parts.*

As pointed out to us (the authors of [74]) by Jacob Fox, there is no graphon (finitely forcible or not) matching the upper bound for infinitely many values of ε tending to 0. In light of this, Theorem 8.4 is almost the best possible.

Proposition 8.5. *There exist no graphon W, positive real c and positive reals ε_i tending to 0 such that every weak ε_i-regular partition of W has at least $2^{c\varepsilon_i^{-2}}$ parts.*

The proof of this proposition is sketched at the end of Section 8.2.

We will refer to the graphon W from Theorem 8.4 as the Švejk graphon. Švejk is the name of a famous (and fictitious) brave Czech soldier and, more importantly for us, it is the name of the restaurant where we (the authors of [74]) usually ate lunch during our work on this subject while three of us were visiting the University of West Bohemia in Pilsen.

8.1 Graph limits

We now introduce notions related to graphons and convergent sequences of graphs. The *density* of a graph H in G, which is denoted by $d(H, G)$, is the probability that $|H|$ randomly chosen vertices of G induce a subgraph isomorphic to H, where $|H|$ is the order (the number of vertices) of H. A sequence of graphs $(G_n)_{n\in\mathbb{N}}$ with the number of their vertices tending to infinity is *convergent* if the sequence $d(H, G_n)$ converges for every graph H. Note that if G_n has $o(|G_n|^2)$ edges, then the sequence $(G_n)_{n\in\mathbb{N}}$ is convergent for trivial reasons. Hence, this notion of graph convergence is of interest for sequences of dense graphs, i.e. graphs with $\Omega(|G_n|^2)$ edges.

A convergent sequence of dense graphs can be represented by an analytic object called a graphon. A *graphon* W is a measurable function from $[0, 1]^2$ to $[0, 1]$ that is symmetric, i.e. it holds that $W(x, y) = W(y, x)$ for every $x, y \in [0, 1]$. The points in $[0, 1]$ are often referred to as the *vertices* of the graphon W.

If W is a graphon, then a *W-random graph* of order k is obtained by sampling k random points $x_1, \ldots, x_k \in [0, 1]$ uniformly and independently, and joining the ith and the jth vertex by an edge with probability $W(x_i, x_j)$. The *density* of a graph H in a graphon W is the probability that a W-random graph of order $|H|$ is isomorphic to H. If $(G_n)_{n\in\mathbb{N}}$ is a convergent sequence of graphs, then there exists a graphon W such that $d(H, W) = \lim_{n\to\infty} d(H, G_n)$ for every graph H [157]. This graphon can be viewed as the limit of the sequence $(G_n)_{n\in\mathbb{N}}$. On the other hand, a sequence of W-random graphs of increasing orders is convergent with probability one and its limit is the graphon W.

Two graphons W_1 and W_2 are *weakly isomorphic* if there exist measure preserving maps φ_1 and φ_2 from $[0, 1]$ to $[0, 1]$ such that $W_1(\varphi_1(x), \varphi_1(y)) = W_2(\varphi_2(x), \varphi_2(y))$ for almost every pair $(x, y) \in [0, 1]^2$. If two graphons W_1 and W_2 are weakly isomorphic, then $d(H, W_1) = d(H, W_2)$ for every graph H. The converse is also true [49]: if two graphons W_1 and W_2 satisfy that $d(H, W_1) = d(H, W_2)$ for every graph H, then W_1 and W_2 are weakly isomorphic. Hence, the limit of a convergent sequence of graphs is unique up to weak isomorphism. We finish with giving a formal definition of a finitely forcible graphon: a graphon W is *finitely forcible* if there exist graphs $H_1, \ldots, H_k$ such that if a graphon W' satisfies that $d(H_i, W') = d(H_i, W)$ for $i = 1, \ldots, k$, then $d(H, W') = d(H, W)$ for every graph H.

8.2 Weak regular partitions

In this section, we recall some basic concepts related to weak regularity for graphs and graphons and cast the lower bound construction of Conlon and Fox from [72] in the language of graphons. Since we do not use any other type of regularity partition, we will just say "regular" instead of "weak regular" in what follows.

We start with defining the notion for graphs. If G is a graph and A and B two subsets of its vertices, let $e(A, B)$ be the number of edges uv such that $u \in A$ and $v \in B$. A partition of a vertex set $V(G)$ of a graph G into subsets $V_1, \ldots, V_k$ is said to be *ε-regular* if it holds that

$$\left| e(A, B) - \sum_{i,j\in[k]} \frac{e(V_i, V_j)}{|V_i||V_j|} |V_i \cap A||V_j \cap B| \right| \leq \varepsilon |V(G)|^2$$

for every two subsets A and B of $V(G)$. It is known that for every ε, there exists $k_0 \leq 2^{O(\varepsilon^{-2})}$ (which depends on ε only) such that every graph has an ε-regular partition with at most k_0 parts [96]. This dependence of k_0 on ε is best possible up to a constant factor in the exponent as shown by Conlon and Fox [72].

We now define the analogous notion for graphons. Let $W : [0,1]^2 \to [0,1]$ be a graphon. If A and B are two measurable subsets of $[0,1]$, then the *density* $d_W(A, B)$ *between* A *and* B is defined to be

$$d_W(A,B) = \int_{A\times B} W(x,y)\;\mathrm{d}x\;\mathrm{d}y\;.$$

We will omit W in the subscript if the graphon W is clear from the context. Note that it always holds that $d(A, B) \leq |A||B|$ where $|X|$ is the measure of a set X. We would like to mention that the density between A and B is often defined in a normalised way, i.e. it is defined to be $\frac{d(A,B)}{|A||B|}$, but this is not the case in this thesis.

A partition $[0, 1]$ into measurable non-null sets $U_1, \ldots, U_k$ is said to be ε-*regular* if it holds that

$$\left| d(A,B) - \sum_{i,j\in[k]} \frac{d(U_i,U_j)}{|U_i||U_j|} |U_i \cap A||U_j \cap B| \right| \leq \varepsilon$$

for every two measurable subsets A and B of $[0, 1]$. The upper bound proof translates directly from graphs to graphons and so we get that for every ε, there exists $k_0 \leq 2^{O\left(\varepsilon^{-2}\right)}$ such that every graphon has an ε-regular partition with at most k_0 parts. Likewise, the example of Conlon and Fox from [72] can be used to obtain a step-graphon W_ε such that every ε-regular partition of W_ε has at least $2^{\Omega\left(\varepsilon^{-2}\right)}$ parts. However, the construction is probabilistic and the description of W_ε is thus not explicit. Based on this construction, we will define an explicit graphon W^m_{CF} which has similar properties as W_ε for $\varepsilon \approx m^{-1/2}$. In fact, a W^m_{CF}-random graph of order $2^{\alpha m}$ for some α close to 0 is the graph constructed by Conlon and Fox in [72].

Fix an integer m. The graphon W^m_{CF} is a step graphon that consists of 2^m parts of equal size. Each of the parts is associated with a vector $u \in \{-1, +1\}^m$. The part of the graphon W^m_{CF} corresponding to vectors u and u' is constantly equal to $\operatorname{trunc}\left(\frac{1}{2} + \frac{\langle u,u'\rangle}{4m^{1/2}}\right)$ where $\operatorname{trunc}(x)$ is equal to x if $x \in [0, 1]$, it is equal to 0 if $x < 0$ and to 1 if $x > 1$. In other words, the operator $\operatorname{trunc}(\cdot)$ replaces values smaller than 0 or larger than 1 with 0 and 1, respectively. Observe that $d([0, 1], [0, 1]) = 1/2$ by symmetry. Using the Chernoff bound, one can show that the measure of the points (x, y) with $0 < W(x, y) < 1$ is at least $1 - 2e^{-2} > 1/2$.

It would be possible to relate the proof presented in [72] to arguments on regular partitions of graphons. However, the probabilistic nature of the construction would make this technical and obfuscate some simple ideas. Because of this, and to keep the chapter self-contained, we decided to present a direct proof following the lines of the reasoning given in [72].

Theorem 8.6. *If $m \geq 25$ and $\varepsilon < \frac{1}{2^{14}m^{1/2}}$, then every ε-regular partition of the graphon W_{CF}^m has at least $2^{m/4}$ parts.*

Proof. Fix an integer $m \geq 25$. Let V_i^- be the vertices of W_{CF}^m in the parts associated with vectors u whose ith coordinate equals -1. Similarly, V_i^+ are the vertices of W_{CF}^m in the parts associated with vectors u whose ith coordinate equals $+1$.

Suppose that W_{CF}^m has an ε-regular partition $U_1, \ldots, U_k$ with $\varepsilon < 2^{-14}m^{-1/2}$ and $k < 2^{m/4}$. We say that a part U_t is *small* if $|U_t| \leq 2^{-m/3}$. Note that the sum of the measures of the small parts is at most $k \cdot 2^{-m/3} \leq 1/2$. For every $t \in [k]$, set

$$S_t = \sum_{i \in [m]} |U_t \cap V_i^-| \cdot |U_t \cap V_i^+| \; . \tag{8.7}$$

If $v \in \{-1, +1\}^m$, then the number of vectors $v' \in \{-1, +1\}^m$ such that v and v' differ in at most $m/16$ coordinates is at most $2^{m - \frac{49}{128}m} \leq 2^{2m/3-1}$ using the Chernoff bound and the fact that $m \geq 25$. Hence, for every $v \in \{-1, +1\}^m$, the measure of the vertices in the parts associated with vectors that differ from v in at most $m/16$ coordinates is at most $2^{-m/3}/2$. Consequently, if U_t is not small, then each vertex of U_t contributes to the sum (8.7) by at least $\frac{m}{16}\left(|U_t| - 2^{-m/3}/2\right) \geq m|U_t|/32$. We conclude that $S_t \geq |U_t|^2 m/32$ if U_t is not small.

We say that the pair $(i,t) \in [m] \times [k]$ is *useful* if $\min\{|U_t \cap V_i^-|, |U_t \cap V_i^+|\} \geq |U_t|/64$. Let M_t, $t \in [k]$, be the number of indices $i \in [m]$ such that the pair (i,t) is useful. Since each term in the sum (8.7) is at most $|U_t|^2/4$ and it is at most $|U_t|^2/64$ if (i,t) is not useful, it follows that $S_t \leq |U_t|^2(M_t/4 + m/64)$. We conclude that $M_t \geq m/16$ unless U_t is small (recall that $S_t \geq |U_t|^2 m/32$ if U_t is not small). Since the sum of the measures of the parts that are not small is at least $1/2$, we obtain that

$$\sum_{t \in [k]} M_t |U_t| \geq m/32 \; . \tag{8.8}$$

In particular, there exists $i_0 \in [m]$ such that the sum of the measure of parts U_t such that the pair (i_0, t) is useful is at least $1/32$. Fix such an index i_0 for the rest of the proof.

Let A^- be any measurable subset of $V_{i_0}^-$ such that $|A^- \cap U_t| = |U_t|/64$ if (i_0, t) is useful, and $|A^- \cap U_t| = 0$ otherwise. Similarly, let A^+ be any measurable subset of $V_{i_0}^+$ such that $|A^+ \cap U_t| = |U_t|/64$ if (i_0, t) is useful, and $|A^+ \cap U_t| = 0$ otherwise. Such sets A^- and A^+ exist because $\min\{|U_t \cap V_{i_0}^-|, |U_t \cap V_{i_0}^+|\} \geq |U_t|/64$ for every t such that (i_0, t) is useful. Note that the sets A^- and A^+ have the same measure and the choice of i_0 implies this measure is at least $1/2048 = 2^{-11}$.

Let $B = V_{i_0}^-$. Since the partition $U_1, \dots, U_k$ is ε-regular, we get that

$$\left| d(A^+, B) - \sum_{i,j \in [k]} \frac{d(U_i, U_j)}{|U_i||U_j|} |U_i \cap A^+||U_j \cap B| \right| \leq \varepsilon$$

and that

$$\left| d(A^-, B) - \sum_{i,j \in [k]} \frac{d(U_i, U_j)}{|U_i||U_j|} |U_i \cap A^-||U_j \cap B| \right| \leq \varepsilon.$$

Since $|A^- \cap U_t| = |A^+ \cap U_t|$ for every $t \in [k]$, we infer that $|d(A^-, B) - d(A^+, B)| \leq 2\varepsilon < 2^{-13} m^{-1/2}$. On the other hand, the choices of A^-, A^+ and B imply that $d(A^-, B) = (1/2 + m^{-1/2}/4)|A^-||B|$ and $d(A^+, B) = (1/2 - m^{-1/2}/4)|A^+||B|$. In particular, it holds that

$$|d(A^-, B) - d(A^+, B)| = \frac{m^{-1/2}(|A^-| + |A^+|)|B|}{4} \geq 2^{-13} \cdot m^{-1/2} .$$

This contradicts the fact that $U_1, \dots, U_k$ is an ε-regular partition of W_{CF}^m with $\varepsilon < 2^{-14} m^{-1/2}$ and $k < 2^{m/4}$. $\square$

We finish this section by sketching a short argument explaining why Proposition 8.5 is true. The proof of the weak regularity lemma in [96] is based on iterative refinements of a partition of a graph(on), at each step doubling the number of parts and increasing the "mean square density" by at least ε^2 until the partition becomes weakly ε-regular. Suppose that there exists a graphon W and $\varepsilon_i \to 0$ as in the proposition. We can assume that $\varepsilon_{i+1} \leq \varepsilon_i/2$ for every $i \in \mathbb{N}$. We start with a trivial partition into a single part and keep refining it until it becomes ε_1-regular; let k_1 be the number of steps made. We then continue with refining until it becomes ε_2-regular and let $k_2 \geq k_1$ be the total number of steps made till this point. We continue this procedure and define k_i, $i \geq 3$, in the analogous way. Setting $k_0 = 0$, we conclude that the mean square density after k_m steps is at least

$$\sum_{i=1}^{m} (k_i - k_{i-1})\varepsilon_i^2 > \sum_{i=1}^{m} k_i \left(\varepsilon_i^2 - \varepsilon_{i+1}^2\right) \geq \sum_{i=1}^{m} \frac{3}{4} k_i \varepsilon_i^2 \, .$$

However, each k_i must be at least $c\varepsilon_i^{-2}$ by the assumption of the proposition (otherwise, W would have a weak ε_i-regular partition with fewer than $2^{c\varepsilon_i^{-2}}$ parts). Consequently, after $m = \left\lceil \frac{4}{3c} \right\rceil + 1$ steps, the mean square density exceeds one, which is impossible.

8.3 Definition of the Švejk graphon

8.3.1 Some ideas behind the construction

At first sight, it might seem natural to consider the limit of the sequence of the graphons W_{CF}^m, $m \in \mathbb{N}$, as a candidate for a finitely forcible graphon with no regular partitions with few parts. It can be shown that the sequence of W_{CF}^m, $m \in \mathbb{N}$, is convergent, however, its limit is (somewhat surprisingly) the graphon that maps every $(x, y) \in [0, 1]^2$ to $1/2$. Clearly this graphon has an ε-regular partition with one part for every $\varepsilon > 0$.

Another natural candidate is the graphon $W_{\mathrm{CF}}^{(m_i)_{i=1}^{\infty}}$ consisting of a disjoint union of copies of the graphons $W_{\mathrm{CF}}^{m_1}, W_{\mathrm{CF}}^{m_2}, \ldots$ where $(m_i)_{i=1}^{\infty}$ is a sequence tending to infinity and the vertex set of the copy of $W_{\mathrm{CF}}^{m_i}$ is the interval $\left[\frac{2^{i-1}-1}{2^{i-1}}, \frac{2^i-1}{2^i}\right) \subseteq [0, 1]$. Provided that $(m_i)_{i=0}^{\infty}$ increases sufficiently rapidly, one can use Theorem 8.6 to show that, for infinitely many ε, this graphon does not have an ε-regular partition with few parts. The problem with this construction, however, is that it is unlikely to be finitely forcible. Strong evidence for this is given by a result of [105] which says that the much simpler graphon consisting of a disjoint union of cliques (i.e. constant 1 graphons) where the vertex set of the ith clique is the interval $\left[\frac{2^{i-1}-1}{2^{i-1}}, \frac{2^i-1}{2^i}\right)$ is not finitely forcible.

A remarkable fact, used in [107, 108], is that it can sometimes be possible to construct a finitely forcible graphon W containing (up to weak isomorphism) a copy of a non-finitely forcible graphon W' on a positive proportion of $[0, 1]$. For instance, the finitely forcible graphons constructed in [107, 108] and in this thesis all contain the non-finitely forcible graphon consisting of a disjoint union of cliques described in the previous paragraph. Our goal in this chapter is to construct a finitely forcible graphon containing the graphon $W_{\mathrm{CF}}^{(m_i)_{i=1}^{\infty}}$ for some rapidly increasing sequence $(m_i)_{i=1}^{\infty}$ on a positive proportion of $[0, 1]$. Again, by Theorem 8.6, for any such graphon there are infinitely many ε for which it does not have an ε-regular partition with few parts (see Proposition 8.9).

8.3.2 Formal definition

The Švejk graphon W_S has ten parts A, B, C, D, E, F, G, P, Q and R. The most important part is $E \times E$ which consists of disjoint copies of the Conlon–Fox graphons, as described in the previous subsection. For simplicity, we will define the graphon W_S as a function W_{13} from $[0, 13) \times [0, 13)$ to $[0, 1]$, and we set $W_S(x, y) = W_{13}(13x, 13y)$. All parts of W_{13} except for Q have measure one and we associate each of them with the unit interval $[0, 1)$, i.e. we view the points of those parts as points in $[0, 1)$. The remaining part Q is associated with $[0, 4)$.

Before formally defining the graphon, we need to introduce a fair bit of notation. We start with defining a tower function $t(n) : \mathbb{N} \to \mathbb{N}$ as follows:

$$t(n) = \begin{cases} 1 & \text{if } n = 0, \text{ and} \\ 2^{t(n-1)} & \text{otherwise.} \end{cases}$$

Note that $t(0) = 1$, $t(1) = 2$, $t(2) = 2^2 = 4$, $t(3) = 2^{2^2} = 16$, $t(4) = 2^{2^{2^2}} = 65536$, etc.

The notation that we define next is summarised in Table 8.1. For $x \in [0, 1)$, we define $[x]_1$ to be the smallest integer k such that $x < 1 - 2^{-k}$. In particular, $[x]_1 = 1$ iff $x \in [0, 1/2)$, $[x]_1 = 2$ iff $x \in [1/2, 3/4)$, $[x]_1 = 3$ iff $x \in [3/4, 7/8)$, etc. This allows us to view the interval $[0, 1)$ as split into *segments* $[0, 1/2)$, $[1/2, 3/4)$, $[3/4, 7/8)$, etc. and $[x]_1$ is the index of the segment containing x (numbered from one). We then define $[\![x]\!]_1$ to be $(x + 2^{1-[x]_1} - 1) \cdot 2^{[x]_1}$, i.e. $[\![x]\!]_1$ is the position of x in the $[x]_1$th segment if the $[x]_1$th segment is scaled to the unit interval.

Next, we let $[x]_2$ equal $\lfloor [\![x]\!]_1 \cdot t([x]_1) \rfloor$. In other words, if the $[x]_1$th segment of $[0, 1)$ is divided into $t([x]_1)$ parts of the same length $2^{-[x]_1}/t([x]_1)$, then $[x]_2$ is the index of the part containing x if the parts are numbered from 0. We refer to these parts of the segments as *subsegments*. Analogously, we let $[x]_3$ be the index of the

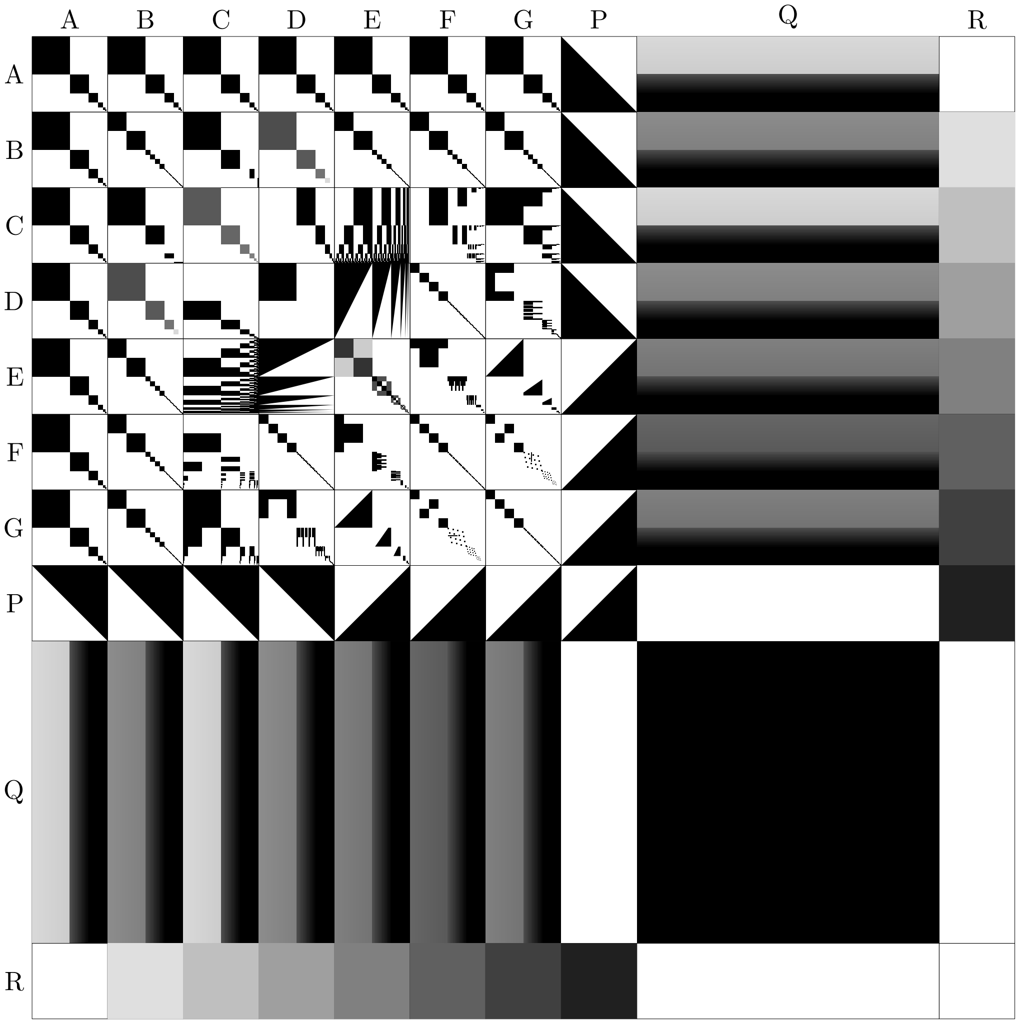

Figure 8.1: The Švejk graphon.

$[x]_1$	index of the segment $[0,1/2)$, $[1/2,3/4)$, $[3/4,7/8)$, etc, containing x
$[x]_2$	index of the subsegment of the $[x]_1$th segment containing x (the segment is divided into $t\left([x]_1\right)$ subsegments)
$[x]_3$	index of the part of the $[x]_2$th subsegment of the $[x]_1$th segment containing x (the subsegment is divided into $t\left([x]_1\right)$ parts)
$[x]_{2,3}$	$[x]_3 + t\left([x]_1\right)[x]_2$
$[\![x]\!]_1$	the position of x in its segment scaled to be in $[0,1)$
$[\![x]\!]_{1,i}^{01}$	the ith bit of the binary representation of $[\![x]\!]_1$
$[\![x]\!]_1^{\pm 1}$	a vector in $\{\pm 1\}^{t([x]_1-1)}$ whose ith coordinate is $2[\![x]\!]_{1,i}^{01} - 1$
$[x]_{j,i}^{01}$	the ith bit of the binary representation of $[x]_j$
$[x]_j^{\pm 1}$	a vector in $\{\pm 1\}^{t([x]_1-1)}$ whose $(i+1)$th coordinate is $2[x]_{j,i}^{01} - 1$

Table 8.1: The notation used in the definition of the Švejk graphon.

part containing x when the $[x]_2$th subsegment of the $[x]_1$th segment is divided into $t([x]_1)$ parts of length $2^{-[x]_1}/t([x]_1)^2$, where the parts are numbered from 0. Define $[x]_{2,3}$ to be $[x]_3+t([x]_1)[x]_2$. Note that $[x]_{2,3}$ can also be viewed as the part containing x when the $[x]_1$th segment is divided into $t([x]_1)^2$ parts of length $2^{-[x]_1}/t([x]_1)^2$, and that $[x]_3$ is equal to $[x]_{2,3}$ reduced modulo $t([x]_1)$.

For $i \geq 1$, let $[\![x]\!]^{01}_{1,i}$ denote the ith bit in the binary representation of $[\![x]\!]_1$. For example, if $[\![x]\!]_1 = 0.375 = .011$ in binary, then $[\![x]\!]^{01}_{1,1} = 0$, $[\![x]\!]^{01}_{1,2} = [\![x]\!]^{01}_{1,3} = 1$ and $[\![x]\!]^{01}_{1,i} = 0$ for $i \geq 4$. We let $[\![x]\!]^{\pm 1}_1$ denote the vector in $\{\pm 1\}^{t([x]_1-1)}$ whose ith coordinate is equal to $2[\![x]\!]^{01}_{1,i} - 1$ (i.e. 1 is mapped to $+1$ and 0 is mapped to -1).

For $j \in \{2,3\}$ and $i \geq 0$, let $[x]^{01}_{j,i}$ denote the ith bit in the binary representation of $[x]_j$. For example, if $[x]_j = 5 = 2^0 + 2^2$, then $[x]^{01}_{j,0} = [x]^{01}_{j,2} = 1$, $[x]^{01}_{j,1} = 0$ and $[x]^{01}_{j,i} = 0$ for $i \geq 3$. We let $[x]^{\pm 1}_j$ denote the vector in $\{\pm 1\}^{t([x]_1-1)}$ whose $(i+1)$th coordinate is equal to $2[x]^{01}_{j,i} - 1$.

A pair (x,y) belongs to	The value of $W_{13}(x,y)$ is 1 if and only if
$A \times (A \cup B \cup \cdots \cup G)$	$[x]_1 = [y]_1$
$B \times (B \cup E \cup F \cup G)$	$[x]_1 = [y]_1$ and $[x]_2 = [y]_2$
$B \times C$	$t([x]_1 - 1) = [y]_1$
$D \times C$	$[x]_1 = [y]_1 + 1$
$D \times D$	$[x]_1 = [y]_1 = 1$
$D \times G$	$[x]_1 = [y]_1$ and $[\![y]\!]_1 \leq \operatorname{trunc}\left(\frac{1}{2} + \frac{\left\langle [x]^{\pm1}_2, [x]^{\pm1}_3 \right\rangle}{4t([x]_1-1)^{1/2}}\right)$
$E \times C$	$[\![x]\!]^{01}_{1,[y]_1} = 1$
$E \times D$	$y \leq 1 - [\![x]\!]_1$
$F \times C$	$[y]_1 \leq t([x]_1 - 1)$, $[\![x]\!]^{01}_{1,[y]_1} = 1$ and $[\![y]\!]_1 \leq t([x]_1)^{-1} 2^{[y]_1}$
$F \times E$	$[x]_1 = [y]_1$ and $[\![y]\!]_1 \leq \frac{1}{2} - \frac{\left\langle [x]^{\pm1}_2, [x]^{\pm1}_3 \right\rangle}{4t([x]_1-1)}$
$F \times (D \cup F)$ or $G \times G$	$[x]_{2,3} = [y]_{2,3}$
$F \times G$	$[x]_3 = [y]_2$
$G \times C$	$[y]_1 \leq t([x]_1 - 1)$ and $[\![y]\!]_1 \leq t([x]_1)^{-1} 2^{[y]_1}$
$G \times E$	$[x]_1 = [y]_1$ and $1 - [\![x]\!]_1 \leq \operatorname{trunc}\left(\frac{1}{2} + t([x]_1 - 1)^{1/2}\left([\![y]\!]_1 - \frac{1}{2}\right)\right)$
$P \times (A \cup B \cup C \cup D)$	$x \leq y$
$P \times (E \cup F \cup G \cup P)$	$x \geq 1 - y$

Table 8.2: The definition of the Švejk graphon on $(A \cup \cdots \cup G \cup P)^2$ except on C^2, E^2, $B \times D$ and $D \times B$.

We will first define the values of the graphon W_{13} between the pairs of the parts not involving Q and R. The graphon W_{13} has values zero and one on $(A\cup\dots\cup G\cup P)^2$ except on C^2, E^2, $B\times D$ and $D\times B$. Table 8.2 determines the values of W_{13} in this zero-one case. The values of W_{13} on C^2, E^2 and $B\times D$ (by symmetry, this also determines the values on $D\times B$) are defined as follows.

$$W_{13}(x,y)=\begin{cases} 2^{-2^{[x]_1-1}} & \text{if } [x]_1=[y]_1\text{, and} \\ 0 & \text{otherwise,}\end{cases} \quad \text{for } (x,y)\in C^2,$$

$$W_{13}(x,y)=\begin{cases} W_{\mathrm{CF}}^{t([x]_1-1)}(\llbracket x\rrbracket_1,\llbracket y\rrbracket_1) & \text{if } [x]_1=[y]_1\text{, and} \\ 0 & \text{otherwise,}\end{cases} \quad \text{for } (x,y)\in E^2\text{, and}$$

$$W_{13}(x,y)=\begin{cases} t([x]_1)^{-1} & \text{if } [x]_1=[y]_1\text{, and} \\ 0 & \text{otherwise,}\end{cases} \quad \text{for } (x,y)\in B\times D.$$

We have defined the values of the graphon W_{13} on $(A\cup\dots\cup G\cup P)^2$, i.e. between all pairs of its parts not involving Q and R.

The part Q is used to equalise degrees of the vertices in the parts $A,\dots,G,P$ (see Section 8.4 for the definition of the degree of a vertex in a graphon). If $x\in A\cup\dots\cup G\cup P=\overline{Q\cup R}$ and $y\in Q$, then

$$W_{13}(x,y)=\frac{1}{4}\left(4-\int\limits_{\overline{Q\cup R}} W_{13}(x,z)\,\mathrm{d}z\right).$$

It is straightforward to verify that $W_{13}(x,y)\in[0,1]$ for every $(x,y)\in\overline{Q\cup R}\times Q$.

The part R distinguishes the parts by vertex degrees. If $y\in R$, then

$$W_{13}(x,y)=\begin{cases} 1/8 & \text{if } x\in B, \\ 2/8 & \text{if } x\in C, \\ 3/8 & \text{if } x\in D, \\ 4/8 & \text{if } x\in E, \\ 5/8 & \text{if } x\in F, \\ 6/8 & \text{if } x\in G, \\ 7/8 & \text{if } x\in P\text{, and} \\ 0 & \text{otherwise.}\end{cases}$$

Finally, the graphon W_{13} is equal to 1 on $Q\times Q$.

The vertices in each of the ten parts of the Švejk graphon have the same degree. This degree is given in Table 8.3. We have not computed the degree of the vertices in

Part	A	B	C	D	E	F	G	P	Q	R
Degree	$\frac{32}{104}$	$\frac{33}{104}$	$\frac{34}{104}$	$\frac{35}{104}$	$\frac{36}{104}$	$\frac{37}{104}$	$\frac{38}{104}$	$\frac{39}{104}$	$\geq \frac{40}{104}$	$\frac{28}{104}$

Table 8.3: The degrees of the vertices in each part of the Švejk graphon.

the part Q exactly since it is enough to establish that this degree is larger than (and thus distinct from) the degrees of the vertices in the other parts.

We finish this section by establishing that the Švejk graphon has no weak regular partitions with few parts.

Proposition 8.9. *The Švejk graphon W_S has no weak ε-regular partition with fewer than $2^{t(n)/4}$ parts if $\varepsilon < \frac{1}{2^{24+2n}t(n)^{1/2}}$ and $n \geq 4$. In particular, there exists a sequence of positive reals ε_i tending to 0 such that every weak ε_i-regular partition of W_S has at least $2^{\Omega\left(\varepsilon_i^{-2}/2^{5\log^* \varepsilon_i^{-2}}\right)}$ parts.*

Proof. The graphon W_S contains a copy $W_{\mathrm{CF}}^{t(n)}$ scaled by $2^{-n-1}/13$ for every $n \in \mathbb{N}$. Note that a weak ε-regular partition of W_S yields a weak $(\varepsilon 2^{-2n}/676)$-regular partition of $W_{\mathrm{CF}}^{t(n)}$ with fewer or the same number of parts. It follows that W_S cannot have a weak ε-regular partition with fewer than $2^{t(n)/4}$ parts for $\varepsilon < \frac{1}{676 \cdot 2^{14+2n} \cdot t(n)^{1/2}}$ and $n \geq 4$ by Theorem 8.6.

Setting $\varepsilon_i = \frac{1}{2^{25+2i}t(i)^{1/2}}$, we obtain the desired sequence of ε_i's. Note that

$$\lim_{i\to\infty} \frac{\log^* \varepsilon_i^{-2}}{i} = \lim_{i\to\infty} \frac{\log^* \left(2^{4i+50}t(i)\right)}{i} = 1$$

and so $\frac{t(i)}{4} \in \Omega\left(\varepsilon_i^{-2}/2^{5\log^* \varepsilon_i^{-2}}\right)$ as desired. □

8.4 Constraints

The proof that the Švejk graphon is finitely forcible uses the notion of decorated constraints, which was introduced in [108] and further developed in [107]. We now present the notion following the lines of [107].

A *constraint* is an equality between two density expressions where a *density expression* is recursively defined as follows: a real number or a graph H are density expressions, and if D_1 and D_2 are two density expressions, then the sum $D_1 + D_2$ and the product $D_1 \cdot D_2$ are also density expressions. The value of the density expression for a graphon W is the value obtained by substituting for each graph H its density in W.

As observed in [108], if W is the unique graphon (up to weak isomorphism) that satisfies a finite set $\mathcal{C}$ of constraints, then it is finitely forcible. In particular, W is the unique graphon with densities of subgraphs appearing in $\mathcal{C}$ equal to their densities in W. Hence, a possible way of establishing that a graphon W is finitely forcible is providing a finite set of constraints $\mathcal{C}$ such that the graphon W is the unique graphon up to weak isomorphism that satisfies these constraints.

If W is a graphon, then the points of $[0,1]$ can be viewed as vertices and we can also speak of the *degree* of a vertex $x \in [0,1]$, defined as

$$\deg_W(x) = \int_{[0,1]} W(x,y)\,\mathrm{d}y\,.$$

Note that the degree is well-defined for almost every vertex of W. We will omit the superscript W when the graphon is clear from the context.

A graphon W is *partitioned* if there exist $k \in \mathbb{N}$ and positive reals $a_1,\dots,a_k$ summing to one and distinct reals $d_1,\dots,d_k$ between 0 and 1 such that the set of vertices of W with degree d_i (referred to as a *part* of the partitioned graphon) has measure a_i. The following lemma was proved in [108].

Lemma 8.10. *Let $a_1,\dots,a_k$ be positive reals summing to one and let $d_1,\dots,d_k$ be distinct reals between 0 and 1. There exists a finite set of constraints $\mathcal{C}$ such that any graphon W satisfying $\mathcal{C}$ must be a partitioned graphon with parts of sizes $a_1,\dots,a_k$ and degrees $d_1,\dots,d_k$.*

We now introduce a stronger type of constraints, which was also used in [107, 108]. We will refer to the constraints introduced earlier as *ordinary constraints* if a distinction needs to be made. Suppose that W is a partitioned graphon with parts $A_i \subseteq [0,1]$, $i \in [k]$, where the part A_i has measure a_i and it contains vertices of degrees d_i. A *decorated graph* is a graph with some vertices distinguished as *roots* and each vertex labeled with one of the parts $A_1,\dots,A_k$; the roots of a decorated graph come with a fixed order. Two decorated graphs are *compatible* if the subgraphs induced by their roots are isomorphic through the isomorphism mapping the ith root of one of them to the ith root of the other; this isomorphism must preserve both the vertex labels and the (non-)edges between the vertices. In particular, two compatible decorated graphs have the same number of roots. A *decorated constraint* is a constraint where all graphs appearing in the density expressions are compatible decorated graphs. Note that decorated graphs and constraints are always defined with a particular type of a partition of a graphon (i.e. names of the parts) in mind.

We now define when a graphon W satisfies a decorated constraint. Fix a decorated constraint C. Let H_0 be the decorated graph induced by the roots of the decorated graphs appearing in the constraint C. The graph H_0 is well-defined since all decorated graphs appearing in C are compatible. All the vertices of H_0 are roots; let n_0 be the number of these vertices. If the probability that a W-random graph of order n_0 is isomorphic to H_0 by mapping its ith vertex to the ith root of H_0 (the isomorphism must preserve both (non-)edges and the vertex labels, i.e. the ith vertex must belong to the part that the ith root is labeled with) is zero, then the decorated constraint C is satisfied. If C is satisfied in this way, we will say that C is *null-satisfied*.

The decorated constraint C that is not null-satisfied is *satisfied* if the following holds for almost every choice of n_0 vertices in $[0, 1]$ such that the ith vertex belongs to the part that the ith root of H_0 is labelled with and the ith and jth vertices are adjacent if and only if the ith and jth roots of H_0 are. Each decorated graph is evaluated to the probability that the random choice of its non-rooted vertices conditioned on each of them chosen from a part that it is labeled with gives the decorated graph (here, we do not allow a permutation of vertices to make the randomly chosen graph isomorphic to the decorated graph). Both sides of the constraint C are required to evaluate to the same value with probability one with respect to a random choice of the roots.

The following lemma was proved in [108], also see [107].

Lemma 8.11. *Let $k \in \mathbb{N}$, let $a_1, \ldots, a_k$ be positive real numbers summing to one, and let $d_1, \ldots, d_k$ be distinct reals between zero and one. If W is a partitioned graphon with k parts formed by vertices of degree d_i and measure a_i each, then any decorated constraint can be expressed as a single ordinary constraint, i.e. W satisfies the decorated constraint if and only if it satisfies the ordinary constraint.*

By Lemma 8.11, we can equivalently work with (formally stronger) decorated constraints instead of ordinary constraints.

It is useful to fix some notation for visualising decorated constraints. We write the decorated constrains as expressions involving decorated graphs where the roots are depicted by squares and non-root vertices by circles, and each vertex is labeled with the name of the respective part of a graphon. The solid lines connecting vertices correspond to the edges and dashed lines to the non-edges. No connection between two vertices means that both edge or non-edge are allowed between the vertices, i.e. the picture should be interpreted as the sum of two graphs, one with an edge and with a non-edge. If more than a single pair of vertices is not joined, the picture should be interpreted as the multiple sum over all non-joined pairs of vertices. To avoid possible ambiguity, the drawing of the graph on the roots is identical for all decorated graphs in each constraint, which makes clear which roots correspond to each other.

We finish this section with the following lemma, which is an easy corollary of Lemma 8.11. In essence, it says that if a graphon W_0 can be finitely forced in its own right, then it can be forced on a single part of a partitioned graphon W without affecting the structure of the other parts.

Lemma 8.12. *Let W_0 be a finitely forcible graphon, let $a_1, \ldots, a_k$ be positive reals summing to one and let $d_1, \ldots, d_k$ be distinct reals between zero and one. Then there exists a finite set $\mathcal{C}$ of decorated constraints such that a partitioned graphon W with k parts formed by vertices of degree d_i and measure a_i each satisfies $\mathcal{C}$ if and only if the subgraphon of W induced by the mth part is weakly isomorphic to W_0. In other words, if the mth part is denoted A_m, then W satisfies $\mathcal{C}$ if and only if there exist measure preserving maps $\varphi : [0, a_m] \to A_m$ and $\varphi_0 : [0,1] \to [0,1]$ such that $W(\varphi(xa_m), \varphi(ya_m)) = W_0(\varphi_0(x), \varphi_0(y))$ for almost every pair $(x, y) \in [0,1]^2$.*

Proof. Let $H_1, \ldots, H_\ell$ and $d_1, \ldots, d_\ell$ be the subgraphs and their densities such that W_0 is the unique graphon (up to weak isomorphism) with these densities. The set $\mathcal{C}$ is formed by ℓ decorated constraints: the left side of the ith constraint contains H_i with all its vertices labelled by A_m and the right side is d_i. If the subgraphon of W induced by A_m is weakly isomorphic to W_0, then clearly these constraints are satisfied. On the other hand, since W_0 is forced by setting the densities of H_i to d_i for every $i \in [\ell]$, the converse is true as well. $\square$

	A	B	C	D	E	F	G	P	Q	R
A	8.5.2	8.5.2	8.5.2	8.5.2	8.5.2	8.5.2	8.5.2	8.5.1	8.5.9	8.5.9
B	8.5.2	8.5.4	8.5.3	8.5.3	8.5.4	8.5.4	8.5.4	8.5.1	8.5.9	8.5.9
C	8.5.2	8.5.3	8.5.3	8.5.2	8.5.5	8.5.5	8.5.5	8.5.1	8.5.9	8.5.9
D	8.5.2	8.5.3	8.5.2	8.5.2	8.5.5	8.5.4	8.5.7	8.5.1	8.5.9	8.5.9
E	8.5.2	8.5.4	8.5.5	8.5.5	8.5.8	8.5.7	8.5.6	8.5.1	8.5.9	8.5.9
F	8.5.2	8.5.4	8.5.5	8.5.4	8.5.7	8.5.4	8.5.4	8.5.1	8.5.9	8.5.9
G	8.5.2	8.5.4	8.5.5	8.5.7	8.5.6	8.5.4	8.5.4	8.5.1	8.5.9	8.5.9
P	8.5.1	8.5.1	8.5.1	8.5.1	8.5.1	8.5.1	8.5.1	8.5.1	8.5.9	8.5.9
Q	8.5.9	8.5.9	8.5.9	8.5.9	8.5.9	8.5.9	8.5.9	8.5.9	8.5.9	8.5.9
R	8.5.9	8.5.9	8.5.9	8.5.9	8.5.9	8.5.9	8.5.9	8.5.9	8.5.9	8.5.9

Table 8.4: The subsections of Section 8.5 where the structure of the Švejk graphon between the corresponding pairs of parts is forced.

8.5 Finite forcibility of the Švejk graphon

Our final and longest section is devoted to proving that the Švejk graphon is finitely forcible. We will prove this by exhibiting a finite set of constraints that the Švejk graphon satisfies and showing that the Švejk graphon is the only graphon up to weak isomorphism that satisfies this set of constraints. By Lemma 8.10, there exists a finite set of constraints such that any graphon that satisfies them is a partitioned graphon with ten parts of the sizes as in the Švejk graphon and degrees of vertices in these parts as in Table 8.3. Hence, we can work with decorated constraints with vertices labeled by the parts $A, \dots, G$, P, Q and R (see Lemma 8.11). We will use decorated constraints to enforce the structure of the graphon between pairs of its parts, one pair after another, often building on the structure enforced by earlier constraints. Table 8.4 gives references to subsections where the structure between the particular pairs of parts is forced.

Fix a graphon W_0 that satisfies all the constraints presented in this section. In particular, W_0 satisfies the constraints given by Lemma 8.10 and it is a partitioned graphon with ten parts of the sizes as in the Švejk graphon and degrees of vertices in these parts as in Table 8.3. These ten parts of W_0 will be denoted by $A_0, \dots, G_0$, P_0, Q_0 and R_0 in correspondence with the parts of the Švejk graphon. We will show that W_0 is weakly isomorphic to the Švejk graphon.

8.5.1 Coordinate system

The half-graphon $W_\triangle$, i.e. the zero-one graphon defined by $W_\triangle(x, y) = 1$ iff $x + y \geq 1$, is finitely forcible [78], also see [159]. By Lemma 8.12, there exists a finite set of decorated constraints such that W_0 satisfies these constraints if and only if the subgraphon induced by the part P_0 is weakly isomorphic to the half-graphon $W_\triangle$. We insist that W_0 satisfies these constraints.

Let $X \in \{A, \ldots, G, P\}$. We use the symbol X_0 to refer to the corresponding element of $\{A_0, \ldots, G_0, P_0\}$. By the Monotone Reordering Theorem (see [153, Proposition A.19] for more details), there exist measure preserving maps $\varphi_X : X_0 \to [0, |X_0|)$ and non-decreasing functions $f_X : [0, |X_0|) \to [0, 1)$, such that

$$f_X(\varphi_X(x)) = 13 \int\limits_{P_0} W_0(x, z) \,\mathrm{d}z$$

for almost every $x \in X_0$. Since we already know that the subgraphon of W_0 induced by P is weakly isomorphic to $W_\triangle$, we must have $W_0(x, y) = 1$ for almost every pair $(x, y) \in P_0^2$ with $f_P(\varphi_P(x)) + f_P(\varphi_P(y)) \geq 1$, $W_0(x, y) = 0$ for almost every pair $(x, y) \in P_0^2$ with $f_P(\varphi_P(x)) + f_P(\varphi_P(y)) < 1$, and $f_P(z) = 13z$ for almost every $z \in [0, 1/13)$.

Set $g_X(x) = f_X(\varphi_X(x))$ for $x \in X_0$ and $X \in \{A, \ldots, G, P\}$. For completeness, let g_Q and g_R be any measurable maps from Q_0 and R_0 to $Q \cong [0, 4)$ and $R \cong [0, 1)$ such that for any measurable subset Z of $[0, 4)$ and $[0, 1)$, respectively, we have $|g_Q^{-1}(Z)| = |Z|/13$ and $|g_R^{-1}(Z)| = |Z|/13$, respectively. Each g_X can be viewed as a map from X_0 to the part X of the graphon W_{13}. The maps $g_A, \ldots, g_G$, g_P, g_Q and g_R constitute a map g from the vertices of W_0 to the vertices of W_{13} and so to those of W_S.

We will argue that the map g as a map from the vertices W_0 to the vertices of W_S is measure preserving and we will show that $W_0(x, y) = W_S(g(x), g(y))$ for almost every pair $(x, y) \in [0, 1]^2$. This will prove that the graphons W_0 and W_S are weakly isomorphic. So far, we have established that $W_0(p, p') = W_S(g(p), g(p'))$ for almost every pair $(p, p') \in P_0^2$ and the map g is measure preserving when restricted to $P_0 \cup Q_0 \cup R_0$.

Let us consider the decorated constraints depicted in Figure 8.2. Let $N_Y(x) = \{y \in Y_0 \mid W_0(x, y) = 1\}$ for $x \in P_0$ and $Y \in \{A, \ldots, G\}$. The first constraint implies that the graphon W_0 is zero-one valued almost everywhere on $P_0 \times (A_0 \cup \cdots \cup G_0)$

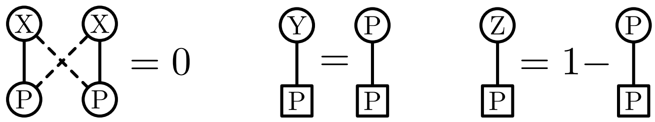

Figure 8.2: Decorated constraints used in Subsection 8.5.1 where X is one of A, B, C, D, E, F or G, Y is E, F or G and Z is A, B, C or D.

and that $N_Y(x) \setminus N_Y(x')$ or $N_Y(x') \setminus N_Y(x)$ has measure zero for almost every pair $(x, x') \in P_0^2$ and for every $Y \in \{A, \ldots, G\}$. The second constraint in Figure 8.2 implies for $Y \in \{E, F, G\}$ that the measure of $N_Y(x)$ is $g_P(x)$ for almost every $x \in P_0$. Hence, it must hold that $f_Y(y) = 13y$ for $y \in [0, 1/13)$ and $W_0(x, y) = 1$ for almost every $(x, y) \in P_0 \times Y_0$ with $g_P(x) + g_Y(y) \geq 1$, where $Y \in \{E, F, G\}$. Similarly, the third constraint in Figure 8.2 implies for $Y \in \{A, B, C, D\}$ that the measure of $N_Y(x)$ is $1 - g_P(x)$ for almost every $x \in P_0$. Consequently, it holds that $f_Y(y) = 13y$ for $y \in [0, 1/13)$ and $W_0(x, y) = 1$ for almost every $(x, y) \in P_0 \times Y_0$ with $g_P(x) \geq g_Y(x)$ for $Y \in \{A, B, C, D\}$. We conclude that g is a measure preserving map on the whole domain and $W_0(x, y) = W_S(g(x), g(y))$ for almost every pair $(x, y) \in P_0 \times \overline{(Q_0 \cup R_0)}$.

The values of the functions $g_A, \ldots, g_G$ can be understood to be the coordinates of the vertices in $A_0, \ldots, G_0$, respectively, and the coordinate of a vertex $x \in A_0 \cup \cdots \cup G_0$ is

$$g_Y(x) = 13 \int\limits_{P_0} W_0(x, z) \,\mathrm{d}z$$

for $Y \in \{A, \ldots, G\}$. This integral is easily expressible as a decorated density expression since it is just the relative edge density (degree) of x to P_0. This view allows us to speak about segments and subsegments of the parts $A_0, \ldots, G_0$. The kth segment of X_0, $X \in \{A, \ldots, G\}$, is formed by those $x \in X_0$ such that $[g_X(x)]_1 = k$. Analogously, the values $[g_X(x)]_2$ determine the subsegments.

8.5.2 Segmenting

We now force that the parts $A_0, \ldots, G_0$ of W_0 are split into segments as in W_S. We also force the structure to recognise the first segment through the clique inside D_0^2 and to have the "successor" relation on the segments through the structure inside $C_0 \times D_0$. All three of these aims will be achieved by the decorated constraints given in Figure 8.3.

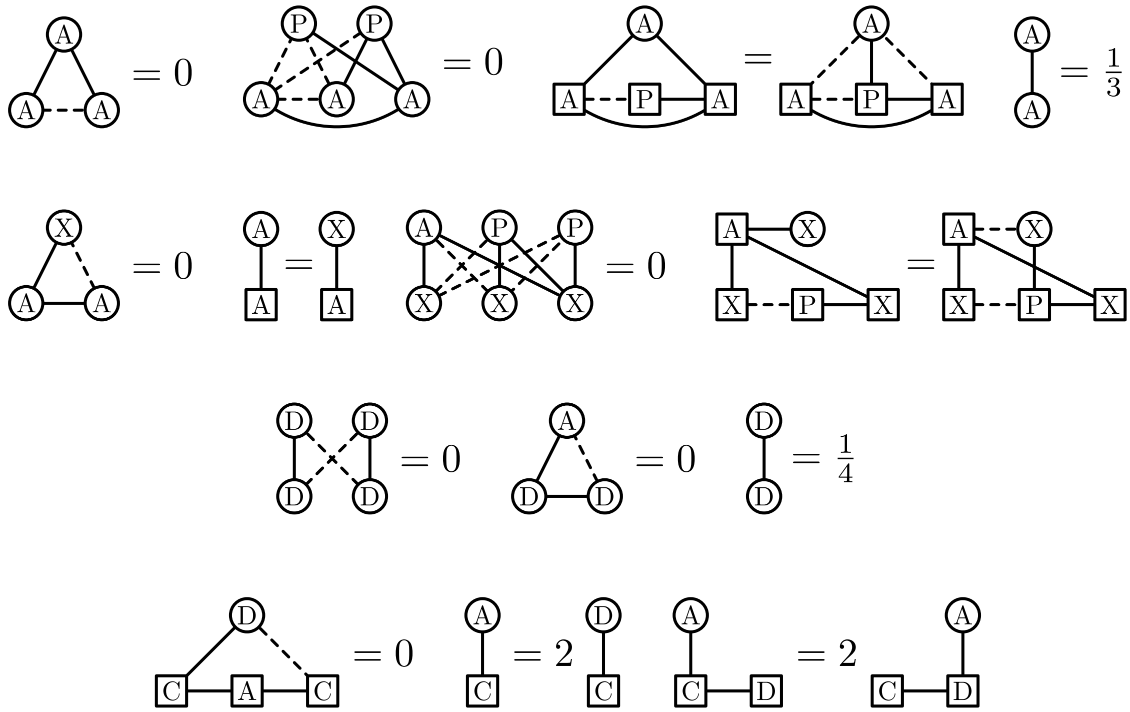

Figure 8.3: Decorated constraints used in Subsection 8.5.2 where X is one of B, C, D, E, F or G.

The four constraints on the first line in Figure 8.3 force the structure on A_0^2. The line of arguments in this subsection follows those given in [107, 108]. The first constraint implies that there exists a set $\mathcal{J}$ of disjoint measurable subsets of A_0 such that for almost every $x \in A_0$, there exists $J \in \mathcal{J}$ such that $W_0(x,y) = 1$ for almost every $x, y \in J$ and $W_0(x,y) = 0$ for almost every $x \in J$ and $y \notin J$. Hence, $W_0(x,y) = 1$ for almost every $(x,y) \in \bigcup_{J \in \mathcal{J}} J^2$ and $W_0(x,y) = 0$ for almost every $(x,y) \in A_0^2 \setminus \bigcup_{J \in \mathcal{J}} J^2$.

We claim that the second constraint together with the structure on $A_0 \times P_0$ yields that for every set $J \in \mathcal{J}$ there exists an open interval $J' \subseteq [0,1)$ such that J and $g_A^{-1}(J')$ differ on a set of measure zero. Note that such an open interval J' might be empty. Since we use an argument of this kind for the first time in this thesis, we give more details. If one of the (non-null) sets J did not have the property, then a random sampling of three points $x, x', x'' \in J \subseteq A_0$ with $g_A(x) < g_A(x') < g_A(x'')$ would satisfy $W_0(x,x') = 0$ and $W_0(x,x'') = 1$ with positive probability. For such three points, the probability of sampling the additional two points from P_0 is

$g_A(x') - g_A(x)$ and $g_A(x'') - g_A(x')$ and the triples of points x, x', x'' such that the differences $g_A(x') - g_A(x)$ and $g_A(x'') - g_A(x')$ would be bounded away from zero have positive measure. Let $\mathcal{J}'$ be the set of open intervals $J' \subseteq [0,1)$ such that $g_A^{-1}(J')$ and J differ on a set of measure zero for some $J \in \mathcal{J}$. Since the sets in $\mathcal{J}$ are disjoint and the sets in $\mathcal{J}'$ are open, the intervals of $\mathcal{J}'$ are also disjoint.

The third constraint implies that if $x, x' \in g_A^{-1}(J')$ for some $J' \in \mathcal{J}'$, then the measure of the interval J', assuming it is non-empty, and the measure of the interval $(\sup J', 1)$ are the same. Again, we provide a detailed justification since we use an argument of this kind for the first time. Almost every choice of the three roots $x \in A_0$, $x' \in P_0$ and $x'' \in A_0$ (the order follows that in the figure) satisfies that $g_A(x) < g_P(x') < g_A(x'')$ (because of the non-edge between x and x' and the edge between x' and x'') and that there exists $J' \in \mathcal{J}'$ such that $x, x'' \in g_A^{-1}(J')$ (because of the edge between x and x''). The left side is then equal to the measure of $g_A^{-1}(J')$, which is the measure of J'. The right side is equal to the measure of those $z \in A_0$ such that $z \notin g_A^{-1}(J')$ and $g_A(z) > g_P(x')$. Hence, the right side is equal to $1 - \sup J'$. Since this holds for almost every triple x, x' and x'', we conclude that the measure of each non-empty interval $J' \in \mathcal{J}'$ is $1 - \sup J'$. Consequently, each non-empty interval J' must be of the form $(1 - 2\alpha, 1 - \alpha)$ for some $\alpha \in [0,1)$. Since the intervals of $\mathcal{J}'$ are disjoint, there can only be a finite number of intervals to the left of each interval of $\mathcal{J}'$. This implies that the set $\mathcal{J}'$ is countable.

Finally, the last constraint on the first line yields that

$$\int\limits_{A_0^2} W(x,y)\,\mathrm{d}x\,\mathrm{d}y = \sum_{J' \in \mathcal{J}'} (\sup J' - \inf J')^2 = \frac{1}{3}\,.$$

However, this equality can hold only if the intervals contained in $\mathcal{J}'$ are exactly the intervals $\left(1 - 2^{1-k}, 1 - 2^{-k}\right)$, $k \in \mathbb{N}$. We conclude that $W_0(x,y)$ and $W_S(g(x), g(y))$ are equal for almost every pair $(x,y) \in A_0^2$.

We now analyse the four constraints on the second line in Figure 8.3. Fix $X \in \{B, \ldots, G\}$. The first constraint implies that for every $J \in \mathcal{J}$, there exists $Z_J \subseteq X_0$ such that $W_0(x,y) = 1$ for almost every $(x,y) \in J \times Z_J$ and $W_0(x,y) = 0$ for almost every $(x,y) \in J \times (X_0 \setminus Z_J)$. The second constraint yields that the measure of Z_J is the same as the measure of J. The third constraint implies that there exists an open interval Z'_J such that Z_J and $g_X^{-1}(Z'_J)$ differ on a set of measure zero. Finally, the last constraint on the second line yields that each of the intervals Z'_J is of the form $(1-2\alpha, 1-\alpha)$ for some $\alpha \in [0,1)$. Since the length of Z'_J is the same as the measure of J, we conclude that if $J = g_A^{-1}((1-2^{1-k}, 1-2^{-k}))$, then $Z'_J = (1-2^{1-k}, 1-2^{-k})$. Hence, $W_0(x,y)$ and $W_S(g(x), g(y))$ are equal for almost every pair $(x,y) \in A_0 \times X_0$, $X \in \{B, \ldots, G\}$.

Let us turn our attention to the three constraints on the third line in Figure 8.3. The first constraint implies that there exists a subset Z_D of D_0 such that $W_0(x,y) = 1$ for almost every $(x,y) \in Z_D^2$ and $W_0(x,y) = 0$ for almost every $(x,y) \in D_0^2 \setminus Z_D^2$. The second constraint yields that Z_D is a subset of Z_J for some $J \in \mathcal{J}$. Finally, the third constraint says that the square of the measure of Z_D is $1/4$, i.e. the measure of Z_D is $1/2$. However, this is only possible if Z_D and $g_D^{-1}((0,1/2))$ differ on a set of measure zero. We conclude that $W_0(x,y)$ and $W_S(g(x), g(y))$ are equal for almost every pair $(x,y) \in D_0^2$.

It remains to analyse the three constraints on the last line in Figure 8.3. The first constraint yields that for every $k \in \mathbb{N}$, there exists $Z_k \subseteq D_0$ such that $W(x,y) = 1$ for almost every $(x,y) \in g_C^{-1}((1-2^{1-k}, 1-2^{-k})) \times Z_k$ and $W(x,y) = 0$ for almost every $(x,y) \in g_C^{-1}((1-2^{1-k}, 1-2^{-k})) \times (D_0 \setminus Z_k)$. The second constraint yields that the measure of Z_k is 2^{-k-1}, and the third constraint that Z_k is a subset of $g_D^{-1}((1-2^{-k}, 1-2^{-k-1}))$ except for a set of measure zero. Hence, Z_k and $g_D^{-1}((1-2^{-k}, 1-2^{-k-1}))$ differ on a set of measure zero, and we can conclude that $W_0(x,y)$ and $W_S(g(x), g(y))$ are equal for almost every pair $(x,y) \in C_0 \times D_0$.

8.5.3 Tower function

In this subsection, we will force a representation of the tower inside $B_0 \times D_0$. This is achieved using the constraints depicted in Figure 8.4. Before analyzing these constraints, we give an analytic observation based on [159, proof of Lemma 3.3].

Lemma 8.13. *Let $F : [0,1)^2 \to [0,1)$ be a measurable function. If*

$$\int_{[0,1)} F(x,z)F(y,z)\ dz = C$$

for almost every $(x,y) \in [0,1)^2$, then

$$\int_{[0,1)} F(x,z)^2\ dz = C$$

for almost every $x \in [0,1)$.

The constraint on the first line in Figure 8.4 yields that

$$\left(\int_{C_0} W_0(x,z)W_0(y,z)\,\mathrm{d}z\right)^2$$

$$= \left(\int_{C_0} W_0(x',z)W_0(x'',z)\,\mathrm{d}z\right)\left(\int_{C_0} W_0(y',z)W_0(y'',z)\,\mathrm{d}z\right)$$

for almost every $x, x', x'' \in C_0$ and $y, y', y'' \in A_0$ all in the same segment (this is implied by the presence of the edges between the roots). By Lemma 8.13, we get that

$$\left(\int_{C_0} W_0(x,z)W_0(y,z)\,\mathrm{d}z\right)^2 = \left(\int_{C_0} W_0(x',z)^2\,\mathrm{d}z\right)\left(\int_{C_0} W_0(y',z)^2\,\mathrm{d}z\right)$$

for almost every $x, x' \in C_0$ and $y, y' \in A_0$ such that $[g_C(x)]_1 = [g_A(y)]_1$. Since the equality holds for almost every $x, x' \in C_0$ and $y, y' \in A_0$, it actually holds that

$$\left(\int_{C_0} W_0(x,z)W_0(y,z)\,\mathrm{d}z\right)^2 = \left(\int_{C_0} W_0(x,z)^2\,\mathrm{d}z\right)\left(\int_{C_0} W_0(y,z)^2\,\mathrm{d}z\right)$$

for almost every $x \in C_0$ and $y \in A_0$ such that $[g_C(x)]_1 = [g_A(y)]_1$. The Cauchy-Schwarz Inequality yields that there exist ξ_k, $k \in \mathbb{N}$, such that $W_0(x,z) = \xi_k \cdot W_0(y,z)$ for almost every $x \in C_0$, $y \in A_0$ and $z \in C_0$ with $[g_C(x)]_1 = [g_A(y)]_1 = k$. Hence, $W_0(x,z) = \xi_k$ for almost every $x \in C_0$ and $z \in C_0$ with $[g_C(x)]_1 = [g_C(z)]_1 = k$ and $W_0(x,z) = 0$ for almost every $x \in C_0$ and $z \in C_0$ with $[g_C(x)]_1 \neq [g_C(z)]_1$. Along the same line, the constraint on the second line implies that there exist ξ'_k, $k \in \mathbb{N}$, such that $W_0(x,z) = \xi'_k \cdot W_0(y,z)$ for almost every $x \in B_0$, $y \in A_0$ and $z \in D_0$ with $[g_B(x)]_1 = [g_A(y)]_1 = k$. Consequently, $W_0(x,z) = \xi'_k$ for almost every $x \in D_0$ and $z \in B_0$ with $[g_D(x)]_1 = [g_B(z)]_1 = k$ and $W_0(x,z) = 0$ for almost every $x \in D_0$ and $z \in B_0$ with $[g_D(x)]_1 \neq [g_B(z)]_1$.

Almost every choice of the roots in the first constraint on the third line satisfies that all the roots belong to the same segment and this segment must be the first segment because of the edge between the two roots from D_0. Hence, this constraint implies that $\xi_1 |g_C^{-1}((0,1/2))| = 1/4$, i.e. $\xi_1 = 1/2$ as desired.

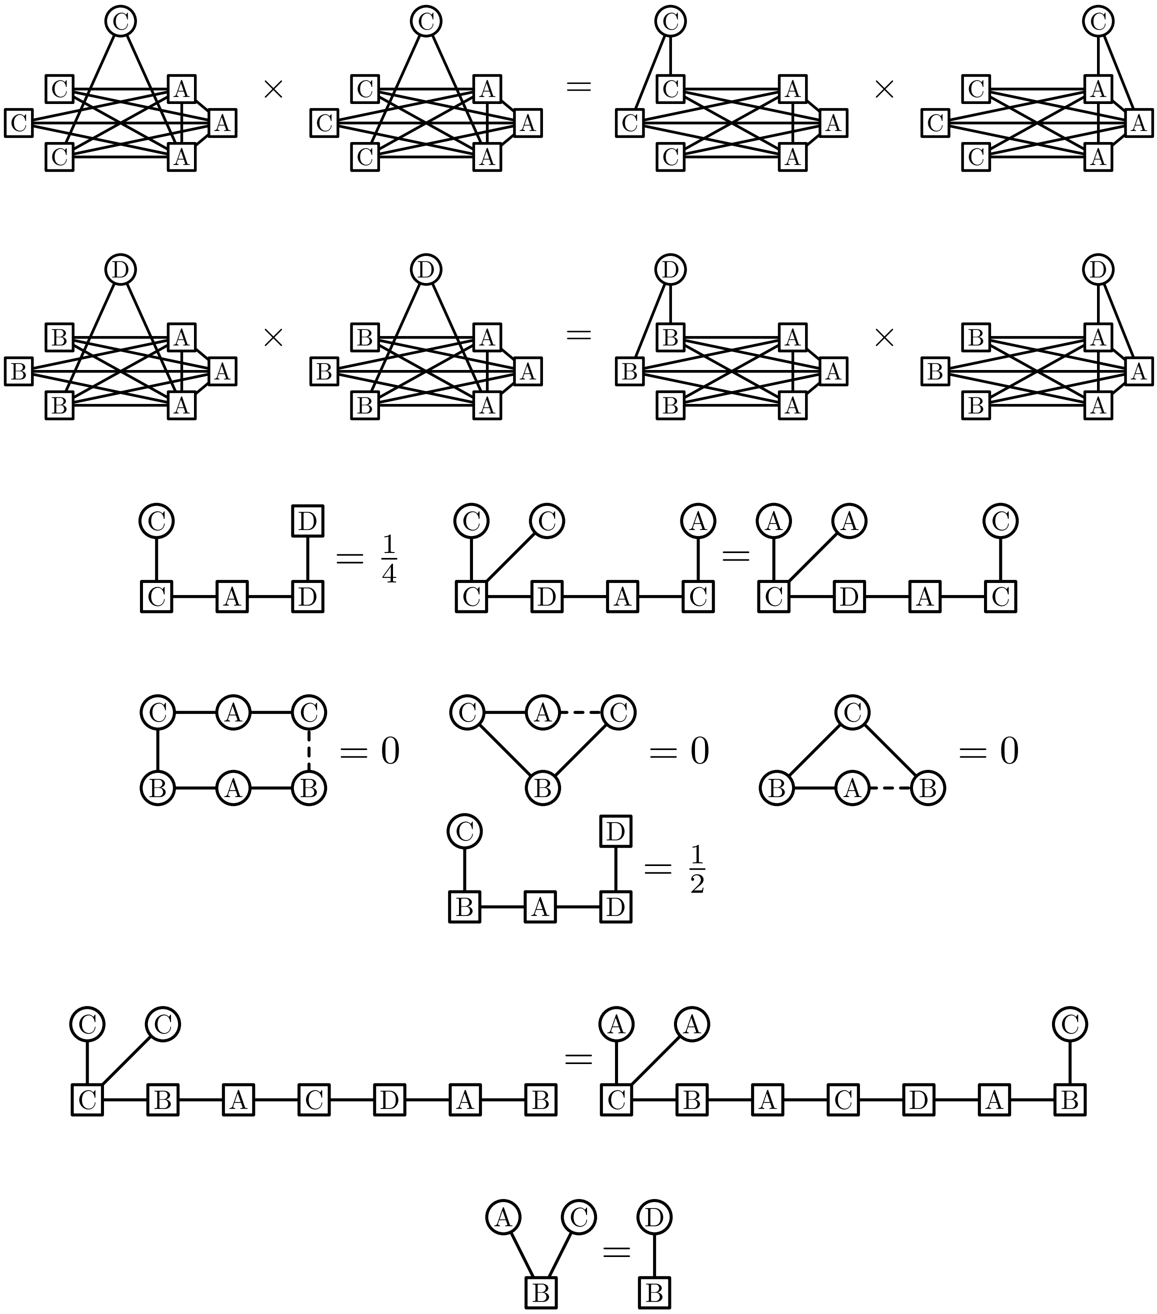

Figure 8.4: Decorated constraints used in Subsection 8.5.3.

Let us now look at the second constraint. Almost every choice of the roots satisfies that if the right root from C_0 is in the kth segment, then the roots from A_0 and D_0 are also in the kth segment and the left root from C_0 is in the $(k-1)$th segment. Since for every k the choice of such roots has positive probability, the constraint implies that the following holds for every $k \in \mathbb{N}$:

$$\left(\xi_k \left|g_C^{-1}\left(\left(1-2^{-k+1}, 1-2^{-k}\right)\right)\right|\right)^2 \left|g_A^{-1}\left(\left(1-2^{-k}, 1-2^{-k-1}\right)\right)\right|$$
$$= \left(\left|g_A^{-1}\left(\left(1-2^{-k+1}, 1-2^{-k}\right)\right)\right|\right)^2 \xi_{k+1} \left|g_C^{-1}\left(\left(1-2^{-k}, 1-2^{-k-1}\right)\right)\right| .$$

Hence, $\xi_{k+1} = \xi_k^2 = 2^{-2^{k-1}}$. We conclude that $W_0(x,y)$ and $W_S(g(x),g(y))$ are equal for almost every pair $(x,y) \in C_0 \times C_0$.

The first three constraints on the fourth line in Figure 8.4 yield that for every $k \in \mathbb{N}$ either $W(x,y) = 0$ for almost every $x \in B_0$ in the kth segment and almost every $y \in C_0$ or there exists m_k such that $W(x,y) = 1$ for almost every $x \in B_0$ in the kth segment and almost every $y \in C_0$ in the m_kth segment and $W(x,y) = 0$ for almost every $y \in C_0$ not in the m_kth segment. The last constraint on the fourth line yields that $m_1 = 1$.

We now show that the constraint on the fifth line implies that m_k exists and $m_k = t(k-1)$ for every $k \in \mathbb{N}$. For almost every choice of the roots in the constraint on the fifth line, if the right root from B_0 belongs to the kth segment, the left root from B_0 belongs to the $(k-1)$th segment and the left root from C_0 belongs to the m_{k-1}th segment. We derive that this constraint implies that

$$\left(2^{-2^{m_{k-1}-1}} \cdot 2^{-m_{k-1}}\right)^2 = \left(2^{-m_{k-1}}\right)^2 \cdot 2^{-m_k} .$$

We conclude that $m_k = 2^{m_{k-1}}$ and so $m_k = t(k-1)$. Consequently, $W_0(x,y)$ and $W_S(g(x),g(y))$ are equal for almost every pair $(x,y) \in B_0 \times C_0$.

The constraint on the last line in Figure 8.4 yields by considering a choice of the root in the kth segment of B_0 that

$$2^{-k} \cdot 2^{-m_k} = \xi'_k \cdot 2^{-k} .$$

We conclude that $\xi'_k = 2^{-m_k} = 2^{-t(k-1)} = t(k)^{-1}$. Hence, $W_0(x,y)$ and $W_S(g(x),g(y))$ are equal for almost every pair $(x,y) \in B_0 \times D_0$.

8.5.4 Subsegmenting

In analogy to Subsection 8.5.2, the first two constraints in Figure 8.5 for $X = B$ yield that there exists a set $\mathcal{J}$ of disjoint open subintervals of $[0, 1)$ such that $W_0(x, y) = 1$ for almost every $(x, y) \in g_B^{-1}(J)^2$ for some $J \in \mathcal{J}$ and $W_0(x, y) = 0$ for almost all other pairs $(x, y) \in B_0^2$. The third constraint implies that every set $g_B^{-1}(J)$ is a subset of $g_A^{-1}((1 - 2^{-k+1}, 1 - 2^{-k}))$ except for a set of measure zero for some $k \in \mathbb{N}$. Hence, each interval $J \in \mathcal{J}$ is a subinterval of $(1-2^{-k+1}, 1-2^{-k})$ for some $k \in \mathbb{N}$. The fourth constraint with $X = B$ yields that the length of each interval J is $2^{-k}t(k)^{-1}$. Finally, the first constraint on the second line can hold only if each interval $(1-2^{-k+1}, 1-2^{-k})$ contains $t(k)$ such intervals J. We conclude that $W_0(x, y)$ and $W_S(g(x), g(y))$ are equal for almost every pair $(x, y) \in B_0^2$.

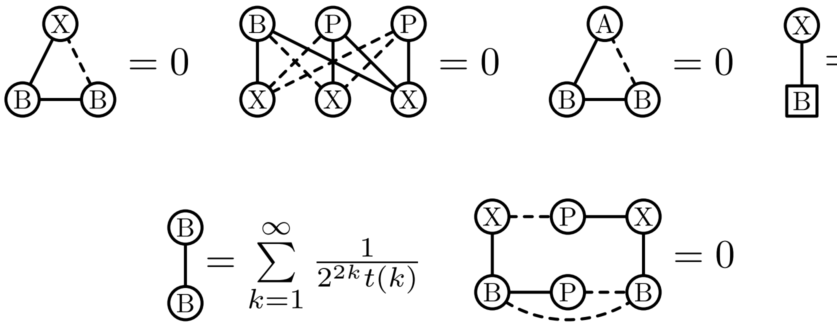

Figure 8.5: The first set of decorated constraints used in Subsection 8.5.4 where $X \in \{B, E, F, G\}$.

For $X \in \{E, F, G\}$, the first, second and fourth constraints on the first line yield that for each $J \in \mathcal{J}$ there exists an open interval J' of the same length as J such that $W_0(x, y) = 1$ for almost every $(x, y) \in J \times J'$ and $W_0(x, y) = 0$ for almost every $(x, y) \in J \times (X \setminus J')$. The last constraint on the second line in Figure 8.5 gives that the intervals J' follow in the same order as the intervals J. Hence, $W_0(x, y)$ and $W_S(g(x), g(y))$ are equal for almost every pair $(x, y) \in B_0 \times (E_0 \cup F_0 \cup G_0)$.

The set of constraints in Figure 8.6 is analogous to those in Figure 8.5. The main difference is the third constraint, which forces that if an interval J is a subinterval of an interval $(1 - 2^{-k+1}, 1 - 2^{-k})$, then $\left(2^{-k}t(k)^{-1}\right)^2 = 2^{-k} \cdot |J|$. Hence, the length of J must be $2^{-k}t(k)^{-2}$. The first constraint on the second line then forces that the interval $(1 - 2^{-k+1}, 1 - 2^{-k})$ must contain $t(k)^2$ such intervals J and the order

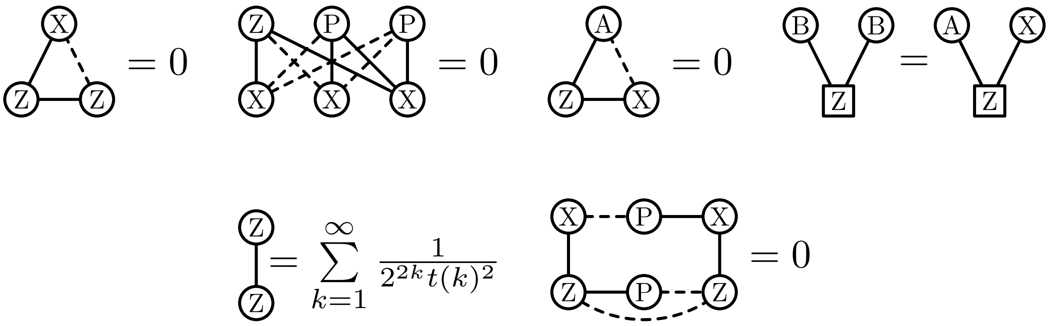

Figure 8.6: The second set of decorated constraints used in Subsection 8.5.4 where $(Z, X) \in \{(F, F), (F, D), (G, G)\}$.

of the corresponding pairs of intervals is forced by the last constraint. We can now conclude that $W_0(x, y)$ and $W_S(g(x), g(y))$ are equal for almost every pair $(x, y) \in F_0^2 \cup G_0^2 \cup (F_0 \times D_0)$.

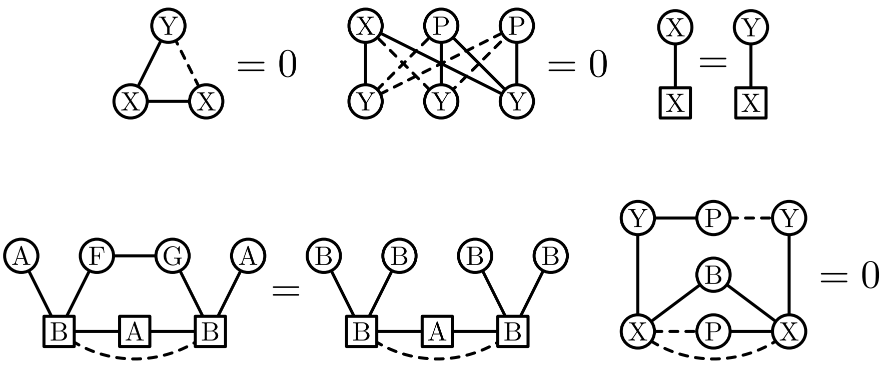

Figure 8.7: The third set of decorated constraints used in Subsection 8.5.4 where $(X, Y) \in \{(F, G), (G, F)\}$.

We now analyse the constraints from Figure 8.7. As in the previous cases, the three constraints on the first line force that each interval $(1 - 2^{-k+1}, 1 - 2^{-k})$, $k \in \mathbb{N}$, contain disjoint open intervals $I_1, \ldots, I_{t(k)^2}$ and $J_1, \ldots, J_{t(k)^2}$, each of length $2^{-k} t(k)^{-2}$, such that $W_0(x, y) = 1$ for almost every $(x, y) \in g_F^{-1}(I_i) \times g_F^{-1}(J_i)$ and $W_0(x, y) = 0$ for almost every $(x, y) \in g_F^{-1}(I_i) \times (F_0 \setminus g_F^{-1}(J_i))$.

Fix a choice of the roots in the first constraint on the second line in Figure 8.7; in almost every choice of the roots, all the three roots belong to the same segment. Suppose that they belong to the kth segment. The right side is equal to $\left(2^{-k} t(k)^{-1}\right)^4$ for almost all choices of the roots (since the structure of B_0^2 has already been forced) and the left side is equal to $2^{-2k} \cdot \left(2^{-k}t(k)^{-2}\right)^2$ multiplied by the number of choices of I_i and J_i such that I_i is contained in the subsegment of the left root and J_i in the subsegment of the right root. Since the left side and the right side must be equal for almost all choices of the roots, we conclude that for any pair of subsegments S and S' of the kth segment there exists a unique index i such that I_i is contained in S and J_i in S'. The last constraint in Figure 8.7 enforces that for any fixed subsegment S of the kth segment and any two subsegments S' and S'' such that S' precedes S'', the pairs $I_i \times J_i \subseteq S \times S'$ and $I_{i'} \times J_{i'} \subseteq S \times S''$ satisfy that the interval I_i precedes the interval $I_{i'}$. This implies that $W_0(x,y)$ and $W_S(g(x),g(y))$ are equal for almost every pair $(x,y) \in F_0 \times G_0$.

8.5.5 Binary expansions

In this section, we force the structure of the graphon inside $E_0 \times D_0$, $E_0 \times C_0$, $G_0 \times C_0$ and $F_0 \times C_0$. This will be achieved using the constraints depicted in Figure 8.8.

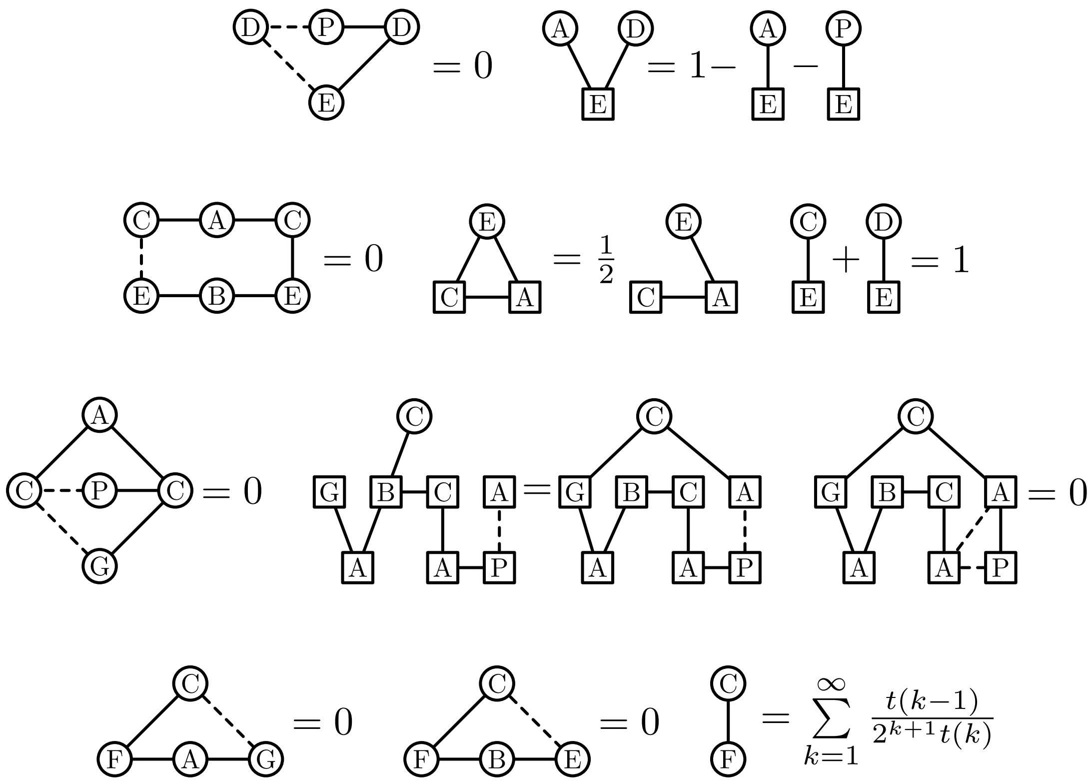

Figure 8.8: The decorated constraints used in Subsection 8.5.5.

The first constraint on the first line in the figure causes that for almost every $x \in$

E_0, there exists ξ_x such that $W_0(x,y) = 1$ for almost every $y \in D_0$ with $g_D(y) \le \xi_x$ and $W_0(x,y)$ for almost every $y \in D_0$ with $g_D(y) > \xi_x$. The second constraint on the line causes that for almost every $x \in E_0$, it holds that $2^{-[g_E(x)]_1}\xi_x = 1 - 2^{-[g_E(x)]_1} - g_E(x)$. It follows that

$$\xi_x = \frac{1 - 2^{-[g_E(x)]_1} - g_E(x)}{2^{-[g_E(x)]_1}} = 1 - [\![g_E(x)]\!]_1$$

for almost every $x \in E_0$. We conclude that $W_0(x,y)$ and $W_S(g(x), g(y))$ are equal for almost every pair $(x,y) \in E_0 \times D_0$.

The first constraint on the second line forces that W_0 is almost everywhere 0 or almost everywhere 1 on each product of a subsegment of E_0 and a segment of C_0. Fix a segment S_E of E_0. The second constraint forces that for almost every $y \in C_0$ the measure of $x \in S_E$ such that $W_0(x,y) = 1$ is exactly half of the measure of S_E. Since the measure of $y \in C_0$ such that $W_0(x,y) = 1$ is $1 - \xi_x = [g_E(x)]_1$ for almost every $x \in E_0$ because of the third constraint on the second line, the choice of the segments and subsegments where W_0 is one almost everywhere is unique. So, we get that $W_0(x,y)$ and $W_S(g(x), g(y))$ are equal for almost every pair $(x,y) \in E_0 \times C_0$.

The first constraint on the third line implies that for almost every $x \in G_0$ there exist $\xi_{x,k}$, $k \in \mathbb{N}$, such that $W_0(x,y) = 1$ for almost every y with $[g_C(y)]_1 = k$ and $[\![g_C(y)]\!]_1 \le \xi_{x,k}$ and $W_0(x,y) = 0$ for almost every other $y \in C_0$. Consider a possible choice of the roots in the second constraint on the third line. For almost every such choice of the roots, if the leftmost root from A_0 lies in the kth segment, then the middle root from A_0 is in the $t(k-1)$th segment (because of the already enforced structure of W_0 on $B_0 \times C_0$ in particular) and the rightmost root from A_0 is in the ℓth segment where $\ell \le t(k-1)$. The left side of the constraint is equal to $2^{-t(k-1)} = t(k)^{-1}$. The right side of the constraint is equal to $\xi_{x,\ell}2^{-\ell}$ where x is the root from G_0. This implies that $\xi_{x,\ell} = 2^{\ell}/t(k)$ where for almost every x from the kth segment of G_0 and $\ell \le t(k-1)$. Finally, almost every choice of the roots in the last constraint on the third line satisfies that if the leftmost root from A_0 lies in the kth segment, then the middle root from A_0 is in the $t(k-1)$th segment and the rightmost root from A_0 is in the ℓth segment where $\ell > t(k-1)$. Hence, $\xi_{x,\ell} = 0$ for almost every x from the kth segment of G_0 and $\ell > t(k-1)$. We conclude that $W_0(x,y)$ and $W_S(g(x), g(y))$ are equal for almost every pair $(x,y) \in G_0 \times C_0$.

The first constraint on the fourth line implies that for almost every $y \in C_0$, if the measure of x with $W_0(x,y) > 0$ from the kth segment of F_0 is positive, then $W_0(x,y) = 1$ for almost every x from the kth segment of G_0. Analogously, the second constraint yields that for almost every $y \in C_0$, if the measure of x with $W_0(x,y) > 0$ from a certain subsegment of F_0 is positive, then $W_0(x,y) = 1$ for almost every x from the corresponding subsegment of G_0. Consequently, $W_0(x,y) = 0$ for almost every pair $(x,y) \in F_0 \times C_0$ such that $W_S(g(x), g(y) = 0$. Since the last constraint implies that the integral of W_0 over $F_0 \times C_0$ is the same as the integral of W_S over $F \times C$, it holds that $W_0(x,y)$ and $W_S(g(x), g(y))$ are equal for almost every pair $(x,y) \in F_0 \times C_0$.

8.5.6 Linear transformation

In this subsection, we focus on the pair G_0 and E_0 of the parts. The first constraint in Figure 8.9 yields that $W_0(x,y) = 0$ for almost every $(x,y) \in G_0 \times E_0$ such that the segments of $[g_G(x)]_1 \neq [g_E(y)]_1$, i.e. the segments of x and y are different. The second constraint implies that for almost every $x \in G_0$ there exists ξ_x such that $W_0(x,y) = 1$ for almost every $y \in E_0$ such that $[g_G(x)]_1 = [g_E(y)]_1$ and $[\![g_E(y)]\!]_1 \geq \xi_x$ and $W_0(x,y) = 0$ for almost every $y \in E_0$ such that $[g_G(x)]_1 = [g_E(y)]_1$ and $[\![g_E(y)]\!]_1 < \xi_x$. The third constraint implies that almost every pair of x and x' from the same segment of G_0 such that $g_G(x) < g_G(x')$ satisfies that $\xi_x \geq \xi_{x'}$. In order to show that $W_0(x,y)$ and $W_S(g(x), g(y))$ are equal for almost every pair $(x,y) \in G_0 \times E_0$, it is enough to show that

$$\xi_x = \frac{1}{2} + t([g_G(x)]_1 - 1)^{1/2} \left(\frac{1}{2} - [\![g_G(x)]\!]_1 \right) \tag{8.14}$$

for almost every $x \in G_0$.

Almost every choice of the roots in the last constraint on the first line in Figure 8.9 satisfies that the root from G is from the first segment. Hence, this constraint implies that almost every $x \in G_0$ with $[g_G(x)]_1 = 1$ satisfies that $\frac{1-\xi_x}{2} = g_G(x)$. Since $t(0) = 1$ and $[\![g_G(x)]\!]_1 = 2g_G(x)$ for such $x \in G_0$, we obtain that (8.14) holds for almost every x from the first segment of G_0.

We now analyse the constraint on the second and third lines in Figure 8.9. Note that almost every choice of the roots satisfies that if the left root from A is in the kth segment, then the root from G is also in the kth segment and the right root from

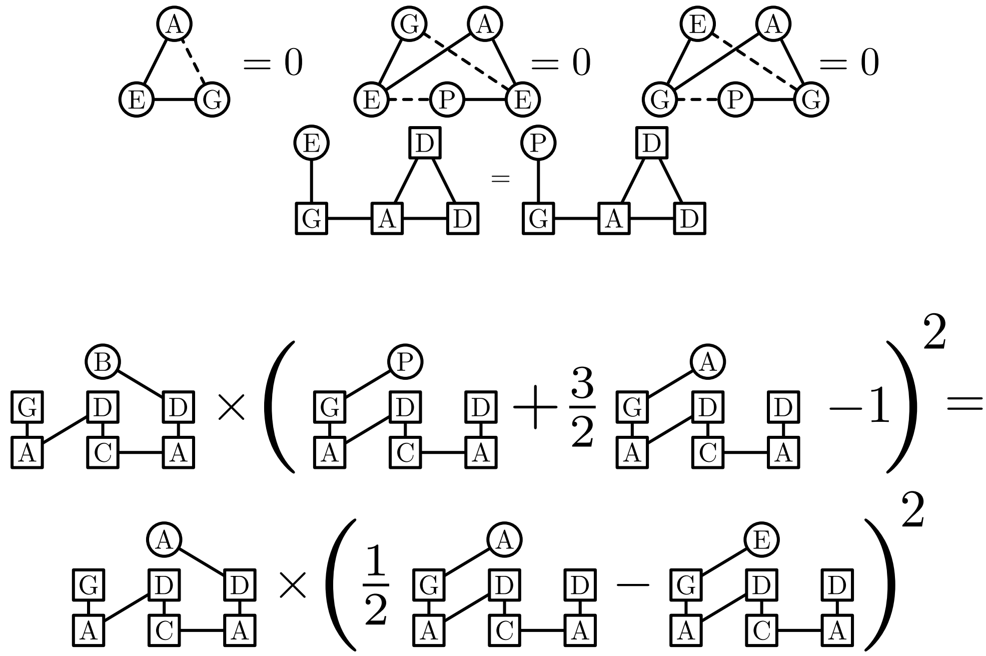

Figure 8.9: The decorated constraints used in Subsection 8.5.6. Note that one of the constraints is on both the second and third lines.

A and the right root from D are in the $(k-1)$th segment. In particular, $k \geq 2$ for almost every choice of the roots. Rewriting the densities using the already established structure of the graphon, we obtain that almost every $x \in G_0$ with $[g_G(x)]_1 = k \geq 2$ satisfies that

$$\frac{2^{-(k-1)}}{t(k-1)} \left(g_G(x) + \frac{3}{2} \cdot 2^{-k} - 1 \right)^2 = 2^{-(k-1)} \left(\frac{1}{2} \cdot 2^{-k} - (1-\xi_x) 2^{-k} \right)^2 .$$

Since it holds that $[\![g_G(x)]\!]_1 = 2^k \cdot (g_G(x) - 1 + 2^{-(k-1)})$, we can rewrite the right side as follows

$$\frac{2^{-(k-1)}}{t(k-1)} \left(2^{-k} [\![g_G(x)]\!]_1 - \frac{1}{2} \cdot 2^{-k} \right)^2 = 2^{-(k-1)} \left(\frac{1}{2} \cdot 2^{-k} - (1-\xi_x) 2^{-k} \right)^2 ,$$

which can be transformed to

$$\left([\![g_G(x)]\!]_1 - \frac{1}{2} \right)^2 = t(k-1) \left(\xi_x - \frac{1}{2} \right)^2 .$$

This implies that ξ_x is equal to

$$\frac{1}{2} + t([g_G(x)]_1 - 1)^{1/2} \left(\frac{1}{2} - [\![g_G(x)]\!]_1 \right) \text{ or } \frac{1}{2} - t([g_G(x)]_1 - 1)^{1/2} \left(\frac{1}{2} - [\![g_G(x)]\!]_1 \right)$$

for almost every $x \in G_0$ not contained in the first segment of G_0. Recall that almost every pair of x and x' from the same segment of G_0 such that $g_G(x) < g_G(x')$, which is equivalent to $[\![g_G(x)]\!]_1 < [\![g_G(x')]\!]_1$, satisfies that $\xi_x \geq \xi_{x'}$. This implies that the latter of two options for the values ξ_x holds for almost no $x \in G_0$ and thus (8.14) also holds for almost every $x \in G_0$ not contained in the first segment of G_0. We conclude that $W_0(x, y)$ and $W_S(g(x), g(y))$ are equal for almost every pair $(x, y) \in G_0 \times E_0$.

8.5.7 Dot product

We now enforce the unique structure inside the pairs $F_0 \times E_0$ and $D_0 \times G_0$. The first constraint in Figure 8.10 yields that $W_0(x, y) = 0$ for almost every $(x, y) \in F_0 \times E_0$ with $[g_F(x)]_1 \neq [g_E(y)]_1$, i.e. the segments of x and y are different. The second constraint implies that for almost every $x \in F_0$ there exists ξ_x such that $W_0(x, y) = 1$ for almost every $y \in E_0$ with $[g_F(x)]_1 = [g_E(y)]_1$ and $[\![g_E(y)]\!]_1 \leq \xi_x$ and $W_0(x, y) = 0$ for almost every $y \in E_0$ with $[g_F(x)]_1 = [g_E(y)]_1$ and $[\![g_E(y)]\!]_1 > \xi_x$. In order to show that $W_0(x, y)$ and $W_S(g(x), g(y))$ are equal for almost every pair $(x, y) \in F_0 \times E_0$, it is enough to show that

$$\xi_x = \frac{1}{2} - \frac{\left\langle [g_F(x)]_2^{\pm 1}, [g_F(x)]_3^{\pm 1} \right\rangle}{4t([g_F(x)]_1 - 1)} \tag{8.15}$$

for almost every $x \in F_0$.

Let x and x' be two elements of F_0 from the same segment, say, the kth segment. By the structure on $F_0 \times C_0$, the measure of $y \in C_0$ such that either $W_0(x, y) = 1$ and $W_0(x', y) = 0$ or $W_0(x, y) = 0$ and $W_0(x', y) = 1$ is equal to the number of different pairs of coordinates in $[\![g_F(x)]\!]_1^{\pm 1}$ and $[\![g_F(x')]\!]_1^{\pm 1}$ multiplied by $t(k)^{-1}$ for almost any pair x and x'. This measure can be rewritten as

$$\frac{t(k-1) - \left\langle [\![g_F(x)]\!]_1^{\pm 1}, [\![g_F(x')]\!]_1^{\pm 1} \right\rangle}{2t(k)} . \tag{8.16}$$

Consider now the constraint on the second line in Figure 8.10. For almost every choice of the roots, their segments are the same and so they all belong to the first segment of their parts (because of the edge between the two roots from D_0). Let x be the right root that belongs to F_0 and x' the left one. Because of the already enforced structure of the graphon, the subsegment of x' is $[g_F(x)]_3$ in almost every choice of

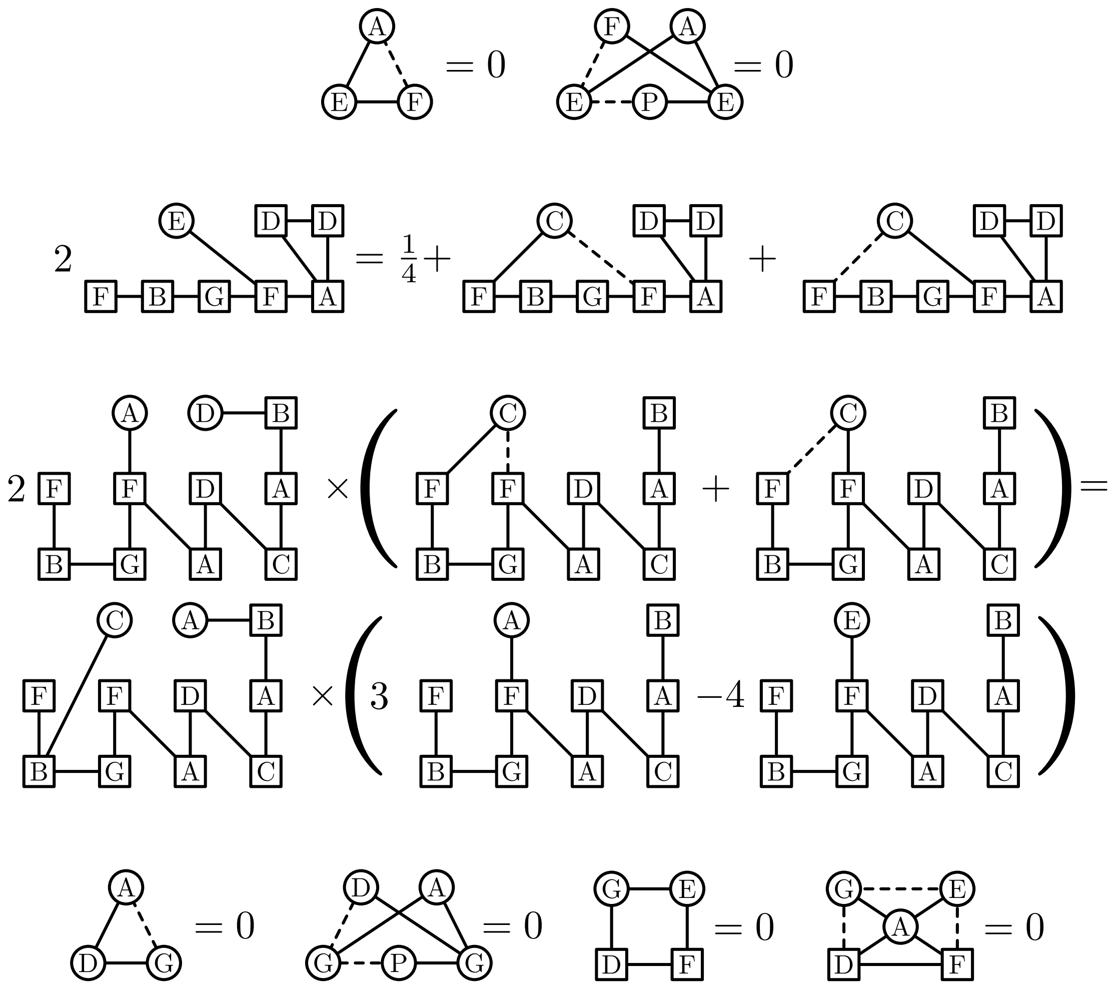

Figure 8.10: The set of decorated constraints used in Subsection 8.5.7. Note that one of the constraints is on both the third and fourth lines.

the roots. Using (8.16) with $t(0) = 1$ and $t(1) = 2$, we derive that it holds for almost every $x \in F_0$ that belongs to the first segment that

$$2 \cdot \frac{\xi_x}{2} = \frac{1}{4} + \frac{1 - \left\langle [g_F(x)]_2^{\pm 1}, [g_F(x)]_3^{\pm 1} \right\rangle}{4} \; .$$

We conclude that the equation (8.15) holds for almost every $x \in F_0$ from the first segment.

We now consider the constraint on the third and fourth lines in Figure 8.10. In almost every choice of the roots, the segment of the two roots from F_0 is one higher than the segment of the root from A_0 adjacent to the root from C_0. Also note that for almost every $x \in F_0$ that belongs to the second or higher segment, there is a set of positive measure of possible choices of the roots. Let x be the right root from F_0, x' the left one and k the common index of their segment. The constraint implies for almost every choice of the roots that

$$2 \cdot \frac{2^{-k} \cdot 2^{-(k-1)}}{t(k-1)} \cdot \frac{t(k-1) - \left\langle [\![g_F(x)]\!]_1^{\pm 1}, [\![g_F(x')]\!]_1^{\pm 1} \right\rangle}{2t(k)} = \frac{2^{-(k-1)}}{2^{t(k-1)}} \left(3 \cdot 2^{-k} - 4\xi_x \cdot 2^{-k}\right) \; .$$

This expression readily transforms to

$$1 - \frac{\left\langle [\![g_F(x)]\!]_1^{\pm 1}, [\![g_F(x')]\!]_1^{\pm 1} \right\rangle}{t(k-1)} = 3 - 4\xi_x \; .$$

Since we can choose the roots for almost every $x \in F_0$ that belongs to the second or higher segment with positive probability, almost every such $x \in F_0$ satisfies (8.15). We conclude that $W_0(x, y)$ and $W_S(g(x), g(y))$ are equal for almost every pair $(x, y) \in F_0 \times E_0$.

In the analogy to the first two constraints on the first last in Figure 8.10, the first two constraints on the last line in the figure yield that for almost every $z \in D_0$, there exists λ_z such that $W_0(z, y) = 1$ for almost every $y \in G_0$ with $[g_D(z)]_1 = [g_G(y)]_1$ and $[\![g_G(y)]\!]_1 \leq \lambda_z$ and $W_0(z, y) = 0$ for almost every $y \in G_0$ with $[g_D(z)]_1 \neq [g_G(y)]_1$ or $[\![g_G(y)]\!]_1 > \lambda_z$. We now consider the last two constraints on the last line in Figure 8.10. Fix a choice of roots $x \in F_0$ and $z \in D_0$. Almost every such choice of roots satisfies $[g_F(x)]_{2,3} = [g_D(z)]_{2,3}$, which implies that $[g_F(x)]_1 = [g_D(z)]_1$. The last but one constraint yields that $W_0(y, y') = 0$ for almost every $y \in E_0$, $y' \in G_0$, $[g_E(y)]_1 = [g_F(x)]_1$, $[g_G(y)]_1 = [g_D(z)]_1$, $[\![g_E(y)]\!]_1 \leq \xi_x$ and $[\![g_G(y')]\!]_1 \leq \lambda_z$. Similarly, the last constraint yields that $W_0(y, y') = 1$ for almost every $y \in E_0$, $y' \in G_0$,

$[g_E(y)]_1 = [g_F(x)]_1$, $[g_G(y)]_1 = [g_D(z)]_1$, $[\![g_E(y)]\!]_1 > \xi_x$ and $[\![g_G(y')]\!]_1 > \lambda_z$. The structure of the graphon W_0 on $E_0 \times G_0$ implies that

$$\lambda_z = \operatorname{trunc}\left(t(k-1)^{1/2}\left(\frac{1}{2} - \xi_x\right) + \frac{1}{2}\right) \tag{8.17}$$

for almost every pair of $x \in F_0$ and $z \in D_0$ such that $[g_F(x)]_{2,3} = [g_D(z)]_{2,3}$ and $[g_F(x)]_1 = [g_D(z)]_1 = k$. The expression (8.15) can be rewritten (for a particular choice of x and z) as

$$\xi_x = \frac{1}{2} - \frac{\left\langle [g_D(z)]_2^{\pm 1}, [g_D(z)]_3^{\pm 1} \right\rangle}{4t(k-1)} . \tag{8.18}$$

We get by substituting (8.18) in (8.17) that

$$\lambda_z = \operatorname{trunc}\left(\frac{1}{2} + \frac{\left\langle [g_D(z)]_2^{\pm 1}, [g_D(z)]_3^{\pm 1} \right\rangle}{4t(k-1)^{1/2}}\right)$$

for almost every $z \in D_0$ that belongs to the kth segment. We conclude that $W_0(z, y)$ and $W_S(g(z), g(y))$ are equal for almost every pair $(z, y) \in D_0 \times G_0$.

8.5.8 Main part

We now focus on the values of W_0 on E_0^2. The first constraint in Figure 8.11 implies that $W_0(x, y) = 0$ for almost every $x, y \in E_0^2$ with $[g_E(x)]_1 \neq [g_E(y)]_1$. Let us consider the second constraint in the figure. For almost every choice of the roots, all the roots belong to the same segment in their respective parts. Let k be the index of this segment and let z and z' be the choices of the bottom and the top roots that belong to B_0. Further, let S and S' be the subsegments of E_0 corresponding to the subsegments of z and z' in B_0, respectively. For almost every choice of the roots, the root from E_0 belongs to S. The second constraint yields that it holds for almost all $x \in S$ that

$$2^{-k} \int\limits_{S'} W(x, y) \, \mathrm{d}y = |S'| \cdot \frac{\operatorname{trunc}\left(\frac{1}{2} + \frac{\left\langle [\![g_B(z)]\!]_1^{\pm 1}, [\![g_B(z')]\!]_1^{\pm 1} \right\rangle}{4t(k-1)^{1/2}}\right)}{2^k} . \tag{8.19}$$

Along the same lines, the third constraint yields that it holds for almost all $x, x' \in S$ that

$$2^{-2k} \int\limits_{S'} W(x, y) W(x', y) \, \mathrm{d}y = |S'| \frac{\left(\operatorname{trunc}\left(\frac{1}{2} + \frac{\left\langle [\![g_B(z)]\!]_1^{\pm 1}, [\![g_B(z')]\!]_1^{\pm 1} \right\rangle}{4t(k-1)^{1/2}}\right)\right)^2}{2^{2k}} . \tag{8.20}$$

By Lemma 8.13, the identity (8.20) implies that

$$\frac{1}{|S'|}\int\limits_{S'} W(x,y)^2 \,\mathrm{d}y = \left(\operatorname{trunc}\left(\frac{1}{2}+\frac{\langle [\![g_B(z)]\!]_1^{\pm1}, [\![g_B(z')]\!]_1^{\pm1}\rangle}{4t(k-1)^{1/2}}\right)\right)^2 . \tag{8.21}$$

for almost every $x \in S$.

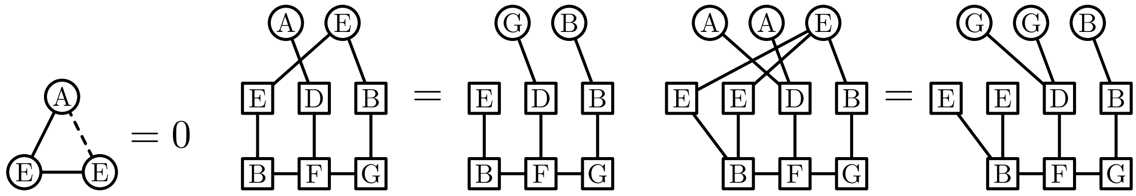

Figure 8.11: The set of decorated constraints used in Subsection 8.5.8.

We derive from (8.19) and (8.21) using Cauchy-Schwarz inequality that

$$W(x,y) = \operatorname{trunc}\left(\frac{1}{2}+\frac{\langle [\![g_B(z)]\!]_1^{\pm1}, [\![g_B(z')]\!]_1^{\pm1}\rangle}{4t(k-1)^{1/2}}\right)$$

for almost every $x \in S$ and every $y \in S'$. Since $[\![g_B(z)]\!]_1^{\pm1} = [\![g_E(x)]\!]_1^{\pm1}$ and $[\![g_B(z')]\!]_1^{\pm1} = [\![g_E(y)]\!]_1^{\pm1}$ for almost every such pair x and y, we conclude that $W_0(x,y)$ and $W_S(g(x),g(y))$ are equal for almost every pair $(x,y) \in E_0^2$.

8.5.9 Degree balancing

Fix $X \in \{A,\ldots,G,P\}$. The first constraint in Figure 8.12 implies that

$$\int\limits_{X_0} W_0(x,y)W_0(x',y)\,\mathrm{d}y = K_X$$

for almost every $x,x' \in Q_0$ for some K_X. By Lemma 8.13, this also implies that

$$\int\limits_{X_0} W_0(x,y)^2\,\mathrm{d}y = K_X$$

for almost every $x \in Q_0$. Hence, there exists a function $h_X : Q_0 \to [0,1]$ such that $W_0(x,y) = h_X(y)$ for almost every $x \in Q_0$ and $y \in X_0$. Since the last constraint on the first line yields that

$$\int\limits_{\overline{R_0}} W_0(x,y)\,\mathrm{d}y = \frac{4}{13}$$

for almost every $x \in X_0$ and it holds that

$$\int\limits_{Q_0 \cup R_0} W_0(x,y) \, \mathrm{d}y = \int\limits_{Q \cup R} W_S(g_X(x), y) \, \mathrm{d}y$$

for almost every $x \in X_0$ because of the already enforced structure, we conclude that $W_0(x,y)$ and $W_S(g(x), g(y))$ are equal for almost every pair $(x,y) \in X_0 \times Q_0$.

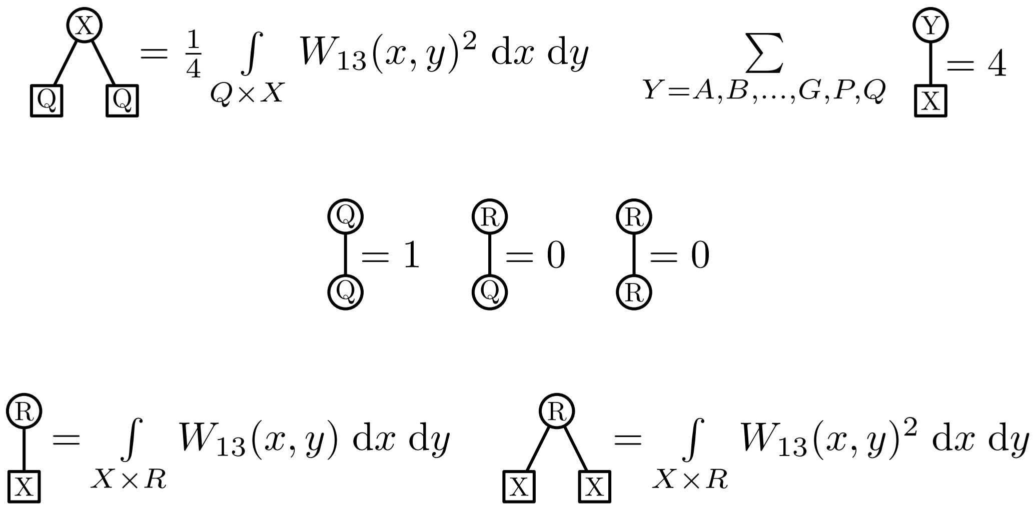

Figure 8.12: The set of decorated constraints used in Subsection 8.5.9 where $X \in \{A, B, C, D, E, F, G, P\}$.

The three constraints on the second line in Figure 8.12 clearly implies that $W_0(x,y)$ and $W_S(g(x), g(y))$ are equal for almost every pair $(x,y) \in (Q_0 \cup R_0)^2$.

Again fix $X \in \{A, \ldots, G, P\}$. The two constraints on the last line in Figure 8.12 give using Lemma 8.13 that

$$\int_{R_0} W_0(x,y) \, \mathrm{d}y = K_X |R_0| \text{ and } \int_{R_0} W_0(x,y)^2 \, \mathrm{d}y = K_X^2 |R_0|$$

for almost every $x \in X_0$ for some K_X. However, this is only possible if $W_0(x,y) = K_X$ for almost every $x \in X_0$ and almost every $y \in R_0$. The right side values of the two constraints yield that the values of K_X matches the corresponding values in W_S. So, we can conclude that $W_0(x,y)$ and $W_S(g(x), g(y))$ are equal for almost every pair $(x,y) \in X_0 \times R_0$. Since we have shown that the constraints in Figures 8.2–8.12 imply that $W_0(x,y)$ and $W_S(g(x), g(y))$ are equal for almost every pair x and y, we have established that the Švejk graphon is finitely forcible. This completes the proof of Theorem 8.4.

8.6 Concluding remark

Proposition 8.5 implies that it is not possible to remove $2^{5\log^* \varepsilon_i^{-2}}$ completely from the denominator in the exponent in Theorem 8.4. However, our construction can be modified to replace $t(n)$ with a faster growing function of n, e.g. with $t(t(n))$, which would replace the function $2^{5\log^* \varepsilon_i^{-2}}$ with a slower growing function of ε^{-1}.

Chapter 9

Complexity of Circular Colouring Reconfiguration

In recent years, a large body of research has emerged concerning reconfiguration variants of combinatorial problems; see the survey of van den Heuvel [122] and the paper of Ito et al. [127] for background. We mainly focus on reconfiguration problems for graph colourings and, more generally, graph homomorphisms (defined below). Problems of this type are of interest in statistical physics. The set of k-colourings of a graph G can be viewed as the space of possible global states of the anti-ferromagnetic Potts model at zero temperature (see Sokal [209]) and the reconfiguration problem corresponds precisely to the Glauber dynamics for this model.

Classical graph colourings and circular colourings are both special cases of graph homomorphisms. A *homomorphism* from a graph G to a graph H (also called an H-*colouring* of G) is a mapping $f : V(G) \to V(H)$ such that $f(u)f(v) \in E(H)$ whenever $uv \in E(G)$. The notation $f : G \to H$ indicates that f is a homomorphism from G to H. In this language, a k-colouring is simply a homomorphism to a complete graph on k vertices. A (p, q)-colouring of G is equivalent to a homomorphism from G to the graph $G_{p,q}$ which has vertex set $\{0, \dots, p-1\}$ and edge set $\{ij : q \leq |i-j| \leq p-q\}$. The graph $G_{p,q}$ is called a *circular clique*.

Remark 9.1. It is well known that $G_{p,q}$ admits a homomorphism to $G_{p',q'}$ if and only if $p/q \leq p'/q'$ [219]. Therefore, since the composition of two homomorphisms is a homomorphism, a graph G admits a (p, q)-colouring if and only if it admits a (p', q')-colouring for all $p'/q' \geq p/q$.

Given two homomorphisms $f, g : G \to H$, we say f *reconfigures* to g if there a sequence $(f = f_0), f_1, f_2, \dots, (f_n = g)$ of homomorphisms from G to H such that f_i and f_{i+1} differ on only one vertex. The sequence is referred to as a *reconfiguration sequence*. Clearly the existence of a reconfiguration sequence from f to g can be determined independently for each component of G, so we may assume that G is connected. We define the general homomorphism reconfiguration problem as follows. Let H be a fixed graph.

***H*-Recolouring**

Instance: A connected graph G, and two homomorphisms $f, g : G \to H$.

Question: Does f reconfigure to g?

When $H = K_k$ or $H = G_{p,q}$ we will call the problem k-Recolouring and (p, q)-Recolouring respectively. The complexity of H-Recolouring is only known for a handful of families of targets. Theorem 1.22 is a dichotomy theorem for the family of complete graphs and Theorem 1.24 is a dichotomy theorem for the family of circular cliques. Recently, Wrochna [224] (see also [223]) proved that H-Recolouring is polynomial whenever H does not contain a 4-cycle. In contrast, one can observe that $G_{p,q}$ contains 4-cycles whenever $p > 2q+1$ and so the polynomial side of Theorem 1.24 does not follow directly from the result of Wrochna.

The rest of the chapter is outlined as follows. In Section 9.1, we provide an explicit polynomial-time algorithm for deciding the (p, q)-Recolouring problem when $2 \leq p/q < 4$. In Section 9.2, we show that, when $p/q \geq 4$, the reconfiguration problem for $\lfloor p/q \rfloor$-colourings can be reduced to the reconfiguration problem for (p, q)-colourings, thereby completing the proof of Theorem 1.24 (via Theorem 1.22).

9.1 The Polynomial Cases: $2 \leq p/q < 4$

We now extend the ideas of [58] to study the complexity of (p, q)-Recolouring for $2 \leq p/q < 4$. Given two 3-colourings f and g of a graph G, the algorithm in [58] consists of two phases. The first phase tests whether the so-called "fixed vertices" (vertices that cannot be recoloured under *any* sequence of recolourings) are assigned the same colours under f and g. If not, then we know that f does not reconfigure to g and the algorithm terminates.

If the algorithm does not terminate at this point, then it enters the second phase. In this phase, the algorithm obtains two edge labellings of G from the vertex colourings f and g. The key property of these labellings is that, given that the algorithm did not terminate after the first phase, one can show that f reconfigures to g if and only if every cycle of G has the same "weight" under both edge labellings. After a polynomial number of steps, the algorithm either produces a reconfiguration sequence or discovers a cycle with different weights under f and g. We follow a similar approach, but we begin by examining the edge relabelling problem which turns out to depend only on cycle weights.

Remark 9.2. The work in [58] has been extended in [134] where the authors find shortest paths between 3-colourings. In our work, we do not consider shortest paths, but rather are simply concerned with testing the existence of paths in polynomial time.

9.1.1 Edge Relabellings

Let f be a (p,q)-colouring of a graph G. Let $\overrightarrow{G}$ be an oriented graph obtained by arbitrarily orienting each edge of G. (This orientation is required to define the weighting of paths, cycles, and cuts below. Thus throughout this section we will assume a graph G has an orientation $\overrightarrow{G}$.) The *edge-labelling induced by* f is defined as $\varphi_f : E(G) \to \{0, 1, \ldots, p-1\}$ by $\varphi_f(e) = f(v) - f(u) \bmod p$ where $e = \overrightarrow{uv}$. Let C be a cycle of G and arbitrarily choose a direction of traversal. Let $E^+(C)$ be those edges of C whose orientation in $\overrightarrow{G}$ agrees with the direction of traversal (forward arcs) and $E^-(C)$ be those edges of C whose orientation in $\overrightarrow{G}$ is reversed to the direction of traversal (backward arcs). Define

$$\varphi_f(C) = \sum_{e \in E^+(C)} \varphi_f(e) + \sum_{e \in E^-(C)} (p - \varphi_f(e)). \tag{9.3}$$

Note that the summation above is in $\mathbb{Z}$, i.e. it is not reduced modulo p. The following properties of φ_f are immediate from the fact that f is a (p,q)-colouring and the sum $\varphi_f(C)$ is a telescoping series in the values of f on C.

P1 $q \le \varphi_f(e) \le p - q$ for all edges e.

P2 $\varphi_f(C) \equiv 0 \pmod p$ for all cycles C.

Given a graph G, an edge-labelling $\psi : E(G) \to \{0, 1, \ldots, p-1\}$ is a (p,q)*-labelling* if it satisfies properties **P1** and **P2** above. Similar to the weighting function for cycles in (9.3), we define a weighting for paths: $\varphi(P)$ is the sum of $\varphi(e)$ and $p - \varphi(e)$ for forward and backward arcs e in P, respectively, according to some chose direction of traversal on P.

We now introduce the reconfiguration process for edge-labellings. Let G be a graph and $\overrightarrow{G}$ an orientation of G. For $\emptyset \neq X \subsetneq V(G)$, the *edge cut* $\partial(X)$ is the set of edges with one end in X and the other in $\overline{X}$. Let $\partial^+(X)$ (resp. $\partial^-(X)$) be those edges oriented from X to $\overline{X}$ (resp. $\overline{X}$ to X) in $\overrightarrow{G}$. Given a (p,q)-labelling $\varphi : E(G) \to \{0, 1, \ldots, p-1\}$, let α be an integer $1 \le \alpha \le p-1$, where $\varphi(e) \ge q + \alpha$ for $e \in \partial^+(X)$ and $\varphi(e) \le p - q - \alpha$ for $e \in \partial^-(X)$. The labelling φ' obtained from φ by *relabelling on* $\partial(X)$ *by* α is defined by:

$$\varphi'(e) = \begin{cases} \varphi(e) - \alpha & \text{if } e \in \partial^+(X), \\ \varphi(e) + \alpha & \text{if } e \in \partial^-(X) \end{cases}$$

This process is called an *edge cut relabelling*. Clearly if φ is a (p,q)-labelling, then φ' is as well. It is easy to verify that for any cycle C, $\varphi(C) = \varphi'(C)$.

We remark an alternative relabelling definition is to simply require that if φ is a (p,q)-labelling, then φ' is as well (with no requirement that $\varphi(e) \ge q + \alpha$ for $e \in \partial^+(X)$ and $\varphi(e) \le p - q - \alpha$ for $e \in \partial^-(X)$). In this case one would naturally reduce φ' modulo p after shifting by α. With this more general definition the weight of a cycle can change after relabelling. For example, consider a directed 4-cycle with $p = 4$ and $q = 1$. The constant labelling of 1 on each edge is a $(4,1)$-labelling. The weight of this cycle is 4 (when traversed in the direction of the edges). Using $\alpha = 2$, one can relabel to obtain the constant labelling $\varphi'(e) = 3$ for all edges: simply add 2 to the first and third edges, and -2 to the second and fourth edges. After reducing the weights modulo 4, we see $\varphi'(e) = 3$ for all edges. The weight of the cycle is now 12. Nonetheless, the two notions of relabelling are equivalent with the assumption $p/q < 4$ and the natural restriction $-p/2 \le \alpha \le p/2$. We will use the first definition as it eases the analysis below.

Theorem 9.4. *Let G be a connected graph and φ, ψ be two (p, q)-edge-labellings of G. Then φ reconfigures to ψ through a sequence of edge cut relabellings if and only if $\varphi(C) = \psi(C)$ for all cycles C in G.*

Proof. Relabelling on an edge cut does not change the weight of any cycles. Thus, if φ reconfigures to ψ, then $\varphi(C) = \psi(C)$ for all cycles C is a necessary condition. To prove sufficiency, assume $\varphi(C) = \psi(C)$ for each cycle C. We may assume $\varphi \neq \psi$. Let $\overrightarrow{wu}$ be an arc such that $\varphi(wu) < \psi(wu)$. (The case where $\varphi(wu) > \psi(wu)$ is analogous.) Let T be a spanning tree of G rooted at u with all arcs directed away from u. For each $v \in V(G)$ define $\mathrm{wt}(v, \varphi) = \varphi(P)$ where P is the (u, v) path in T, traversed from u to v. Similarly define $\mathrm{wt}(v, \psi)$.

Let $X = \{v : \mathrm{wt}(v, \varphi) \leq \mathrm{wt}(v, \psi)\}$. Thus, $\overline{X} = \{v : \mathrm{wt}(v, \varphi) > \mathrm{wt}(v, \psi)\}$. Clearly $u \in X$ and it is easy to see $w \in \overline{X}$. Suppose to the contrary $w \in X$. Then uw is not an edge of T and thus the uw path in T plus the edge wu is a cycle C. Moreover this cycle must have $\varphi(C) < \psi(C)$ when traversed with the direction of the path P, contrary to our assumption. Hence $\emptyset \neq X \subsetneq V(G)$ and $\partial(X)$ is a well defined edge cut.

We claim that $\partial^+(X)$ must only contain arcs where $\varphi(e) > \psi(e)$ and arcs in $\partial^-(X)$ must have $\varphi(e) < \psi(e)$. Suppose to the contrary there is an edge $e = \overrightarrow{xy}, x \in X, y \in \overline{X}$ such that $\varphi(e) \leq \psi(e)$. (The proof for arcs in $\partial^-(X)$ analogous.) By construction there is a (u, x)-path P_1 such that $\varphi(P_1) \leq \psi(P_1)$. Thus $P_1 + y$ is an (u, y)-path such that $\varphi(P_1 + y) \leq \psi(P_1 + y)$. On the other hand, since $y \in \overline{X}$, the path (u, y) path in T, say P_2, has the property that $\varphi(P_2) > \psi(P_2)$. Hence, reverse of the path is a (y, u)-path, P_2^R such that $\varphi(P_2^R) < \psi(P_2^R)$. The concatenation of $P_1 + y$ and P_2^R is a closed (u, u)-walk with less weight under φ than ψ. In particular, it must contain a cycle C such that $\varphi(C) \neq \psi(C)$, a contradiction. Let $\alpha = \min_{e \in \partial(X)}\{|\varphi(e) - \psi(e)|\}$. Relabel $\partial(X)$ by α to obtain φ'. Now φ' agrees with ψ on more edges than φ. The result follows. $\square$

Observe the second half of the proof of Theorem 9.4 shows that if there is a cycle C with different weights under φ and ψ, then C is the fundamental cycle in $T + e$ where $e \in \partial(X)$. The discovery of such a cycle through examining fundamental cycles plays a key role in our polynomial time algorithm.

Corollary 9.5. *Let G be a graph together with two (p, q)-edge-labellings φ and ψ. Either φ reconfigures to ψ through a sequence of edge cut relabellings or there is a cycle C for which $\varphi(C) \neq \psi(C)$. Morever, in the latter case C may be taken to be the fundamental cycle in $T + e$ where T is a fixed spanning tree of G, and $e \in \partial(X)$ as defined in the proof of Theorem 9.4. Determining which of the two cases (exclusively) holds can be determined in polynomial time.*

9.1.2 Vertex Recolouring

We now return to recolouring vertices. We have already observed that given a (p, q)-colouring f of a graph, we can construct a (p, q)-edge labelling φ_f. Conversely given φ, a (p, q)-edge labelling, using the ideas from the proof of Theorem 9.4, we can generate a weight, $\mathrm{wt}(v, \varphi)$, for each vertex v. Reducing these weights modulo p yields a (p, q)-colouring f where $\varphi = \varphi_f$. The following proposition is immediate. Throughout this section addition (for shifting colours) is done modulo p.

Proposition 9.6. *Let G be a connected graph. An edge-labelling $\psi : E(G) \to \{0, 1, \ldots, p-1\}$ is a (p, q)-labelling if and only if $\psi = \varphi_f$ for some $f : G \to G_{p,q}$. Moreover, if $\varphi_f = \varphi_g$ for $f, g : G \to G_{p,q}$, then there exists $k \in \{0, 1, \ldots, p-1\}$ such that $f(v) = g(v) + k$ for all $v \in V(G)$.*

To connect our work on edge cut relabellings back to vertex recolourings, we must realise edge cut relabelling through vertex recolourings where we recolour one vertex at a time. First, we observe that relabelling φ on $\partial(X)$ by α corresponds to the following vertex recolouring where we recolour an entire set of vertices (simultaneously). Let f' be the colouring obtained from f by:

$$f'(v) = \begin{cases} f(v) + \alpha & \text{if } v \in X, \\ f(v) & \text{otherwise.} \end{cases}$$

Then $\varphi_{f'}$ is the edge-labelling obtained from φ_f by relabelling on $\partial(X)$ by α. We call this process *recolouring the vertex set X by α*. Provided $\varphi(e) \geq q + \alpha$ for all $e \in \partial^+(X)$ and $\varphi(e) \leq p - q - \alpha$ for all $e \in \partial^-(X)$, the new colouring f' is a proper (p, q)-colouring. The following corollary is immediate from Proposition 9.6 and Theorem 9.4.

Corollary 9.7. *Let f and g be two (p,q)-colourings of a graph G. Then there is a sequence of colourings $f, f_1, f_2, \ldots, f_n$, each obtained from its predecessor by a vertex set recolouring, where $f_n(v) = g(v) + k$ for all vertices v and some constant k if and only if $\varphi_f(C) = \varphi_g(C)$ for all cycles C in G.*

To complete our algorithm, we need to determine when recolouring a vertex set X by α can be realised by a sequence of single vertex recolourings. We begin by proving that for $p/q < 4$, recolouring by α can always be achieved through recolouring X by 1 (α times). Note this not the case when $p/q \geq 4$. For example, consider a $(4,1)$-colouring of C_4 where the vertices are coloured $0, 1, 2, 3$ (in their cyclic order). The vertex with colour 0 can be recoloured to 2. However, the recolouring cannot be done by increasing the colour by 1 (twice).

Given a pair $a, b \in \{0, \ldots, p-1\}$, the *interval* $[a,b]$ is defined to be the set $\{a, a+1, \ldots, b-1, b\}$ (where addition is modulo p). An important property of $G_{p,q}$ used in the proof below is that, when $2 \leq p/q < 4$, the common neighbours of any two vertices form an interval. This fact, and its proof, appear in [56]. (Roughly, if x and y are common neighbours in two non-contiguous intervals, then the non-neighbours of x and the non-neighbours of y are disjoint intervals. This requires $p \geq 2 + 2\cdot(2q-1)$ or $p/q \geq 4$.)

Proposition 9.8. *For $2 \leq p/q < 4$, let f, g be two (p,q)-colourings of a graph G where g is obtained from f by recolouring the vertex set X by α. Then there is a sequence of colourings $(f = f_0), f_1, f_2, \ldots, (f_\alpha = g)$ where each colouring is obtained from its predecessor by recolouring X by 1 (for the entire sequence) or each is obtained from its predecessor by recolouring X by -1 (for the entire sequence).*

Proof. Let $v \in X$ and vu be an edge. If $u \in X$, then $(f(v)+1)(f(u)+1)$ and $(f(v)-1)(f(u)-1)$ are both edges of $G_{p,q}$. Consequently we can increment the colour by 1 of each vertex in X, or decrement the colour by 1 of each vertex in X, and maintain a proper (p,q)-colouring on X.

On the other hand, consider a neighbour u of v such that $u \notin X$. Note $f(u)f(v)$ and $f(u)(f(v)+\alpha)$ are both edges of $G_{p,q}$. If $\alpha = p/2$, then $f(u)$ is distance at least q from both $f(v)$ and $f(v)+p/2$ which implies $p/2 \geq 2q$, contrary to our assumption. If $\alpha < p/2 < 2q$, then $f(u)$ must belong to the interval $[f(v)+\alpha, f(v)]$. Thus $f(u)$

is adjacent in $G_{p,q}$ to every vertex of $[f(v), f(v) + \alpha]$ and in particular to $f(v) + 1$. Note this argument depends only on $\alpha < p/2$. Hence the colouring f_1 obtained from f by increasing the colour of each vertex in X by 1 is a (p, q)-colouring of G.

A similar argument shows for $\alpha > p/2$, the colouring f_1 obtained from f by decreasing the colour of each vertex in X by 1 is a (p, q)-colouring. The result follows. □

Given G, an orientation $\overrightarrow{G}$, and an (p, q)-edge-labelling, say φ, let D_φ be the subdigraph of $\overrightarrow{G}$ induced by edges with label q or $p - q$. Without loss of generality we may assumed D_φ is oriented so that all edges have label q. (To simplify notation we shall write D_f instead of D_{φ_f} when the edge-labelling φ_f comes from some vertex colouring f.) In the case $p/q = 2$, each edge will be oriented both ways, thus we have a directed 2-cycle on each edge.

Lemma 9.9. *Let f be a (p, q)-colouring of a graph G. Let f' be the (p, q)-colouring obtained from f by recolouring some vertex set X by 1. Then f' can obtain from f by a sequence of recolourings of single vertices if and only if induced subdigraph $D_f[X]$ contains no directed cycles.*

Proof. Suppose $D_f[X]$ contains no directed cycles. We topologically sort the vertices of X to obtain an ordering $x_1, x_2, \ldots, x_t$ where $e = \overrightarrow{x_i x_j} \in E(D_f[X])$ implies $i < j$. For each $i = t, t - 1, \ldots, 1$ increase $f(x_i)$ by 1.

On the other hand, suppose $D_f[X]$ contains a directed cycle. If one could sequentially increase the colour of each vertex in X by 1, then there would be a first vertex c_i in the cycle to have its colour changed. However, c_i has an out neighbour c_{i+1} (in $D_f[X]$) meaning $f(c_{i+1}) - f(c_i) = q$. The resulting colouring gives the edge $c_i c_{i+1}$ a weight of $q - 1$, i.e. the resulting colouring is a not a proper (p, q)-colouring. □

9.1.3 Polynomial Time Algorithm

We now describe the polynomial time algorithm for recolouring. We present a proof of correctness for each step below. The algorithm itself is in Figure 9.1.3.

Let f and g be two (p, q)-colourings of a graph G. As defined in [58], vertices whose colours can never change under any sequence of recolourings are called *fixed*. Following the ideas of [58] and using the results above, we identify (the only) three obstructions preventing the recolouring of f to g: fixed vertices, cycle weights, and weights of paths between fixed vertices.

The first step of the algorithm is to identify fixed vertices and verify $f(v) = g(v)$ for all fixed vertices.

Lemma 9.10. *Let f be a (p, q)-colouring of a graph G where $2 \leq p/q < 4$. Suppose v is a vertex in a strongly connected component of D_f. Then the colour of v cannot be changed under any sequence of recolourings.*

Proof. Since v belongs to a strongly connected component of D_f, v must belong to a directed cycle in D_f. By definition of D_f, the predecessor and successor of v on this cycle receive colours $f(v) - q$ and $f(v) + q$ respectively. As noted above, and proved in [56], the common neighbours of two vertices in $G_{p,q}$ form an interval. As the colour of v must be a common neighbour of $f(v) - q$ and $f(v) + q$, the only choice for v is $f(v)$. That is, the colour of v cannot change until the colour of one its neighbours changes. However, this is true for every vertex on the directed cycle. All the vertices in the strongly connected component are fixed. □

The strongly connected components of D_f can be found in linear time; hence, we can find the fixed vertices and verify equality of f and g on the fixed vertices in linear time. (See, e.g., [2].)

In Step 2 (a) of the algorithm we construct a spanning tree T rooted a vertex u, and construct the cut $\partial(X)$ as in the proof of Theorem 9.4. Corollary 9.5 shows if some cycle has different weight under φ_f and φ_g, then we will discover a fundamental cycle in $T + e$ for some edge $e \in \partial(X)$.

In Step 2 (b) we calculate the shift α that will increase the number of edges upon which φ_f and φ_g agree. Proposition 9.8, Lemma 9.9, and Lemma 9.11 show that the vertices of X or $\overline{X}$ can be recoloured 1 at a time, or there is a path between fixed vertices certifying that f does not recolour to g. Specifically, in the case we wish to increase the vertex set X by 1 and $D_f[X]$ contains a directed cycle, we can achieve the same change to the edge labelling on $\partial(X)$ by decreasing $\overline{X}$ by 1 provided $D_f[\overline{X}]$ does not contain a directed cycle. In the event both subgraphs contain a directed cycle, we have an obstruction to recolouring as described in the following lemma.

Lemma 9.11. *Let G be a graph with two (p,q)-colourings f and g. Suppose T is a spanning tree rooted at u. Let $X = \{v : \mathrm{wt}(v,\varphi_f) \leq \mathrm{wt}(v,\varphi_g)\}$ (as defined in Theorem 9.4). If $D_f[X]$ and $D_f[\overline{X}]$ both contain a directed cycle, then f does not recolour to g.*

Proof. Let x belong to a directed cycle in X and y belong to a directed cycle in $\overline{X}$. Then there are paths in T from u to x, say P_1, and from u to y, say P_2 such that $\varphi_f(P_1) \leq \varphi_g(P_1)$ and $\varphi_f(P_2) > \varphi_g(P_2)$. Thus the reverse of P_1 concatenated with P_2 is an x,y-path such that $\varphi_f(P_1^R P_2) > \varphi_g(P_1^R P_2)$. Since the end points of this path belong to directed cycles in D_f, their colours are fixed. Let z be an internal vertex of the path. If z can be recoloured by α it is easy to check the sum of the weight of the two path edges incident with z does not change. (For example, if both edges are directed in the direction of traversal of the path, one increases by α and the other decreases by α. The analyses for the other possible orientations of the two edges are similar.) $\square$

In [58], fixed vertices are found by successively deleting sources and sinks from D_f. Thus vertices belonging to a directed cycle will be fixed, but also vertices on a directed path between two directed cycles are also fixed. In our work, we identify fixed vertices of the former type using strongly connected components of D_f and vertices of the latter type by the weight of paths between strongly connected components (our third obstruction) as directed paths in D_f have the smallest possible weight over all edge labellings.

In Step 3 of the algorithm, we have recoloured vertices to obtain a colouring f_n such that $\varphi_{f_n} = \varphi_g$. By Corollary 9.7, $f_n(v) = g(v) + k$ for some constant k. If D_f contains any directed cycles, then there are fixed vertices under f. In Step 1 we have checked that $f(v) = g(v)$ for fixed vertices, so we can conclude $k = 0$ and $f_n = g$. On the other hand, if $k \neq 0$, then there are no directed cycles in D_f. Applying Proposition 9.8 with $X = V(G)$ shows we can recolour the vertices one at a time to reconfigure f_n to g.

The following theorem is immediate from the algorithm.

Theorem 9.12. *Let $2 \leq p/q < 4$. Suppose f and g are (p,q)-colourings of a graph G, $\overrightarrow{G}$ is an orientation of G and φ_f, φ_g are the edge labellings of G obtained from f and g. Then f recolours to g if and only if:*

The Recolouring Algorithm

Let $2 \leq p/q < 4$.

Input: A graph G and two (p, q)-colourings f, g.

Output: A recolouring sequence from f to g or an obstruction to recolouring.

1. Find the strongly connected components of D_f. For each vertex v belonging to a nontrivial strongly connected component, verify $f(v) = g(v)$. If for some such v, $f(v) \neq g(v)$, then answer NO, return a directed cycle to which v belongs, and STOP.

2. Construct a spanning tree T rooted at a vertex u. Partition $V(G)$ into three sets: X_ℓ, X_e, X_s consisting of those vertices v whose (u, v)-path in T has weight larger, equal, smaller under φ_f versus φ_g respectively. Repeat until $X_\ell \cup X_s = \emptyset$:

   (a) If $X_\ell \neq \emptyset$, then let $X = X_e \cup X_s$ and $\overline{X} = X_\ell$. (If $X_\ell = \emptyset$ and $X_s \neq \emptyset$ apply an analogous process reversing their roles.) For each $e \in \partial^+(X)$ (resp. $\partial^-(X)$), verify $\varphi_f(e) > \varphi_g(e)$ (resp. $\varphi_f(e) < \varphi_g(e)$). If some edge fails this test, then answer NO, return a cycle with different weights under φ_f and φ_g, and STOP.

   (b) If $D_f[X]$ and $D_f[\overline{X}]$ both contain nontrivial strongly connected components, then answer NO, return a path P between two fixed vertices with $\varphi_f(P) \neq \varphi_g(P)$, and STOP.

   (c) Let $\alpha = \min_{e \in \partial(X)} |\varphi_f(e) - \varphi_g(e)|$. Recolour the vertices of X (or $\overline{X}$) by 1 using a sequence of single vertex recolourings. Repeat this recolouring by X (by 1) α times.

   (d) Update the path weights in T and the sets X_ℓ, X_e, X_s.

3. At this point $\varphi_f = \varphi_g$. If $f(u) = g(u)$, then answer YES, return the sequence of recolourings, and STOP. Otherwise, we know G contains no directed cycles and we can topologically sort $V(G) : v_1, v_2, \ldots v_n$. Increase (or decrease) the colour of all the vertices by 1 in the order v_n to v_1. Repeat this until $f(u) = g(u)$. Answer YES, return the recolouring sequence, and STOP.

Figure 9.1: The Recolouring Algorithm

1. *the same set of vertices are fixed under f and g and $f(v) = g(v)$ for all fixed vertices;*

2. *for each cycle C of G, $\varphi_f(C) = \varphi_g(C)$; and*

3. *for each path P whose end points are fixed, $\varphi_f(P) = \varphi_g(P)$.*

Moreover, one can find a recolouring sequence from f to g or an obstruction of the above type in polynomial time.

At each iteration of the algorithm we recolour $O(|V(G)|)$ vertices and increase X_e by at least one vertex. Once a vertex is in X_e it never leaves. Thus we do $O(|V(G)|)$ recolouring steps.

Corollary 9.13. *Let $2 \leq p/q < 4$. Suppose f and g are (p, q)-colourings of a graph G. If f recolours to g, then there is a reconfiguration sequence of length $O(|V(G)|^2)$ which certifies this.*

9.2 The PSPACE-complete Cases: $p/q \geq 4$

9.2.1 Overview

As in the previous section, addition and subtraction of in $\{0, \ldots, p-1\}$ is viewed modulo p. We say that $i, j \in \{0, \ldots, p-1\}$ are *compatible* if they are adjacent in $G_{p,q}$. For notational convenience, let us define $k := \lfloor p/q \rfloor$ and $r := p - kq$.

Consider, for a moment, the case $r = 0$. If $q = 1$, then the complexity is PSPACE-complete by Theorem 1.22. Otherwise, let $\gamma : \{0, \ldots, k-1\} \to \{0, \ldots, p-1\}$ be defined by $\gamma(i) := qi$ and $\phi : \{0, \ldots, p-1\} \to \{0, \ldots, k-1\}$ be defined by $\phi(j) := \lfloor j/q \rfloor$. It is not hard to see that γ is a homomorphism from K_k to $G_{p,q}$ and that ϕ is a homomorphism from $G_{p,q}$ to K_k. Also, given an instance (G, f, g) of the k-Recolouring problem, we have that f reconfigures to g if and only if γf reconfigures to γg as (p, q)-colourings. Therefore, the (p, q)-Recolouring problem is PSPACE-complete if $p/q \geq 4$ and $r = 0$. Thus while the proofs below work when $r = 0$, to avoid trivialities we assume that $r \geq 1$.

Given an instance (G, f, g) of the k-Recolouring problem, we will construct an instance (G', α, β) of the (p, q)-Recolouring problem in polynomial time such that:

- $|V(G')| = O(|V(G)| + |E(G)|)$ and
- f reconfigures to g if and only if α reconfigures to β.

It will follow from Theorem 1.22 that the reconfiguration problem for (p,q)-colourings is PSPACE-complete, thereby completing the proof of Theorem 1.24.

We give a brief outline of the ideas behind the construction before moving into the finer details. The first step is to divide the set $\{0,\dots,p-1\}$ into k intervals which we will, in some sense, treat as k separate colours. Define

$$S_0 := [0, q+r-1], \text{ and}$$

$$S_i := [iq + r, (i+1)q + r - 1] \text{ for } 1 \le i \le k-1.$$

It is clear that S_i and S_j are disjoint for $i \neq j$ and that $\bigcup_{i=0}^{k-1} S_i = \{0,\dots,p-1\}$. Let $\gamma : \{0,\dots,k-1\} \to \{0,\dots,p-1\}$ be the function which maps i to the left endpoint of S_i for $0 \le i \le k-1$. For $i \neq j$, it is easily observed that the left endpoint of S_i is adjacent to the left endpoint of S_j in $G_{p,q}$. Therefore, given any k-colouring $f : V(G) \to \{0,\dots,k-1\}$, the composition γf is a (p,q)-colouring.

This observation guides part of our construction. The graph G' will contain G as an induced subgraph and the (p,q)-colourings α and β will be defined so that their restrictions to G will be equal to γf and γg, respectively. Now, if it is possible to reconfigure α to β in such a way that none of the intermediate colourings map a pair of adjacent vertices of G to the same set S_i, then we immediately obtain a reconfiguration sequence taking f to g. That is, to obtain a sequence of k-colourings taking f to g, one could simply compose each (p,q)-colouring of the sequence with the function $\varphi : \{0,\dots,p-1\} \to \{0,\dots,k-1\}$ which maps every vertex of S_i to i for $0 \le i \le k-1$. Notice that, for $i \neq 0$, the set S_i is a stable set in $G_{p,q}$ and therefore no adjacent pair is ever mapped to S_i. Thus, all that we need to worry about is that some of the intermediate colourings may map two adjacent vertices of G to S_0. To remedy this, we will add some structures (or "gadgets") to G' which will forbid adjacent vertices of G from mapping to S_0.

To this end, we start by adding a copy of $G_{p,q}$ disjoint from G with vertex set $\{y_0,\dots,y_{p-1}\}$, where y_iy_j is an edge whenever $q \le |i-j| \le p-q$. We extend the (p,q)-colourings α and β to $\{y_0,\dots,y_{p-1}\}$ by setting $\alpha(y_i) = \beta(y_i) = i$ for all i. It is clear that each vertex y_i is fixed in α and β and, moreover, in any (p, q)-colouring which reconfigures to α or β.

Now, for each edge uv of G, we will add two *forbidding paths* P_{uv} and P_{vu} from u to v to G' which are internally disjoint from $V(G) \cup \{y_0, \ldots, y_{p-1}\}$ and one another, as well as from every other such path. These paths restrict the colour pairs which can appear on u and v during a reconfiguration process. We achieve this by assigning to each interval vertex of the path, a list of allowed colours for that vertex, i.e. we use a list colouring. This is accomplished by joining the internal vertices of the paths to specific subsets of $\{y_0, \ldots, y_{p-1}\}$. For example consider the case $p = 18, q = 4$ which we explore below. Suppose we want a path vertex x to always be coloured with an element of the list $\{4, 5, \ldots, 11\}$. Then it suffices to join x to the vertices $\{y_{15}, y_{16}, y_{17}, y_0\}$. As we will show, the colours in the lists forbids u and v from mapping to S_0. Also, once we describe our construction in detail, it will be clear that the length of the forbidding paths depends only on p and q, and so $|V(G')| = O(|V(G)| + |E(G)|)$.

The difficulty now comes in proving that if f reconfigures to g, then α reconfigures to β. Specifically, given a reconfiguration sequence $(h_i)_{i=1}^s$ taking f to g, we need that the lists assigned to the internal vertices of the forbidding paths are flexible enough that we can use $(h_i)_{i=1}^s$ to obtain a reconfiguration sequence taking α to β. This will be proved using a moderate amount of case analysis at the end of the section. Before moving on, we remark that the general strategy of using some sort of forbidding paths in which the internal vertices are confined to lists was also used by Bonsma and Cereceda [48] (in a somewhat different manner) to prove that the reconfiguration problem for k-colourings is PSPACE-complete for $k \geq 4$.

9.2.2 Defining the Forbidding Paths

Let uv be an edge of G. We will now describe our construction of the forbidding path $P_{uv} = ux_0^{uv}x_1^{uv} \ldots x_t^{uv}v$ from u to v, including the definition of the lists assigned to the internal vertices $x_0^{uv}, \ldots, x_t^{uv}$. As mentioned before, the construction of G' also includes a path P_{vu} from v to u which is defined similarly.

First, the construction of the lists depends on the sequence $q, 2q, 3q, \ldots, r$. Define t to be the smallest positive integer such that $(t+1)q \equiv r \pmod{p}$. Now, the *lists* for

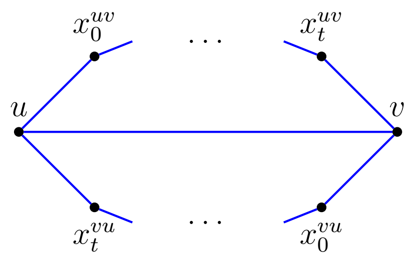

Figure 9.2: Forbidding paths attached to each edge uv

the internal vertices of P_{uv} are intervals of $G_{p,q}$ and are defined as follows:

$$L\left(x_0^{uv}\right) = L\left(x_t^{uv}\right) := [p-1, 2q-1],$$

$$L\left(x_i^{uv}\right) := [iq, (i+2)q-1] \text{ for } 1 \le i \le t-1.$$

As stated in the outline, in order to enforce these lists we add a copy of $G_{p,q}$ on vertex set $\{y_0, \dots, y_{p-1}\}$ and join the vertices of $P_{uv} \setminus \{u, v\}$ to the appropriate vertices of this set. Specifically, we join the vertices x_0^{uv} and x_t^{uv} to $\{y_j : j \in [3q-1, p-q-1]\}$ and join x_i^{uv} to $\{y_j : j \in [(i+3)q-1, (i-1)q]\}$ for $1 \le i \le t-1$.

Figure 9.3 contains a useful tabular representation of the lists, and a specific case $(p, q) = (18, 4)$ is depicted in Figure 9.4. In both of these figures, the colours are laid out so that below colour c in the table is the colour $c + q$. This gives the following proposition.

Proposition 9.14. *When colouring the path $x_0^{uv}, x_1^{uv}, \dots, x_t^{uv}$ the colours used (one per row) must be directly below or to the right of the preceding vertex, and any such sequence of colours from the table gives a good colouring of the path with the exception of $L(x_{t-1}^{uv})$ and $L(x_t^{uv})$. In this case colours in the interval $[r-2q, r-1]$ of $L(x_{t-1}^{uv})$ are compatible with colours below and to the right in the interval $[p-1, q+r]$ of $L(x_t^{uv})$. Similarly, the interval $[-q-1, r-1]$ is compatible with colours below and to the right of $[p-1, 2q-1]$.*

Proof. For $1 \le i \le t-2$ vertex x_i^{uv} receives a colour from the interval $[iq, (i+2)q-1]$ and x_{i+1}^{uv} receives a colour from $[(i+1)q, (i+3)q-1]$. Observe iq is compatible with all colours in the interval $[(i+1)q, (i-1)q]$. Since $p + (i-1)q \ge (i+3)q - 1$, we have iq is compatible with every colour on the list for x_{i+1}^{uv}. The colour $iq+1$ is compatible with $[(i+1)q+1, (i-1)q+1]$ which include all colours below and to the right of $iq+1$.

Continuing one sees that $(i+2)q-1$ is compatible with $[(i+3)q-1, (i+1)q-1]$ which includes only the last vertex of the list for x_{i+1}^{uv}. Thus the proposition holds for $1 \leq i \leq t-2$. It is straightforward to verify the case $i=0$ as well. Consider now $i = t-1$. The argument is similar; however, we need to consider the two cases in the statement of the proposition. First for a colour $[r-2q, r-1]$, a colour in $[p-1, q+r]$ is compatible if and only if it is below and to the right. (Note the right hand end point $q+r$ is compatible with $r-2q$ if and only if $p+r-2q-q \geq q+r$ or $p \geq 4q$.) A similar statement holds for $[-q-1, r-1]$ and $[p-1, 2q-1]$. The result follows. $\square$

Finally the vertical bar in the table provides the following information: if vertex u receives colour 0, then only colours to the right of the vertical bar can be used to colour the path P_{uv}.

Vertex	Lists										
x_0^{uv}	$p-1$	0	$\ldots$	$q-r-1$	$\ldots$	$q-1$	q	$\ldots$	$2q-1$		
x_1^{uv}		q	$\ldots$	$2q-r-1$	$\ldots$	$2q-1$	$2q$	$\ldots$	$3q-1$		
x_2^{uv}		$2q$	$\ldots$	$3q-r-1$	$\ldots$	$3q-1$	$3q$	$\ldots$	$4q-1$		
$\vdots$			$\vdots$		$\vdots$			$\vdots$			
x_i^{uv}		iq	$\ldots$	$(i+1)q-r-1$	$\ldots$	$(i+1)q-1$	$(i+1)q$	$\ldots$	$(i+2)q-1$		
$\vdots$			$\vdots$		$\vdots$			$\vdots$			
x_{t-1}^{uv}		$r-2q$	$\ldots$	$-q-1$	$\ldots$	$r-q-1$	$r-q$	$\ldots$	$r-1$		
x_t^{uv}				$p-1$	$\ldots$	$r-1$	r	$\ldots$	$q+r-1$	$\ldots$	$2q-1$

Figure 9.3: Assignment of lists to the internal vertices of the forbidding path.

Vertex	Lists									
x_0^{uv}	17	0	1	2	3	4	5	6	7		
x_1^{uv}		4	5	6	7	8	9	10	11		
x_2^{uv}		8	9	10	11	12	13	14	15		
x_3^{uv}		12	13	14	15	16	17	0	1		
x_4^{uv}			17	0	1	2	3	4	5	6	7

Figure 9.4: Assignment of lists to the internal vertices of the forbidding path for the $(p,q) = (18,4)$ case.

Before moving on, let us check that the paths P_{uv} and P_{vu} actually do the job that they are meant to do; namely, that they forbid u and v from both being mapped to S_0.

Proposition 9.15. *Let uv be an edge of G. If ψ is a (p, q)-colouring of $P_{uv} \cup P_{vu}$ such that the internal vertices of each path are coloured from their lists, then $\psi(u)$ and $\psi(v)$ cannot both be contained in S_0.*

Proof. Suppose to the contrary that $\psi(u), \psi(v) \in S_0$. In $G_{p,q}$, the only edges contained in S_0 are between vertices in $[0, r-1]$ and vertices in $[q, q+r-1]$. So, without loss of generality, we can assume that

$$\psi(u) \in [0, r-1] \text{ and } \psi(v) \in [q, q+r-1]. \tag{9.16}$$

However, the fact that $\psi(u) \in [0, r-1]$ implies that $\psi\left(x_0^{uv}\right) \in [q, 2q-1]$ (Note that $r-1 < q-1$ and thus $(r-1, p-1)$ is not an edge of $G_{p,q}$.) By our observations above, we see that only colours to the right of the vertical line in Figure 9.3 can be used to colour P_{uv}. In particular $\psi\left(x_t^{uv}\right) \in [r, 2q-1]$. This implies that $\psi(v) \notin [q, q+r-1]$, which contradicts (9.16) and completes the proof. □

9.2.3 Defining α and β on the Forbidding Paths

As we have already mentioned, α and β are defined in such a way that their restrictions to G are equal to γf and γg, respectively, and $\alpha(y_i) = \beta(y_i) = i$ for all i. Next, we will describe the way in which we extend α and β to the vertices of the forbidding paths.

Recall that the vertices of G each receive a colour from $\{0, q+r, 2q+r, \ldots, (k-1)q+r\}$ under α. Given an edge uv of G, the internal vertices of P_{uv} are coloured as follows. If $\alpha(u) \neq 0$ and $\alpha(v) \neq 0$, then colour x_i^{uv} with colour iq for $i = 0, 1, 2, \ldots, t-1$ and colour x_t^{uv} with 0. (These colours correspond to the left hand column in Figure 9.3 for $1 \leq i \leq t-1$.) The path P^{vu} is similarly coloured so that $\alpha(x_i^{vu}) = iq$ and $\alpha(x_t^{vu}) = 0$. In particular, $\alpha(x_0^{uv}) = \alpha(x_t^{vu}) = 0$ for all $v \in N(u)$. Note that $(t-1)q \equiv r - 2q \pmod{p}$ which is compatible with 0. On the other hand, if $\alpha(u) = 0$, then set $\alpha(x_i^{uv}) = (i+1)q$ for $i = 0, 1, \ldots, t$. For P_{vu}, set $\alpha(x_i^{vu}) = iq$ for $i = 0, 1, \ldots, t-1$ and $\alpha(x_t^{vu}) = q$. Observe in this case, the internal vertices around u have colour q, i.e. $\alpha(x_0^{uv}) = \alpha(x_t^{vu}) = q$ for all $v \in N(u)$, and the internal vertices around v have colour r or 0, i.e. $\alpha(x_t^{uv}) = r$ and $\alpha(x_0^{vu}) = 0$ for all $v \in N(u)$.

We similarly extend β using g. A technical detail we will exploit below is that in extending α and β the colour for x_t^{uv} belongs to $\{0, r, q\}$. These colours are compatible with all colours above and to the left in the row for x_{t-1}^{uv} by Proposition 9.14. We call these colourings of P_{uv} the *standard path colourings*. At the start of each reconfiguration step we assume the paths have a standard path colouring, and the end of each reconfiguration step we ensure the paths have a standard path colouring. The standard path colourings correspond to coloumns in Figure 9.3 with possible changes at x_t^{uv}.

By Proposition 9.15, given any sequence of (p, q)-colourings which reconfigures α to β, we can compose each of these colourings with γ to obtain a sequence of k-colourings taking f to g. This proves the following proposition.

Proposition 9.17. *Suppose (G', α, β) is an instance of (p, q)-*Recolouring *obtained from (G, f, g), an instance of k-*Recolouring*, as described above. If α reconfigures to β, then f reconfigures to g.*

9.2.4 Recolouring G'

To complete the reduction we need to prove that if f reconfigures to g, then α reconfigures to β. Let $(h_i)_{i=1}^s$ be any reconfiguration sequence taking f to g. We show that there is a sequence $(\eta_i)_{i=1}^s$ of (p, q)-colourings of G' such that

- $\eta_1 = \alpha$ and $\eta_s = \beta$,
- the restriction of η_i to G is γh_i for $1 \leq i \leq s$, and
- η_i reconfigures to η_{i+1} for $1 \leq i \leq s-1$.

Clearly this will prove that α reconfigures to β and complete the proof of Theorem 1.24.

Before tackling the general case we illustrate our method with the example in Figure 9.5, in which $(p, q) = (18, 4)$. In this case, the colourings η_i map $V(G)$ to the colours $\{0, 6, 10, 14\}$. Let $1 \leq j \leq s-1$ be fixed and suppose that $(\eta_i)_{i=1}^j$ have been constructed to satisfy the conditions above; our goal is to construct the colouring η_{j+1} and a reconfiguration sequence taking η_j to η_{j+1}.

By definition, h_j and h_{j+1} differ on at most one vertex, say $u \in V(G)$. Suppose, for example that $h_j(u) = 0$ and $h_{j+1}(u) = 1$; this requires us change the colour of u from 0 (its current colour under η_j) to 6 (its desired colour under η_{j+1}) without changing the colour of any other vertex of G. We remark that this is the most involved

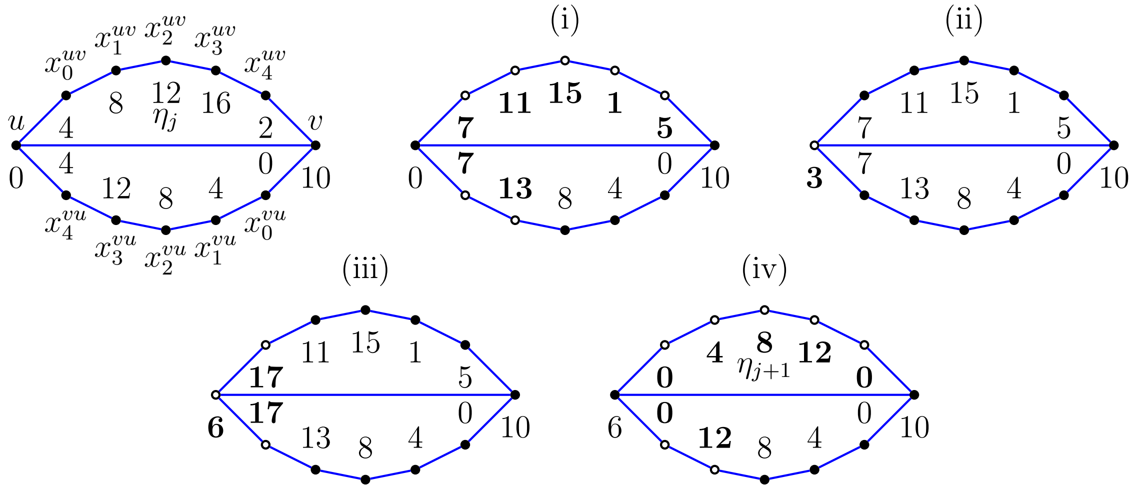

Figure 9.5: Recolouring u from 0 to 6 in the $(18, 4)$ case.

case, and corresponds to Case (c) in Lemma 9.18 below. The other cases use similar ideas but are more straightforward. For each vertex $v \in V(G)$ to which u is adjacent, v must receive a colour from $\{2, 3\}$ under h_j to facilitate the recolouring of u, which implies that $\eta_j(v) \in \{10, 14\}$. Each step below is applied for each neighbour v of u, one by one, before moving on to the subsequent step. The recoloured vertices at each step are shown as white vertices in Figure 9.5 with the new colours indicated in boldface.

(i) First, recolour x_3^{vu} to 13 and x_4^{vu} to 7. Then, recolour P_{uv} by recolouring x_4^{uv} to 5, x_3^{uv} to 1, x_2^{uv} to 15, x_1^{uv} to 11 and x_0^{uv} to 7. At this point all vertices of $\{x_0^{uv}, x_t^{vu} : v \in N_G(u)\}$ have colour 7.

(ii) Now, change the colour of u to 3 (temporarily).

(iii) Next, change the colour of x_0^{uv} to 17 and x_t^{vu} to 17. At this point all vertices of $\{x_0^{uv}, x_t^{vu} : v \in N_G(u)\}$ have colour 17. Recolour u to 6.

(iv) Recolour the paths to the standard colourings: x_i^{uv} is coloured iq for $i = 0, 1, \ldots, t-1$ and x_t^{uv} is coloured 0; x_{t-1}^{vu} is coloured 12 and x_t^{vu} is coloured 0.

We now describe the recolouring technique in general.

Lemma 9.18. *Let h and h' be k-colourings of G which differ on a unique vertex u and let η be a (p, q)-colouring of G' such that*

- *η is the standard path colouring for all paths P_{uv},*
- *$\eta(y_i) = i$ for all i, and*
- *the restriction of η to G is γh.*

Then there exists a (p, q)-colouring η' of G' such that η' is the standard path colouring for all paths P_{uv}, the restriction of η' to G is $\gamma h'$, and η reconfigures to η'.

Proof. We provide a reconfiguration sequence which takes η to a colouring η' with the desired properties. The proof is divided into cases. In each case, each of the reconfiguration steps is done for every neighbour v of u, one by one, before moving on to the subsequent step.

As mentioned above, when extending α from $V(G)$ to $V(G')$, the colour for x_t^{uv} is selected from the interval $[p-1, q+r]$. In some cases below we may assign x_t^{uv} a colour outside of this interval (but still on its list). However, the assignment will be compatible with the current colour of x_{t-1}^{uv} and at the end of the case the colour of x_t^{uv} will be in $[p-1, q+r]$.

Let uv be an edge in G and assume the paths P_{uv} and P_{vu} have standard path colourings.

Case (a): $\eta(u), \eta'(u) \in \{q+r, 2q+r, \ldots, (k-1)q+r\}$ In standard path colourings $\eta(x_0^{uv}), \eta(x_t^{vu}) \in \{0, r\}$ for each $v \in N(u)$. The colour $\eta'(u)$ is compatible with both 0 and r, so we may change the colour of u to $\eta'(u)$. The paths still have standard path colourings.

Case (b): $\eta(u) = 0, \eta'(u) \in \{2q+r, 3q+r, \ldots, (k-1)q+r\}$

(i) By the definition of standard path colourings, $\eta(x_0^{uv}) = \eta(x_t^{vu}) = q$. Change the colour of u to $\eta'(u)$.

(ii) Recolour x_t^{vu} to 0. The path P_{vu} now has a standard path colouring.

(iii) Recolour x_i^{uv} to iq for $i = 0, 1, \ldots, t-1$ and x_t^{uv} to 0. The path P_{uv} now has a standard path colouring.

Case (c): $\eta(u) = 0, \eta'(u) = q+r$

(i) We begin observing that $h(u) = 0$ and $h'(u) = 1$, which implies that $h(v) \in \{2, \ldots, k-1\}$. Thus, $\eta(v) \in \{2q+r, 3q+r, \ldots, (k-1)q+r\}$. For $i = t, \ldots, 0$, in order, recolour x_i^{uv} to $(i+2)q-1$.

(ii) Change the colours of x_{t-1}^{vu} to $-q-1$, and x_t^{vu} to $2q-1$. These colours are compatible with each other, and by Propsoition 9.14, the colour of x_{t-1}^{vu} and x_{t-2}^{vu} are compatible as well.

(iii) At this point all vertices of $\{x_0^{uv}, x_t^{vu} : v \in N_G(u)\}$ have colour $2q-1$. Recolour u to $q-1$ (temporarily). Note that this colour is compatible with all neighbours v of u in G since $\eta(v) \in \{2q+r, 3q+r, \ldots, (k-1)q+r\}$.

(iv) Recolour x_0^{uv} to $p-1$ and x_t^{vu} to $p-1$.

(v) At this point all vertices of $\{x_0^{uv}, x_t^{vu} : v \in N_G(u)\}$ are coloured $p-1$. Change the colour of u to $q+r$.

(vi) Recolour x_{t-1}^{vu} to $r-2q$ and x_t^{vu} to 0. Recolour x_i^{uv} to iq for $i = 0, 1, \ldots, t-1$ and x_t^{uv} to 0. The paths now have the standard path colourings.

Case (d): $\eta(u) \in \{2q+r, 3q+r, \ldots (k-1)q+r\}, \eta'(u) = 0$

(i) Change the colour of x_t^{vu} to q, and recolour x_i^{uv} to $(i+1)q$ for $i = t, t-1, \ldots, 0$.

(ii) At this point all vertices of $\{x_0^{uv}, x_t^{vu} : v \in N_G(u)\}$ have colour q. Now change the colour of u to 0.

(iii) The paths have standard path colourings.

Case (e): $\eta(u) = q+r, \eta'(u) = 0$

(i) Recolour x_t^{vu} to $p-1$ and x_0^{uv} to $p-1$.

(ii) At this point all vertices of $\{x_0^{uv}, x_t^{vu} : v \in N_G(u)\}$ have colour $p-1$. Change the colour of u to $q-1$ (temporarily). Note that this colour is compatible with all neighbours v of u in G since $\eta(v) \in \{2q+r, 3q+r, \ldots, (k-1)q+r\}$.

(iii) Since $\eta(v) \in \{2q+r, 3q+r, \ldots, (k-1)q+r\}$, for $i = t, \ldots, 0$ we can recolour x_i^{uv} to $(i+2)q-1$.

(iv) Change the colour of x_{t-1}^{vu} to $-q-1$, x_t^{vu} to $2q-1$ and u to 0.

(v) Finally, recolour x_i^{uv} to $(i+1)q$ for $i = 0, 1, \dots, t$. Recolour x_t^{vu} to q, and x_{t-1}^{vu} to $r - 2q$. The paths have the standard path colourings.

This completes the proof of the lemma, and of Theorem 1.24. $\square$

References

[1] J. Adler and U. Lev. Bootstrap percolation: Visualizations and applications. *Braz. J. Phys.*, 33:641–644, 2003.

[2] A. V. Aho, J. E. Hopcroft, and J. D. Ullman. *Data Structures and Algorithms.* Addison-Wesley, 1983.

[3] M. Aizenman and J. L. Lebowitz. Metastability effects in bootstrap percolation. *J. Phys. A*, 21(19):3801–3813, 1988.

[4] N. Alon. An extremal problem for sets with applications to graph theory. *J. Combin. Theory Ser. A*, 40(1):82–89, 1985.

[5] N. Alon, P. Erdős, R. Holzman, and M. Krivelevich. On k-saturated graphs with restrictions on the degrees. *J. Graph Theory*, 23(1):1–20, 1996.

[6] H. Amini. Bootstrap percolation in living neural networks. *J. Stat. Phys.*, 141(3):459–475, 2010.

[7] H. Amini and N. Fountoulakis. Bootstrap percolation in power-law random graphs. *J. Stat. Phys.*, 155(1):72–92, 2014.

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

[9] R. Baber. Turán densities of hypercubes. Preprint available as arXiv: 1201.3587v2.

[10] R. Baber and J. Talbot. Hypergraphs do jump. *Combin. Probab. Comput.*, 20(2):161–171, 2011.

[11] R. Baber and J. Talbot. A solution to the 2/3 conjecture. *SIAM J. Discrete Math.*, 28(2):756–766, 2014.

[12] D. B. Bahr and E. Passerini. Statistical mechanics of collective behavior: Macro-sociology. *J. Math. Sociol.*, 23(1):29–49, 1998.

[13] P. Balister, B. Bollobás, J. R. Johnson, and M. Walters. Random majority percolation. *Random Structures Algorithms*, 36(3):315–340, 2010.

[14] P. Balister, B. Bollobás, M. Przykucki, and P. Smith. Subcritical $\mathcal{U}$-bootstrap percolation models have non-trivial phase transitions. *Trans. Amer. Math. Soc.*, 368(10):7385–7411, 2016.

[15] J. Balogh and B. Bollobás. Bootstrap percolation on the hypercube. *Probab. Theory Related Fields*, 134(4):624–648, 2006.

[16] J. Balogh, B. Bollobás, H. Duminil-Copin, and R. Morris. The sharp threshold for bootstrap percolation in all dimensions. *Trans. Amer. Math. Soc.*, 364(5):2667–2701, 2012.

[17] J. Balogh, B. Bollobás, and R. Morris. Bootstrap percolation in three dimensions. *Ann. Probab.*, 37(4):1329–1380, 2009.

[18] J. Balogh, B. Bollobás, and R. Morris. Majority bootstrap percolation on the hypercube. *Combin. Probab. Comput.*, 18(1-2):17–51, 2009.

[19] J. Balogh, B. Bollobás, and R. Morris. Bootstrap percolation in high dimensions. *Combin. Probab. Comput.*, 19(5-6):643–692, 2010.

[20] J. Balogh, B. Bollobás, and R. Morris. Graph bootstrap percolation. *Random Structures Algorithms*, 41(4):413–440, 2012.

[21] J. Balogh, B. Bollobás, R. Morris, and O. Riordan. Linear algebra and bootstrap percolation. *J. Combin. Theory Ser. A*, 119(6):1328–1335, 2012.

[22] J. Balogh, P. Hu, B. Lidický, and H. Liu. Upper bounds on the size of 4- and 6-cycle-free subgraphs of the hypercube. *European J. Combin.*, 35:75–85, 2014.

[23] J. Balogh, P. Hu, B. Lidický, O. Pikhurko, B. Udvari, and J. Volec. Minimum number of monotone subsequences of length 4 in permutations. *Combin. Probab. Comput.*, 24(4):658–679, 2015.

[24] J. Balogh, H. Liu, Š. Petříčková, and M. Sharifzadeh. The typical structure of maximal triangle-free graphs. *Forum Math. Sigma*, 3:e20, 19, 2015.

[25] J. Balogh, H. Liu, M. Sharifzadeh, and A. Treglown. The number of maximal sum-free subsets of integers. *Proc. Amer. Math. Soc.*, 143(11):4713–4721, 2015.

[26] J. Balogh, R. Morris, and W. Samotij. Independent sets in hypergraphs. *J. Amer. Math. Soc.*, 28(3):669–709, 2015.

[27] J. Balogh, R. Mycroft, and A. Treglown. A random version of Sperner's theorem. *J. Combin. Theory Ser. A*, 128:104–110, 2014.

[28] J. Balogh, Y. Peres, and G. Pete. Bootstrap percolation on infinite trees and non-amenable groups. *Combin. Probab. Comput.*, 15(5):715–730, 2006.

[29] J. Balogh and Š. Petříčková. The number of the maximal triangle-free graphs. *Bull. Lond. Math. Soc.*, 46(5):1003–1006, 2014.

[30] J. Balogh and B. G. Pittel. Bootstrap percolation on the random regular graph. *Random Structures Algorithms*, 30(1-2):257–286, 2007.

[31] J. Balogh, A. Treglown, and A. Z. Wagner. Applications of graph containers in the boolean lattice. Preprint available as arXiv:1602.05870v1.

[32] J. Balogh and A. Z. Wagner. Further applications of the container method. Preprint available as arXiv:1601.07809v1.

[33] C. A. Barefoot, L. H. Clark, R. C. Entringer, T. D. Porter, L. A. Székely, and Zs. Tuza. Cycle-saturated graphs of minimum size. *Discrete Math.*, 150(1-3):31–48, 1996.

[34] M. Biskup and R. H. Schonmann. Metastable behavior for bootstrap percolation on regular trees. *J. Stat. Phys.*, 136(4):667–676, 2009.

[35] B. Bollobás. On generalized graphs. *Acta Math. Acad. Sci. Hungar*, 16:447–452, 1965.

[36] B. Bollobás. On a conjecture of Erdős, Hajnal and Moon. *Amer. Math. Monthly*, 74:178–179, 1967.

[37] B. Bollobás. Weakly k-saturated graphs. In *Beiträge zur Graphentheorie (Kolloquium, Manebach, 1967)*, pages 25–31. Teubner, Leipzig, 1968.

[38] B. Bollobás. Relations between sets of complete subgraphs. In *Proceedings of the Fifth British Combinatorial Conference (Univ. Aberdeen, Aberdeen, 1975)*, pages 79–84. Congressus Numerantium, No. XV, Utilitas Math., Winnipeg, Man., 1976.

[39] B. Bollobás. *Extremal graph theory*, volume 11 of *London Mathematical Society Monographs*. Academic Press, Inc., London-New York, 1978.

[40] B. Bollobás. *Combinatorics*. Cambridge University Press, Cambridge, 1986. Set systems, hypergraphs, families of vectors and combinatorial probability.

[41] B. Bollobás. Extremal graph theory. In *Handbook of combinatorics, Vol. 2*, pages 1231–1292. Elsevier, Amsterdam, 1995.

[42] B. Bollobás. *Random graphs*, volume 73 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, second edition, 2001.

[43] B. Bollobás. *The art of mathematics: Coffee time in Memphis*. Cambridge University Press, New York, 2006.

[44] B. Bollobás, K. Gunderson, C. Holmgren, S. Janson, and M. Przykucki. Bootstrap percolation on Galton-Watson trees. *Electron. J. Probab.*, 19:no. 13, 27, 2014.

[45] B. Bollobás, M. Przykucki, O. Riordan, and J. Sahasrabudhe. On the maximum running time in graph bootstrap percolation. Preprint available as arXiv:1510.07096v1.

[46] B. Bollobás and O. Riordan. Sparse graphs: metrics and random models. *Random Structures Algorithms*, 39(1):1–38, 2011.

[47] B. Bollobás, P. Smith, and A. Uzzell. Monotone cellular automata in a random environment. *Combin. Probab. Comput.*, 24(4):687–722, 2015.

[48] P. Bonsma and L. Cereceda. Finding paths between graph colourings: PSPACE-completeness and superpolynomial distances. *Theoret. Comput. Sci.*, 410(50):5215–5226, 2009.

[49] C. Borgs, J. Chayes, and L. Lovász. Moments of two-variable functions and the uniqueness of graph limits. *Geom. Funct. Anal.*, 19(6):1597–1619, 2010.

[50] C. Borgs, J. Chayes, L. Lovász, V. T. Sós, B. Szegedy, and K. Vesztergombi. Graph limits and parameter testing. In *STOC'06: Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, pages 261–270. ACM, New York, 2006.

[51] C. Borgs, J. T. Chayes, L. Lovász, V. T. Sós, and K. Vesztergombi. Convergent sequences of dense graphs. I. Subgraph frequencies, metric properties and testing. *Adv. Math.*, 219(6):1801–1851, 2008.

[52] C. Borgs, J. T. Chayes, L. Lovász, V. T. Sós, and K. Vesztergombi. Convergent sequences of dense graphs II. Multiway cuts and statistical physics. *Ann. of Math. (2)*, 176(1):151–219, 2012.

[53] M. Borowiecki and E. Sidorowicz. Weakly $\mathcal{P}$-saturated graphs. *Discuss. Math. Graph Theory*, 22(1):17–29, 2002. Conference on Graph Theory (Elgersburg, 2000).

[54] P. Brass, H. Harborth, and H. Nienborg. On the maximum number of edges in a C_4-free subgraph of Q_n. *J. Graph Theory*, 19(1):17–23, 1995.

[55] R. C. Brewster, S. McGuinness, B. Moore, and J. A. Noel. A dichotomy theorem for circular colouring reconfiguration. *Theoret. Comput. Sci.*, 639:1–13, 2016.

[56] R. C. Brewster and J. A. Noel. Mixing homomorphisms, recolorings, and extending circular precolorings. *J. Graph Theory*, 80(3):173–198, 2015.

[57] E. Candellero and N. Fountoulakis. Bootstrap percolation and the geometry of complex networks. *Stochastic Process. Appl.*, 126(1):234–264, 2016.

[58] L. Cereceda, J. van den Heuvel, and M. Johnson. Finding paths between 3-colorings. *J. Graph Theory*, 67(1):69–82, 2011.

[59] R. Cerf and E. N. M. Cirillo. Finite size scaling in three-dimensional bootstrap percolation. *Ann. Probab.*, 27(4):1837–1850, 1999.

[60] R. Cerf and F. Manzo. The threshold regime of finite volume bootstrap percolation. *Stochastic Process. Appl.*, 101(1):69–82, 2002.

[61] R. Cerf and F. Manzo. A d-dimensional nucleation and growth model. *Probab. Theory Related Fields*, 155(1-2):427–449, 2013.

[62] J. Chalupa, P. L. Leath, and G. R Reich. Bootstrap percolation on a Bethe latice. *J. Phys. C*, 12:L31–L35, 1979.

[63] Y.-C. Chen. Minimum C_5-saturated graphs. *J. Graph Theory*, 61(2):111–126, 2009.

[64] Y.-C. Chen. All minimum C_5-saturated graphs. *J. Graph Theory*, 67(1):9–26, 2011.

[65] S. Choi and P. Guan. Minimum critical squarefree subgraph of a hypercube. In *Proceedings of the Thirty-Ninth Southeastern International Conference on Combinatorics, Graph Theory and Computing*, volume 189, pages 57–64, 2008.

[66] F. R. K. Chung. Subgraphs of a hypercube containing no small even cycles. *J. Graph Theory*, 16(3):273–286, 1992.

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

[68] T. Coker and K. Gunderson. A sharp threshold for a modified bootstrap percolation with recovery. *J. Stat. Phys.*, 157(3):531–570, 2014.

[69] M. Collares Neto and R. Morris. Maximum-size antichains in random sets. Preprint available as arXiv:1404.5258v3.

[70] M. Conder. Hexagon-free subgraphs of hypercubes. *J. Graph Theory*, 17(4):477–479, 1993.

[71] D. Conlon. An extremal theorem in the hypercube. *Electron. J. Combin.*, 17(1):Research Paper 111, 7 pages, 2010.

[72] D. Conlon and J. Fox. Bounds for graph regularity and removal lemmas. *Geom. Funct. Anal.*, 22(5):1191–1256, 2012.

[73] D. Conlon and W. T. Gowers. Combinatorial theorems in sparse random sets. Preprint available as arXiv:1011.4310v2.

[74] J. W. Cooper, T. Kaiser, D. Kráľ, and J. A. Noel. Weak regularity and finitely forcible graph limits. Preprint available as arXiv:1507.00067v2.

[75] S. Das, W. Gan, and B. Sudakov. The minimum number of disjoint pairs in set systems and related problems. Preprint available as arXiv:1305.6715v1. To appear in Combinatorica.

[76] S. Das, W. Gan, and B. Sudakov. Sperner's theorem and a problem of Erdős, Katona and Kleitman. *Combin. Probab. Comput.*, 24(4):585–608, 2015.

[77] A. N. Day. Saturated graphs of prescribed minimum degree. Preprint available as arXiv:1407.6664v1.

[78] P. Diaconis, S. Holmes, and S. Janson. Threshold graph limits and random threshold graphs. *Internet Math.*, 5(3):267–320 (2009), 2008.

[79] A. P. Dove, J. R. Griggs, R. J. Kang, and J.-S. Sereni. Supersaturation in the Boolean lattice. *Integers*, 14A:Paper No. A4, 7, 2014.

[80] D. A. Duffus and D. Hanson. Minimal k-saturated and color critical graphs of prescribed minimum degree. *J. Graph Theory*, 10(1):55–67, 1986.

[81] G. Elek. On limits of finite graphs. *Combinatorica*, 27(4):503–507, 2007.

[82] A. C. D. van Enter. Proof of Straley's argument for bootstrap percolation. *J. Statist. Phys.*, 48(3-4):943–945, 1987.

[83] P. Erdős. On a lemma of Littlewood and Offord. *Bull. Amer. Math. Soc.*, 51:898–902, 1945.

[84] P. Erdős and A. H. Stone. On the structure of linear graphs. *Bull. Amer. Math. Soc.*, 52:1087–1091, 1946.

[85] P. Erdős. On a theorem of Rademacher-Turán. *Illinois J. Math.*, 6:122–127, 1962.

[86] P. Erdős. Some recent results on extremal problems in graph theory. Results. In *Theory of Graphs (Internat. Sympos., Rome, 1966)*, pages 117–123 (English); pp. 124–130 (French). Gordon and Breach, New York; Dunod, Paris, 1967.

[87] P. Erdős. On the number of complete subgraphs and circuits contained in graphs. *Časopis Pěst. Mat.*, 94:290–296, 1969.

[88] P. Erdős. On some problems in graph theory, combinatorial analysis and combinatorial number theory. In *Graph theory and combinatorics (Cambridge, 1983)*, pages 1–17. Academic Press, London, 1984.

[89] P. Erdős, A. Hajnal, and J. W. Moon. A problem in graph theory. *Amer. Math. Monthly*, 71:1107–1110, 1964.

[90] P. Erdős and R. Holzman. On maximal triangle-free graphs. *J. Graph Theory*, 18(6):585–594, 1994.

[91] P. Erdős and A. Rényi. On random graphs. I. *Publ. Math. Debrecen*, 6:290–297, 1959.

[92] P. Erdős and M. Simonovits. Supersaturated graphs and hypergraphs. *Combinatorica*, 3(2):181–192, 1983.

[93] J. R. Faudree, R. J. Faudree, and J. R. Schmitt. A survey of minimum saturated graphs. *Electron. J. Combin.*, 18:Dynamic Survey 19, 36 pp. (electronic), 2011.

[94] L. R. Fontes, R. H. Schonmann, and V. Sidoravicius. Stretched exponential fixation in stochastic Ising models at zero temperature. *Comm. Math. Phys.*, 228(3):495–518, 2002.

[95] P. Frankl. An extremal problem for two families of sets. *European J. Combin.*, 3(2):125–127, 1982.

[96] A. Frieze and R. Kannan. Quick approximation to matrices and applications. *Combinatorica*, 19(2):175–220, 1999.

[97] Z. Füredi. On a theorem of Erdős and Simonovits on graphs not containing the cube. In *Number theory, analysis, and combinatorics*, De Gruyter Proc. Math., pages 113–125. De Gruyter, Berlin, 2014.

[98] Z. Füredi. A proof of the stability of extremal graphs, Simonovits' stability from Szemerédi's regularity. *J. Combin. Theory Ser. B*, 115:66–71, 2015.

[99] Z. Füredi and Y. Kim. Cycle-saturated graphs with minimum number of edges. *J. Graph Theory*, 73(2):203–215, 2013.

[100] Z. Füredi and L. Özkahya. On 14-cycle-free subgraphs of the hypercube. *Combin. Probab. Comput.*, 18(5):725–729, 2009.

[101] Z. Füredi and L. Özkahya. On even-cycle-free subgraphs of the hypercube. *J. Combin. Theory Ser. A*, 118(6):1816–1819, 2011.

[102] Z. Füredi and Á. Seress. Maximal triangle-free graphs with restrictions on the degrees. *J. Graph Theory*, 18(1):11–24, 1994.

[103] W. Gan, D. Korándi, and B. Sudakov. $K_{s,t}$-saturated bipartite graphs. *European J. Combin.*, 45:12–20, 2015.

[104] D. Gerbner, B. Keszegh, N. Lemons, C. Palmer, D. Pálvölgyi, and B. Patkós. Saturating Sperner Families. *Graphs Combin.*, 29(5):1355–1364, 2013.

[105] R. Glebov, A. Grzesik, T. Klimošová, and D. Král'. Finitely forcible graphons and permutons. *J. Combin. Theory Ser. B*, 110:112–135, 2015.

[106] R. Glebov, C. Hoppen, T. Klimošová, Y. Kohayakawa, D. Kráľ, and H. Liu. Large permutations and parameter testing. Preprint available as arXiv: 1412.5622v2.

[107] R. Glebov, T. Klimošová, and D. Král'. Infinite dimensional finitely forcible graphon. Preprint available as arXiv:1404.2743v1.

[108] R. Glebov, D. Král', and J. Volec. Compactness and finite forcibility of graphons. Preprint available as arXiv:1309.6695v6.

[109] A. W. Goodman. On sets of acquaintances and strangers at any party. *Amer. Math. Monthly*, 66:778–783, 1959.

[110] R. Gould, T. Łuczak, and J. Schmitt. Constructive upper bounds for cycle-saturated graphs of minimum size. *Electron. J. Combin.*, 13(1):Research Paper 29, 19 pages, 2006.

[111] J. Gravner, A. E. Holroyd, and R. Morris. A sharper threshold for bootstrap percolation in two dimensions. *Probab. Theory Related Fields*, 153(1-2):1–23, 2012.

[112] B. Green. The Cameron-Erdős conjecture. *Bull. London Math. Soc.*, 36(6):769–778, 2004.

[113] A. Grzesik. On the maximum number of five-cycles in a triangle-free graph. *J. Combin. Theory Ser. B*, 102(5):1061–1066, 2012.

[114] K. Gunderson, S. Koch, and M. Przykucki. The time of graph bootstrap percolation. Preprint available as arXiv:1503.01454v2.

[115] K. Gunderson and M. Przykucki. Lower bounds for bootstrap percolation on Galton-Watson trees. *Electron. Commun. Probab.*, 19:no. 42, 7, 2014.

[116] N. G. Hadžiivanov and V. S. Nikiforov. The Nordhaus-Stewart-Moon-Moser inequality. *Serdica*, 4(4):344–350, 1978.

[117] A. Hajnal. A theorem on k-saturated graphs. *Canad. J. Math.*, 17:720–724, 1965.

[118] R. W. Hamming. Error detecting and error correcting codes. *Bell System Tech. J.*, 29:147–160, 1950.

[119] H. Hatami, J. Hladký, D. Král, S. Norine, and A. Razborov. Non-three-colourable common graphs exist. *Combin. Probab. Comput.*, 21(5):734–742, 2012.

[120] H. Hatami, J. Hladký, D. Kráľ, S. Norine, and A. Razborov. On the number of pentagons in triangle-free graphs. *J. Combin. Theory Ser. A*, 120(3):722–732, 2013.

[121] P. E. Haxell, Y. Kohayakawa, and T. Łuczak. Turán's extremal problem in random graphs: forbidding even cycles. *J. Combin. Theory Ser. B*, 64(2):273–287, 1995.

[122] J. van den Heuvel. The complexity of change. In *Surveys in combinatorics 2013*, volume 409 of *London Math. Soc. Lecture Note Ser.*, pages 127–160. Cambridge Univ. Press, Cambridge, 2013.

[123] A. E. Holroyd. Sharp metastability threshold for two-dimensional bootstrap percolation. *Probab. Theory Related Fields*, 125(2):195–224, 2003.

[124] C. Hoppen, Y. Kohayakawa, C. G. Moreira, B. Ráth, and R. M. Sampaio. Limits of permutation sequences. *J. Combin. Theory Ser. B*, 103(1):93–113, 2013.

[125] C. Hoppen, Y. Kohayakawa, C. G. Moreira, and R. M. Sampaio. Limits of permutation sequences through permutation regularity. Preprint available as arXiv:1106.1663v1.

[126] C. Hoppen, Y. Kohayakawa, C. G. Moreira, and R. M. Sampaio. Testing permutation properties through subpermutations. *Theoret. Comput. Sci.*, 412(29):3555–3567, 2011.

[127] T. Ito, E. D. Demaine, N. J. A. Harvey, C. H. Papadimitriou, M. Sideri, R. Uehara, and Y. Uno. On the complexity of reconfiguration problems. *Theoret. Comput. Sci.*, 412(12-14):1054–1065, 2011.

[128] F. Jaeger and C. Payan. Nombre maximal d'arêtes d'un hypergraphe τ-critique de rang h. *C. R. Acad. Sci. Paris Sér. A-B*, 273:A221–A223, 1971.

[129] S. Janson. On percolation in random graphs with given vertex degrees. *Electron. J. Probab.*, 14:no. 5, 87–118, 2009.

[130] S. Janson. Poset limits and exchangeable random posets. *Combinatorica*, 31(5):529–563, 2011.

[131] S. Janson, T. Łuczak, and A. Rucinski. *Random graphs.* Wiley-Interscience Series in Discrete Mathematics and Optimization. Wiley-Interscience, New York, 2000.

[132] S. Janson, T. Łuczak, T. Turova, and T. Vallier. Bootstrap percolation on the random graph $G_{n,p}$. *Ann. Appl. Probab.*, 22(5):1989–2047, 2012.

[133] J. R. Johnson and T. Pinto. Saturated subgraphs of the hypercube. Preprint available as arXiv:1406.1766v1.

[134] M. Johnson, D. Kratsch, S. Kratsch, V. Patel, and D. Paulusma. Finding shortest paths between graph colourings. In *Parameterized and exact computation*, volume 8894 of *Lecture Notes in Comput. Sci.*, pages 221–233. Springer, Cham, 2014.

[135] J. Kahn. Entropy, independent sets and antichains: a new approach to Dedekind's problem. *Proc. Amer. Math. Soc.*, 130(2):371–378, 2002.

[136] G. Kalai. Hyperconnectivity of graphs. *Graphs Combin.*, 1(1):65–79, 1985.

[137] A. R. Kansal, S. Torquato, G. R. Harsh IV, E. A. Chiocca, and T. S. Deisboeck. Cellular automaton of idealized brain tumor growth dynamics. *BioSystems*, 55:119–127, 2000.

[138] L. Kászonyi and Z. Tuza. Saturated graphs with minimal number of edges. *J. Graph Theory*, 10(2):203–210, 1986.

[139] G. O. H. Katona. Solution of a problem of A. Ehrenfeucht and J. Mycielski. *J. Combinatorial Theory Ser. A*, 17:265–266, 1974.

[140] P. Keevash. The existence of designs. Preprint available as arXiv:1401.3665v1.

[141] D. Kleitman. A conjecture of Erdős-Katona on commensurable pairs among subsets of an n-set. In *Theory of Graphs (Proc. Colloq., Tihany, 1966)*, pages 215–218. Academic Press, New York, 1968.

[142] D. Kleitman. On Dedekind's problem: The number of monotone Boolean functions. *Proc. Amer. Math. Soc.*, 21:677–682, 1969.

[143] D. Kleitman and G. Markowsky. On Dedekind's problem: the number of isotone Boolean functions. II. *Trans. Amer. Math. Soc.*, 213:373–390, 1975.

[144] D. J. Kleitman and K. J. Winston. On the number of graphs without 4-cycles. *Discrete Math.*, 41(2):167–172, 1982.

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

[146] D. Korándi and B. Sudakov. Saturation in random graphs. Preprint available as arXiv:1510.09187v2.

[147] A. D. Korshunov. The number of monotone Boolean functions. *Problemy Kibernet.*, 38:5–108, 272, 1981.

[148] D. Král', C.-H. Liu, J.-S. Sereni, P. Whalen, and Z. B. Yilma. A new bound for the 2/3 conjecture. *Combin. Probab. Comput.*, 22(3):384–393, 2013.

[149] D. Král', L. Mach, and J.-S. Sereni. A new lower bound based on Gromov's method of selecting heavily covered points. *Discrete Comput. Geom.*, 48(2):487–498, 2012.

[150] D. Král' and O. Pikhurko. Quasirandom permutations are characterized by 4-point densities. *Geom. Funct. Anal.*, 23(2):570–579, 2013.

[151] I. Leader and E. Long. Tilted Sperner families. *Discrete Appl. Math.*, 163(part 2):194–198, 2014.

[152] L. Lovász. Flats in matroids and geometric graphs. In *Combinatorial surveys (Proc. Sixth British Combinatorial Conf., Royal Holloway Coll., Egham, 1977)*, pages 45–86. Academic Press, London, 1977.

[153] L. Lovász. *Large networks and graph limits*, volume 60 of *American Mathematical Society Colloquium Publications*. American Mathematical Society, Providence, RI, 2012.

[154] L. Lovász and M. Simonovits. On the number of complete subgraphs of a graph. In *Proceedings of the Fifth British Combinatorial Conference (Univ. Aberdeen, Aberdeen, 1975)*, pages 431–441. Congressus Numerantium, No. XV. Utilitas Math., Winnipeg, Man., 1976.

[155] L. Lovász and M. Simonovits. On the number of complete subgraphs of a graph. II. In *Studies in pure mathematics*, pages 459–495. Birkhäuser, Basel, 1983.

[156] L. Lovász and V. T. Sós. Generalized quasirandom graphs. *J. Combin. Theory Ser. B*, 98(1):146–163, 2008.

[157] L. Lovász and B. Szegedy. Limits of dense graph sequences. *J. Combin. Theory Ser. B*, 96(6):933–957, 2006.

[158] L. Lovász and B. Szegedy. Testing properties of graphs and functions. *Israel J. Math.*, 178:113–156, 2010.

[159] L. Lovász and B. Szegedy. Finitely forcible graphons. *J. Combin. Theory Ser. B*, 101(5):269–301, 2011.

[160] D. Lubell. A short proof of Sperner's lemma. *J. Combinatorial Theory*, 1:299, 1966.

[161] F. J. MacWilliams and N. J. A. Sloane. *The theory of error-correcting codes. I.* North-Holland Publishing Co., Amsterdam-New York-Oxford, 1977. North-Holland Mathematical Library, Vol. 16.

[162] F. J. MacWilliams and N. J. A. Sloane. *The theory of error-correcting codes. II.* North-Holland Publishing Co., Amsterdam-New York-Oxford, 1977. North-Holland Mathematical Library, Vol. 16.

[163] S. Maerivoet and B. De Moor. Cellular automata models of road traffic. *Phys. Rep.*, 419(1):1–64, 2005.

[164] W. Mantel. Problem 28. *Wiskundige Opgaven*, 10:60–61, 1907.

[165] K. Matzke. The saturation time of graph bootstrap percolation. Preprint available as arXiv:1510.06156v2.

[166] L. D. Mešalkin. A generalization of Sperner's theorem on the number of subsets of a finite set. *Teor. Verojatnost. i Primenen*, 8:219–220, 1963.

[167] J. W. Moon and L. Moser. On a problem of Turán. *Magyar Tud. Akad. Mat. Kutató Int. Közl.*, 7:283–286, 1962.

[168] R. Morris. Zero-temperature Glauber dynamics on $\mathbb{Z}^d$. *Probab. Theory Related Fields*, 149(3-4):417–434, 2011.

[169] N. Morrison and J. A. Noel. Extremal bounds for bootstrap percolation in the hypercube. Preprint available as arXiv:1506.04686v1.

[170] N. Morrison, J. A. Noel, and A. Scott. Saturation in the hypercube and bootstrap percolation. Preprint available as arXiv:1408.5488v2. To appear in Combin. Probab. Comput.

[171] N. Morrison, J. A. Noel, and A. Scott. On saturated k-Sperner systems. *Electron. J. Combin.*, 21(3):Paper 3.22, 17 pages, 2014.

[172] G. Moshkovitz and A. Shapira. Exact bounds for some hypergraph saturation problems. *J. Combin. Theory Ser. B*, 111:242–248, 2015.

[173] J. S. Myers. The minimum number of monotone subsequences. *Electron. J. Combin.*, 9(2):Research paper 4, 17 pp. (electronic), 2002/03. Permutation patterns (Otago, 2003).

[174] J. von Neumann. *Theory of Self-Reproducing Automata*. University of Illinois Press, 1966. A. W. Burks (ed.).

[175] V. Nikiforov. The number of cliques in graphs of given order and size. *Trans. Amer. Math. Soc.*, 363(3):1599–1618, 2011.

[176] J. A. Noel, A. Scott, and B. Sudakov. Supersaturation in posets and applications involving the container method. In preparation.

[177] E. A. Nordhaus and B. M. Stewart. Triangles in an ordinary graph. *Canad. J. Math.*, 15:33–41, 1963.

[178] L. T. Ollmann. $K_{2,2}$ saturated graphs with a minimal number of edges. In *Proceedings of the Third Southeastern Conference on Combinatorics, Graph Theory, and Computing (Florida Atlantic Univ., Boca Raton, Fla., 1972)*, pages 367–392. Florida Atlantic Univ., Boca Raton, Fla., 1972.

[179] D. Osthus. Maximum antichains in random subsets of a finite set. *J. Combin. Theory Ser. A*, 90(2):336–346, 2000.

[180] B. Patkós. Supersaturation and stability for forbidden subposet problems. *J. Combin. Theory Ser. A*, 136:220–237, 2015.

[181] O. Pikhurko. *Extremal Hypergraphs.* PhD thesis, University of Cambridge, October 1999.

[182] O. Pikhurko. The minimum size of saturated hypergraphs. *Combin. Probab. Comput.*, 8(5):483–492, 1999.

[183] O. Pikhurko. Weakly saturated hypergraphs and exterior algebra. *Combin. Probab. Comput.*, 10(5):435–451, 2001.

[184] O. Pikhurko. Results and open problems on minimum saturated hypergraphs. *Ars Combin.*, 72:111–127, 2004.

[185] O. Pikhurko and A. Razborov. Asymptotic structure of graphs with the minimum number of triangles. Preprint available as arXiv:1204.2846v2.

[186] O. Pikhurko and E. R. Vaughan. Minimum number of k-cliques in graphs with bounded independence number. *Combin. Probab. Comput.*, 22(6):910–934, 2013.

[187] A. A. Razborov. Flag algebras. *J. Symbolic Logic*, 72(4):1239–1282, 2007.

[188] A. A. Razborov. On the minimal density of triangles in graphs. *Combin. Probab. Comput.*, 17(4):603–618, 2008.

[189] A. A. Razborov. On the minimal density of triangles in graphs. *Combin. Probab. Comput.*, 17(4):603–618, 2008.

[190] A. A. Razborov. On 3-hypergraphs with forbidden 4-vertex configurations. *SIAM J. Discrete Math.*, 24(3):946–963, 2010.

[191] C. Reiher. The clique density theorem. Preprint available as arXiv:1212.2454v2.

[192] V. Rödl. On a packing and covering problem. *European J. Combin.*, 6(1):69–78, 1985.

[193] V. Rödl. On universality of graphs with uniformly distributed edges. *Discrete Math.*, 59(1-2):125–134, 1986.

[194] G.-C. Rota. *Indiscrete thoughts.* Birkhäuser Boston, Inc., Boston, MA, 1997. With forewords by Reuben Hersh and Robert Sokolowski, Edited with notes and an epilogue by Fabrizio Palombi.

[195] W. Samotij. Stability results for random discrete structures. *Random Structures Algorithms*, 44(3):269–289, 2014.

[196] W. Samotij. Counting independent sets in graphs. *European J. Combin.*, 48:5–18, 2015.

[197] W. Samotij and B. Sudakov. On the number of monotone sequences. *J. Combin. Theory Ser. B*, 115:132–163, 2015.

[198] W. Samotij and B. Sudakov. The number of additive triples in subsets of abelian groups. *Math. Proc. Cambridge Philos. Soc.*, 160(3):495–512, 2016.

[199] A. A. Sapozhenko. On the number of connected subsets with given cardinality of the boundary in bipartite graphs. *Metody Diskret. Analiz.*, 45:42–70, 96, 1987.

[200] A. A. Sapozhenko. The number of antichains in multilayered ranked sets. *Diskret. Mat.*, 1(2):110–128, 1989.

[201] A. A. Sapozhenko. On the number of independent sets in extenders. *Diskret. Mat.*, 13(1):56–62, 2001.

[202] A. A. Sapozhenko. Independent sets in quasi-regular graphs. *European J. Combin.*, 27(7):1206–1210, 2006.

[203] F. Sausset, C. Toninelli, G. Biroli, and G. Tarjus. Bootstrap percolation and kinetically constrained models on hyperbolic lattices. *J. Stat. Phys.*, 138(1-3):411–430, 2010.

[204] D. Saxton and A. Thomason. List colourings of regular hypergraphs. *Combin. Probab. Comput.*, 21(1-2):315–322, 2012.

[205] D. Saxton and A. Thomason. Hypergraph containers. *Invent. Math.*, 201(3):925–992, 2015.

[206] M. Schacht. Extremal results for random discrete structures. preprint, November 2014.

[207] R. H. Schonmann. On the behavior of some cellular automata related to bootstrap percolation. *Ann. Probab.*, 20(1):174–193, 1992.

[208] M. Simonovits. A method for solving extremal problems in graph theory, stability problems. In *Theory of Graphs (Proc. Colloq., Tihany, 1966)*, pages 279–319. Academic Press, New York, 1968.

[209] A. D. Sokal. The multivariate Tutte polynomial (alias Potts model) for graphs and matroids. In *Surveys in combinatorics 2005*, volume 327 of *London Math. Soc. Lecture Note Ser.*, pages 173–226. Cambridge Univ. Press, Cambridge, 2005.

[210] E. Sperner. Ein Satz über Untermengen einer endlichen Menge. *Math. Z.*, 27(1):544–548, 1928.

[211] T. G. Tarján. Complexity of lattice-configurations. *Studia Sci. Math. Hungar.*, 10(1–2):203–211, 1975.

[212] A. Thomason. Pseudo-random graphs. In A. Barlotti, M. Biliotti, A. Cossu, G. Korchmaros, and G. Tallini, editors, *North-Holland Mathematics Studies*, volume 144, pages 307–331. North-Holland, 1987.

[213] A. Thomason. Random graphs, strongly regular graphs and pseudorandom graphs. In *Surveys in combinatorics 1987 (New Cross, 1987)*, volume 123 of *London Math. Soc. Lecture Note Ser.*, pages 173–195. Cambridge Univ. Press, Cambridge, 1987.

[214] M. Truszczyński and Z. Tuza. Asymptotic results on saturated graphs. *Discrete Math.*, 87(3):309–314, 1991.

[215] P. Turán. Eine Extremalaufgabe aus der Graphentheorie. *Mat. Fiz. Lapok*, 48:436–452, 1941.

[216] Z. Tuza. Extremal problems on saturated graphs and hypergraphs. *Ars Combin.*, 25(B):105–113, 1988. Eleventh British Combinatorial Conference (London, 1987).

[217] Z. Tuza. C_4-saturated graphs of minimum size. *Acta Univ. Carolin. Math. Phys.*, 30(2):161–167, 1989. 17th Winter School on Abstract Analysis (Srní, 1989).

[218] S. Ulam. Random processes and transformations. In *Proceedings of the International Congress of Mathematicians, Vol. 2, Cambridge, Mass., 1950*, pages 264–275. Amer. Math. Soc., Providence, R. I., 1952.

[219] A. Vince. Star chromatic number. *J. Graph Theory*, 12(4):551–559, 1988.

[220] W. Wessel. Über eine Klasse paarer Graphen. I. Beweis einer Vermutung von Erdős, Hajnal und Moon. *Wiss. Z. Techn. Hochsch. Ilmenau*, 12:253–256, 1966.

[221] W. Wessel. Über eine Klasse paarer Graphen. II. Bestimmung der Minimalgraphen. *Wiss. Z. Techn. Hochsch. Ilmenau*, 13:423–426, 1967.

[222] H. S. Wilf. Three problems in combinatorial asymptotics. *J. Combin. Theory Ser. A*, 35(2):199–207, 1983.

[223] M. Wrochna. Reconfiguration and structural graph theory. Master's thesis, University of Warsaw, 2014.

[224] M. Wrochna. Homomorphism reconfiguration via homotopy. In *32nd International Symposium on Theoretical Aspects of Computer Science*, volume 30 of *LIPIcs. Leibniz Int. Proc. Inform.*, pages 730–742. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2015.

[225] K. Yamamoto. Logarithmic order of free distributive lattice. *J. Math. Soc. Japan*, 6:343–353, 1954.

[226] X. Zhu. Circular chromatic number: a survey. *Discrete Math.*, 229(1-3):371–410, 2001. Combinatorics, graph theory, algorithms and applications.

[227] A. A. Zykov. On some properties of linear complexes. *Mat. Sbornik N.S.*, 24(66):163–188, 1949.