Problems on Colourings and Reconstruction



Youri Tamitegama

Keble College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2025

To Joe

Acknowledgements

Miniseminars, problems, meetings, workshops, late nights, papers,
Varsity matches, events, hangouts, long hikes, rowing, dinners.

My thanks, to Alex for making this possible.
To the group for their warmth and camaraderie.
To my Keble friends for their solidarity.
To my family, they showed me I'm capable.

Finally, to Jane for bearing with me
and supporting me through it all.

Abstract

This thesis addresses problems in the areas of graph colourings, reconstruction, and permutations.

We consider three topics related to graph colourings. First, we investigate dominating K_t-models introduced by Illingworth and Wood in 2024 which strengthens the notion of containing a clique minor. They asked whether the corresponding strengthening of Hadwiger's conjecture could hold. Surprisingly, we provide evidence towards a positive answer to this stronger conjecture. Notably, we show that graphs with no K_t-dominating models have average degree $O(t \log^3 t)$, which significantly improves the previous bound.

Next, motivated by recent breakthroughs in graph product structure theory, we investigate how useful such results can be to prove general bounds in the context of clustered colourings. Given a (non-proper) colouring of a graph, its clustering number is the largest size of a monochromatic component. For colourings of various strong products of graphs that use two or three colours, we show asymptotically tight bounds on their clustering numbers.

Thirdly, the interval thickness of a graph G is the smallest number of parts E_i in an edge-partition where each E_i has a colouring $c_i \colon E_i \to \mathbb{N}$ so that the colours of edges at any vertex form an interval of $\mathbb{N}$. Asratian, Casselgren and Petrosyan introduced this notion in the context of scheduling problems. We provide the first non-constant lower bounds on interval thickness, and improve upper bounds.

Our first reconstruction problem is probabilistic. A binary matrix is said to be in shredded form if we are given its multiset of unlabelled rows and its multiset of unlabelled columns. Let M be a binary matrix where each entry is randomly, independently chosen to be 1 with probability $p \leqslant 1/2$. Atamanchuk, Devroye and Vicenzo introduced the problem of reconstructing such a matrix M from its shredded form. For any $\varepsilon > 0$ they proved that with high probability M can be reconstructed when $p \geqslant$

$(2+\varepsilon)\frac{1}{n}\log n$. We show that the sharp threshold for reconstruction is at $p \sim \frac{1}{2n}\log n$.

Next, we examine reconstruction of CAT(0) cube complexes. These are cube complexes which satisfy a certain local non-positive curvature condition. Inspired by a result of Haslegrave who proved that disc quadrangulations with internal degrees at least 4 can be reconstructed from the restriction of their distance matrix to boundary vertices, we generalise this result to 3-dimensional CAT(0) cube complexes. This can be thought of as a discrete analogue of the boundary rigidity conjecture of Michel from 1981 regarding simple compact Riemannian manifolds.

Switching gears to permutations, we say that a family of permutations t-shatters a set if it induces at least t distinct permutations on that set. What is the minimum number $f_k(n,t)$ of permutations of $\{1,\dots,n\}$ that t-shatter all subsets of size k? For $t \leqslant 2$, $f_k(n,t) = \Theta(1)$. Spencer showed that $f_k(n,t) = \Theta(\log\log n)$ for $3 \leqslant t \leqslant k$ and $f_k(n,k!) = \Theta(\log n)$. In 1996, Füredi asked whether partial shattering with permutations must always fall into one of these three regimes. Johnson and Wickes recently settled the case $k = 3$ affirmatively and proved that $f_k(n,t) = \Theta(\log n)$ for $t > 2(k-1)!$. We give a surprising negative answer to the question of Füredi by showing that a fourth regime exists for $k \geqslant 4$, and settle the asymptotic behaviour of $f_k(n,t)$ for a large range of values of t.

Statement of Originality

The material in this thesis has not been submitted for any other degree or qualification at any university.

Chapter 2 is based on joint unpublished work with António Girão, Sergey Norin and Jane Tan. Chapter 3 is based on joint work with Campbell, Gollin, Hendrey, Lesgourges, Mohar, Tan and Wood [30]. Chapter 4 is based on joint work with Axenovich, Girão, Powierski and Savery [15]. Chapter 5 is based on joint work with Balister, Kronenberg and Scott [18]. Chapter 6 is based on joint work with Haslegrave, Tan and Scott [64]. Chapter 7 is based on joint work with Girão and Michel [58].

Contents

1 Introduction **1**

1.1 Graph colourings and structure 1
1.2 Reconstruction 4
1.3 Families of shattering permutations 6
1.4 General definitions and results 6

2 A dominating variant of Hadwiger's conjecture **8**

2.1 Preliminaries 11
2.2 Bounds on average degree of graphs excluding K_t-dominating models . 12
2.3 Dominating K_5-models 19
2.4 Dominating $K^*_{s,t}$-models 22
2.5 Deficiency colourings 26

3 Clustered colouring of graph products **29**

3.1 Definitions and tools 34
3.2 Two colours 35
3.3 Three colours 37
3.4 Arbitrary number of colours 45
3.5 Products of bounded degree graphs 46
3.6 Planar graphs 47

4 Improved bounds on the interval thickness of graphs **48**

4.1 Proof of the lower bounds 49
4.2 Proof of the upper bounds 54
4.3 Concluding remarks 56

5 Reconstructing matrices in shredded form **58**

5.1 Preliminaries 63
5.2 Reconstructing almost all rows and columns 68

5.3 Full reconstruction 74
5.4 A fast algorithm for reconstruction 77
5.5 Conclusion 78

6 Reconstructing CAT(0) cube complexes from their boundary metric 79
6.1 Definitions 84
6.2 Proof overview and discussion 90
6.3 Technical toolbox 93
6.4 Cleaning 97
6.5 Rows of cubes 102
6.6 Rows of Cubes II: The Thickening 106
6.7 Existence of rows of cubes on the boundary 109
6.8 Finite CAT(0) cube complexes of dimensions 1 and 2 114

7 Small families of partially shattering permutations 116
7.1 A first approach 119
7.2 Lower bounds 120
7.3 Upper bounds 124
7.4 Open problems and discussion 128

Bibliography 130

List of Figures

2.1 A triad with apex a and with leg L_1 in blue. 20
2.3 A sufficient subgraph for a $K^*_{2,3}$ dominating minor of type I. In blue the two most dominating branch sets, and in red the other three. 23

3.1 A fan F_{10}. 30
3.2 $H_n = n^{\widehat{2\ F_{n^4}}}$, where dotted edges are in H_n as well. 39
3.3 Subgraph $G[J]$ and the subgrid of L bounded by R_s and R_t. 41
3.4 Subgrid of L bounded by R_s and R_t with dummy vertices. 42
3.5 Proof of Claim 3.23 43
3.6 Paths P_1 and P_2 44

5.1 Isolated ones are the main obstruction to weak reconstructibility. Permuting rows r_1, r_2 as shown produces two different matrices. 65
5.2 The matrix M to be reconstructed, where $*$ represents an entry yet to be revealed. c_1, c_2 are not admissible due to conditions (i) and (ii) respectively, c_3 is admissible. In our regime and for large values of n, almost all columns in $N(r)$ will be admissible. 72

6.1 A cube hidden within another. Links of vertices of the 'hidden' cube are not flag: four cubes meet at each such vertex, but the complex has no cells of dimension 4. 81
6.2 Two pure cube complexes with the same boundary distances. The dark shaded cubes are removed. 82
6.3 Two embeddings of the same (infinite) complex with differing geometric boundary (in blue). 85
6.4 From left to right: a nonogon, a bigon and a self-intersecting hyperplane. The relevant immersed hyperplanes are represented in dashed lines. . . 90

6.5 CAT(0) cube complex where removing a row of cubes gives a complex which is not homeomorphic to a ball: the dark blue cube is a row of cubes on the boundary (its red edge is a path of length 1 with endpoints and internal vertices of boundary degrees 3 and 4 respectively), yet removing it leaves two cubes sharing an edge. 92
6.6 CAT(0) cube complex where removing the face bounded by the central blue square gives a complex which is not contractible. 92
6.7 Pattern corresponding to a vertex v of degree 3 not in a cube. 100
6.8 Row configuration. 103
6.9 Sketch of proof of Lemma 6.18. If the dashed path γ_* is a geodesic, we can project its intersection with the tubular neighbourhood of H to the side of A, giving a subpath γ_0. This produces a (p, x)-path of length at most $d_X(p, y) - 1$. 105
6.10 Example of gluing faces of cubes C_F together. 108
6.11 Two row configurations sharing endpoint x. In fact, F_0 and F_2 must share an edge. 112

7.1 A subdivision of S inside T. Edges of T are dashed for better visibility. . 122

Chapter 1

Introduction

The content of this thesis is based on material from six papers, each with its own dedicated chapter. In this introduction, we set the scene by giving a brief overview of the literature that they fit in and introduce important definitions.

1.1 Graph colourings and structure

Graph colourings is an old, central topic of graph theory with significant connections to computer science. The *chromatic number* $\chi(G)$ of a graph G is the smallest number of colours required to colour its vertices so that no neighbours share the same colour. This is perhaps the most studied graph invariant, partly due to its natural definition but also because of its historical role in computer science. Famously, determining the chromatic number of a graph was one of the first major examples of a so-called NP-complete problem.

It is interesting to reflect on what causes high chromaticity in graphs. An immediate example is containing a large clique, but there are many instances of triangle-free graphs with high chromatic number (e.g. see Section 2 of [110]). This raises the following natural question.

Question 1.1. *What substructures must a graph of chromatic number at least k contain?*

The study of this question under different definitions of 'containment' of substructures has given rise to different research areas. For example, one direction is to study chi-bounded graph classes, namely classes $\mathcal{G}$ of graphs whose chromatic number is bounded as a function of the size of the largest cliques they contain.

We discuss three topics in graph colourings.

Hadwiger's conjecture

Here, we are interested in containment of substructures as minors. Say that G contains an *H-minor* if the vertex set of G contains a collection of disjoint sets V_x indexed by vertices x of H such that each V_x induces a connected subgraph of G, and for any edge xy of H, there is an edge between V_x and V_y in G.

In light of Question 1.1, the iconic Four Colour Theorem of Appel and Haken [7] from the 1970s is a partial answer for graphs of chromatic number at least five; they must be non-planar and therefore contain either a K_5-minor or a $K_{3,3}$-minor by a result of Wagner from the 1930s. In fact, Wagner showed that graphs without K_5-minors are either planar or isomorphic to certain 4-colourable graphs, so this together with the Four Colour Theorem implies that graphs of chromatic number at least 5 must contain K_5-minors. Hadwiger's conjecture attempts to generalise this result by characterising graphs whose minors have chromatic number less than k as being precisely those which do not contain a clique of size k as a minor.

Conjecture 1.2 (Hadwiger's conjecture)**.** *If a graph does not contain a K_k-minor then it has chromatic number less than k.*

If true, this conjecture would provide a complete answer to Question 1.1 for minors.

In Chapter 2, we discuss this in more detail and explore a strengthening of this conjecture proposed by Illingworth and Wood [72] where the usual notion of minor is tweaked to require domination between adjacent vertex sets.

Clustering number and treewidth

Motivated by a problem related to dynamically evolving databases, Kleinberg, Motwani, Raghavan and Venkatasubramanian [79] first introduced the notion of clustered colourings under the language of 'fragmented colourings' in 1997. Around the same time, Ding, Oporowski, Sanders and Vertigan [40] reframed this notion within the realm of graph partitions: a k-colouring of a graph G with *clustering number* c is a partition of the vertices of G into k parts $V_1, \dots, V_k$ such that the graph induced by each V_i has connected components of size at most c. Specifically they explored the relationship between bounded treewidth, bounded degree and 2-colourings of bounded clustering.

Treewidth is an important structural invariant; it gives an indication of how 'tree-like' a graph is. We give a formal description of this notion in Section 1.4, but roughly speaking a graph G has treewidth at most k if we can view its vertices as subtrees of some tree T where adjacent vertices in G have overlapping subtrees, and at most $k+1$ overlap

in any given node of T. First introduced by Bertelè and Brioschi [20] under the name of 'dimension', they observed that many otherwise difficult decision problems can be solved efficiently using dynamic programming, as long as this parameter is guaranteed to be bounded. This property makes graphs with bounded treewidth particularly attractive for computational applications.

The Grid Minor Theorem of Robertson and Seymour [106] from 1986 gives a rough characterisation of bounded treewidth in terms of forbidden minors: a graph has treewidth at most t if it does not contain an $n \times n$ grid as a minor for some large t depending on n, and conversely large grids have large treewidth. The former part of this statement is now known to hold for t polynomial in n, with Chuzhoy and Tan [34] proving in 2021 that $t = O(n^9 \operatorname{polylog} n)$ suffices.

Having bounded treewidth is tightly related to containing good separators. A *separator* in a graph G on n vertices is a set S of vertices such that removing S from G leaves connected components of restricted size. This is usually referred to as a *balanced separator* in the literature when the bound restricting the sizes of these components is linear in n. Planar graphs provide a classical example of a class of graphs with small balanced separators: Lipton and Tarjan [87] showed that a planar graph on n vertices contains a balanced separator of size $O(\sqrt{n})$. Having small balanced separators is appealing from an algorithmic point of view as these are helpful for divide and conquer strategies.

In another of their seminal papers on minors published in 1986, Robertson and Seymour [105] showed that graphs have balanced separators of size linear in their treewidth. This result was recently extended by Dvořák and Wood [47] who showed that graphs of bounded treewidth have efficient separators in a strong sense. Specifically, they show that for any graph G on n vertices and $p \in \mathbb{N}$ there is a set of at most p vertices which, when removed from G, leaves connected components of size at most $(t+1)n/p$, where t is the treewidth of G. Conversely, Dvořák and Norin [46] proved a linear bound on treewidth in terms of the balanced separators of its subgraphs.

In [86], Linial, Matoušek, Sheffet and Tardos used separators to prove general bounds on the clustering number of graphs of bounded treewidth and minor closed families of graphs. Bearing in mind recent breakthroughs which used graph product structure results, we revisit these ideas in Chapter 3 from the perspective of graph product structure theory.

Interval colourings

Consider the following problem. Say we are hosting a conference and some of the attendees wish to have one on one meetings of fixed duration with some other attendees, but no one is willing to wait between two of their meetings on any given day. Assuming that we can host as many meetings in a day as we like, what is the minimum number of days that the conference needs to run for so that all meetings take place?

Viewing attendees as a set of vertices and meetings as edges between pairs of attendees, the answer is precisely the interval thickness of the graph. Given a graph G, we say it has *interval thickness* at most k if its edges can be decomposed into k sets $E_1, \dots, E_k$ so that there are colourings $c_i \colon E_i \to \mathbb{N}$ where for any vertex v, the numbers $c_i(vx)$ for $vx \in E_i$ are distinct and form an interval in $\mathbb{N}$. A graph with an interval thickness of one is said to be *interval colourable*.

Interval colourable graphs were initially introduced by Asratian and Kamalian [11] in 1987, and have applications in the context of scheduling problems such as the one above. In Chapter 4 we present the first non-trivial lower bound on the interval thickness of general graphs, and improve upper bounds.

1.2 Reconstruction

The problem of reconstructing a graph from partial information dates back as far as 1941, when Kelly and Ulam conjectured that any finite simple graph G on at least 3 vertices can be identified up to isomorphism from the collection of its maximal proper induced subgraphs, namely the multiset $\{G \setminus \{v\} : v \in V(G)\}$.[1]

While still open in general form, the Reconstruction Conjecture is much easier to approach in a random setting. Müller [96] proved in 1976 that it holds with high probability for the random graph $G(n, \frac{1}{2})$ and Bollobás [23] later showed that here reconstruction is even possible from much less information: one only requires three of the subgraphs $G \setminus \{v\}$.

Shotgun reconstruction

One can think of the Reconstruction Conjecture as asking if graphs can be reconstructed from the most structural information we can possibly have about G, without knowing G already. In contrast, the 'shotgun reconstruction' problem introduced by Mossel

[1] This is widely known as the 'Reconstruction Conjecture'. The date of 1941 comes from a 1977 paper of Bondy and Hemminger [27], who attribute it to "reliable sources".

and Ross [93] in 2019 asks about reconstructibility when only provided with local information. Specifically, for a graph G we are provided with subgraphs $N_G^r(v)$ induced by the vertices at distance at most r from v for each vertex v of G, where *only* the label of the centre vertex is known. We say that G is 'reconstructible' if we can identify it up to isomorphism from the multiset of partially labelled r-balls $\{N_G^r(v) : v \in V(G)\}$.

This question in the random graph setting has sparked intense interest, with several papers [41, 55, 70, 76] sharpening the initial results of Mossel and Ross over a very short span of time between 2019 and 2021.

In the Mossel–Ross reconstruction problem, only the central vertex is labelled in each provided r-ball. Atamanchuk, Devroye and Vicenzo [13] considered a problem regarding 'shredded binary matrices' that, in the language of random graphs, essentially asks about reconstruction when all vertices of r-balls are labelled *except* for the centre. We discuss this in Chapter 5.

Boundary rigidity

Isoperimetric inequalities are a class of results which generally speaking describe the extremal area or volume of objects with fixed size boundary in given host spaces. The aforementioned planar separator theorem of Lipton and Tarjan [87] can be thought of as an isoperimetric inequality: it says that in any planar graph on n vertices there is a set of $\sqrt{n}$ vertices which separates the graph in two parts where the smaller has size linear in n. Here the number of vertices corresponds to the volume, the set of $\sqrt{n}$ vertices is the fixed size boundary and the whole graph the host space. This has a similar flavour to the well-known result that discs maximize surface area for a given boundary length, yet it is not quite a discrete analogue of it.

Angel, Benjamini and Horesh [6] gave a discrete analogue of this result: disc triangulations with n vertices have boundary of size $\Omega(\sqrt{n})$, so long as internal vertices have minimum degree 6. One should note that this minimum degree requirement is not benign as it corresponds to a non-positive curvature condition through the combinatorial analogue of the Gauss–Bonnet theorem.[2] Spaces with non-positive curvature tend to have very rigid metrics; for instance, Pestov and Uhlmann [99] showed that

[2]For the reader's convenience, we recall one version of this result. Consider a finite 2-dimensional simplex X. Then

$$\sum_{v \notin \partial X} (6 - t(v)) + \sum_{v \in \partial X} (3 - t(v)) = 6 \cdot \chi(X),$$

where $t(v)$ denotes the number of triangles that v is incident to, and $\chi(X)$ is the Euler characteristic of X. The two terms on the left hand side represent the local curvature at internal and boundary points of X respectively. So the minimum degree 6 condition corresponds to non-positive internal curvature.

two dimensional simple compact Riemannian manifolds with boundary are uniquely determined by their metric on the boundary. A discrete analogue was proven by Halsegrave [63] who showed that 2-dimensional triangulations and quadrangulations can be reconstructed from boundary distances, as long as they satisfy certain internal degree conditions. In Chapter 6, we generalise this result to 2-dimensional CAT(0) cube complexes and 3-dimensional CAT(0) cube complexes with an embedding in $\mathbb{R}^3$.

1.3 Families of shattering permutations

Finally, we consider a problem which is different in flavour from the topics above. It concerns a variant of the notion of 'shattering' which underlies the well-studied VC dimension invariant. Specifically, we say that a family of orders $\mathcal{P}$ on $[n] := \{1, \dots, n\}$ *t-shatters* a set $S \subseteq [n]$ if it induces at least t distinct orderings on the elements of S. This concept can be traced back to a paper of Spencer [113] from the 1970s where he considered the minimal size of families shattering all sets of a given size.

In Chapter 7, we disprove a conjecture of Füredi [54] from 1996 concerning the possible orders of growth of minimal families which t-shatter all subsets of a fixed size. Perhaps surprisingly, ideas and constructions borrowed from Ramsey-type problems turn out to be effective in this setting.

1.4 General definitions and results

Let G be a graph on n vertices. All graphs considered in this thesis are simple unless otherwise stated. We write $e(G) = |E(G)|$. Say G is *d-regular* if all its vertices have degree d, and *d-degenerate* if any induced subgraph contains a vertex of degree at most d. We say G is *k-colourable* if it has chromatic number at most k.

Let d_G denote the shortest path metric on G. For $k \in \mathbb{N}$, the *kth neighbourhood* of v is $N_k(v) = \{u \in V \setminus v : d_G(u, v) = k\}$. We usually write $N(v)$ for its first neighbourhood, and $\deg(v) = |N(v)|$ for its *degree*. For a set $S \subseteq V$, the *subgraph induced by S*, denoted $G[S]$, is the subgraph on vertex set S and edges of G with both endpoints in S. A graph H is *contained* in a graph G if a subgraph of G is isomorphic to H.

For a graph G and a tree T, a *T-decomposition* for G is a collection $\mathcal{W} = (W_x : x \in V(T))$ such that

1. for each $v \in V(G)$, the subgraph of T induced by $\{x : v \in W_x\}$ is a tree,
2. for each $uv \in E(G)$, there is an $x \in V(T)$ such that $u, v \in V_x$.

We sometimes write $(T, \mathcal{W})$ for a T-decomposition. Each set W_x of $\mathcal{W}$ is called a *branch set*. The *width* of this decomposition is $\max\{|W_x| : x \in V(T)\}$, a *tree decomposition* is a T-decomposition for any tree T and the *treewidth* $\mathrm{tw}(G)$ is the smallest width of a tree decomposition of G, minus one.

For graphs G and H, an *H-minor model* in G is a tuple $(V_x)_{x \in V(H)}$ of disjoint vertex sets of G such that each $G[V_x]$ is connected, and there is an edge between V_x, V_y whenever $xy \in E(H)$. A *subdivision of H* is a graph that can be obtained from H by replacing its edges with vertex disjoint paths.

We write P_n for the path on n vertices, K_n for the complete graph on n vertices and $K_{n,m}$ for the complete bipartite graph with parts of sizes n and m repectively.

A *colouring* of a graph G is a function that assigns a 'colour' to each vertex of G. A *c-colouring* is a colouring that uses at most c colours. We say that this colouring is *proper* if it assigns distinct colours to adjacent vertices.

Lastly, in Chapters 2, 5 and 4 we use the well-known bounds on the tails of binomial distributions known as the Chernoff bounds (see e.g. [92], Theorem 4.4).

Lemma 1.3 (Chernoff bounds)**.** *Let $0 < p \leqslant \frac{1}{2}$, $X \sim \mathrm{Bin}(n, p)$ and $\varepsilon > 0$. Then,*

$$\mathbb{P}\big(X \geqslant (1+\varepsilon)np\big) \leqslant \exp\Big(-\tfrac{\varepsilon^2 np}{2+\varepsilon}\Big),$$
$$\mathbb{P}\big(X \leqslant (1-\varepsilon)np\big) \leqslant \exp\Big(-\tfrac{\varepsilon^2 np}{2}\Big).$$

Chapter 2

A dominating variant of Hadwiger's conjecture

This chapter is based on unpublished joint work with Girão, Norin and Tan.

Graph minors, and more specifically forbidden minors, are cornerstone notions of modern graph theory. For instance, a classical result of Wagner from the 1930s states that planar graphs are precisely those which do not contain K_5 or $K_{3,3}$ as a minor. Kuratowski independently proved the stronger result that forbidding subdivisions of K_5 or $K_{3,3}$ suffices.

Much like planarity, many natural graph properties are preserved under taking minors. The celebrated Robertson–Seymour Graph Minor Theorem cemented the role of forbidden minors in graph theory; it states that any minor-closed graph family $\mathcal{F}$ is characterised by a finite set of forbidden minors. In other words, there is a finite set of forbidden minors which acts as a certificate for membership of $\mathcal{F}$.

The famous conjecture of Hadwiger bounds the chromatic number of graph families which exclude a clique minor: graphs which contain no K_t-minor can be properly coloured with less than t colours. This conjecture remains open as stated or even up to a multiplicative constant, but significant progress has been made towards it over the past decade. For many years the best asymptotic bound was due to Kostochka [82, 83][1] and Thomason [115] who showed in the 1980s that graphs with no K_t-minors are $O(t\sqrt{\log t})$-degenerate, which moreover is tight. In 2019, a preprint by Norin and Song [98] gave the first significant improvement on this degeneracy bound by showing that K_t-minor-free graphs can be properly coloured with $O(t \log^{\beta} t)$ colours for any $\beta \geqslant 0.354$. This breakthrough was due to the use of a novel 'density increment'

[1] These two references in fact point to the same paper, with one version in Russian and the other in English. They are published in different journals, likely due to the language difference.

argument, which they sharpened together with Postle [97] in a subsequent version of the same paper. With an independent approach, Delcourt and Postle [37] later showed the currently best known bound of $O(t \log \log t)$.

While a proof of Hadwiger's conjecture seems currently intractable, a great deal of intuition has been gained from understanding how the picture changes when terms in the statement of the conjecture are tweaked. For example, a conjecture attributed to Hajós [59][2] replaces 'minor' with 'subdivision'. This is long known to be false for almost all graphs [51]. Another example is if we replace 'chromatic number' with 'fractional chromatic number'. The analogue of Hadwiger's conjecture in this setting is known to hold, up to a factor of two [102]. It is interesting to note that many such variants have been shown to either be far too restrictive, or already proven true. The present chapter deals with a variant of Hadwiger that is unusually robust.

First introduced by Illingworth and Wood [72], K_t-dominating models are a simple strengthening of the usual notion of minor, where the adjacency condition between branch sets is replaced by domination. For a graph G, we say that it contains a K_t-*dominating model* if there are disjoint subsets $T_1, \dots, T_t$ of its vertex set such that each $G[T_i]$ is connected and each $v \in T_j$ admits a neighbour in T_i for every $i < j$. Although this seems far more restrictive than the traditional notion of minor and one therefore expects K_t-minor free graphs to be a much narrower family than graphs with no K_t-dominating minors, proving this in a quantitative manner is surprisingly elusive.

A natural avenue to distinguish between K_t-minor free and K_t-dominating minor free graphs would be to show that Hadwiger's conjecture cannot be strengthened by replacing minors with dominating minors. Yet this seems far more difficult than initially expected, and much evidence supports otherwise.

Conjecture 2.1. *[Dominating Hadwiger] Let t be a positive integer and G a graph with no dominating K_t-model. Then G has a proper colouring with $t-1$ colours.*

A notable difference is in the smallest average degree forcing K_t-minors. As mentioned earlier, Kostochka [82, 83] and Thomason [115] showed in the 1980s that graphs without K_t-minors are $O(t\sqrt{\log t})$-degenerate, meaning that their average degree cannot exceed this bound. On the other hand, Illingworth and Wood [72] showed that random graphs with average degree $\Theta(t \log t)$ whp do not contain dominating K_t-models. They also provided an upper bound of 2^{t-2} on the smallest average degree forcing dominating

[2]The exact attribution of this conjecture is uncertain, as Thomassen writes in [116]: "Sometimes the paper [59] is cited, but B. Toft (oral communication) points out that Hajós did not mention the conjecture in that paper and that G. A. Dirac knew the conjecture in the 1950s".

K_t-models. In this chapter, we narrow this gap by showing that graphs of average degree $\Omega(t \log^3 t)$ contain dominating K_t-models, and show that several techniques used to prove special cases of Hadwiger's conjecture can be adapted to the dominating case.

Theorem 2.2. *There is an absolute constant $c > 0$ such that for any positive integer t, if G is a graph with $d(G) \geqslant ct \log^3 t$ then G contains a K_t-dominating minor.*

Illingworth and Wood proved that Conjecture 2.1 holds for $t \leqslant 4$. This leaves the cases $t = 5, 6$ where Hadwiger is known to hold by the Four Colour theorem and results of Robertson, Seymour and Thomas [104]. We investigate the case $t = 5$.

We first obtain a bound on the number of edges of graph with no dominating K_5-models which is achieved by planar graphs.

Theorem 2.3. *For $n \geqslant 4$, every n-vertex graph with no dominating K_5-model has at most $3n - 6$ edges.*

It is not hard to see that 4-connected graphs with a $K_{3,3}$-minor contain a K_5-minor, so we naturally consider non-planar 4-connected graphs next. Using similar techniques to those used by Robertson, Seymour and Thomas [104] in their proof of Hadwiger for $t = 6$, we show that these contain dominating K_5-models provided they have a triangle.

Theorem 2.4. *Let G be 4-connected non-planar and suppose it has at least one triangle. Then G contains a dominating K_5-model.*

Our argument uses results from [104] guaranteeing existence of substructures called 'legless tripods' with specified 'feet'. The same argument shows that 4-connected graphs contain dominating models of K_5 if we allow one edge to be missing from the model, and our hope is that a similar approach can be extended to general graphs.

Let $K^*_{s,t}$ be the graph whose vertex set has two parts A, B of sizes s and t respectively, all edges with both endpoints in A as well as all edges across A and B, and no edges within B. We will formally define what we mean by '$K^*_{s,t}$-dominating model of type I' in Section 2.4, but it essentially is the same definition as for K_t-dominating models, except that we need to specify a fixed ordering of its vertices as not all edges are present.

Corollary 2.5. *Let G be a simple graph that is 4-connected and not 4-colourable. Then for any $v \in V(G)$, G contains a dominating $K^*_{3,2}$-model of Type I where v is in the most dominating branch set.*

Motivated by results of Kostochka [81, 84] that show graphs with chromatic number $s + t$ have a $K^*_{s,t}$-minor when t is much larger than s, we prove the following result.

Theorem 2.6. *Let G be a graph which is not $(t+1)$-colourable. Then G contains a $K^*_{2,t}$ dominating model of Type I.*

Say that a graph G admits a *c-defective* colouring with t colours if there is a partition $X_1, \dots, X_t$ of its vertex set so that the maximum degree of each $G[X_i]$ is at most c. Edwards, Kang, Kim, Oum, Seymour [49] proved that for every $t > 0$ there is a constant c_t such that if a graph G has no K_t-minor then it has a c_t-defective colouring with $t-1$ colours, and van den Heuvel and Wood [118] showed that the defect c_t can be chosen to be $t-2$. We show that a similar partitioning can be obtained for graphs without K_t-dominating models.

Theorem 2.7. *For every $t \in \mathbb{N}$ there is a constant $c_t > 0$ such that the vertex set of any graph G that does not contain a dominating K_t-model can be partitioned into $t-1$ parts $X_1, \dots, X_{t-1}$ with $\Delta(G[X_i]) \leqslant c_t$ for each i.*

The chapter is organised as follows. In Section 2.1 we define general notation specific to this chapter. Section 2.2 is dedicated to the proof of Theorem 2.2. Section 2.3 focuses on finding dominating K_5-models and contains the proof of Theorem 2.3 and Theorem 2.4. In Section 2.4 we prove Corollary 2.5 and Theorem 2.6. Finally, Section 2.5 treats Theorem 2.7.

2.1 Preliminaries

Let G be a graph. In this chapter we denote the degree of a vertex v of G by $d(v)$. If $S \subseteq V(G)$ let $N(S) = \{u \in V(G) \setminus S : uv \in E(G) \text{ for some } v \in S\}$ be the *outer neighbourhood* of S. In this chapter we use 'neighbourhood' to mean '*outer* neighbourhood' unless otherwise stated. Given a set F of edges (possibly not in $E(G)$), write $G + F$ for the graph obtained from G by adding edges of F (where we ignore duplicate edges), and $G - F$ the one obtained by removing them. For a set $S \subseteq V(G)$, let $G \setminus S$ be the graph obtained from G by removing vertices of S. When S or F is a singleton we sometimes omit braces. For two subgraphs G_1, G_2 of a host graph G, we write $G_1 \cup G_2$ for the subgraph of G on vertex set $V(G_1) \cup V(G_2) \subseteq V(G)$ and edge set $E(G_1) \cup E(G_2) \subseteq E(G)$. For an edge uv of G, we write G/uv for the graph obtained by contracting uv, namely the graph on vertex set $V(G) \setminus \{u\}$ with all edges they induce in G, together with all non-duplicate edges from $\{xv : x \in N(u) \setminus \{v\}\}$.

A K_t-*dominating model* in G is a tuple $(T_1, \dots, T_t)$ of disjoint sets of vertices, where any $v \in T_j$ admits a neighbour in T_i for every $i < j$, and such that $G[T_i]$ is connected. We

commonly refer to these T_is as being the *branch sets* of the model, and to T_1 as the *most dominating* branch set.

2.2 Bounds on average degree of graphs excluding K_t-dominating models

First, we treat the almost-regular case (Corollary 2.9) which is then used in the proof of the general case.

In both cases, arguments rely on technical lemmas which find a connected subgraph that dominates a set on which we recurse. In the almost-regular case, this role is filled by Lemma 2.8 and in the general case by Lemma 2.11.

Lemma 2.8. *Let c, d be constants with $d \geqslant 3$, and suppose that G is a graph on n vertices with average degree at least d and maximum degree $\Delta(G) \leqslant cd$. Then, G contains a connected set X whose neighbourhood induces a graph H with average degree at least $d(1 - \frac{51c \log d}{d})$.*

Proof. As removing vertices with degree less than $d/2$ does not decrease the average degree, we may assume without loss of generality that G has minimum degree at least $d/2$.

Let S_0 be obtained by sampling each vertex of G independently with probability $16 \log d/d$. Say that a vertex is *unusual* if it fails to have at least $4 \log d \geqslant 1$ neighbours in S_0. By the standard Chernoff bounds (Lemma 1.3), a vertex is unusual with probability at most d^{-1}. Let S_1 be the union of S_0 and unusual vertices. Since $\mathbb{E}|S_1| \leqslant \frac{17 \log d}{d} n$, we may fix an outcome where this bound is realised.

Let $X_1, \dots, X_t$ be the connected components of S_1, and define $E_{in} = \{uv \in E(G) : u, v \in N(X_i)$, for some $i\}$, i.e. the edges contained inside their neighbourhoods, and $E_{ext} = E(G) \setminus E_{in}$, the edges not covered by neighbourhoods. Starting with $S = S_1$, repeat the following two steps until neither condition is satisfied.

1. If a vertex $v \in V \setminus S$ has neighbours in X_i and X_j for $i \neq j$, then add v to S.
2. If $uv \in E_{ext}$, add both u and v to S.

Continue until this process stops, so that neither condition is satisfied for S. Since the number of connected components is $t \leqslant |S_1| \leqslant \frac{17 \log d}{d} n$ and each step of the process reduces it by at least one, in $G[V \setminus S]$ there are at most $\Delta(2t + |S_1|) \leqslant 51c \log d \cdot n$ fewer edges than in G. In particular, the average degree in $G[V \setminus S]$ is at least $d(1 - \frac{51c \log d}{d})$.

Finally, as every edge in $G[V \setminus S]$ is contained in the neighbourhood of some connected component of $G[S]$, the desired connected dominating set X may be chosen to be the connected component of S whose neighbourhood induces a subgraph of largest average degree. □

As a consequence, we derive the result for almost-regular graphs.

Corollary 2.9. *Suppose that G is a graph on n vertices with average degree d and let c be such that $\Delta(G) \leqslant cd$. If $t \geqslant \max(2, (1000c)^{1/17})$ and $d \geqslant 1000ct \log t$ then G contains a K_t-dominating model.*

Moreover since $t \geqslant 2$, the inequality $t \geqslant (1000c)^{1/17}$ holds when $c \leqslant 10^{10}t$.

Proof. It suffices to check that Lemma 2.8 may be applied t times to G, each time repeating with the subgraph H obtained from the lemma. Every application decreases the average degree by at most $51c \log(1000ct \log t)$, so we can repeat at least

$$\frac{1000ct \log t}{51c(\log 1000c + \log t + \log\log t)} \geqslant t$$

times. Since $t \geqslant 2$ we have $\log t \geqslant \log\log t$, so it suffices to show $898 \log t \geqslant 51 \log 1000c$, which holds by choice of t. □

For the general case, we use the following dichotomy lemma of Girão and Hunter [57] that either finds an almost-regular subgraph where Corollary 2.9 applies, or a highly imbalanced bipartition. The bounds stated below differ slightly from those recorded in [57], but these can be obtained by a careful reading of their proof.

For a graph H of average degree d, say it is *ℓ-almost-regular* if $\Delta(H) \leqslant \ell d$.

Lemma 2.10 ([57], Lemma 3.4)**.** *There is an absolute constant $1200 \geqslant c > 0$ such that if $d \geqslant 2$ and $L \geqslant 6$ are constants and G a d-degenerate graph on n vertices of average degree d, then G contains one of:*

1) *a $(c \log L)$-almost-regular graph of average degree at least $d/c \log^2 L$;*
2) *a bipartition (A, B) such that $|A| \geqslant L|B|$ and $e(A, B) \geqslant dn/8$.*

Proof sketch of Lemma 2.10. Consider the partition of G into A, B where $B = \{v \in V(G) : d(v) > Ld\}$. If at least $1/4$ of all edges have endpoints in both A and B, then G is highly imbalanced, as in the second case.

Otherwise, by d-degeneracy, a constant proportion of edges of G have both endpoints in A. The idea here is to find a bipartite subgraph $H = (A', B')$ inside A which has at least

a roughly $1/(\log L)^2$ proportion of all edges, and which satisfies a certain regularity property. That is, we want degrees in A' and B' not to exceed their expected degrees by more than a factor of $\log L$: for each $a \in A'$, $d_H(a) \leqslant \log L \cdot e(A,B)/|A|$, and similarly for each $b \in B'$, $d_H(b) \leqslant \log L \cdot e(A,B)/|B|$.

Say we found such a bipartite subgraph and $|A'| \geqslant |B'|$. If A' and B' are of comparable size then we are done. Otherwise, we expect to be close to the desired almost-regular graph after sampling A' with probability $|B'|/|A'|$ as this shrinks most degrees in B' to roughly match those in A'. Up to minor details, this yields the desired almost-regular subgraph.

To find such a bipartite subgraph we partition vertices in A dyadically according to their degrees into about $\log L$ sets. That is, partition vertices into sets so that degrees in the same set differ by a factor of at most 2. By averaging, at least a $1/\log L$ proportion of all edges have an endpoint in one of these sets X. Then, technical details aside, we roughly repeat this process with the neighbourhood of X to obtain another set Y such that (X,Y) has at least a $1/\log^2 L$ proportion of all edges, and degrees in Y differ by a factor of at most 2. When repeating this process, the average degree of vertices in X drops by at most $1/\log L$, so the resulting bipartite subgraph (X,Y) satisfies the advertised regularity property. □

Choosing L to be polynomial in the average degree d, case 1) essentially reduces to Lemma 2.8, so only the imbalanced bipartite case is interesting. We sketch the argument to find good connected dominating sets in this setting, where $G = (A,B)$ with high average degree and $|A| \geqslant L|B|$. This is the content of Lemma 2.11.

Following the outline of Lemma 2.8, sample vertices in A and B to obtain a set S_0 which dominates most of the graph. We may assume vertices in A have degree d to start with and want to preserve some form of regularity in A, so whenever we add a vertex v to S_0 we add $N_A(v)$ as well. This also allows us to exert control over the number of edges incident to S_0 by controlling the number of vertices in A added to S_0 throughout. After sampling, vertices in A have around their expected number of neighbours in S_0, so we can assume that $A \setminus S_0$ is $(1-\varepsilon)d$-regular in the rest of the graph. Then, it suffices to find a connected component X_i of $G[S_0]$ such that we can induct on $N(X_i)$, namely that $N(X_i)$ is a subgraph (A',B') with $|A'| \geqslant (1-\alpha)L|B'|$ and vertices of A' have at least $(1-\varepsilon)d$ neighbours in B', for small $\alpha, \varepsilon > 0$.

In an ideal scenario, each vertex in $V(G) \setminus S_0$ sees at most one connected component of S_0, and there are no 'external' edges, namely edges between neighbourhoods of connected components. An averaging argument would then give the desired X_i, $N(X_i)$.

In reality this may not be true at this stage, but we hope to pay a small number of vertices of A to make the situation as close as possible to the ideal one. We add vertices to S_0 which decrease the number of connected components, so long as the cost per connected component does not exceed a threshold cost σ, and we remove external edges by adding the endpoint in A to S_0. When this stops, there are no external edges, each vertex in A sees a single connected component of $G[S_0]$ and if $\sigma \gg L$ then vertices of B do not see too many components overall. The picture is then good enough for an averaging argument to work.

A last technical point regarding the choice of $\sigma \gg L$ is that we need the number of components in S_0 to be much smaller than $|B|$ after sampling in order to counteract the fact that σ is so large. We do this at the sampling stage by showing that second neighbourhoods of vertices of B are large, as otherwise there would be an almost-regular subgraph.

Lemma 2.11. *Let d be sufficiently large and set $q = d^8$, $r = d^7$. There is an absolute constant $\eta > 0$ such that any $f > d^{17}$ the following holds. If $G = (A, B)$ is a non-empty bipartite graph with $|A| = f|B|$ where each vertex in A has degree $\delta \geqslant d/2$, then G contains one of:*

(i) a bipartite subgraph (A', B') where vertices of A' have at least $d/2$ neighbours and $|A'| = q$, $|B'| \leqslant r$;

(ii) a connected set S of vertices which dominates a subgraph (A', B') where $|A'| \geqslant f(1 - \delta^{-2})|B'|$ and each vertex of A' has degree $d \cdot (1 - \frac{1}{2}\eta \log d/d)$.

Proof. After throwing away vertices from A if necessary, we may assume that $f \leqslant d^{100}$.

First, we find a set S_0 which dominates most of the graph by random sampling. Sample vertices of A and B independently with probabilities $p_A = \delta^{-5}$ and $p_B = \eta \log \delta/\delta$ to form sets S_A and S_B respectively, and let $S_0 = S_A \cup S_B$.

Let $g = \eta \log \delta$ and $h = f\delta \cdot p_A$ be the expected degrees of vertices in A and B to S_0 respectively. Call a vertex $a \in A$ *unusual* if either $d_{S_0}(a) < \frac{1}{2}g$ or $d_{S_0}(a) > \frac{3}{2}g$, and a vertex $b \in B$ *unusual* if $d_{S_0}(b) < \frac{1}{2}h$. Here, we write $d_X(y)$ for the number of edges $xy \in E(G)$ with $x \in X$, where $X \subset V(G)$ and y is a vertex of G.

We make an observation which is useful to bound the number of connected components in S_0. Fix a sampling S_A of A and for a vertex $b \in B$ which is not unusual, say that b is *bad* if its second neighbourhood in the sampled vertices is too sparse, specifically if $|N(N(b) \cap S_A) \cap S_B| < \frac{1}{2}p_B \cdot r$. Notice that if $|N(N(b) \cap S_A)| \leqslant r$ then as $\frac{1}{2}h \geqslant q$, the subgraph $G[N(b) \cap S_A, N(N(b) \cap S_A)]$ is as described in Lemma 2.11 (i). We may therefore assume that all such $b \in B$ satisfy $|N(N(b) \cap S_A)| > r$, and so by the standard

Chernoff bounds (Lemma 1.3) we expect their second neighbourhood to contain a large number of sampled vertices:

$$\mathbb{P}\left(b \text{ is bad} \mid S_A,\ b \text{ not unusual}\right) \leqslant \exp(-\frac{1}{8}p_B \cdot r) \leqslant \delta^{-\eta/12}.$$

Lemma 1.3 bounds the probability of unusual vertices occuring unconditionally:

$$\mathbb{P}(a \text{ is unusual}) \leqslant 2 \cdot \delta^{-\eta/12};$$
$$\mathbb{P}(b \text{ is unusual}) \leqslant e^{-h/12}.$$

Hence a vertex $a \in A$ is adjacent to an unusual vertex with probability at most $\delta \cdot e^{-h/12}$, and to a bad vertex with probability at most $\delta \cdot \delta^{-\eta/12} \leqslant \delta^{-\eta/15}$ for sufficiently large η. Remove all unusual and bad vertices in $A \cup B \setminus S_0$ as well as their neighbours in $A \setminus S_0$, if any. Then all remaining vertices in A still have degree δ in G, and the expected number of vertices we removed from A is

$$\begin{aligned}\mathbb{E}[\#\text{ vertices removed from } A] &\leqslant |A| \cdot (2 \cdot \delta^{-\eta/12} + \delta^{-\eta/15} + \delta e^{-h/12}) \\ &\leqslant |A| \cdot \delta^{-4}\end{aligned}$$

for large enough η. Since further, $\mathbb{E}|S_A| = p_A|A| = |A| \cdot \delta^{-5}$ and

$$\frac{|A|}{|B|} \cdot \mathbb{E}[\#\text{ bad vertices in } B] \leqslant |A| \cdot \delta^{-\eta/8},$$

we may therefore fix an outcome where the number of vertices removed from A is at most $|A|/\delta^3$, where $|S_A| \leqslant |A|/\delta^3$ and where the number of bad usual vertices in B is at most $|B|/\delta^3$.[3]

Let A_0, B_0 be the remaining vertices in $A \setminus S_A$, $B \setminus S_B$ respectively after the removal step. Then S_0 dominates A_0 and B_0, vertices in A_0 have between $\delta - 3g/2$ and $\delta - g/2$ neighbours in B_0 and $|A_0| \geqslant |A|(1 - \delta^{-3})$. Moreover we may remove excess edges to assume that every vertex in A_0 sends exactly $\delta(1 - \frac{\eta \log \delta}{2\delta})$ edges to B_0.

Let $X_1, \ldots, X_t$ be the connected components of $G[S_0]$.

Claim 2.12. *$G[S_0]$ contains at most $3|A|/(f\delta^7)$ connected components.*

[3]Here we are applying the elemetary fact that if $\mathbb{E}[a] + \mathbb{E}[b] + \mathbb{E}[c] = \mathbb{E}[a + b + c] \leqslant \delta^{-3}$ for random variables a, b, c then there exists a probabilistic outcome where a, b and c are at most δ^{-3} simultaneously. For us, a is the number of vertices removed from A, b is $|S_A|$ and c is $\frac{|A|}{|B|}$ times the number of bad vertices in B.

Proof. For sufficiently large η, there are at most $|B|/\delta^7$ bad or unusual vertices in B, and at most $|A|/f\delta^7$ unusual vertices in A. As any component X_i must either contain a single unusual vertex of A, a bad or unusual vertex of B, or at least $\frac{1}{2}p_B r \geqslant \frac{1}{2}\eta\delta^6$ vertices from B, we may bound the number of such components by

$$t \leqslant |B|/\delta^7 + |A|/f\delta^7 + 2|B|/\eta\delta^6 \leqslant 3|A|/f\delta^7. \qquad \blacksquare$$

Next, we massage our set S_0 by adding to it special vertices from A_0 and B_0 until every edge in $G[N(S_0)]$ has endpoints covered by some $N(X_i)$ and some additional conditions hold. In the latter case, averaging produces a pair X_i, $N(X_i)$ which satisfies the second case of the lemma statement.

Let $E_{in} = \{ab \in E(G[N(S_0)]) : a, b \in N(X_i) \text{ for some } i\}$ be the set of *internal* edges, i.e. those that are covered by neighbourhoods of X_is, and $E_{ext} = E(G[N(S_0)]) \setminus E_{in}$ the *external* edges. For a vertex v, let $t(v) = |\{i : v \in N(X_i)\}|$ be the number of X_is it is adjacent to, and $m(v) = |\{i : av \in E_{ext} \text{ for some } a \in N(X_i)\}|$ the number of neighbourhoods to which it has external edges.

Call $\sigma = \delta^4 f$ the *cost threshold*.

To simplify notation in the conditions below, the degree $d(v)$ of a vertex v refers to its degree in $G[A_0, B_0]$ at that particular step. Similarly, $t(v)$ and 'external' edges are defined in terms of the components of S_0 at that particular step. Perform the following operations until no condition in satisfied:

1. If $a \in A \setminus S_0$ and $t(a) \geqslant 2$, remove a from $A_0 \cup B_0$ and add it to S_0.
2. If $b \in B \setminus S_0$ and $d(b) \leqslant \sigma t(b)$, remove b and $N_A(b)$ from $A_0 \cup B_0$ and add them to S_0.
3. If ab is external with $b \in B$ and b has external edges to, say, $N(X_1), \ldots, N(X_k)$, then choose representatives $v_1, \ldots, v_k$ such that $v_i \in N(X_i)$ and bv_i is external. Remove $v_1, \ldots, v_k$ from $A_0 \cup B_0$ and add them to S_0.

Let A_1, B_1, S_1 be the sets obtained after the process stops.

Claim 2.13. *There are no external edges after the process stops.*

Proof. Since Item 3 does not hold, there cannot be any external edges. $\blacksquare$

Claim 2.14. $|A_1| \geqslant |A|(1 - 9\delta^{-3})$ *and each* $a \in A_1$ *has* $\delta(1 - \eta \log \delta/2\delta)$ *neighbours in* B_1.

Proof. Items 1 and 2 decrease the number of components by precisely $t(a) - 1$ and $t(b) - 1$, while removing at most 1 and $\sigma t(b)$ vertices from A_0, respectively. Additionally, each time Item 3 is performed it reduces the number of connected components by at least k so in total this operation removes at most t vertices from A_0. Therefore, in total the process removes at most $2(t + t \cdot \sigma) < 8|A|/\delta^3$ vertices from A_0. The desired bound follows as $|A_0| \geqslant |A|(1 - \delta^{-3})$.

As we only ever remove vertices from A_0 or vertices in B_0 with all their neighbours, the degree condition is satisfied. ■

Claim 2.15. *Once the process stops,* $\sum_{v:t(v)\geqslant 2} t(v) \leqslant |A|/\delta^3 f$.

Proof. Since Item 2 does not hold, any vertex $v \in B_1$ satisfies $d(v) > \sigma t(v)$ and thus

$$\sum_{v:t(v)\geqslant 2} t(v) \leqslant \frac{1}{\sigma} \sum_{b\in B} d(b) = \frac{\delta|A|}{\delta^4 f}.$$ ■

Finally, we seek a connected component X_i of S_1 whose neighbourhood satisfies $|N(X_i) \cap A_1| \geqslant f(1 - \delta^{-2})|N(X_i) \cap B_1|$. Once this is found we are done by taking $S = X_i$, $A' = N(X_i) \cap A_1$, $B' = N(X_i) \cap B_1$ as all vertices in $N(X_i) \cap A_1$ have $(1 - \eta \log \delta/2\delta)\delta$ neighbours in $N(X_i) \cap B_1$. We find this X_i via an averaging argument.

Suppose no such X_i exists, then

$$\begin{aligned}\sum |N(X_i) \cap A_1| &< f(1 - \delta^{-2}) \sum |N(X_i) \cap B_1| \\ &\leqslant f(1 - \delta^{-2}) \sum_{v\in B_1} t(v) \\ &\leqslant f(1 - \delta^{-2})(\sum_{v:t(v)\geqslant 2} t(v) + |B|).\end{aligned}$$

By Claim 2.14 the expression on the left hand side is at least $|A|(1 - 9\delta^{-3})$ and Claim 2.15 bounds the expression on the right hand side by at most

$$f(1 - \delta^{-2})(|A|/\delta^3 f + |B|) \leqslant (1 - \delta^{-2})(2 + |A|),$$

since $|A| = f|B|$. Now, $|A| > 2(\delta^2 + 9\delta + 7\frac{\delta}{\delta-9})$, so $|A|(1 - 9\delta^{-3}) > (|A| + 2)(1 - \delta^{-2})$, which is a contradiction. □

We now prove the main theorem. Note that the statement holds with $c \geqslant 1200$.

Theorem 2.2. *There is an absolute constant $c > 0$ such that for any positive integer t, if G is a graph with $d(G) \geqslant ct \log^3 t$ then G contains a K_t-dominating minor.*

In the proof of Lemma 2.10, the highly imbalanced bipartition is obtained by considering the partition (A, B) where $B = \{v \in V(G) : d(v) \geqslant Ld\}$. In particular, we use below that if $\Delta(G) < L$ then the first case of the lemma must hold.

Proof. By Lemma 2.10 applied with $L = d^{100}$, either G contains an $O(\log d)$-almost-regular graph of average degree $\Omega(d/\log^2 d)$, or there is a bipartite subgraph (A, B) where $|A| \geqslant d^{100}|B|$ and $e(A, B) \geqslant e(G)/8$. In the first case, apply Corollary 2.9 with $d = d/\log^2 d$ and $c = 1200 \log d \leqslant 10^{10} t$ to obtain a K_t-dominating model.

In the second case, we repeatedly apply Lemma 2.11. If Lemma 2.11 (i) ever occurs after applying the lemma then G contains a subgraph (A', B') with average degree at least $d \cdot d^8/(2d^8 + 2d^7) \geqslant d/4$ and maximum degree at most d^8, so by Lemma 2.10 applied with $L = d^{100}$ it contains an $O(\log d)$-almost-regular graph of average degree at least $\Omega(d/\log^2 d)$ and therefore contains a K_t-dominating model, as in the previous paragraph.

Otherwise, let S_i be the connected dominating set obtained after the ith iteration, and (A_i, B_i) the corresponding bipartite subgraph it dominates. Each successive application of Lemma 2.11 reduces the average degree by an additive term of at most $\eta \log d$, therefore since for c large enough we have $d - t \cdot \eta \log d \geqslant d/2$. Hence, for $i \leqslant t$ we have $|A_i| \geqslant (1 - 4d^{-2})^t d^{100}|B_i| > d^{17}|B_i|$, so we can apply Lemma 2.11 at least t times to obtain connected sets $S_1, \ldots, S_t$ such that S_i dominates S_j for $j > i$, as required. □

2.3 Dominating K_5-models

The following result shows that graphs with no dominating K_5-model have minimum degree at most 5 and therefore are 6-colourable.

Theorem 2.3. *For $n \geqslant 4$, every n-vertex graph with no dominating K_5-model has at most $3n - 6$ edges.*

Proof. For $n = 3, 4$, all n-vertex graphs have at most $3n - 6$ edges and no dominating K_5-model.

So suppose that $n \geqslant 5$, and let G be any n-vertex graph with at least $3n - 5$ edges. By passing to a subgraph if necessary, we may assume without loss of generality that G is connected. Fix a root vertex $r \in V(G)$. We will show by induction that G contains a dominating K_5-model $(T_1, T_2, T_3, T_4, T_5)$ with $r \in T_1$. When $n = 5$, the only graph with more than $3n - 6$ edges is K_5 itself, and this is a dominating K_5-model with singleton branch sets so any vertex can be chosen to be in T_1.

Now let $n > 5$ and consider $N(r)$. If there is a vertex $s \in N(r)$ such that rs is part of at most two triangles in G, then contract the edge rs to form a new vertex $r' \in G/rs$. The graph G/rs has $n-1$ vertices and at least $3n-5-3 = 3(n-1)-5$ edges, so by the inductive hypothesis G/rs has a dominating K_5-model with $r' \in T_1$. We can then obtain the desired dominating K_5-model in G by replacing r' with $\{r, s\}$ in T_1.

Otherwise, if there is no such vertex s, then all vertices in $N(r)$ have degree at least 3. Then by [72, Theorem 7], there is a dominating K_4-model in $G[N(r)]$, say (T_1', T_2', T_3', T_4'). Then $(\{r\}, T_1', T_2', T_3', T_4')$ is a dominating K_5-model in G. □

We now work towards a proof of Theorem 2.4. In the following we make use of results on tripods and triads from [104]. A *triad* T is a tree with three leaves (called the *feet* of T) and a single vertex of degree three, its *apex*. Removing the apex from T produces three disjoint paths, the *legs* of T. We say that a triad T is *lean* if any triad T' with $V(T') \subseteq V(T)$ and same feet has $V(T) = V(T')$. A *legless tripod* with feet v_1, v_2, v_3 is a pair (T_1, T_2) of triads with feet v_1, v_2, v_3 that are otherwise vertex disjoint.

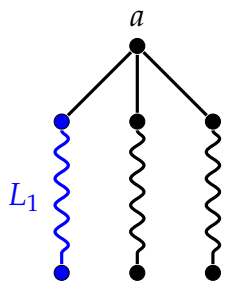


Figure 2.1. A triad with apex a and with leg L_1 in blue.

Lemma 2.16 ([104, 3.1])**.** *Let G be a 4-connected graph and suppose T is a triad in G with feet v_1, v_2, v_3 and W a connected component of $G \setminus T$. Then there is a lean triad T' with feet v_1, v_2, v_3 such that $G \setminus T'$ is connected and $V(T') \cap V(W) = \varnothing$.*

Lemma 2.17 ([104, 3.4-3.5])**.** *Let G be a 4-connected non-planar graph and v_1, v_2, v_3 distinct vertices. Then G has a legless tripod with feet v_1, v_2, v_3.*

It is well known that non-planar graphs contain K_5 or $K_{3,3}$ as a minor, and a short argument shows that they must contain a K_5-minor if they are further 4-connected. We show that such graphs also contain a K_5-dominating model, provided they contain a triangle. It seems difficult to remove this last condition as it is harder to find K_5-dominating models with many singleton branch sets in graphs of high girth.

Theorem 2.4. *Let G be 4-connected non-planar and suppose it has at least one triangle. Then G contains a dominating K_5-model.*

This immediately follows from the lemma below by setting v_1, v_2, v_3 to be vertices of a triangle, as $(G - T_b, T_b \setminus \{v_1, v_2, v_3\}, v_1, v_2, v_3)$ is then a dominating K_5-model.

Lemma 2.18. *Let G be a 4-connected non-planar graph and v_1, v_2, v_3 three of its vertices. Then G contains a legless tripod (T_a, T_b) with feet v_1, v_2, v_3 such that $V(G) \setminus V(T_b)$ is a connected dominating set.*

Moreover, if $v \notin \{v_1, v_2, v_3\}$ is a vertex of G we may assume that $v \notin V(T_b)$.

Proof. Fix v_1, v_2, v_3 distinct vertices in G. Lemma 2.17 gives a legless tripod (T_a, T_b) with feet v_1, v_2, v_3. By Lemma 2.16 we may assume that T_b is lean and $X = G \setminus V(T_b)$ is connected, with $V(T_a) \cap V(T_b) = \{v_1, v_2, v_3\}$. If at this stage $v \in V(T_b)$, we may relabel the legless tripod so that $v \in V(T_a)$, and after applying Lemma 2.16 again we may assume T_b is lean with X connected and $V(T_a) \cap V(T_b) = \{v_1, v_2, v_3\}$. We claim that for every vertex $v \in V(T_b) \setminus \{v_1, v_2, v_3\}$, the degree of v in $G[V(T_b)]$ is at most 3. This suffices to prove the statement as every vertex in T_b then has at least one neighbour in X, and $v \notin V(T_b)$.

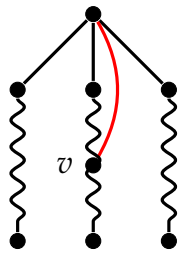
Let b be the apex vertex in T_b and L_1, L_2, L_3 the three possibly empty *legs* of T_b labelled so that $v_i \in L_i$, and u_i the neighbour of b in L_i.

Claim 2.19. *The following hold.*

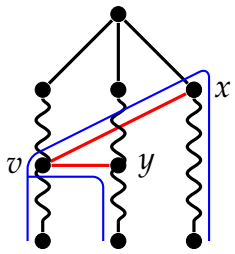
1. *The vertex b has precisely 3 neighbours in $V(T_b)$.*
2. *No $v \in L_1 \setminus v_1$ is adjacent to both a vertex in L_2 and a vertex in L_3.*
3. *No $v \in L_1$ is adjacent to more than one vertex in any L_i.*
4. *No $v \in L_1$ is adjacent to a vertex $w \in L_1$ other than its neighbours in L_1.*

Proof.

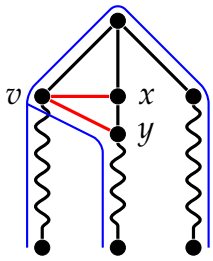
1. Suppose for contradiction that b is adjacent to $v \in V(T_b) \setminus \{u_1, u_2, u_3\}$ in G, say $v \in L_1$. Then the path P from v to v_1 in L_1 is strictly shorther than L_1 and so the triad T' with apex b and legs P, L_2, L_3 contradicts the leanness of T.
2. If not, say $v \neq v_1$ in L_1 is adjacent to x_2 in L_2, x_3 in L_3 and let x_1 be its neighbour closer to v_1 on L_1, then for each i let P_i, be the subpath of L_i from x_i to v_i. Now the triad T' with apex v and legs P_1, P_2, P_3 contradicts the leanness of T.
3. If not, say v in L_1 is adjacent to x, y in L_2, where y is on the path in L_2 from x to v_2. Then let $P_1 \subseteq L_1$ be the subpath from v to v_1, $P_2 \subseteq L_2$ be the subpath from y to v_2 and P_3 be obtained by joining the path from v to u_1, the edges u_1b, bu_3 and L_3. Then the triad with apex v and legs P_1, P_2, P_3 is contained in T_b yet misses x, a contradiction.



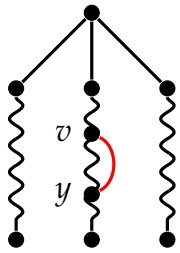
(a) Apex has degree three.



(b) Neighbours in $\leqslant 1$ other leg.



(c) At most one neighbour per leg.



(d) No other neighbours in own leg.

4. If not, say vy is such an edge where y is on the vv_1 path in L_1 and there is a distinct vertex x on the vy path in L_1. Let $P \subseteq L_1$ be the path obtained by joining bu_1, the bv path, vy and the yv_1 path. Then the triad T with apex b and legs P, L_2, L_3 is contained in T_b yet misses x, a contradiction. ■

By symmetry the analogous statements hold for vertices in L_2 and L_3, and therefore each $v \in V(T_b) \setminus \{v_1, v_2, v_3\}$ is adjacent to at most three vertices in $G[V(T_b)]$. □

When G is 4-connected non-planar and has girth at least 4, the above approach fails as we can no longer choose the three most dominated sets to be singleton vertices. The simplest structure for these would consist of a path and two dominated vertices forming a cycle, but it then becomes difficult to find two connected sets dominating this cycle with one further dominating the other.

However, we can still choose v_1, v_2, v_3 in Lemma 2.18 to form a path, thereby producing a dominating $K^*_{3,2}$-model. This allows us to show Corollary 2.5 with a bit more work, as detailed in the next section.

2.4 Dominating $K^*_{s,t}$-models

The definition of a minor model generalises to graphs other than cliques, provided we specify some additional information. Namely, suppose that H is a graph on k vertices with an ordering $\sigma = (\sigma(1), \ldots, \sigma(k))$ of its vertices. We say that G contains

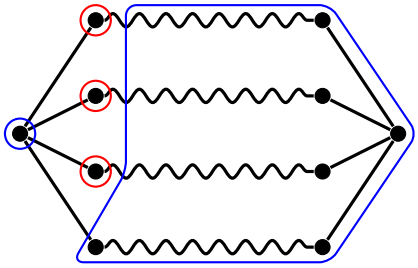


Figure 2.3. A sufficient subgraph for a $K^*_{2,3}$ dominating minor of type I. In blue the two most dominating branch sets, and in red the other three.

a *dominating (H, σ)-model* if there is a sequence of pairwise disjoint non-empty vertex sets $(T_1, \dots, T_k)$ such that

- each T_i induces a connected subgraph of G,
- for all $1 \leqslant i < j \leqslant k$, if $\sigma(i)\sigma(j) \in E(H)$ then every vertex in T_j has a neighbour in T_i.

That is, $\sigma(i)$ is the vertex of H corresponding to the ith branch set T_i in the model, and T_i dominates the sets in $\{T_{i+1}, \dots, T_k\}$ corresponding to vertices $\sigma(j)$ of H adjacent to $\sigma(i)$. Note that in this definition the relative order of consecutive vertices within an independent set or within a clique is not important, in the sense that if σ and τ only differ on such pairs of vertices, then any (H, σ)-model is an (H, τ)-model.

We consider the case $H = K^*_{s,t}$ for some specific values of s and t. Recall that $K^*_{s,t}$ is obtained from the complete bipartite graph $K_{s,t}$ on parts A, B of sizes s, t respectively by adding all edges with both endpoints in A. We will consider two types of orderings of vertices in $K^*_{s,t}$; orderings in which all s vertices occur before all t vertices will be called *Type I*, and those with all t vertices before all s vertices are called *Type II*. Accordingly, we will say that a graph G has a *dominating $K^*_{s,t}$-model of Type I or Type II* if it has a dominating $(K^*_{s,t}, \sigma)$-model where σ is of Type I or Type II respectively.

The result below follows directly from the results of the previous section. One might hope that an inductive argument analysing cuts of different sizes would bypass the connectedness restriction, but this seems non-trivial.

Corollary 2.5. *Let G be a simple graph that is 4-connected and not 4-colourable. Then for any $v \in V(G)$, G contains a dominating $K^*_{3,2}$-model of Type I where v is in the most dominating branch set.*

Proof. Since G is 4-connected, there is a path $v_1v_2v_3$ disjoint from v. By Lemma 2.18, G contains a legless tripod (T_a, T_b) with feet v_1, v_2, v_3 such that $V \setminus V(T_b)$ is a dominating set inducing a connected subgraph and $v \notin V(T_b)$. Then $(G \setminus V(T_b), T_b \setminus \{v_1, v_2, v_3\}$,

$v_2, v_1, v_3)$ is a $K^*_{3,2}$-dominating model of Type I with v in the most dominating branch set, as desired. □

In the arguments below we use a result of Carmesin, Diestel, Hamann and Hundertmark [31] regarding existence of k-lean tree decompositions. A tree decomposition $(T, \mathcal{V})$ is called *k-lean* if none of its bags contain another and for any two sets of vertices $Z_1 \subseteq V_{t_1}, Z_2 \subseteq V_{t_2}$ of size $|Z_1| = |Z_2| = \ell \leqslant k+1$ either there are ℓ vertex disjoint paths from Z_1 to Z_2, or there is an edge tt' along the t_1, t_2 path in T such that $|V_t \cap V_{t'}| < \ell$. We sometimes refer to this last condition as the *leanness property*.

Theorem 2.20 ([31])**.** *Every graph has a k-lean tree decomposition.*

Theorem 2.6. *Let G be a graph which is not $(t+1)$-colourable. Then G contains a $K^*_{2,t}$ dominating model of Type I.*

Proof. Let G be a simple graph which is not $(t+1)$-colourable. It suffices to find distinct vertices x, y with $t+1$ internally vertex disjoint paths connecting them. Indeed, let $P_0, \dots, P_t$ be such paths in increasing order of edge length, and for each $i > 0$ let x_i be the neighbour of x in P_i. Then $(G \setminus \{x, x_1, \dots, x_t\}, x, x_1, \dots, x_t)$ is a $K^*_{2,t}$-dominating model of type I, as desired.

Let G be a minimal counterexample, meaning that any smaller graph is either $(t+1)$-colourable or contains a $K^*_{2,t}$-dominating minor, and let $(T, \mathcal{V})$ be a $(t+1)$-lean tree decomposition of G. We show that the tree T must satisfy the following properties. First, that no leaf is adjacent to a vertex of degree at most 2, and second, that no two leaves are adjacent to the same vertex. No tree can satisfy both conditions unless it consists of a single vertex. This contradicts the fact that G is a counterexample, as shown below.

Every vertex in G has degree at least $t+1$ as otherwise we may $(t+1)$-colour G inductively. Since T consists of a single vertex, for any two vertices $x, y \in V(G)$ we may apply the leanness property to subsets $Z_1 \subseteq N(x), Z_2 \subseteq N(y)$ of size $t+1$ and obtain the desired $t+1$ vertex disjoint paths between x and y. So we assume from now on that T has at least two vertices.

Claim 2.21. *Let t be a leaf of T and p its neighbour. Then $V_t \setminus V_p = \{v\}$ for some vertex $v \in V(G)$.*

Proof. Suppose not and let $v, v' \in V_t \setminus V_p$ be distinct. Then as v, v' are contained in no other bag of T, $N(v), N(v') \subseteq V_t$. The leanness property applied to subsets

$Z_1 \subseteq N(v), Z_2 \subseteq N(v')$ of size $t+1$ gives $t+1$ internally vertex disjoint paths from v to v', a contradiction. ■

Claim 2.22. *Let t be a leaf of T and p its neighbour. Then t is the unique leaf adjacent to p.*

Proof. Suppose not and let t' be another leaf adjacent to p. By Claim 2.21 there are vertices v, v' such that $\{v\} = V_t \setminus V_p, \{v'\} = V_{t'} \setminus V_p$, so that $N(v), N(v') \subset V_p$. Apply leanness to $Z_1 \subseteq N(v)$, $Z_2 \subseteq N(v')$ subsets of size $t+1$ to obtain $t+1$ internally vertex disjoint paths between v and v', as desired. ■

Claim 2.23. *Let t be a leaf of T and p its neighbour. Then p has degree at least three.*

Proof. Suppose p has degree 1 and write $\{v\} = V_t \setminus V_p$ and $\{v'\} = V_p \setminus V_t$ so that $N(v), N(v') \subset V_p$. Apply the leanness property to subsets of these to obtain $t+1$ internally vertex disjoint paths from v to v', a contradiction.

Suppose p has degree 2, let t' be its second neighbour, write $\{v\} = V_t \setminus V_p$ so that $N(v) \subseteq V_p$ and choose $v' \in V_p \setminus V_{t'}$. If $v' \notin V_t$ then $N(v') \subseteq V_p$ and leanness applied to subsets of $N(v), N(v')$ gives $t+1$ internally vertex disjoint paths between v and v', as desired.

If $v' \in V_t$, we may assume without loss of generality that v is adjacent to v' as otherwise by Claim 2.21, $N(v') \subseteq V_p$ which is the same situation as in the previous paragraph.

We claim that $|N(v)| = t+1$. Otherwise, let $Z_1 \subseteq N(v) \setminus v' \subseteq V_p$ and $Z_2 \subseteq v' \cup N(v') \setminus v \subseteq V_p$ be sets of size $t+1$. Leanness applied to Z_1, Z_2 gives $t+1$ vertex disjoint paths between $v \cup N(v)$ and $N(v') \setminus v'$, at most one of which uses the edge vv'. Namely, there are at least t vertex disjoint paths from $N(v) \setminus v'$ to $N(v') \setminus v$ and the edge vv', as desired.

Now let $Z_1 = N(v) \subseteq V_p$ and $Z_2 \subseteq v' \cup N(v') \setminus v \subseteq V_p$ be a set of size $t+1$. Leanness applied to Z_1, Z_2 gives $t+1$ vertex disjoint paths from $N(v)$ to $v' \cup N(v) \setminus v$, and since each path uses exactly one vertex of Z_1, no path can use the vertex v. Hence, we found t vertex disjoint paths from $N(v) \setminus v'$ to $N(v') \setminus v$ which, together with the edge vv', form the desired $t+1$ internally vertex disjoint paths from v to v'. ■

□

To find a $K^*_{2,t}$-dominating minor of Type II, it suffices to find two adjacent vertices with $t+1$ internally vertex disjoint paths (including the edge) between them. It is natural to try a similar approach to the above and while Claim 2.21 and Claim 2.23 can be adapted

by using the fact that tree decomposition with lexicographically minimal fatness are lean, it seems difficult to obtain a result akin to Claim 2.22.

2.5 Deficiency colourings

Lemma 2.24. *Let $t \in \mathbb{N}$ and $G = (A, B)$ be a bipartite graph with at least one vertex such that $|A| \geqslant (t-1)!|B|$ and $\deg(a) = t - 1$ for each $a \in A$. Then G contains a dominating K_t-model.*

Proof. We may assume that G is connected by averaging over its connected components.

We proceed by induction on t. When $t = 1$ or 2 the statement is clear (interpreting $-1! = 0! = 0$) since a single vertex and a single edge are dominating K_1- and K_2-models respectively. So suppose $t \geqslant 3$.

For a subset $B' \subseteq B$, define $A^{\geqslant 2}_{B'}$ to be the set of vertices $a \in A$ with at least 2 neighbours in B'. Choose such a subset $B' \subset B$ maximising $|B'|$ subject to the conditions

1. $G[A^{\geqslant 2}_{B'} \cup B']$ is connected, and
2. $|A^{\geqslant 2}_{B'}| < (t-1)! \cdot |B'|$.

Now define $A^1_{B'}$ to be the set of vertices $a \in A$ with exactly 1 neighbour in B', and let $B'' := N(A^1_{B'}) \setminus B'$ be its neighbourhood outside B'. Since A, B satisfy $|A| \geqslant (t-1)!|B| \geqslant (t-1)!|B'|$, we must have $A^{\geqslant 2}_{B'} \neq A$ and therefore $B' \neq B$, so $A^1_{B'}$ and B'' are non-empty as G is connected. Also, any $a \in A$ and its neighbourhood $N(a)$ satisfy both conditions, so both $|A^{\geqslant 2}_{B'}|$ and $|B'|$ are positive.

By definition $A^1_{B'}$ is dominated by the connected subgraph $G[A^{\geqslant 2}_{B'} \cup B']$. In fact, we show in the claim below that after adding not too many vertices to this subgraph we can assume B'' is also dominated. This allows us to apply induction if $A^1_{B'}$ is large enough to begin with.

Claim 2.25. *If $|A^1_{B'}| \geqslant ((t-2)! + 1) \cdot |B''|$, then G contains a dominating K_t-model.*

Proof. Choose a minimal set of vertices $S \subset A^1_{B'}$ such that for every $b \in B''$ there is a neighbour of b in S. By definition each $b \in B''$ has a neighbour in $A^1_{B'}$, so $|S| \leqslant |B''|$. Suppose we move the vertices in S to $A^{\geqslant 2}_{B'}$; that is, define $A' := A^{\geqslant 2}_{B'} \cup S$, and $A'' := A^1_{B'} \setminus S$.

Since B' dominates $A^1_{B'}$ by definition, it dominates A'' as well. Further, by construction A' dominates B'', so $A' \cup B'$ dominates $A'' \cup B''$. In addition, $G[A' \cup B']$ is connected

since $G[A^{\geqslant 2}_{B'}, B']$ was connected and all vertices in S have a neighbour in B'. We take $A' \cup B'$ to be the most dominating branch set in our dominating model.

To find the rest of the model, observe that each vertex in A'' has at least $t-2$ neighbours in B'', and $|A''| = |A^1_{B'}| - |S| \geqslant |A^1_{B'}| - |B''| \geqslant (t-2)! \cdot |B''|$ using the assumption of the claim. By the induction hypothesis, $G[A'' \cup B'']$ contains a dominating K_{t-1}-model $(F_1, \ldots, F_{t-1})$. Then $(A' \cup B', F_1, \ldots, F_{t-1})$ is a dominating K_t-model, as desired. ■

To conclude the proof, we show that if $|A^1_{B'}| < ((t-2)!+1) \cdot |B''|$ then $|B'|$ is not maximal to reach a contradiction. As vertices in $A^1_{B'}$ have exactly one neighbour in B', we can bound

$$e(A^1_{B'}, B'') = (t-2)|A^1_{B'}| \leqslant (t-2)((t-2)!+1)|B''| \leqslant (t-1)!|B''|.$$

Hence, there is a vertex $b \in B''$ with at most $(t-1)!$ neigbours in $A^1_{B'}$. Call this set of neighbours $X \subseteq A^1_{B'}$. Note that each $a \in X$ has exactly one neighbour in B' and is adjacent to b by definition, so $A^{\geqslant 2}_{B' \cup b}$ is precisely $A^{\geqslant 2}_{B'} \cup X$. Moreover, $G[A^{\geqslant 2}_{B' \cup b} \cup (B' \cup b)]$ is connected.

$$|A^{\geqslant 2}_{B' \cup b}| = |A^{\geqslant 2}_{B'}| + X \leqslant (t-1)! \cdot |B'| + (t-1)! = (t-1)! \cdot |B' \cup b|.$$

Thus, $B' \cup b$ (and the corresponding set $A^{\geqslant 2}_{B' \cup b}$) satisfies conditions 1 and 2 above, which contradicts the maximality of $|B'|$. □

Theorem 2.7. *For every $t \in \mathbb{N}$ there is a constant $c_t > 0$ such that the vertex set of any graph G that does not contain a dominating K_t-model can be partitioned into $t-1$ parts $X_1, \ldots, X_{t-1}$ with $\Delta(G[X_i]) \leqslant c_t$ for each i.*

Proof. Suppose G is a minimal counterexample. Define $B = \{v \in V(G) : \deg(v) > c_t\}$ to be the vertices of degree larger than c_t and $A = V(G) \setminus B$ the remaining vertices. If $B = \varnothing$ then $\Delta(G) \leqslant c_t$ so G is not a counterexample.

Suppose $a \in A$ has at most $t-2$ neighbours in B. Since G is a minimal counterexample the graph $G \setminus a$ may be partitioned into $t-1$ parts $X'_1, \ldots, X'_{t-1}$ with induced maximum degree at most c_t. Then a is not adjacent to any vertex in (say) $X'_{t-1} \cap B$, so the partition $X'_1, \ldots, X'_{t-2}, X'_{t-1} \cup a$ of G has induced maximum degrees at most c_t, a contradiction. Hence we may assume $|\{ab : b \in B\}| \geqslant t-1$ for each $a \in A$.

Without loss of generality we may remove edges from the bipartite graph $G[A, B]$ until each $a \in A$ has *exactly* $t-1$ neighbours in B. As G is a counterexample it must satisfy

$\frac{e(A,B)}{|A|+|B|} \leqslant d(G) = 400t\log^3 t$ as otherwise by Theorem 2.2 it would contain a dominating K_t-model. In particular, $c_t|B| + (t-1)|A| \leqslant 400t\log^3 t \cdot (|A| + |B|)$ and hence

$$|A| \geqslant \frac{c_t - 400t\log^3 t}{400t\log^3 t - t + 1}|B|.$$

The bipartite graph $G[A, B]$ satisfies the conditions of Lemma 2.24 and hence G contains a dominating K_t-model, a contradiction. □

Chapter 3

Clustered colouring of graph products

This chapter is based on joint work with Campbell, Gollin, Hendrey, Lesgourges, Mohar, Tan and Wood [30]. Generally speaking, we are concerned with the growth of clustering numbers of c-colourings of graph products of the form $H_1 \boxtimes H_2$ in terms of their number of vertices, with restrictions on the maximum treewidth t of the H_is, and occasionally on the maximum degree Δ of one or both of them. As such, in this chapter we usually consider t, c to be constants, as well as Δ where relevant.

Recall that a *colouring* of a graph G is a function that assigns a 'colour' to each vertex of G. A *c-colouring* is a colouring that uses at most c colours. We allow adjacent vertices to be assigned the same colour. A colouring of a graph has *clustering* k if the maximum number of vertices in a monochromatic component equals k. Here, a *monochromatic component* is a connected component of the subgraph induced by the vertices assigned a single colour.

Directly applying structural tools from [39], namely that graphs of bounded treewidth and bounded degree have bounded tree-partition-width, Ding, Oporowski, Sanders and Vertigan [40] showed that graphs of bounded treewidth and bounded degree have 2-colourings of bounded clustering.

Theorem 3.1 ([40])**.** *Every graph with maximum degree Δ and treewidth k admits a 2-colouring with clustering at most $24k\Delta$.*

In this statement the bounded degree condition is necessary but not sufficient. Necessity can be seen by considering the fan graph F_n formed by adding a dominating vertex to P_n (see Figure 3.1): any 2-colouring of F_n has a monochromatic component with roughly $\sqrt{n}$ vertices. Examples showing that bounded degree is not sufficient can be obtained by taking the line graph of any 4-regular graph G of large girth, namely viewing edges of G as vertices of a new graph with adjacency between two edges of G whenever they

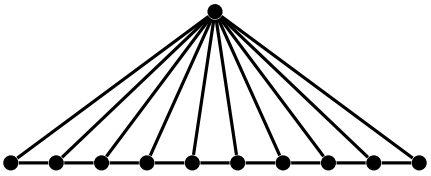


Figure 3.1. A fan F_{10}.

share an endpoint. Such a graph is 6-regular yet any partition of its edges into two has a part with average degree at least 2 that therefore contains a cycle. Since the initial girth is large, the component containing this cycle is also large.

This last example appeared in a later paper of Alon, Ding, Oporowski and Vertigan [4] and is perhaps more concise than the original construction of Ding, Oporowski, Sanders and Vertigan. Another class of examples showing bounded degree is not sufficient is the set of triangulated $n \times n$ grids: every vertex has degree at most 6, but by the Hex Lemma stated below any 2-colouring of its vertices contains a monochromatic path of length at least n.

Lemma 3.2 (Hex Lemma [89, Proposition 6.1.4])**.** *Let G be a plane internal triangulation with outer-cycle (a, b, c, d). In every 2-colouring χ of G with $\chi(a) = \chi(c)$ and $\chi(b) = \chi(d)$, there is a monochromatic path either between a and c or between b and d.*

This lemma is particularly useful for us as several constructions in this chapter involve taking strong products of graphs that contain long paths, producing grid-like structures. It was used in a similar manner by Matoušek and Přívětivý [90] to prove bounds on the clustering number of 2-colourings of d-dimensional grids. We also note that some of our arguments are inspired from the paper of Linial, Matoušek, Sheffet and Tardos [86] who were interested in clustering numbers of colourings of minor-closed classes.

Product structure theorems first came to light in a 2019 breakthrough paper by Dujmović, Joret, Micek, Morin, Ueckerdt and Wood [44] where they showed that any planar graph can be found as a subgraph of $H \boxtimes P$, where H is a graph of bounded treewidth and P is a long enough path in order to prove that planar graphs have bounded queue number. This result naturally sparked two tightly related lines of research: finding similar product structure theorems for natural classes of graphs, and studying the properties of such graph products.[1]

[1]While product structure results have been in the spotlight recently, one can argue that they are not a novel concept. For example, the proof of Theorem 3.1 comes from a fact shown in 1996 by Ding and Oporowski [39] that can be rephrased to say that graphs of maximum degree Δ and treewidth k are isomorphic to subgraphs of $T \boxtimes K_{24\Delta k}$, where T is a tree of maximum degree $O(\Delta)$.

The work presented in this chapter follows the latter line: we study clustering numbers of colourings of graph products with an eye towards graph product structure theorems. In particular, we focus on clustered colouring of strong products of two graphs with bounded treewidth (see Section 3.1 for definitions of these notions). One motivation for considering such graphs is that it has been recently shown that many natural graph classes are subgraphs of the strong product of a graph with bounded treewidth and a path [44, 45, 43, 71, 67] or the strong product of two graphs with bounded treewidth [88, 42] (see Section 3.6 for more on this theme).

We study two variations of these products. In the first, we assume that one of the graphs has bounded maximum degree, generalising the paths in the aforementioned products. We then study the scenario with no restriction on the maximum degree. In each scenario, given an integer $c \geqslant 2$, our goal is to prove tight bounds on the optimal clustering in c-colourings of products $H_1 \boxtimes H_2$, where H_1 and H_2 are graphs of bounded treewidth. Our main results are stated in the next section. They focus on the exponent of $|V(H_1 \boxtimes H_2)|$ in the clustering value.

Main results

First, consider 2-colourings of graph products. We show the following asymptotically tight bound on the optimal clustering in this case.

Theorem 3.3. *For any fixed integer $t \geqslant 1$ and any graphs H_1, H_2 with treewidth at most t, the graph $H_1 \boxtimes H_2$ is 2-colourable with clustering $O(|V(H_1 \boxtimes H_2)|^{2/3})$. Furthermore, there are infinitely many graphs H with treewidth 2 and paths P such that every 2-colouring of $H \boxtimes P$ has clustering $\Omega(|V(H \boxtimes P)|^{2/3})$.*

Theorem 3.3 shows that in the case of two colours, perhaps surprisingly, the maximum degree condition has no significant impact on the optimal clustering value. For more than two colours, the maximum degree condition has a significant impact on the optimal clustering value of colourings of graph products.

Now consider the case of three colours. When one graph has bounded degree, we prove the following asymptotically tight bound.

Theorem 3.4. *For any fixed integers $t, \Delta \geqslant 1$, for any graph H_1 with treewidth at most t, and any graph H_2 with treewidth at most t and maximum degree at most Δ, the graph $H_1 \boxtimes H_2$ is 3-colourable with clustering $O(|V(H_1 \boxtimes H_2)|^{3/7})$. Furthermore, there are infinitely many graphs H with treewidth 3 and paths P such that every 3-colouring of $H \boxtimes P$ has clustering $\Omega(|V(H \boxtimes P)|^{3/7})$.*

The following result shows that dropping the maximum degree restriction impacts the exponent in both upper and lower bounds.

Theorem 3.5. *For any fixed integer $t \geqslant 1$ and any graphs H_1, H_2 with treewidth at most t, the graph $H_1 \boxtimes H_2$ is 3-colourable with clustering $O(|V(H_1 \boxtimes H_2)|^{4/7})$. Furthermore, there exist infinitely many pairs of graphs H_1, H_2 with treewidth 2 such that every 3-colouring of $H_1 \boxtimes H_2$ has clustering $\Omega(|V(H_1 \boxtimes H_2)|^{1/2})$.*

For $c \geqslant 4$ colours, when one graph has bounded degree, we obtain the following general upper bound.

Theorem 3.6. *For any fixed integers $t, c, \Delta \geqslant 1$, for any graph H_1 with treewidth at most t, and any graph H_2 with treewidth at most t and maximum degree at most Δ, the graph $H_1 \boxtimes H_2$ is c-colourable with clustering $O(|V(H_1 \boxtimes H_2)|^{c/(c^2-c+1)})$.*

We use the lower bound from Theorem 3.4 to prove the following general lower bound.

Proposition 3.7. *For any fixed integer $c \geqslant 3$, there are infinitely many graphs H with treewidth c and infinitely many paths P, such that every c-colouring of $H \boxtimes P$ has clustering $\Omega(|V(H \boxtimes P)|^{1/(c-\frac{2}{3})})$.*

There is a sharp change in the asymptotic behaviour of both upper and lower bounds on the clustering in c-colourings of $H_1 \boxtimes H_2$ when we remove the restriction on the maximum degree of H_2. Indeed, while the bounds of Theorem 3.6 and Proposition 3.7 are of the form $n^{1/(c-O(1))}$, in the unrestricted setting we prove the following upper and lower bound of the form $n^{\Theta(1/\sqrt{c})}$.

Theorem 3.8. *For any fixed integers $t \geqslant 1$ and $c \geqslant 4$, for any graphs H_1 and H_2 with treewidth at most t, the graph $H_1 \boxtimes H_2$ is c-colourable with clustering $O(|V(H_1 \boxtimes H_2)|^{1/\lfloor\sqrt{c}\rfloor})$. Furthermore, for any integer $c \geqslant 4$, there are infinitely many pairs of graphs H_1, H_2 with treewidth at most $\sqrt{c}$ such that every c-colouring of $H_1 \boxtimes H_2$ has clustering $\Omega(|V(H_1 \boxtimes H_2)|^{1/2\sqrt{c}})$.*

Of course, Theorems 3.6 and 3.8 are only of interest for $c < (t+1)^2$. Indeed, graphs with treewidth t are t-degenerate and thus properly $(t+1)$-colourable; that is, they are $(t+1)$-colourable with clustering 1. A product colouring then shows that $H_1 \boxtimes H_2$ is properly $(t+1)^2$-colourable.

Summary

Table 3.1 summarises all our results. The notation $\Omega(f(n))$ means that for some $\delta = \delta(t, \Delta, c) > 0$, there are infinitely many graphs H_1, H_2 such that every c-colouring of $H_1 \boxtimes H_2$ has clustering at least $\delta f(|V(H_1 \boxtimes H_2)|)$. The notation $O(f(n))$ means that for

some $\delta = \delta(t, \Delta, c) > 0$ and for any graphs H_1, H_2, there exists a c-colouring of $H_1 \boxtimes H_2$ with clustering at most $\delta f(|V(H_1 \boxtimes H_2)|)$. If $f(n)$ does not depend on n (that is, $c \geqslant 3$ in the first column), the constant δ is absolute (and does not depend on t, Δ, c).

Table 3.1. Clustering in c-colourings of $H_1 \boxtimes H_2$, where $n := |V(H_1 \boxtimes H_2)|$ and $\Delta_i = \Delta(H_i)$, $t_i = \mathrm{tw}(H_i)$ for each $i = 1, 2$.

	$\Delta_1 \leqslant \Delta, t_1 \leqslant t$ $\Delta_2 \leqslant \Delta, t_2 \leqslant t$	$t_1 \leqslant t$ $\Delta_2 \leqslant \Delta, t_2 \leqslant t$	$t_1 \leqslant t$ $t_2 \leqslant t$
$c = 2$	$\Theta(n^{1/2})$	$\Theta(n^{2/3})$	$\Theta(n^{2/3})$
$c = 3$	$O(t^3\Delta^4)$	$\Theta(n^{3/7})$	$\Omega(n^{1/2})$ $O(n^{4/7})$
$c \geqslant 4$	$O(t^2\Delta^2)$	$\Omega(n^{1/(c-2/3)})$ $O(n^{c/(c^2-c+1)})$	$\Omega(n^{1/(2\sqrt{c}))}$ $O(n^{1/\lfloor\sqrt{c}\rfloor})$

For each case in which upper and lower bounds do not match, we conjecture the given upper bound to be asymptotically tight. In particular, we believe the upper bound in Theorem 3.6 is asymptotically tight, and furthermore conjecture that there exists a construction matching this upper bound where the bounded degree graph is a path.

Conjecture 3.9. *For any fixed integer $c \geqslant 2$, there exists an integer t such that for infinitely many graphs H with treewidth t and paths P, every c-colouring of $H \boxtimes P$ has clustering $\Omega(|V(H \boxtimes P)|^{c/(c^2-c+1)})$.*

We also believe that the upper bounds in Theorems 3.8 and 3.5 are asymptotically tight.

Conjecture 3.10. *There exists an integer $t \geqslant 1$ such that for infinitely many graphs H_1, H_2 with treewidth t, every 3-colouring of $H_1 \boxtimes H_2$ has clustering $\Omega(|V(H_1 \boxtimes H_2)|^{4/7})$.*

Conjecture 3.11. *For any fixed integer $c \geqslant 4$ there exists an integer k such that for infinitely many graphs H_1, H_2 with treewidth k, every c-colouring of $H_1 \boxtimes H_2$ has clustering $\Omega(|V(H_1 \boxtimes H_2)|^{1/((1+o_c(1))\sqrt{c})})$.*

In this chapter, we mainly focus on clustering numbers of 2 and 3-colourings, so most proofs of bounds from Table 3.1 in the case $c \geqslant 4$ are omitted. However, we do include a sketch of the upper bound when only Δ_2 is bounded as this provides the asymptotically tight upper bound for the case $c = 3$. We also include the case when both degrees are bounded as it is rather simple.

The rest of this chapter is organised as follows. In Section 3.1 we establish the basic notation and key tools. Our result with two colours, Theorem 3.3, is proved in Section 3.2. The results with three colours, Theorems 3.4 and 3.5, are proved in Section 3.3. We sketch the proof of Theorem 3.6 in Section 3.4. Section 3.5 studies clustered colouring of products of two graphs, both with bounded treewidth and bounded maximum degree.

This setting turns out to be much simpler than the above two scenarios, and the proofs are straightforward. Section 3.6 concludes with a discussion on planar graphs.

3.1 Definitions and tools

Let F_n be the *fan* graph consisting of a path P_n (called the *base*) and one additional dominant vertex adjacent to every vertex along the path. If v is the dominant vertex and $(w_1, \dots, w_n)$ is the base path, then the collection of bags $\{v, w_1, w_2\}, \{v, w_2, w_3\}, \dots, \{v, w_{n-1}, w_n\}$ defines a P_{n-1}-decomposition of F_n with width 2. This demonstrates that $\mathrm{tw}(F_n) \leqslant 2$. The following is another well-known fact about fans.

Lemma 3.12. *Every 2-colouring of F_n has clustering at least $\lfloor \sqrt{n} \rfloor$.*

Proof. Consider any red/blue colouring of F_n with clustering k. Without loss of generality, the dominant vertex of F_n is red. So there are at most $k-1$ red vertices in the base path P_n of F_n. Thus there are at most k blue components in P_n, each with at most k vertices. Hence $n \leqslant (k-1) + k^2 < (k+1)^2$, implying $k > \sqrt{n} - 1$ and $k \geqslant \lfloor \sqrt{n} \rfloor$. □

The *cartesian product* of graphs G and H, denoted by $G \,\square\, H$, is the graph with vertex set $V(G) \times V(H)$, where (u, v) is adjacent to (u', v') in $G \,\square\, H$ if $u = u'$ and $vv' \in E(H)$, or $v = v'$ and $uu' \in E(G)$.

The *strong product* of graphs G and H, denoted by $G \boxtimes H$, is the graph with vertex set $V(G) \times V(H)$, where (u, v) is adjacent to (u', v') in $G \boxtimes H$ if $u = u'$ and $vv' \in E(H)$, or $v = v'$ and $uu' \in E(G)$, or $uu' \in E(G)$ and $vv' \in E(H)$.

Strong products and clustered colouring are inherently related. In particular, it follows from the definitions that a graph G is c-colourable with clustering at most k if and only if G is contained in $H \boxtimes K_k$ for some c-colourable graph H.

For an integer $m \geqslant 1$ and a graph G, define $\widehat{m\,G}$ to be the graph constructed by taking m pairwise disjoint copies of G and adding one additional vertex adjacent to all vertices in each copy. Note that

$$\mathrm{tw}(\widehat{m\,G}) \leqslant \mathrm{tw}(G) + 1. \tag{3.1}$$

To see this, take a tree-decomposition of each copy of G, each with width $\mathrm{tw}(G)$. Add the new dominating vertex to each bag of each decomposition, and add some edges to connect the trees together, forming one large tree. This is a tree-decomposition of $\widehat{m\,G}$ with width $\mathrm{tw}(G) + 1$, hence $\mathrm{tw}(\widehat{m\,G}) \leqslant \mathrm{tw}(G) + 1$.

Several proofs in this article use the following two results of Dvořák and Wood [47]. The first allows us to find separators in graphs with bounded treewidth.

Lemma 3.13 ([47] Lemma 25)**.** *For any integers $n, t \geqslant 1$ and any real number $p > 0$, every graph G on n vertices with treewidth at most t has a set S of at most p vertices such that each component of $G - S$ has at most $\frac{(t+1)n}{p}$ vertices.*

The second lemma is a corollary of the isoperimetric inequality of Bollobás and Leader [25].

Lemma 3.14 ([47] Lemma 13)**.** *If S is any set of at most $\frac{n^2}{2}$ vertices in $G := P_n \square P_n$ and $k \leqslant \frac{n^2}{e^2}$ and each component of $G - S$ has at most k vertices, then $n^2 \leqslant 4|S|k^{1/2}$.*

3.2 Two colours

We start by studying 2-colourings of graph products, proving Theorem 3.3. In this case, we demonstrate that the upper bound with no maximum degree restriction and the lower bound with a bounded degree graph are asymptotically tight up to a multiplicative constant. Theorem 3.3 follows immediately from Lemmas 3.15 and 3.16.

Lemma 3.15. *For any graphs H_1 and H_2 both with treewidth at most t, the graph $G := H_1 \boxtimes H_2$ is 2-colourable with clustering at most $2((t+1)|V(G)|)^{2/3}$.*

Proof. Let $n_i := |V(H_i)|$ for $i \in \{1,2\}$. Let $n := |V(G)| = n_1 n_2$. By Lemma 3.13, for each $i \in \{1,2\}$, there is a set S_i of at most $(t+1)^{2/3} n_i / n^{1/3}$ vertices in H_i such that each component of $H_i - S_i$ has at most $((t+1)n)^{1/3}$ vertices. Let $X := (S_1 \times V(H_2)) \cup (V(H_1) \times S_2)$. So

$$|X| \leqslant |S_1|n_2 + |S_2|n_1 \leqslant (t+1)^{2/3} n_1 n_2 / n^{1/3} + (t+1)^{2/3} n_2 n_1 / n^{1/3} = 2((t+1)n)^{2/3}.$$

Colour each vertex in X blue and colour all other vertices red. Each red component Y is contained in $Y_1 \boxtimes Y_2$ for some component Y_1 of $H_1 - S_1$ and some component Y_2 of $H_2 - S_2$. Thus $|V(Y)| \leqslant |V(Y_1 \boxtimes Y_2)| \leqslant ((t+1)n)^{2/3}$. Hence G is 2-coloured with clustering at most $2((t+1)n)^{2/3}$. □

We now prove that Lemma 3.15 is asymptotically tight, even when H_2 is a path.

Lemma 3.16. *There are infinitely many graphs H with treewidth 2 and paths P such that every 2-colouring of $H \boxtimes P$ has clustering at least $\frac{1}{3}|V(H \boxtimes P)|^{2/3}$.*

Proof. Let n be a sufficiently large integer so that $\frac{1}{2}n^2 \geqslant \frac{1}{3}(n^3+n^2)^{2/3}$, and let $P := P_n$, $H := F_{n^2}$, and $G := H \boxtimes P$ so that $|V(G)| = n^3 + n^2$. With x being the dominating vertex of the fan H, we refer to $\{(x, v) : v \in P\}$, the set of copies of x in G, as the *spine* of G.

Suppose, toward a contradiction, that there exists a 2-colouring of G in red and blue with clustering less than $\frac{1}{2}n^2$. Suppose that two adjacent vertices of the spine are coloured differently. Then there is a red vertex and a blue vertex of G both dominating the same copy of P_{n^2}. Without loss of generality, at least $\frac{1}{2}n^2$ of these vertices are coloured red, and they are all connected into one component by the dominating red vertex, contradicting the fact that the colouring has clustering less than $\frac{1}{2}n^2$.

It follows that the entire spine is monochromatic. Let this colour be red. Since vertices along the spine are dominating in their respective copies of H, this means that all red vertices in the graph G are joined into a single component and therefore there are less than $\frac{1}{2}n^2$ red vertices in G.

Consider the crossed $(n^2 \times n)$-grid, $P_{n^2} \boxtimes P_n$, where the first factor is the base of the fan H. At most $\frac{1}{2}n^2$ columns contain a red vertex. So at least $\frac{1}{2}n^2$ columns are all-blue. Moreover, there is a row R with less than $\frac{1}{2}n$ red vertices. There are at most $\frac{1}{2}n + 1$ blue components in R. At least one of these components intersects at least $(\frac{1}{2}n^2)/(\frac{1}{2}n+1) \geqslant \frac{1}{2}n$ all-blue columns, which are thus contained in a blue component with at least $(\frac{1}{2}n)n = \frac{1}{2}n^2$ vertices. □

Note that the bound $\mathrm{tw}(H) \leqslant 2$ in Lemma 3.16 is best possible. In particular, for any trees T_1 and T_2, by Proposition 3.28 there exists a 2-colouring of $T_1 \boxtimes T_2$ with clustering at most $|V(T_1 \boxtimes T_2)|^{1/2}$.

We present an alternative proof that the upper bound in Lemma 3.15 is asymptotically tight, using a symmetric construction. This construction is perhaps nicer, but the one above is more interesting for the purposes of graph product theory.

Lemma 3.17. *For any integer $n \geqslant 8$, every 2-colouring of $F_n \boxtimes F_n$ has clustering at least $\frac{1}{3}(n+1)^{4/3}$. Namely, setting $m = |V(F_n \boxtimes F_n)|$, such a colouring has clustering at least $\frac{1}{3}m^{2/3}$.*

Proof. Let A and B be copies of F_n, with vertex sets $V(A) = \{a_0, a_1, \ldots, a_n\}$ and $V(B) = \{b_0, b_1, \ldots, b_n\}$, where a_0 dominates the path $a_1, \ldots, a_n$ in A, and b_0 dominates the path $b_1, \ldots, b_n$ in B. The graph $A \boxtimes B$ contains an induced copy of $P_n \boxtimes P_n$ on the vertices (a_i, b_j) with $1 \leqslant i, j \leqslant n$, and consider its subgraph $H \cong P_n \,\square\, P_n$.

Suppose, for the sake of contradiction, that there exists a 2-colouring of $A \boxtimes B$ with clustering less than $k := \frac{1}{3}(n+1)^{4/3}$. Without loss of generality, assume that (a_0, b_0) is red. Note that, since $n \geqslant 8$, we have $k \leqslant \frac{n^2}{2}$ and $k \leqslant \frac{n^2}{e^2}$.

Let S be the set of red vertices in H. Since (a_0, b_0) dominates $A \boxtimes B$, we have $|S| \leqslant k$. Each component of $H - S$ is contained in a monochromatic component of $A \boxtimes B$, and thus has at most k vertices. By Lemma 3.14,

$$n^2 \leqslant 4|S|k^{1/2} \leqslant 4k^{3/2} \leqslant 4(\tfrac{1}{3}(n+1)^{4/3})^{3/2} < n^2,$$

since $n \geqslant 8$, which is the desired contradiction. □

3.3 Three colours

We now move to 3-colourings of graph products, proving Theorems 3.4 and 3.5.

Unrestricted maximum degree

We start by proving Theorem 3.5, where H_1, H_2 both have bounded treewidth. For the upper bound, we build on the 2-colour construction from Lemma 3.15 to prove the following result with three colours. This approach leads to a bound for all c by induction. However, for $c \geqslant 4$, the stronger upper bound in Theorem 3.8 is obtained by colouring each graph separately and taking a product colouring.

Lemma 3.18. *For any graphs H_1 and H_2 both with treewidth at most t, the graph $G := H_1 \boxtimes H_2$ is 3-colourable with clustering at most $2(t+1)^{6/7}|V(G)|^{4/7}$.*

Proof. Let $n_i := |V(H_i)|$ for $i \in \{1,2\}$. Let $n := |V(G)| = n_1 n_2$. By Lemma 3.13, for each $i \in \{1,2\}$, there is a set S_i of at most $(t+1)^{6/7} n_i / n^{3/7}$ vertices in H_i such that each component of $H_i - S_i$ has at most $(t+1)^{1/7} n^{3/7}$ vertices. Let $X := (S_1 \times V(H_2)) \cup (V(H_1) \times S_2)$. Then,

$$|X| \leqslant |S_1|n_2 + |S_2|n_1 \leqslant (t+1)^{6/7} n_1 n_2 / n^{3/7} + (t+1)^{6/7} n_2 n_1 / n^{3/7} = 2(t+1)^{6/7} n^{4/7}.$$

Each component Y of $G - X$ is contained in $Y_1 \boxtimes Y_2$ for some component Y_1 of $H_1 - S_1$ and some component Y_2 of $H_2 - S_2$, implying that $|V(Y_1 \boxtimes Y_2)| \leqslant (t+1)^{2/7} n^{6/7}$. By Lemma 3.15, $Y_1 \boxtimes Y_2$ (and hence also Y) is 2-colourable with clustering at most

$$2((t+1)|V(Y_1 \boxtimes Y_2)|)^{2/3} \leqslant 2((t+1)(t+1)^{2/7} n^{6/7})^{2/3} \leqslant 2(t+1)^{6/7} n^{4/7}.$$

By using a third colour for X, we obtain a 3-colouring of G with clustering at most $2(t+1)^{6/7} n^{4/7}$. □

The next lemma proves the lower bound in Theorem 3.5, and shows that the upper bound in Theorem 3.4 for the bounded degree scenario does not hold in the unrestricted scenario.

Lemma 3.19. *Every 3-colouring of $F_n \boxtimes F_n$ has clustering at least $\left(1 - \frac{1}{\sqrt{2}}\right)n$.*

Proof. Let A and B be copies of F_n with vertex sets $V(A) = \{a_0, a_1, \ldots, a_n\}$ and $V(B) = \{b_0, b_1, \ldots, b_n\}$, where a_0 dominates the path $a_1, \ldots, a_n$ in A, and b_0 dominates the path $b_1, \ldots, b_n$ in B. Suppose there exists a red/blue/green-colouring of $G := A \boxtimes B$ with clustering less than δn, where $\delta := 1 - \frac{1}{\sqrt{2}}$. We may assume without loss of generality that the dominating vertex (a_0, b_0) is red. In particular, there are less than δn red vertices in G.

Let $X = \{(a_0, b_i) : i \in [n]\}$ and $Y = \{(a_i, b_0) : i \in [n]\}$. Note that for any $i, j \in [n]$, the vertices (a_0, b_i) and (a_j, b_0) are adjacent in G. Assume that X contains at least one blue vertex and at least one green vertex. Since each vertex of X dominates Y, it follows that Y contains less than δn blue and less than δn green vertices. This is a contradiction, since there are less than δn red vertices in G and $\delta < 1/3$ while $|Y| = n$. Therefore, without loss of generality, X contains more than $(1 - \delta)n$ blue vertices and no green vertices.

We can apply similar reasoning to Y. Given that X contains at least one blue vertex, Y has less than δn blue vertices, and therefore Y contains more than $(1 - \delta)n$ green vertices and no blue vertices.

Now focus on the subgraph H of G induced by $\{(a_i, b_j) : i, j \in [n]\}$, which forms a copy of $P_n \boxtimes P_n$. Each vertex of X dominates a row of H, and each vertex of Y dominates a column. Since there are less than δn red vertices in G, there exist more than $(1 - \delta)n$ rows with no red vertices, each dominated by a blue vertex from X. Each such row contains less than δn blue vertices, and hence more than $(1 - \delta)n$ green vertices. Therefore, H contains more than $(1 - \delta)^2 n^2$ green vertices. If we consider columns dominated by green vertices, an identical argument yields that there are more than $(1 - \delta)^2 n^2$ blue vertices in H. Since $2(1 - \delta)^2 = 1$, there are more than n^2 vertices in H coloured blue or green, but this is a contradiction since H has only n^2 vertices. □

Note that the clustering value in Lemma 3.19 is best possible (up to the multiplicative constant) for 3-colourings of $F_n \boxtimes F_n$. To see this, colour all the high-degree vertices red. What remains is $P_n \boxtimes P_n$, which can be 2-coloured by alternating rows. So $F_n \boxtimes F_n$ is 3-colourable with clustering $2n$.

Bounded maximum degree

We now focus on Theorem 3.4, when H_2 has bounded treewidth and bounded maximum degree. The upper bound follows directly from the general $c \geqslant 2$ result, Theorem 3.6, presented in Section 3.4. We prove that this upper bound is asymptotically tight (up to the multiplicative factor) when $c = 3$, which concludes the proof of Theorem 3.4.

Recall that F_n is the fan graph with base path P_n. Define $H_n := \widehat{n^2 F_{n^4}}$, and let x be the dominating vertex in H_n (see Figure 3.2). The *fans of* H_n refer to the n^2 pairwise disjoint copies of F_{n^4} in H_n. Note that H_n is planar, and $\mathrm{tw}(H_n) \leqslant \mathrm{tw}(F_{n^4}) + 1 = 3$ by Equation (3.1).

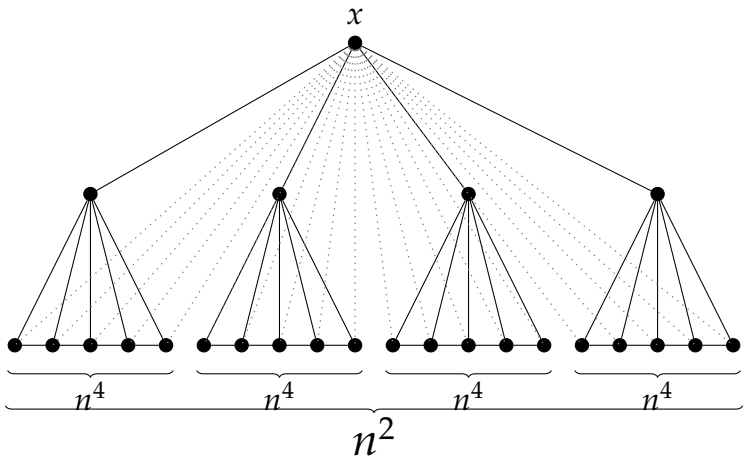


Figure 3.2. $H_n = \widehat{n^2 F_{n^4}}$, where dotted edges are in H_n as well.

Lemma 3.20. *For any $\delta < 1/2$ and any n large enough, every 3-colouring of $G := H_n \boxtimes P_n$ has clustering at least δn^3.*

Since $|V(G)| = n^7 + n^3$ and $\mathrm{tw}(H_n) \leqslant 3$, Lemma 3.20 proves the lower bound of Theorem 3.4.

Proof overview: The proof is by contradiction. We first show in Claim 3.21 that a specific set of vertices of G must be monochromatic, say red. This allows us to bound the total number of red vertices in G. We then restrict our attention to a copy J of $F_{n^4} \boxtimes P_n$ with fewer than δn red vertices. We recognise two subgraphs in J: a *path* induced by the copies of the dominant vertex of F_{n^4} and a *grid* L formed by the strong product of the path of F_{n^4} with P_n. As there are few red vertices, linearly many vertices of the path of J must be of the same non-red colour, say blue. In Claim 3.22 we argue moreover that these blue vertices all lie in the same blue component S of J. We then consider the intersection K of rows with blue associated dominant vertex and columns disjoint from S in L, and note in Claim 3.23 that vertices $u, v \in K$ in the same component

of $L - S$ remain so when deleting all blue vertices. This allows us to conclude the proof as, loosely speaking, the red vertices split K into few components so there must be a large monochromatic component in the last remaining colour.

Proof. Let $H := H_n$. We refer to $\{(x, v) : v \in P_n\}$, the set of copies of x in G, as the *spine* of G. Suppose, toward a contradiction, that the vertices of G are 3-coloured in red, blue and green, with clustering at most δn^3.

Claim 3.21. *The spine of G is monochromatic.*

Proof. Suppose that the spine of G is not monochromatic. Let r, b be two adjacent vertices of P_n such that (x, r) and (x, b) have different colours, say red and blue respectively. In G, every vertex $w \in P_n$ has an associated copy $H^{(w)}$ of H on vertex set $\{(z, w) : z \in V(H)\}$ (with edges between two vertices (y, w) and (z, w) precisely when y and z are adjacent in H). Since r and b are adjacent in P_n, then (x, r) and (x, b) are both dominating vertices for $H^{(r)}$. Therefore, as G contains no monochromatic component of size δn^3, there must be less than δn^3 red and less than δn^3 blue vertices in $H^{(r)}$.

Let $\mathcal{F}^{(r)}$ be the set of fans of $H^{(r)}$. As $|\mathcal{F}^{(r)}| = n^2$, by averaging there is some $F \in \mathcal{F}^{(r)}$ which contains less than $2\delta n < n$ vertices coloured red or blue. Hence, the path of length n^4 in F contains a subpath of green vertices of length at least n^3, a contradiction. □

Without loss of generality, we may assume that all vertices of the spine are coloured red. Since the spine is a dominating set for G, there are less than δn^3 red vertices in G. Let $\mathcal{F}$ be the set of fans in H. By averaging, there is an $F \in \mathcal{F}$ such that the subgraph J of G induced by $V(F) \times V(P_n)$ contains less than δn red vertices.

Consider the restriction of the colouring of G to J. Note that J contains a copy of the grid $P_{n^4} \boxtimes P_n$, formed by the strong product of the path of F with P_n (see the graph $G[J]$ on the left of Figure 3.3). In this grid, rows are copies of the path of F and columns are of the form $\{v\} \boxtimes P_n$, where v is a vertex of the path of F.

Let y be the dominating vertex of the fan F and call a vertex of J of the form (y, v) with $v \in V(P_n)$ a *top vertex*.

As previously, for $v \in V(P_n)$, denote by $F^{(v)}$ the copy of F associated to v in J. Since there are less than δn red vertices in J, we may assume without loss of generality that there are at least $\frac{1-\delta}{2} n$ blue top vertices.

Claim 3.22. *All blue top vertices of J are in the same monochromatic blue component in J.*

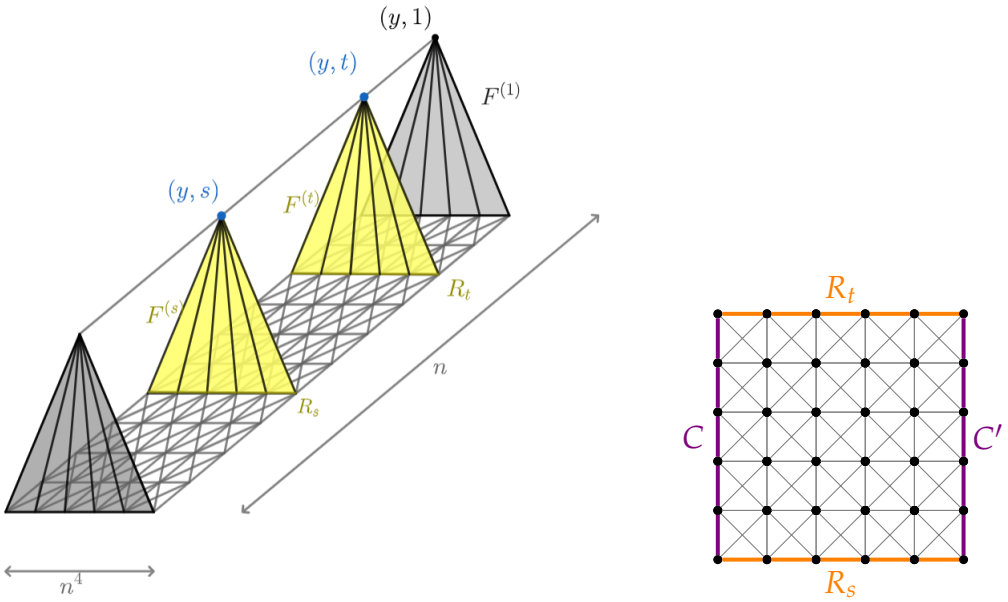


Figure 3.3. Subgraph $G[J]$ and the subgrid of L bounded by R_s and R_t.

Proof. Let (y,s), (y,t) be two blue top vertices in J, for some $s,t \in [n]$. We show that there exists a blue path between (y,s) and (y,t) in J. Let R_s and R_t be the sets of vertices in the paths of the fans $F^{(s)}$ and $F^{(t)}$ respectively, both of size n^4. Note that in J, these two rows bound a subgrid of L. Let C denote the first column of this subgrid, and C' the last.

Remove the top vertices from J and add two dummy vertices u_s, u_t adjacent to all vertices in R_s and R_t (respectively). Colour them blue. Add two dummy vertices v, v' adjacent to all vertices in C and C' (respectively). Colour them red or green . This produces an auxiliary graph as shown in Figure 3.4.

Then by the Hex Lemma (Lemma 3.2) applied to this auxiliary graph, either there is a blue path from u_s to u_t, or there is a path from v to v' with no blue vertices. If the latter occurs, This yields a path of length n^4 with no blue vertices. Since there are less than δn red vertices in J, it contains a subpath of n^3 green vertices and we are done. We may therefore assume that there is a blue path from u_s to u_t. This blue path intersects both R_s and R_t. Since (y,s) (respectively (y,t)) is dominating for R_s (respectively R_t), this extends to a blue path from (y,s) to (y,t) in J, as claimed. □

Let S be the vertices of the monochromatic blue component of J containing all blue top vertices. Note that there may be blue vertices in J which are not in S, and as G contains no monochromatic component of size δn^3, we have that $|S| < \delta n^3$.

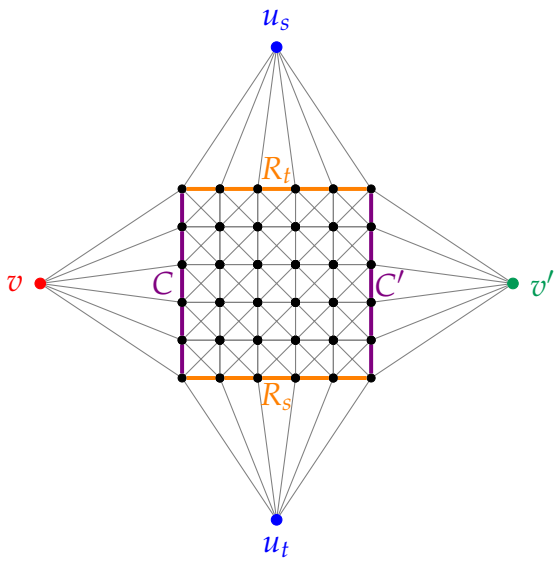


Figure 3.4. Subgrid of L bounded by R_s and R_t with dummy vertices.

Say that a row of L is *blue dominated* if its associated top vertex is blue, and that a column of L is *S-free* if it does not intersect S. In particular, as there are n^4 columns in L and less than δn^3 elements in S, for any $\varepsilon > 0$, for n large enough, there are $(1-\varepsilon)n^4$ columns which are S-free. Define the *boundary* of S in J, denoted by ∂S, to be the set of vertices outside of S that are adjacent to a vertex in S.

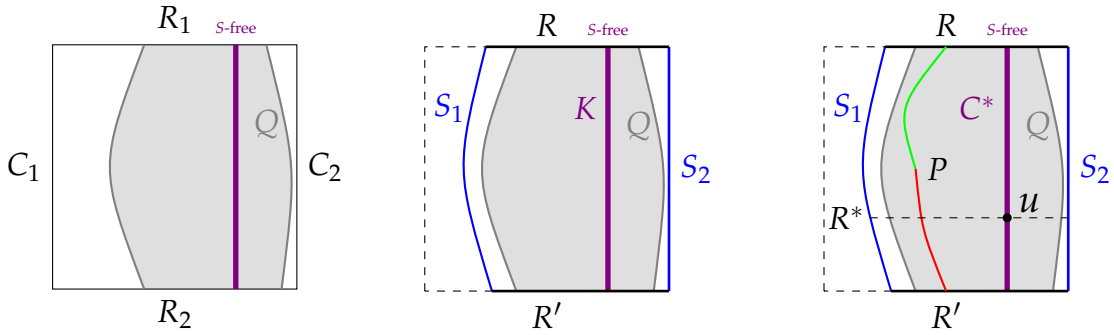
Note that after removing vertices of S from L, there may still be blue vertices remaining.

Claim 3.23. *Let $u, v \in \partial S$ be vertices simultaneously on blue dominated rows and S-free columns. Suppose u, v are in the same component Q of the grid L after the vertices of S have been removed, and that Q contains an S-free column. Then, there is a path consisting of red and green vertices between u and v in L. In other words, u, v remain in the same component of L after removing all blue vertices.*

Proof. Let Q be a component of $L - S$ that contains an S-free column. Let R_1, R_2, C_1, C_2 be the top row, bottom row, leftmost column, and rightmost column of L, respectively, see Figure 3.5a.

Let D be the smallest subgraph of L containing Q, bounded up and down by R_1 and R_2, left by a blue path S_1 from R_1 to R_2 (or $S_1 = C_1$ if such a path does not exist), and right by a blue path S_2 from R_1 to R_2 (or $S_2 = C_2$ if such a path does not exist).

We now restrict our attention to D. Let R and R' denote the vertices of its top and bottom rows, and fix a full S-free column K of Q, see Figure 3.5b. By the Hex Lemma



(a) Q in $L - S$ **(b)** Domain D **(c)** Vertex u

Figure 3.5. Proof of Claim 3.23

(Lemma 3.2), either there is a blue path from S_1 to S_2 or a red/green path from R to R'.

Suppose there is a blue path B from S_1 to S_2. If $S_1 = C_1$ and $S_2 = C_2$, then this path has length at least n^4, and is therefore a monochromatic component of size at least δn^3, a contradiction. Otherwise, at least one of S_1, S_2 is a blue monochromatic path crossing every row in D. Without loss of generality, suppose S_1 is such a path. As some rows of L are blue dominated, this path is entirely in S. Hence, as B is blue and touches S_1, $B \subset S$ as well. But B certainly intersects K, an S-free column, which is impossible.

Hence, there must be a path P from R to R' whose vertices are not blue.

It now suffices to show that for each $u \in Q$ which is the intersection of an S-free column with a blue dominated row, there is a path from u to P whose vertices are not blue. Indeed, any $u, v \in Q$ are then connected with non-blue vertices by concatenating the paths from u to P, from v to P, and P itself.

If $u \in P$ we are done, so we may assume that $u \notin P$. Let C^* and R^* be the column and row of u, respectively, see Figure 3.5c. We let

- P_1 be the shortest path from u to P in D whose vertices consist only of vertices from C^* *above* u and possibly vertices of R, and $x \in P$ the other endpoint of P_1, and
- P_2 be the shortest path from u to P in D whose vertices consist only of vertices from C^* *below* u and possibly vertices of R', and $y \in P$ the other endpoint of P_2,

as in Figure 3.6.

Let A be the subgraph of D induced by the vertices in the region enclosed by P_1, P_2 and the subpath of P between x and y. By the Hex Lemma (Lemma 3.2), either there is a blue path from $P_1 - u$ to $P_2 - u$, or there is a red/green path from u to P.

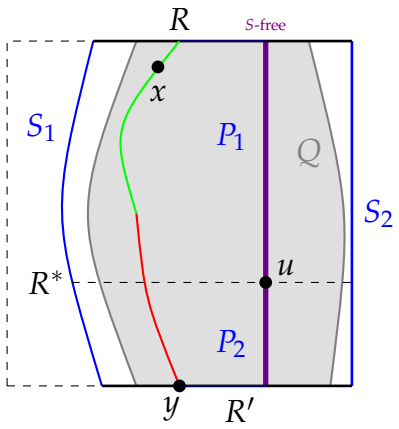


Figure 3.6. Paths P_1 and P_2

Suppose there is a blue path P' from $P_1 - u$ to $P_2 - u$. Note that, as all vertices of P_1 lie above or on row R^* and all vertices of P_2 lie below or on row R^*, any path from P_1 to P_2 must cross row R^*. In particular, P' must be entirely in S as it is blue. Consider the endpoints of P'. If they lie in R and R', then, as P' is blue and in S, both of its endpoints lie to one side of the S-free column $K \subseteq Q$. Thus, P' splits D into two non-trivial parts with one containing Q, contradicting the minimality of D. Hence, by construction, some endpoint of P' must lie on C^*. But C^* is an S-free column, so this is also a contradiction. Therefore, there is a red/green path from u to P, as desired. □

Let $K \subset \partial S$ be the set of vertices of L simultaneously on blue dominated rows and S-free columns. Recall that L contains at least $(1-\varepsilon)n^4$ S-free columns and at least $\frac{1-\delta}{2}n$ blue top vertices, therefore $|K| \geqslant \frac{(1-\varepsilon)(1-\delta)}{2}n^5$.

Since $|S| < \delta n^3$, at least one of the n rows of the grid L contains less than δn^2 vertices from S. It follows that the S-free columns of L live in at most δn^2 components of $L - S$.

As K contains only vertices from S-free columns, the vertices of K live in at most δn^2 components of $L - S$. By Claim 3.23 the vertices of K still live in at most δn^2 components of L after all blue vertices are removed. We now remove the less than δn red vertices. Note that for any $\varepsilon' > 0$, for n large enough, we have at least $(1-\varepsilon')|K|$ green vertices in K. As each vertex of L has degree at most 8, at each removal we add at most 7 components. After removing all red vertices, for any $\varepsilon'' > 0$, for n large enough, we still have less than $(\delta + \varepsilon'')n^2$ components. Therefore, using $\delta < 1/2$ and choosing ε, ε', and ε'' small enough, there is a component of size at least

$$\frac{(1-\varepsilon)(1-\varepsilon')(1-\delta)}{2(\delta+\varepsilon'')}n^3 > \delta n^3$$

containing only green vertices, a contradiction. □

3.4 Arbitrary number of colours

Upper bound for bounded maximum degree

In this section, we sketch the proof of Theorem 3.6. The overall strategy is to view $H_1 \boxtimes H_2$ as a subgraph of $H_1 \boxtimes T \boxtimes K_\ell$ and use known results about clustering numbers of products. We use the following lemmas.

Lemma 3.24 ([86, Theorem 1.2]). *Every graph on n vertices with treewidth at most t has a c-colouring with clustering number $O(n^{1/c})$.*

This lemma is a simple application of Lemma 3.13 regarding sublinear separators. Note that it is only interesting if $c \leqslant t$, since if $c \geqslant t+1$, then G is $(c-1)$-degenerate and thus properly c-colourable.

Lemma 3.25 ([53, Lemma 14]). *If a graph G is c-colourable with clustering k, then $G \boxtimes K_\ell$ is c-colourable with clustering $k\ell$.*

This can be seen by taking a c-colouring of G with clustering number k and colouring every vertex (x, y) of $G \boxtimes K_\ell$ with the colour it assigns to x.

Lemma 3.26 ([53, Lemma 32]). *If a graph G is c-colourable with clustering k and T is a tree with maximum degree at most $\Delta \geqslant 3$, then $G \boxtimes T$ is $(c+1)$-colourable with clustering less than $2k(\Delta-1)^{c-1}$.*

A sufficient bound for our purposes can easily be seen as follows. Choose a root vertex y_0 in T and fix a c-colouring of G with clustering k using colours from $[c]$. For each (x, y) in $G \boxtimes T$, if x has colour $i \in [c]$, we assign (x, y) colour $c+1$ if $d_T(y_0, y) = 2i \bmod 2c$, and i otherwise. Then, two vertices (x, y), (x', y') in a same monochromatic component have $d_T(y, y') \leqslant 4c$, so a bound of the form $k\Delta^{4c}$ follows as non-adjacent vertices in G remain non-adjacent across copies of G in $G \boxtimes T$.

Lemma 3.27 (Ding, Oporowski [38]). *Every graph H with treewidth at most t and maximum degree at most Δ is contained in $T \boxtimes K_{24t\Delta}$ for some tree T with maximum degree $O(\Delta)$.*

Proof sketch of Theorem 3.6. Let H_1, H_2 be graphs of treewidth at most t with $\Delta(H_2) \leqslant \Delta$ and write $a = |V(H_1)|$, $n = |V(H_1 \boxtimes H_2)|$. By Lemma 3.27, there is a tree T of maximum degree $O(\Delta)$ such that $H_1 \boxtimes H_2 \subseteq H_1 \boxtimes T \boxtimes K_{24t\Delta}$. By Lemma 3.25, it suffices to show that $H_1 \boxtimes T$ has c-clustering number $O(n^{c/(c^2-c+1)})$. We proceed according to the relative sizes of a and $b := |V(T)|$.

If $a \geqslant b^{c(c-1)}$, fix a c-colouring of H_1 with clustering $O(a^{1/c})$ and colour each vertex (x, y) in $H_1 \boxtimes T$ by the colour of x. This c-colouring has clustering number $O(a^{1/c} b)$.

If $a \leqslant b^{c(c-1)}$, apply Lemma 3.26 to find a c-colouring of $H_1 \boxtimes T$ with clustering $O(a^{c-1})$. Since $n = O(ab)$, one checks that both bounds are $O(n^{c/(c^2-c+1)})$.[2] □

3.5 Products of bounded degree graphs

This section considers clustered colouring of the strong product of two graphs, both with bounded treewidth and bounded degree. This setting is much simpler than with none or one of the graphs having bounded degree. We only require the following proposition to solve this case.

Proposition 3.28. *For any integer $k \geqslant 1$ and any graphs H_1, H_2 that are c-colourable with clustering k, the graph $H_1 \boxtimes H_2$ is c-colourable with clustering $k\,|V(H_1 \boxtimes H_2)|^{1/2}$.*

Proof. Let $G := H_1 \boxtimes H_2$ and $n := |V(G)|$. Without loss of generality, $|V(H_2)| \leqslant |V(H_1)|$, so $|V(H_2)| \leqslant n^{1/2}$. Fix a c-colouring of H_1 with clustering k. Colour each vertex (x, y) of G by the colour assigned to x. Now G is c-coloured. If $(x_1, y_1)(x_2, y_2)$ is a monochromatic edge in G, then x_1 and x_2 had the same colour in H_1, and $x_1 = x_2$ or $x_1 x_2 \in E(H_1)$. Hence, each monochromatic component M in G projects to a monochromatic component in H_1, implying that $|V(M)| \leqslant k\,|V(H_2)| \leqslant k\,n^{1/2}$, as claimed. □

Now consider graphs H_1, H_2 with treewidth at most t and maximum degree at most Δ. By Lemma 3.27, each H_i is contained in $T_i \boxtimes K_{O(t\Delta)}$ for some tree T_i. As noted by Alon, Ding, Oporowski and Vertigan [5, Theorem 2.2], any proper 2-colouring of T_i determines a 2-colouring of H_i with clustering $O(t\Delta)$. Therefore:

- For $c = 2$, Proposition 3.28 implies that $H_1 \boxtimes H_2$ is 2-colourable with clustering $O(t\Delta\,|V(H_1 \boxtimes H_2)|^{1/2})$. This bound is tight for fixed t and Δ, since the Hex Lemma (Lemma 3.2) implies that any 2-colouring of $P_n \boxtimes P_n$ has clustering at least $n = |V(P_n \boxtimes P_n)|^{1/2}$.
- For $c = 3$, Esperet and Wood [53, Theorems 4 and 33] proved that $H_1 \boxtimes H_2$ is 3-colourable with bounded clustering, at most $O(t^3\Delta^4)$.
- For $c = 4$, using the product of the 2-colourings of each H_i, it is easy to see that $H_1 \boxtimes H_2$ is 4-colourable with bounded clustering, at most $O(t^2\Delta^2)$. This implies bounded clustering for $c > 4$.

[2] Eg. if $n = ab$, when $a = b^{c(c-1)}$ we have $b^c = n^{c/(c^2-c+1)}$. Observe that $\frac{1}{c} \leqslant \frac{c}{c^2-c+1} \leqslant \frac{1}{c-1}$ for $c \geqslant 2$. So for a fixed b, $a^{1/c} \leqslant n^{1/c}$ grows slower than $n^{c/(c^2-c+1)}$, and $a^{1/(c-1)} = O((n/b)^{1/c})$ decreases faster than $n^{c/(c^2-c+1)}$.

3.6 Planar graphs

Now consider clustered colouring of planar graphs. Linial, Matoušek, Sheffet and Tardos [86] introduced the following definition. Let $f_c(n)$ be the minimum integer such that every n-vertex planar graph is c-colourable with clustering $f_c(n)$. The 4-Colour Theorem says that $f_c(n) = 1$ for every $c \geqslant 4$. Linial, Matoušek and Tardos [86] showed that $f_2(n) = \Theta(n^{2/3})$ and $\Omega(n^{1/3}) \leqslant f_3(n) \leqslant O(n^{1/2})$.

One approach for closing the gap in these bounds on $f_3(n)$ is to apply the Planar Graph Product Structure Theorem of Dujmović, Joret, Micek, Morin, Ueckerdt and Wood [44], which says that every planar graph is contained in $H \boxtimes P$ for some graph H of treewidth at most 8 and some path P. The idea is to 3-colour $H \boxtimes P$ with small clustering, thus determining a 3-colouring of G. Two issues arise, however. First, it may be that $|V(H \boxtimes P)|$ is significantly larger than $|V(G)|$ (the best bounds are $|V(H)| \leqslant |V(G)|$ and $|V(P)| \leqslant O(|V(G)|^{(1+\epsilon)/2})$ by Hickingbotham and Wood [68]). So clustering $O(|V(H \boxtimes P)|^{\beta})$ does not immediately imply clustering $O(|V(G)|^{\beta})$. Second, by Lemma 3.20, there are graphs H with treewidth 3 and there are paths P such that every 3-colouring of $H \boxtimes P$ has clustering $\Omega(|V(H \boxtimes P)|^{3/7})$. So the best upper bound on $f_3(n)$ that one could hope for using this method is $f_3(n) \leqslant O(n^{3/7})$, which is between the known bounds mentioned above. Determining $f_3(n)$ is a tantalising open problem.

Chapter 4

Improved bounds on the interval thickness of graphs

This chapter is based on joint work with Axenovich, Girão, Powierski and Savery [15].

We say that a graph G is *interval colourable* if it has an *interval colouring* c, that is, a proper edge-colouring $c\colon E(G) \to \mathbb{N}$ such that the set of colours incident to each vertex $v \in V(G)$, $\{c(vw)\colon w \in N(v)\}$, consists of consecutive integers. Since these notions were introduced by Asratian and Kamalian [11] in 1987, interval colourable graphs and their properties have been studied extensively (see for example, [10, 11, 12, 28, 32, 56, 61, 62, 100, 122, 16, 111]). Examples of interval colourable graphs include trees and regular class 1 graphs (and hence, in particular, regular bipartite graphs), while non-interval colourable graphs include odd cycles and complete graphs with an odd number of vertices (see [111] for a further, bipartite, example).

In the present chapter, we study a parameter introduced recently by Asratian, Casselgren, and Petrosyan [10] to quantify how far a graph is from being interval colourable. The *interval colouring thickness* of a graph G, denoted $\theta(G)$, is the minimum number k such that G can be edge-decomposed into k interval colourable subgraphs. We write $\theta(n)$ for the maximum interval colouring thickness of a graph on n vertices, and $\theta'(m)$ for the maximum interval colouring thickness of a graph with m edges. Upper bounds on $\theta(G)$ in terms of the edge chromatic number of G obtained by Asratian, Casselgren, and Petrosyan [10] imply that $\theta(n) \leqslant 2\lceil n/5 \rceil$, and the same authors observe that an arboricity result of Dean, Hutchinson, and Scheinerman [36] gives $\theta'(m) \leqslant \lceil \sqrt{m/2} \rceil$. Axenovich and Zheng [16] improved on the first of these bounds, showing that $\theta(n)$ is sublinear[1]. We show that a result of Rödl and Wysocka [107] implies the following

[1]While Axenovich and Zheng do not state explicit bounds in their paper, their arguments rely on Szemeredi's regularity lemma. If $f(\varepsilon, m)$ is the bounding function for the maximum number of parts in a

polynomial improvement.

Theorem 4.1. *We have* $\theta(n) \leqslant n^{5/6+o(1)}$ *and* $\theta'(m) \leqslant m^{5/11+o(1)}$.

As noted above, various graphs with $\theta(G) \geqslant 2$ are known, but to the best of our knowledge none with $\theta(G) \geqslant 3$ were known prior to [15]. We show that $\theta(n) = \omega(1)$.

Theorem 4.2. *We have* $\theta(n) \geqslant (\log n)^{1/3-o(1)}$, *and thus* $\theta'(m) \geqslant (\log m)^{1/3-o(1)}$.

The first bound in Theorem 4.2 is proved via a random construction, which will, in fact, produce a bipartite graph. The bound on $\theta'(m)$ follows trivially (see Section 4.1). One corollary of this theorem is that for every integer k there exists a graph with interval colouring thickness k, answering a question of Asratian, Casselgren, and Petrosyan [10]. Indeed, given a graph G of some interval colouring thickness K, and an edge-decomposition of G into K interval colourable subgraphs, the union of any $k \leqslant K$ of these subgraphs has interval colouring thickness k.

Additional results

Shortly after the material from this chapter appeared on arXiv as a preprint [15], an independent preprint by Hollom, Portier and Versteegen [69] appeared with similar results. As overlap between results claimed was significant, we joined the two works for the final journal version of the paper [14].

In this final version we improve Theorem 4.2 to a lower bound of $\theta(n) = \Omega(\frac{\log n}{\log\log n})$ and prove the following result on interval colourings of planar graphs.

Theorem 4.3. *Let* G *be a planar graph on* $n \geqslant 2$ *vertices that admits an interval colouring. Then an interval colouring of* G *uses at most* $(3/2)n - 2$ *colours.*

Theorem 4.3 was also proven independently by Hambardzumyan and Muradyan [60].

The improvement on Theorem 4.2 is obtained via a slight tweak to the key Lemma 4.5 and a sharpening of the pigeonhole argument in Theorem 4.2. We give a short overview of how this can be done at the end of Section 4.1.

4.1 Proof of the lower bounds

In order to lower bound the interval colouring thickness of the graph G that we construct below, we will need to show that every edge-decomposition of G into sufficiently

ε-regular partition with at least m parts guaranteed by the regularity lemma, then their arguments seem to produce a bound of the form $f(\varepsilon, m) + \sqrt{100\varepsilon + 1/(2m)} \cdot n$. So they improve on the linear bound by a function which grows slower than an iterated logarithm.

few parts has a part which is not interval colourable. Our approach to disproving interval colourability is based on the following observation, noted by Sevastianov [111].

Observation 4.4. Let G be an interval colourable graph and let $U \subseteq V(G)$. Suppose that there exists $d \in \mathbb{N}$ such that for all distinct $v, w \in U$ there is a path P in G from v to w such that $\sum_{x \in V(P)} d(x) \leqslant d$. Then for all $u \in V(G)$, we have $|N(u) \cap U| \leqslant d$.

Proof. If $|U| \leqslant 1$ then the lemma holds trivially, so assume $|U| \geqslant 2$. Let c be an interval colouring of G, let $u \in V(G)$, and let $v, w \in N(u) \cap U$ be distinct. Fix a path $P = x_1 \dots x_k$ in G from $v = x_1$ to $w = x_k$ such that $\sum_{i=1}^{k} d(x_i) \leqslant d$. The colours under c of any two edges incident to a vertex x of G differ by at most $d(x) - 1$, so it follows from the existence of P that any edge incident to v and any edge incident to w have colours differing by at most $\sum_{i=1}^{k} (d(x_i) - 1) \leqslant d - 1$. Observe that we can choose v and w such that $c(uv) - c(uw) \geqslant |N(u) \cap U| - 1$ which gives $|N(u) \cap U| \leqslant d$. □

We shall construct a bipartite graph of large interval colouring thickness by patching together many bipartite graphs of the form given by Lemma 4.5.

Lemma 4.5. *Fix $\alpha, \delta \in (0, 1/3)$ such that $\delta \geqslant 2\alpha$, and let $a \geqslant 8/\alpha\delta^2$ be an integer. Then for all integers $n \geqslant 1000(\log(a) + 1)/\alpha^2$ there is a bipartite graph G on parts A and B of sizes a and n respectively satisfying the following.*

(a) *For all $x \in A$, $d(x) = \lfloor \alpha n \rfloor$.*

(b) *Let H be a subgraph of G with at least $a\delta\alpha n$ edges. Then there exist $A' \subseteq A$ and $B' \subseteq B$ with $|A'| \leqslant \delta/\alpha$ and $|B'| \geqslant \delta^2 n/32$ such that the induced subgraph $H[A' \cup B']$ has diameter at most 6.*

Proof. Let A and B be vertex sets of the desired sizes and consider the random bipartite graph between them in which every edge appears independently with probability $4\alpha/3$. Drawing on $n \geqslant 1000(\log(a) + 1)/\alpha^2$, a standard application of a Chernoff bound (Lemma 1.3) and a union bound over all vertices $x \in A$ yields that with failure probability at most $1/e$ we have $d(x) \geqslant \alpha n$ for all $x \in A$. Similarly, we can show that, with failure probability at most $1/e$, we have $|N(x) \cap N(y)| \leqslant 2\alpha^2 n$ for all distinct $x, y \in A$, a property which we will call $(\star)$. Thus, there is a positive probability that in this random graph we have both $d(x) \geqslant \alpha n$ for all $x \in A$, and property $(\star)$. Given an outcome with these properties, we can delete edges where necessary to obtain a graph G satisfying (a) and $(\star)$. We will now show that G also satisfies (b).

Fix a subgraph H of G with at least $a\delta\alpha n$ edges. Let $\{s_1, \dots, s_k\} \subseteq A$ be a maximal set with the property that there is a collection of pairwise disjoint sets $\{S_1, \dots, S_k\}$ such that $S_i \subseteq N_H(s_i)$ and $|S_i| \geqslant \delta\alpha n/4$ for each $i \in [k]$. Let $S := \bigcup_i S_i$. Note that $k \leqslant K := 4/\alpha\delta$ (indeed, otherwise $n = |B| \geqslant \sum_i |S_i| > K\delta\alpha n/4 = n$). By (a) and the fact that $a \geqslant 8/\alpha\delta^2$, there exists $z \in A \setminus \{s_1, \dots, s_k\}$ with

$$d_H(z) \geqslant \frac{e(H) - k\lfloor \alpha n \rfloor}{a} \geqslant \delta\alpha n - \frac{4n}{a\delta} \geqslant \frac{\delta\alpha n}{2}.$$

Now, if $|N_H(z) \cap (B \setminus S)| \geqslant \delta\alpha n/4$, then we could add z to S, giving a contradiction. Hence, $|N_H(z) \cap S| \geqslant \delta\alpha n/4$. By property $(\star)$, we have $|N_H(z) \cap S_i| \leqslant 2\alpha^2 n$ for all $i \in [k]$, so $N_H(z)$ intersects at least $\frac{\delta\alpha n/4}{2\alpha^2 n} = \frac{\delta}{8\alpha}$ of the S_i. Denote by B' the union of the first $\left\lceil \frac{\delta}{8\alpha} \right\rceil$ of the S_i which intersect $N_H(z)$ and let A' be the set containing z and the corresponding s_i. By construction $H[A' \cup B']$ has diameter at most 6, and since $\delta \geqslant 2\alpha$ we have $|A'| = \lceil \delta/8\alpha \rceil + 1 \leqslant \delta/\alpha$ and $|B'| \geqslant \lceil \delta/8\alpha \rceil \cdot \delta\alpha n/4 \geqslant \delta^2 n/32$, as required. □

We now give an overview of the proof of Theorem 4.2. We shall construct a graph G that is a union of $t' \approx \log(n)$ bipartite graphs G_i with bipartitions (A_i, B) such that $|A_i| = \sqrt{n}$ and $|B| = n$, and where the degrees of vertices in A_{i-1} are twice those of vertices in A_i. We apply Lemma 4.5 to obtain such G_i that also satisfy (b) with $\delta = 1/\ell$, where $\ell \approx (\log n)^{1/3}$. Suppose that G can be edge-partitioned into ℓ interval colourable graphs, and pick one of these graphs H such that there are at least $t = t'/\ell \approx (\log n)^{2/3}$ many i where $H \cap G_i$ has at least a δ proportion of the edges of G_i. Henceforth, we restrict to indices i that have this property.

By the lemma, for each i there is a subgraph G_i' of $H \cap G_i$ on $A_i' \cup B_i$ satisfying the properties in (b), i.e. G_i' has small diameter, and $|A_i'|$ is small while $|B_i| \gtrsim n/\ell^2$. For indices $j \ll k$, since the diameter of G_k' is small, we can find a short, low degree path in H between any two vertices in B_k. Thus Observation 4.4 yields that vertices in A_j' (which have much larger degrees) can't send a large proportion of their edges into B_k. Since A_j' is small, it follows that $B_j \cap B_k$ is small. Hence, by slightly thinning our index set, we have that $B_j \cap B_k$ is small for all $j \neq k$. But each B_i individually has size at least n/ℓ^2 so that the union of the t many B_i is larger than $n = |B|$, which is a contradiction.

Proof of Theorem 4.2. Fix $\varepsilon > 0$ and let n be large in terms of ε. Since n is large, in what follows we will not concern ourselves with whether or not expressions are integers. Let $\ell = (\log n)^{1/3-\varepsilon}$. To prove the theorem we will construct a graph G on at most $2n$ vertices which we will show has interval colouring thickness greater than ℓ.

Let $t = 64\ell^2$, $d = \log(7t^2/\ell) = \log(7 \cdot 2^{12}\ell^3)$, and $t' = d\ell t = 64d\ell^3$, and note that

$$t' = 64 \cdot (\log n)^{1-3\varepsilon} \cdot \log(7 \cdot 2^{12} \cdot (\log n)^{1-3\varepsilon}) \leqslant \frac{\log n}{3},$$

since n is large. Let B be a set of n vertices. For each $i \in \left\{\frac{\log n}{9}, \dots, t' + \frac{\log n}{9}\right\}$, let A_i be a set of $\sqrt{n}$ vertices (disjoint from vertices in other A_js and B) and let G_i be a bipartite graph with bipartition (A_i, B) of the form given by Lemma 4.5, where α is taken to be $\alpha_i = 2^{-i}$ and $\delta = 1/\ell$. To see that we may apply the lemma in this setting, note that $n^{1/9} \leqslant \alpha_i^{-1} \leqslant n^{4/9}$ and n is large, so $\alpha_i, \delta \in (0, 1/3)$ with $\delta \geqslant 2\alpha$,

$$8/\alpha_i\delta^2 \leqslant 8 \cdot n^{4/9} \cdot (\log n)^{2/3-2\varepsilon} \leqslant \sqrt{n} = a,$$

and

$$\frac{1000(\log(a)+1)}{\alpha_i^2} \leqslant 1000 \cdot n^{8/9} \cdot \left(\frac{\log n}{2} + 1\right) \leqslant n.$$

Let G be the union of the G_i, so that G is a bipartite graph with bipartition $(\bigcup_i A_i, B)$. Note that $|V(G)| \leqslant n + \log(n)\sqrt{n}/3 \leqslant 2n$, and suppose for a contradiction that G has an edge-decomposition into ℓ interval colourable subgraphs. For each subgraph H in such a decomposition, let I_H be the set of those i for which $|E(H) \cap E(G_i)| \geqslant e(G_i)/\ell$. By two applications of the pigeonhole principle, we first find for each i a part H such that $i \in I_H$, and then a part H with $|I_H| \geqslant t'/\ell$. Finally, we may obtain a sequence $i_1, \dots, i_t \in I_H$ such that $i_{j+1} \geqslant i_j + d$ for all $j \in [t-1]$. Now, by (b) from Lemma 4.5, for each $j \in [t]$ there are $A'_{i_j} \subseteq A_{i_j}$ and $B_{i_j} \subseteq B$ with $|A'_{i_j}| \leqslant 2^{i_j}/\ell$, $|B_{i_j}| \geqslant n/(32\ell^2)$ such that $G'_{i_j} = (G_{i_j} \cap H)[A'_{i_j} \cup B_{i_j}]$ has diameter at most 6.

We next argue that for any fixed $j, k \in [t]$ with $j < k$ the intersection $B_{i_j} \cap B_{i_k}$ is small. The degrees in G of vertices in A'_{i_k} are $2^{-i_k}n$ and those of vertices in $B_{i_k} \subseteq B$ are at most $t'\sqrt{n} \leqslant \log(n)\sqrt{n} \leqslant 2^{-i_k}n$ since $i_k \leqslant 4\log(n)/9$. Combining this with the fact that G'_{i_k} has diameter at most 6 gives that for any two vertices $v, w \in B_{i_k}$ there is a path P in $G'_{i_k} \subseteq H$ from v to w with $\sum_{x \in V(P)} d_G(x) \leqslant 7 \cdot 2^{-i_k}n$. Thus Observation 4.4 yields that for all $u \in A_{i_j}$ we have $|N_H(u) \cap B_{i_k}| \leqslant 7 \cdot 2^{-i_k}n$.

Since $B_{i_j} \subseteq \bigcup_{u \in A'_{i_j}} N_H(u)$, this gives

$$|B_{i_j} \cap B_{i_k}| \leqslant \sum_{u \in A'_{i_j}} |N_H(u) \cap B_{i_k}| \leqslant \frac{2^{i_j}}{\ell} \cdot 7 \cdot 2^{-i_k}n.$$

Thus

$$\Big|\bigcup_{j \in [t]} B_{i_j}\Big| \geqslant \sum_j |B_{i_j}| - \sum_{j<k} |B_{i_j} \cap B_{i_k}| \geqslant \frac{tn}{32\ell^2} - \sum_{j<k} \frac{7 \cdot 2^{i_j - i_k}n}{\ell} \geqslant \frac{tn}{32\ell^2} - \binom{t}{2}\frac{7n}{\ell \cdot 2^d}.$$

Since $t \geqslant 64\ell^2$ and $d > \log(7\binom{t}{2}/\ell)$, this yields

$$n \geqslant |\bigcup_{j\in[t]} B_{i_j}| > 2n - n = n,$$

giving the required contradiction. This completes the proof of the bound on $\theta(n)$. It is straightforward to deduce the bound on $\theta'(m)$ by noting that any n-vertex graph contains at most n^2 edges and θ' is an increasing function, so $\theta'(n^2) \geqslant \theta(n)$ for all n. □

The lower bound revisited

In fact, the following improvement on Theorem 4.2 holds essentially by making small modifications to the analysis of the construction for that theorem.

Theorem 4.6 ([14]). *We have* $\theta(n) = \Omega(\frac{\log n}{\log\log n})$.

Remark 4.7. It seems that this is the best possible bound that can be obtained by this construction. Indeed, the final step of the proof of Theorem 4.2 is to bound intersections between suitably chosen B_is. We do so by showing that each $|B_i \cap B_j|$ is bounded by an expression of the form $2^{i-j}n$ (for $i < j$). Since we take a sum over at most $\log n$ such pairs, we require that 2^{i-j} counteract the number $(\log n)^2$ of summands. This is achieved when $j - i > 2\log\log n$, so we require a subsequence of subgraphs (A_i, B_i) whose indices differ by at least $2\log\log n$, where each (A_i, B_i) is dense enough after edge-partitioning to apply Lemma 4.5.

If we think of the initial total 'density' as being about $\log n$ (density 1 for each (A_i, B_i)), partitioning edges into ℓ parts decreases this 'density' by $1/\ell$. After further reducing to a subsequence with gaps at least $2\log\log n$ the remaining 'density' is $\delta = \frac{1}{\ell}\frac{\log n}{2\log\log n}$.

At this point, the best possible case for us is that all B_is in this subsequence are disjoint. Assume this and write δ_i for the density of (A_i, B_i), so that $\sum \delta_i = \delta$. Assume moreover that, improving on (b) from Lemma 4.5, we can show $|B_i| \geqslant c\delta_i n$ for some constant c. Then to derive a contradiction we still need that $\sum |B_i| = c\delta n > n$, namely that $\delta > 1/c$. In particular, this implies $\ell = O(\log n/\log\log n)$, so an edge-decomposition cannot have more parts.

Overview of Theorem 4.6. The improved bound can seen by making two small modifications to the arguments in Section 4.1.

First, in (b) of Lemma 4.5 we can guarantee that the subgraph (A', B') of H that we find has $|A'| \leqslant 1/\alpha$ and $|B'| \geqslant c\delta n$ for some constant c. To do so, it suffices to replace property $(\star)$ in the proof by the following. For every $U \subseteq B$, all but a negligible fraction

(about $n/\alpha|U|$) of vertices $y \in A$ satisfy $|N(y) \cap U| \leqslant 2\alpha|U|$. A simple application of Chernoff bounds shows that this holds for all U with constant probability. This improves on property $(\star)$ by a factor of α, and therefore using this we are able to shave off a factor of δ in the final bound on $|B'|$.

Next, we tighten the pigeonhole argument in the proof of Theorem 4.2. This has essentially already been described in the above remark, up to a technical detail.

After partitioning into $\ell = \log n / \log\log n$ parts, we restrict our attention to a part with remaining 'edge density' of order at least $\log\log n$. Here, we think of 'density' in the same way as in the above remark, where the whole graph starts with $\log n$ 'density'. Let δ_i be the density of (A_i, B_i) in this part. We can restrict to a subsequence (A_i, B_i) with gaps of length at least $2\log\log n$, such that $\sum \delta_i > k$ for some large constant k.

Now our hope is to be able to apply the strengthened version of Lemma 4.5, but unfortunately some subgraphs (A_i, B_i) may have extremely low density δ_i. Thankfully these contribute very little to the total sum of densities: there are at most $\log n$ subgraphs (A_i, B_i) so removing each subgraph with $\delta_i < 1/\log n$ only decreases $\sum \delta_i$ by one.

With additional technical details, these two observations suffice to prove the stronger Theorem 4.6. □

4.2 Proof of the upper bounds

We prove Theorem 4.1 by decomposing into forests and regular bipartite subgraphs. Such graphs are interval colourable (regular bipartite graphs are edge-decomposable into perfect matchings which can be taken as colour classes). To find large regular subgraphs of relatively dense graphs we rely on the following result of Rödl and Wysocka [107].

Theorem 4.8 ([107]). *Let $\gamma\colon \mathbb{N} \to [0, 1/2)$ satisfy $\gamma(n) = \omega(n^{-1/3})$ as $n \to \infty$. Then every n-vertex graph with at least γn^2 edges contains an $\Omega(\gamma^3 n)$-regular subgraph.*

Before proving Theorem 4.1, we note that it is relatively straightforward to use Theorem 4.8 and the bound $\theta'(m) = O\left(\sqrt{m}\right)$ [10, 36] mentioned in the introduction to prove that $\theta(n) = n^{1-\Omega(1)}$. Indeed, let G be an n-vertex graph and suppose that $e(G) = \Omega(n^{2-2/13})$. Then G contains a bipartite subgraph with $\Omega(n^{2-2/13})$ edges, which in turn contains an $\Omega(n^{1-6/13})$-regular (bipartite) subgraph by Theorem 4.8. Such a subgraph contains $\Omega(n^{2-12/13})$ edges. Delete these edges from G, and repeat until the remaining graph has fewer than $n^{2-2/13}$ edges. This process terminates after $O(n^{2-2+12/13}) = O(n^{12/13})$ steps, at which point we have partitioned

G into $O(n^{12/13})$ regular bipartite (and hence interval colourable) subgraphs, and a subgraph with $O(n^{2-2/13})$ edges. This final portion has interval colouring thickness $O(\sqrt{n^{2-2/13}}) = O(n^{12/13})$ by the bound on $\theta'(m)$, so $\theta(G) = O(n^{12/13})$.

When upper bounding the number of steps in this process we assumed the worst: that there were on the order of n^2 edges in G to begin with. If this were the case, however, then we could (and should) have taken the first few regular bipartite subgraphs to have more edges than we did. Thus, to make optimal use of Theorem 4.8 we must adjust the sizes of the regular bipartite subgraphs as we go through. Note that at the end of the process we appeal to a bound on $\theta'(m)$, so the quality of the bound on $\theta(n)$ that we can obtain from this argument will depend on the quality of our $\theta'(m)$ bound. The next lemma optimises this approach for any given polynomial upper bound on $\theta'(m)$.

Lemma 4.9. *Suppose that $\theta'(m) = O(m^\beta)$ for some $\beta \in (0, 1/2]$. Then $\theta(n) \leqslant n^{\alpha+o(1)}$, where $\alpha = \frac{10\beta}{5+\beta}$.*

Proof. Fix $\varepsilon > 0$. It is sufficient to show that every n-vertex graph has interval colouring thickness $O(n^{\alpha+\varepsilon})$. Let $\alpha_i = \frac{\alpha}{5}\left(1 - \frac{1}{6^i}\right)$ for each $i \in \mathbb{N}$, let $\alpha_0 = 0$, and suppose that G is an n-vertex graph with $n^{2-\alpha_{i+1}} \leqslant e(G) \leqslant n^{2-\alpha_i}$ for some $i \in \mathbb{N}_0$. Then G contains a bipartite subgraph with at least $n^{2-\alpha_{i+1}}/2$ edges which in turn has an $\Omega(n^{1-3\alpha_{i+1}})$-regular (bipartite) subgraph by Theorem 4.8 (note that $\alpha_{i+1} < 1/3$). This subgraph has $\Omega(n^{2-6\alpha_{i+1}})$ edges, so if we delete these edges from G and repeat until the remaining graph has fewer than $n^{2-\alpha_{i+1}}$ edges, then the process will terminate after $O(n^{2-\alpha_i}/n^{2-6\alpha_{i+1}}) = O(n^{6\alpha_{i+1}-\alpha_i}) = O(n^\alpha)$ steps.

Let $C \in \mathbb{N}$ be large enough that $\alpha_C \geqslant \frac{\alpha}{5} - \beta^{-1}\varepsilon$. Every n-vertex graph G has at most $n^2 = n^{2-\alpha_0}$ edges, so by repeatedly applying the above, we can remove at most $C \cdot O(n^\alpha) = O(n^\alpha)$ regular bipartite graphs from G so that fewer than $n^{2-\alpha_C}$ edges remain. By the assumed bound on $\theta'(m)$, this remaining portion has interval colouring thickness $O(n^{\beta(2-\alpha_C)}) = O(n^{\beta(2-\alpha/5+\beta^{-1}\varepsilon)}) = O(n^{\alpha+\varepsilon})$. □

The next lemma converts an upper bound on $\theta(n)$ to one on $\theta'(m)$. The proof is based on the idea that every connected graph with m edges is either dense enough that we can apply the bound on $\theta(n)$, or sparse enough that a spanning tree contains a large proportion of its edges.

Lemma 4.10. *Suppose that $\theta(n) = O(n^\alpha)$ for some $\alpha \in (0, 1]$. Then $\theta'(m) = O(m^\beta)$ where $\beta = \frac{\alpha}{1+\alpha}$.*

Proof. Let $\lambda \geqslant 1$ be such that $\theta\left(\left\lceil \frac{1}{\beta}x \right\rceil\right) \leqslant \lambda x^{\alpha}$ for all $x \in [1,\infty)$. We will show that $\theta(G) \leqslant \lambda m^{\beta}$ for all graphs G with $e(G) = m$. We proceed by induction on m. First, observe that if G has connected components $C_1, \dots, C_k$, then $\theta(G) = \max_i \theta(C_i)$, so we may assume that G is connected. If $|V(G)| > m' := \left\lceil \frac{1}{\beta}m^{1-\beta} \right\rceil$, then G contains a tree T with m' edges. By induction and applying the mean value theorem to the function $x \mapsto x^{\beta}$, we obtain

$$\theta(G-T) \leqslant \lambda(m-m')^{\beta} \leqslant \lambda(m^{\beta} - m'\beta m^{\beta-1}) \leqslant \lambda(m^{\beta}-1) \leqslant \lambda m^{\beta} - 1$$

since $\lambda \geqslant 1$. As T is interval colourable we have $\theta(G) \leqslant \lambda m^{\beta}$ as required. On the other hand, if $|V(G)| \leqslant m'$, then since θ is monotonic we have $\theta(G) \leqslant \theta(m') \leqslant \lambda m^{\alpha-\beta\alpha} = \lambda m^{\beta}$ by our choice of λ. □

We can now prove Theorem 4.1 by playing the preceding lemmas off against each other.

Proof of Theorem 4.1. Let $\alpha_0 = 1$, then for each $i \in \mathbb{N}$ let $\beta_i = \frac{\alpha_{i-1}}{1+\alpha_{i-1}}$ and $\alpha_i = \frac{10\beta_i}{5+\beta_i} = \frac{10\alpha_{i-1}}{5+6\alpha_{i-1}}$. Noting that $\beta_i = \frac{10\beta_{i-1}}{5+11\beta_{i-1}}$ for $i \geqslant 2$, we see that $\alpha_i \in (0,1)$ and $\beta_i \in (0,1/2]$ for all $i \in \mathbb{N}$. Hence, if $\theta(n) \leqslant n^{\alpha_i+o(1)}$ for some $i \in \mathbb{N}_0$, then Lemma 4.10 yields $\theta'(m) \leqslant m^{\beta_{i+1}+o(1)}$. This allows us to deduce from Lemma 4.9 that $\theta(n) \leqslant n^{\alpha_{i+1}+o(1)}$. Clearly $\theta(n) \leqslant n = n^{\alpha_0}$, so $\theta(n) \leqslant n^{\alpha_i+o(1)}$ and $\theta'(m) \leqslant m^{\beta_i+o(1)}$ for all $i \in \mathbb{N}$. The theorem now follows from the fact that $\alpha_i \to \frac{5}{6}$ and $\beta_i \to \frac{5}{11}$ as $i \to \infty$. □

4.3 Concluding remarks

We have improved the bounds on the maximum interval colouring thickness of a graph on n vertices, showing that $(\log n)^{1/3-o(1)} \leqslant \theta(n) \leqslant n^{5/6+o(1)}$. The current best known lower bound is $\theta(n) = \Omega(\log n / \log\log n)$. Clearly there remains a large gap between the bounds, which would be interesting to narrow. We are particularly interested to know whether $\theta(n)$ is polynomial in n, and we believe this is not the case.

Conjecture 4.11. *We have* $\theta(n) = n^{o(1)}$.

We proved our upper bound on $\theta(n)$ by showing that any graph could be edge-decomposed into the desired number of regular bipartite graphs and forests (which are all interval colourable). However, no bound of the form $\theta(n) = o(\sqrt{n})$ can be obtained by decomposing only into these two types of graphs, as the following example shows.

Observation 4.12. Every edge-decomposition of $K_{n,\sqrt{n}}$ into forests and regular graphs has at least $(1-o(1))\sqrt{n}$ parts.

This suggests that to further improve the upper bound on $\theta(n)$, it might be useful to prove an upper bound on the interval colouring thickness of biregular graphs (where a graph is *biregular* if it is bipartite and all vertices in the same part have the same degree). It is a longstanding conjecture of Hansen [61] (see also [74, 114]) that such graphs are, in fact, interval colourable.

Conjecture 4.13 ([61])**.** *Every biregular graph is interval colourable.*

We pose a weaker conjecture.

Conjecture 4.14. *There exists an absolute constant C such that $\theta(G) \leqslant C$ for every biregular graph G.*

A proof of Conjecture 4.14 might pave the way to an improved bound on $\theta(n)$ via decompositions into forests and biregular graphs. We ask whether this strategy has the potential to prove Conjecture 4.11.

Question 4.15. *Can every n-vertex graph be edge-decomposed into $n^{o(1)}$ forests and biregular graphs?*

Chapter 5

Reconstructing matrices in shredded form

This chapter is based on joint work with Balister, Kronenberg and Scott [18].

Let M be an $n \times n$ matrix with entries all either 0 or 1 and let $\mathcal{R}$ and $\mathcal{C}$ be the collections (multisets) of the n binary strings of length n representing the rows and columns of M respectively. When is it possible to reconstruct M just from the knowledge of $\mathcal{R}$ and $\mathcal{C}$?

The matrix M is *reconstructible* (or *weakly reconstructible*) if M is uniquely determined by the multisets $\mathcal{R}$ and $\mathcal{C}$ of its rows and columns. We say M is *strongly reconstructible* if the positions of all rows and columns are determined by $\mathcal{R}$ and $\mathcal{C}$, that is, for each row $r = (r_1, \dots, r_n) \in \mathcal{R}$ we can determine a unique i such that $M_{i,j} = r_j$ for all j, and similarly for columns. Clearly a weakly reconstructible matrix M is strongly reconstructible if and only if there are no two identical rows and no two identical columns, i.e., $\mathcal{R}$ and $\mathcal{C}$ are actually sets.

We study the threshold for reconstructibility when entries of M are random, independently chosen to be 1 with probability $p \leqslant \frac{1}{2}$. Note that corresponding results for $p \geqslant \frac{1}{2}$ can be obtained by exchanging 0 and 1; so we will always assume $p \leqslant \frac{1}{2}$. In [13], Atamanchuk, Devroye and Vicenzo showed that for any $\varepsilon > 0$, M is strongly reconstructible with high probability whenever $(2+\varepsilon)\frac{1}{n}\log n \leqslant p \leqslant \frac{1}{2}$. (We say that an event E happens *with high probability* (*w.h.p.*) if $\mathbb{P}(E) \to 1$ as $n \to \infty$.)

As mentioned above, a simple obstruction to strong reconstructibility is when two rows or two columns are equal, and $p \sim \frac{1}{n}\log n$ is a sharp threshold for such an event (see Lemma 5.6). Here we show that this is the main obstacle, locating the threshold for strong reconstructibility at $p \sim \frac{1}{n}\log n$. In fact, we show a stronger statement. We prove that the threshold for (weak) reconstructibility is $p \sim \frac{1}{2n}\log n$, so above this value

of p, with high probability, the only likely obstruction to strong reconstructibility is the presence of duplicate rows and/or columns. Moreover, we also identify the main obstacle to weak reconstructibility as a pair of 1s in M, each of which is the only 1 in its row and column. The threshold $p \sim \frac{1}{2n} \log n$ for weak reconstructibility is simply the threshold for the disappearance of these obstacles (see Lemma 5.7).

Our main result is the following.

Theorem 5.1. *Suppose that* $p = \frac{1}{2n}(\log n + \log\log n + c_n) \leqslant \frac{1}{2}$.

(a) If $c_n \to \infty$ *as* $n \to \infty$, *then with high probability* M *is reconstructible from the collection of its rows and columns.*

(b) If $c_n \to -\infty$ *as* $n \to \infty$, *then with high probability either* M *has at most one* 1 *or* M *is not reconstructible from the collection of its rows and columns.*

(c) If $c_n \to c$ *as* $n \to \infty$, *then the probability that* M *is reconstructible tends to an explicit constant depending on* c *that is strictly between* 0 *and* 1.

The proof of Theorem 5.1 runs in two phases. We begin by knowing the 'row values' $\mathcal{R}$ and 'column values' $\mathcal{C}$, but not where the rows and columns are placed. In the first phase, we attempt to assign most of the rows and columns to the correct 'row position' or 'column position' using 'local' information. It is helpful here to consider the matrix as the adjacency matrix of a bipartite graph G with bipartition $(\mathcal{I}, \mathcal{J})$, where the vertex class $\mathcal{I}$ corresponds to the indices of the rows of M, the vertex class $\mathcal{J}$ corresponds to the indices of the columns of M, and the edges correspond to the 1s in the matrix. Hence for $i \in \mathcal{I}$, $j \in \mathcal{J}$, ij is an edge if and only if $M_{i,j} = 1$. (We assume $\mathcal{I}$ and $\mathcal{J}$ are disjoint, but both have natural bijections to $[n] := \{1, 2, \ldots, n\}$.) We deduce information about the local structure of G in two different ways:

- For row values $r \in \mathcal{R}$, we deduce the structure of the ball of radius 3 around r by looking at the multiset $\mathcal{R}$ of rows of M, which gives the adjacencies in G between the row values in $\mathcal{R}$ and the column indices (positions) in $\mathcal{J}$.
- For row positions $i \in \mathcal{I}$, we deduce the structure of the ball of radius 3 around the ith row of M by looking at the multiset $\mathcal{C}$ of columns of M, which gives the adjacencies between the row indices in $\mathcal{I}$ and the column values in $\mathcal{C}$.

We complete the first part of the argument by matching up most row values with their row positions and most column values with their column positions. To do this we show that with high probability (for $p > \frac{\delta}{n} \log n$ with any $\delta > 0$), the ball of radius 3 is enough to uniquely identify most vertices of G. Here we use a method similar to

Johnston, Kronenberg, Roberts, and Scott [76] for reconstructing Erdős–Renyi random graphs.

Once we have identified the correct position for most rows and columns, we show in the second phase that we now have enough information to fill in the remaining rows and columns with high probability unless certain substructures occur, and identify the thresholds for these substructures occurring.

In [13], Atamanchuk, Devroye and Vicenzo also proved that for $p \geqslant \frac{(16+\varepsilon)\log^2 n}{n(\log\log n)^2}$ there is an algorithm that succeeds in producing a strong reconstruction of the matrix in $O(n^2)$ time with high probability and in expectation. Our proof gives an algorithm that also produces the matrix in $O(n^2)$ time with high probability above the threshold for weak reconstructibility. See Lemma 5.17 for more details.

As mentioned above, in light of Lemma 5.6, we obtain the threshold for strong reconstructibility as a simple corollary of Theorem 5.1.

Corollary 5.2. *Suppose that* $p = \frac{1}{n}(\log n + c_n) \leqslant \frac{1}{2}$.

(a) If $c_n \to \infty$ *as* $n \to \infty$, *then with high probability, M is strongly reconstructible from the collection of its rows and columns.*

(b) If $c_n \to -\infty$ *as* $n \to \infty$, *then with high probability, M is not strongly reconstructible from the collection of its rows and columns.*

(c) If $c_n \to c$ *as* $n \to \infty$, *then the probability that M is strongly reconstructible tends to an explicit constant depending on c that is strictly between* 0 *and* 1.

Discussion and related results

Reconstruction of binary matrices is closely connected to graph reconstruction.

Reconstruction of bipartite graphs. We note that every binary matrix can be viewed as a bipartite graph where one part acts as the rows of the graph, and the other as the columns. There is an edge ij in this graph if there is 1 at the (i, j) entry of the matrix. Thus, reconstruction of a binary matrix by the collection of its rows and columns, can be achieved by the reconstruction of the corresponding balanced bipartite graph from the collection of its 1-balls, where the centre vertex is unlabelled, but the other vertices are labelled. We will use this connection in our proofs. Our main theorem thus also says the following.

Corollary 5.3. *Suppose that* $p = \frac{1}{2n}(\log n + \log\log n + c_n) \leqslant \frac{1}{2}$ *and let G be a random subgraph of* $K_{n,n}$ *obtained by keeping each edge independently with probability p.*

(a) *If $c_n \to \infty$ as $n \to \infty$, then with high probability G is reconstructible from the collection of its* 1*-balls with unlabelled centres.*

(b) *If $c_n \to -\infty$ as $n \to \infty$, and assuming that G has at least two edges, then with high probability M is not reconstructible from the collection of its* 1*-balls with unlabelled centres.*

(c) *If $c_n \to c$ as $n \to \infty$, then the probability that G is reconstructible from the collection of its* 1*-balls with unlabelled centres tends to an explicit constant depending on c that is strictly between* 0 *and* 1*.*

Reconstruction of directed graphs. Every $n \times n$ binary matrix is equivalent to a directed graph on n vertices, where a directed edge $\vec{ij}$ appears if and only if there is 1 in the (i, j) entry of the matrix (possibly with self-loops). Thus, reconstruction of a binary matrix by the collection of its rows and columns, is equivalent to the reconstruction of a directed graph by the collection of its in- and out- neighbourhoods, where the centre vertex is unlabelled, but the other vertices are labelled. As above, it is straightforward to write down a corollary of Theorem 5.1 for random directed graphs.

Reconstruction of (random) graphs.

There is a huge literature on graph reconstruction, and a growing body of work on reconstructing random graphs and other random combinatorial structures. Kelly and Ulam conjectured in 1941 that every finite simple graph on at least 3 vertices can be determined (up to isomorphism) by its collection of vertex-deleted subgraphs, that is the multiset $\{G \setminus \{v\} : v \in V(G)\}$ of subgraphs of G obtaining by deleting one vertex each time [78, 117]. The Reconstruction Conjecture has a long history and was proved in some special cases, but the general statement is still open (see e.g., [27, 26, 9, 85] for surveys and background).

In a probabilistic setting, the Reconstruction Conjecture tends to be far more tractable. Müller [96] proved in 1976 that it holds for almost all graphs, in the sense that it holds with high probability for the binomial random graph $G(n, \frac{1}{2})$. Bollobás [23] subsequently showed a much stronger result: in fact, with high probability, the graph can be reconstructed from any three of the subgraphs $G \setminus \{v\}$. This led to the understanding that for reconstruction of random combinatorial objects, much less information is needed.

A significant line of recent research looks at when graphs can be reconstructed from 'local' information. Mossel and Ross [93] introduced the 'shotgun reconstruction' problem. The terminology was motivated by the shotgun assembly problem for DNA sequences, where the goal is to reconstruct a DNA sequence from random local 'reads',

corresponding to short subsequences (see [48, 8, 95] among many references). In the context of graphs, the goal is to reconstruct the graph from balls of small radius. For a graph G and a vertex v of G, let $N_G^r(v)$ be the induced graph of the vertices of distance at most r from v (where v is the only vertex in this graph whose original label is known). Then G is *r-reconstructible* if every graph with the same multiset of r-balls as G is isomorphic to G. In other words, G can be identified up to isomorphism from the multiset $\{N_G^r(v) : v \in V(G)\}$.

Mossel and Ross [93] proved that if $p = \frac{\lambda}{n}$, for $\lambda > 1$, then $r = \Theta(\log n)$ is enough for r-reconstruction of $G(n,p)$ with high probability. Sharp asymptotics were obtained by Ding, Jiang, and Ma in 2021 [41]. Mossel and Ross also looked at the problem of reconstructing from balls of constant radius. They showed that if $np/\log^2 n \to \infty$, then $r = 3$ is enough for $G(n,p)$ with high probability. This was later improved by Gaudio and Mossel [55] who also obtained bounds on p for the cases $r = 1$ and 2. The case $r = 1$ was improved by Huang and Tikhomorov [70], who showed that there is a phase transition in 1-reconstructibility around $p = n^{-1/2}$, where the upper and lower bounds differ by a polylogarithmic factor. In [76], this problem was settled for $r \geqslant 3$, and improved bounds were obtained for $r = 1$ and 2.

There is also work on other graph models including random regular graphs [94], random geometric graphs [2] and random simplicial complexes [1].

Going back to reconstruction of matrices, one may see this study as the complement of the previous one. Here, the subgraphs given have all vertices labelled except the centre, while in the Mossel-Ross type of random graph reconstruction the only labelled vertex in each given subgraph is the middle one. As we will see later, the fact that in the matrix case the neighbourhood is labelled, will allow us to obtain information on larger and larger balls, and by that to apply some of the techniques that were presented in [76] for r-reconstruction, where $r \geqslant 3$. In this case, however, the notion of r-neighbourhoods is slightly different, as all vertices at odd distance from the root of the ball are labelled but all vertices at even distance are not. But the informative structure of the matrix does allow us to obtain statistics on the vertices of distance at most r of each vertex, which will be sufficient for reconstruction.

Organisation

The chapter is organised as follows. In Section 5.1 we give classical probabilistic results and thresholds for the substructures that provide obstacles to either weak or strong reconstructibility. In Section 5.2 we start the proof of Theorem 5.1 by showing how to

reconstruct almost all rows and columns. In Section 5.3 we show that we can complete this into a full reconstruction unless some specific substructures occur, and complete the proof of Theorem 5.1. Finally in Section 5.4 we show that reconstruction can be achieved w.h.p. in $O(n^2)$ time.

5.1 Preliminaries

In this section we state probabilistic bounds which will be useful later in the chapter.

We make frequent use of the Chernoff bounds on the tails of the binomial distribution as stated in Lemma 1.3. We will also be interested in tail bounds for binomial distributions where $np \to 0$ as $n \to \infty$, for which we use the following simple observation.

Lemma 5.4. *Let $X \sim \mathrm{Bin}(n,p)$ and $k \in \mathbb{N}$. Then*

$$\mathbb{P}(X \geqslant k) \leqslant \tbinom{n}{k} p^k \,.$$

Proof. Indeed, $\mathbb{P}(X \geqslant k)$ is at most the expected number of k-tuples of trials that all succeed, which is $\binom{n}{k}p^k$. □

We now quote the following result on convergence to a Poisson distribution which follows directly from [24, Theorem 1.23]. Here we use the *falling factorial* notation $(n)_r := n(n-1)\cdots(n-r+1)$.

Lemma 5.5. *Let $\lambda_1, \dots, \lambda_k$ be non-negative reals and $X_n^{(i)}$, $i = 1, \dots, k$, be sequences of random variables such that for all k-tuples $(r_1, \dots, r_k)$ of non-negative integers,*

$$\mathbb{E}\big[(X_n^{(1)})_{r_1} \cdots (X_n^{(k)})_{r_k}\big] \to \lambda_1^{r_1} \cdots \lambda_k^{r_k} \qquad \text{as } n \to \infty,$$

Then $(X_n^{(1)}, \dots, X_n^{(k)})$ converges jointly in distribution to independent Poisson random variables with means $\lambda_1, \dots, \lambda_k$. Namely, for any non-negative integers $s_1, \dots, s_k$,

$$\lim_{n\to\infty} \mathbb{P}\big(X_n^{(1)} = s_1, \dots, X_n^{(k)} = s_k\big) = \prod_{i=1}^{k} \frac{e^{-\lambda_i}\lambda_i^{s_i}}{s_i!} \,.$$

Recall that M is an $n \times n$ matrix with i.i.d. Bernoulli random entries that are 1 with probability p. The following lemma gives the probability for the main obstruction to strong reconstructibility, which is the appearance of two identical rows or two identical columns.

Lemma 5.6. *Assume* $p = \frac{1}{n}(\log n + c_n) \leqslant \frac{1}{2}$. *Then, in the matrix* M,

$$\mathbb{P}(\exists \text{ two equal rows or two equal columns}) \to \begin{cases} 1, & \text{if } c_n \to -\infty, \\ 1 - \left((1+e^{-c})e^{-e^{-c}}\right)^2, & \text{if } c_n \to c, \\ 0, & \text{if } c_n \to \infty. \end{cases}$$

Proof. Let N be the number of pairs of identical rows in M that are *not* entirely zero. The probability that two given rows are identical is $((1-p)^2+p^2)^n$ and the probability that the two rows are identically zero is $(1-p)^{2n}$. Hence,

$$\begin{aligned} \mathbb{E}[N] &= \binom{n}{2}\left(\left((1-p)^2+p^2\right)^n - \left((1-p)^2\right)^n\right) \\ &\leqslant n^2 \cdot np^2 \cdot \left((1-p)^2+p^2\right)^{n-1} \\ &\leqslant n^2 \cdot np^2 \cdot e^{-2(n-1)p(1-p)}, \end{aligned}$$

where we have used comparison with a geometric series (or the Mean Value Theorem) in the second line, and the inequality $1-x \leqslant e^{-x}$ in the third line. This last expression tends to zero except in the case $p \leqslant (\frac{1}{2}+o(1))\frac{1}{n}\log n$, in which case $c_n \to -\infty$.

Hence it is enough to consider pairs of identically zero rows or columns. As the number of zero rows and/or zero columns is stochastically decreasing as p increases, it is enough to prove the result just in the case when $c_n \to c$ as the other two cases follow from stochastic domination and taking limits $c \to \pm\infty$.

Let X be the number of zero rows and Y the number of zero columns. Then for $r, s \geqslant 0$, $(X)_r(Y)_s$ counts the number of choices of r-tuples of rows and s-tuples of columns, all filled with zero. Such a configuration consists of a union of r rows and s columns with all $rn+sn-rs$ entries equal to zero. There are $(n)_r$ r-tuples of distinct rows and $(n)_s$ s-tuples of distinct columns so, for fixed r and s,

$$\mathbb{E}\big[(X)_r(Y)_s\big] = (n)_r(n)_s(1-p)^{rn+sn-rs} \to e^{-c(r+s)} \qquad \text{as } n \to \infty,$$

as $(n)_r/n^r, (n)_s/n^s, (1-p)^{-rs} \to 1$ and $n(1-p)^n = \exp(\log n - pn + O(p^2 n)) \to e^{-c}$ as $n \to \infty$. Thus by Lemma 5.5, (X, Y) converges in distribution to i.i.d. $\mathrm{Po}(e^{-c})$ random variables. The probability that there are either two empty rows or two empty columns is $\mathbb{P}(X \geqslant 2 \text{ or } Y \geqslant 2)$, which converges to the expression given in the statement of the lemma. □

Similarly, the following lemma gives the probability of the main obstruction to weak reconstructibility. We say an entry in M is an *isolated* 1 if the entry is a 1 but all other entries in the same row or column are zero. In the graph G this corresponds to an

	$*$	0	$*$	$*$	0	$*$			$*$	0	$*$	$*$	0	$*$
	$*$	0	$*$	$*$	0	$*$			$*$	0	$*$	$*$	0	$*$
r_1	0	1	0	0	0	0		r_2	0	0	0	0	1	0
	$*$	0	$*$	$*$	0	$*$			$*$	0	$*$	$*$	0	$*$
r_2	0	0	0	0	1	0		r_1	0	1	0	0	0	0
	$*$	0	$*$	$*$	0	$*$			$*$	0	$*$	$*$	0	$*$

Figure 5.1. Isolated ones are the main obstruction to weak reconstructibility. Permuting rows r_1, r_2 as shown produces two different matrices.

isolated edge, i.e., a component consisting of a single edge. The main obstruction to weak reconstructibility turns out to be the appearance of two or more isolated 1s in the matrix. Indeed, in such a case the isolated 1s must appear in distinct rows and columns, and any permutation of their rows, say, results in a matrix distinct from M, but with identical multisets of rows and columns, as in Figure 5.1.

Lemma 5.7. *Suppose* $p = \frac{1}{2n}(\log n + \log\log n + c_n) \leqslant \frac{1}{2}$. *Let* X *be the number of* 1s *in the matrix* M *and let* Y *be the number of isolated* 1s *in* M. *Then,*

$$\mathbb{P}(Y \geqslant 2 \text{ or } X < 2) \to \begin{cases} 1, & \text{if } c_n \to -\infty, \\ 1 - (1 + e^{-c}/2)e^{-e^{-c}/2}, & \text{if } c_n \to c, \\ 0, & \text{if } c_n \to \infty. \end{cases}$$

Proof. Note that $X \sim \mathrm{Bin}(n^2, p)$, so the $X < 2$ condition is only significant if $n^2 p$ is bounded, i.e., only in the $c_n \to -\infty$ case.

Now $(Y)_r$ counts the number of r-tuples of isolated 1s, all of which must lie in distinct rows and columns. Hence,

$$\mathbb{E}\big[(Y)_r\big] = (n)_r (n)_r p^r (1-p)^{2nr - r^2 - r}.$$

If $c_n \to \infty$ then when $p \leqslant n^{-1/2}$,

$$\mathbb{E}[Y] = n^2 p (1-p)^{2n-2} = n^2 p e^{-2np - O(np^2 + p)} = \frac{np}{\log n} e^{-c_n + O(1)} \to 0,$$

and clearly $\mathbb{E}[Y] = O(n^2 e^{-2np}) \to 0$ for larger p. Hence $\mathbb{P}(Y \geqslant 2) \to 0$ by Markov and, as noted above, $\mathbb{P}(X < 2) \to 0$ as well.

If $c_n \to c$ then,

$$n^2 p (1-p)^{2n} = n^2 p e^{-2np - O(np^2)} = \frac{np}{\log n} e^{-c_n + o(1)} \to e^{-c}/2.$$

As r is fixed and $p \to 0$, we then have $\mathbb{E}[(Y)_r] \to (e^{-c}/2)^r$. Thus Y tends in distribution to a $\mathrm{Po}(e^{-c}/2)$ random variable. Also $\mathbb{P}(X < 2) \to 0$, so $\mathbb{P}(Y \geqslant 2 \text{ or } X < 2)$ converges to the expression given.

If $c_n \to -\infty$ but $n^2 p \to \infty$ then $\mathbb{E}[Y] \to \infty$. However,

$$\frac{\mathbb{E}[Y(Y-1)]}{\mathbb{E}[Y]^2} = \frac{(n-1)^2}{n^2}(1-p)^{-2} \to 1.$$

Thus $\mathrm{Var}[Y] = \mathbb{E}[Y(Y-1)] + \mathbb{E}[Y] - \mathbb{E}[Y]^2 = o(\mathbb{E}[Y]^2)$. Hence by Chebychev's inequality (i.e., the second moment method),

$$\mathbb{P}(Y < 2) \leqslant \mathbb{P}\big(|Y - \mathbb{E}[Y]| > \mathbb{E}[Y] - 2\big) \leqslant \frac{\mathrm{Var}[Y]}{(\mathbb{E}[Y]-2)^2} \to 0,$$

so $\mathbb{P}(Y \geqslant 2) \to 1$.

Finally we may assume $n^2 p = O(1)$, in which case $Y = X$ w.h.p. as the probability of any row or column containing at least two 1s is $O(n^3 p^2) = o(1)$. Hence in this case $\mathbb{P}(Y \geqslant 2 \text{ or } X < 2) \geqslant \mathbb{P}(X = Y) \to 1$. □

The following lemma will allow us to bound the probability that two multisets of i.i.d. random variables are the same. We shall use the notation $[x_1, \ldots, x_d]$ to denote the multiset consisting of the elements $x_1, \ldots, x_d$.

Lemma 5.8. *Let $X_1, \ldots, X_d$ be i.i.d. discrete random variables with $\mathbb{P}(X_i = x) \leqslant p_0$ for all x. Then for any multiset $\mathcal{M}$ of d possible values of X_i,*

$$\mathbb{P}\big([X_1, \ldots, X_d] = \mathcal{M}\big) \leqslant \frac{(2\pi d + 2)^{1/2}}{(2\pi p_0 d + 1)^{1/(2p_0)}} = O\big(\sqrt{d}(2\pi p_0 d)^{-1/(2p_0)}\big).$$

The proof boils down to bounding an expression of the form $\binom{d}{d_1,\ldots,d_t} \cdot p_0^d$. This is maximised when the product $d_1! \cdots d_t!$ is minimised, which should occur when all d_is are equal. If $p_0 d \leqslant 1$, then the worst case should be when all d_is are 1, in which case we have the simple bound of $d! p_0^d$ obtained by summing over all permutations $\sigma \in S_d$ the probability $\mathbb{P}(X_1 = x_{\sigma(1)}, \ldots, X_d = x_{\sigma(d)})$, where $\mathcal{M} = [x_1, \ldots, x_d]$.

Our bound is weaker however, as we could have $d > 1/p_0$ in general. If we suppose that $d_i = d'$ for each i, then by Stirling's formula the asymptotics of this expression are essentially given by

$$\sqrt{2\pi d} \cdot (2\pi d')^{-d/(2d')} \cdot (d p_0 / d')^d.$$

But at most d/d' distinct x_is can occur with probability p_0, so $dp_0/d' \leqslant 1$. In particular, the above expression is bounded by

$$\sqrt{2\pi d} \cdot (2\pi d p_0)^{-d/(2d')}.$$

In fact, we argue below that this expression is maximised when the equality $dp_0/d' = 1$ holds, so this last exponent can be replaced by $-1/(2p_0)$.

Proof. Let the multiset $\mathcal{M}$ with the highest probability have d_i copies of an element x_i where x_i occurs with probability p_i, $i = 1, 2, \ldots$. We will use the following version of Stirling's formula which holds for all $d \geqslant 0$:

$$(d/e)^d \sqrt{2\pi d + 1} \leqslant d! \leqslant (d/e)^d \sqrt{2\pi d + 2}.$$

We note that this clearly holds for $d = 0$ (with the interpretation that $0^0 = 1$) and follows easily for $d \geqslant 1$ from the explicit bounds

$$(d/e)^d \sqrt{2\pi d}\, e^{1/(12d+1)} \leqslant d! \leqslant (d/e)^d \sqrt{2\pi d}\, e^{1/12d}$$

proved by Robbins [103]. Thus,

$$\mathbb{P}\big([X_1, \ldots, X_d] = \mathcal{M}\big) = \frac{d!}{d_1! \cdots d_t!} p_1^{d_1} \cdots p_t^{d_t} \leqslant \sqrt{2\pi d + 2} \cdot \prod_{i=1}^{t} \Big(\frac{dp_i}{d_i}\Big)^{d_i} \frac{1}{\sqrt{2\pi d_i + 1}},$$

where we reindex the p_i and d_i if necessary so that $d_1, \ldots, d_t > 0$ and $d_i = 0$ for $i > t$. Now set $\alpha_i = p_i/p_0$ and $d_i' = d_i/\alpha_i$. We note that $dp_i/d_i = dp_0/d_i'$ and $2\pi d_i + 1 \geqslant (2\pi d_i' + 1)^{\alpha_i}$, where the last bound follows by Bernouilli's inequality for the range $0 < \alpha_i \leqslant 1$. Hence we can rewrite the bound as

$$\log \mathbb{P}\big([X_1, \ldots, X_d] = \mathcal{M}\big) \leqslant \log \sqrt{2\pi d + 2} + \sum_i \alpha_i \Big(d_i' \log \tfrac{dp_0}{d_i'} - \tfrac{1}{2} \log\big(2\pi d_i' + 1\big)\Big),$$

where we take this last sum over all $i \geqslant 1$, noting that all summands vanish when $i > t$. Now maximise over the d_i', assumed just to be non-negative reals with $\sum \alpha_i d_i' = d$. It is easy to check that $f(x) = x \log \frac{dp_0}{x} - \frac{1}{2}\log(2\pi x + 1)$ is concave, where we set $f(0) = 0$. Indeed $f(x) \to 0$ as $x \to 0^+$ and $f''(x) = -\frac{1}{x} + \frac{2\pi^2}{(2\pi x+1)^2} = -\frac{(2\pi x-1)^2 + (4-\pi)(2\pi x)}{x(2\pi x+1)^2} < 0$ for all $x > 0$. By Jensen's inequality for concave functions, to maximise $\sum \alpha_i f(d_i')$ subject to $\sum \alpha_i d_i' = d$ requires taking all d_i' to be equal, say $d_i' = d'$. We do this for all i, even $i > t$. But then $\frac{d}{d'} = \sum \alpha_i = \sum \frac{p_i}{p_0} = \frac{1}{p_0}$, so all $d_i' = d' = p_0 d$. Thus,

$$\log \mathbb{P}\big([X_1, \ldots, X_d] = \mathcal{M}\big) \leqslant \log \sqrt{2\pi d + 2} - \tfrac{1}{2p_0} \log\big(2\pi p_0 d + 1\big),$$

as required. □

5.2 Reconstructing almost all rows and columns

For this section, let M be a binary matrix whose entries are i.i.d. Bernoulli random variables taking value 1 with probability p, where $\frac{\delta}{n} \log n \leqslant p \leqslant \frac{1}{2}$ for some fixed small constant $\delta > 0$.

Recall that the matrix M corresponds to a bipartite graph G with vertex classes $\mathcal{I}$ and $\mathcal{J}$, both of size n, corresponding to the indices of the rows and columns. For $i \in \mathcal{I}$, $j \in \mathcal{J}$, ij is an edge of G if and only if the matrix M has 1 in the entry (i, j).

Given the multisets $\mathcal{R}$ and $\mathcal{C}$, we can reconstruct two graphs, both isomorphic to G. The first is G_R, which is a bipartite graph with vertex classes $\mathcal{R}$ and $\mathcal{J}$ where the row value $r \in \mathcal{R}$, which is a binary vector $r = (r_1, \ldots, r_n)$, is joined to all columns $j \in \mathcal{J}$ where $r_j = 1$. The second is G_C, which is a bipartite graph with vertex classes $\mathcal{I}$ and $\mathcal{C}$ where the column value $c = (c_1, \ldots, c_n)^T \in \mathcal{C}$ is joined to all rows $i \in \mathcal{I}$ where $c_i = 1$. Clearly G_R and G_C are both isomorphic to G by the correct identification of the row values in $\mathcal{R}$ with their indices in $\mathcal{I}$, and the column values in $\mathcal{C}$ with their indices in $\mathcal{J}$ respectively.

For a vertex v of G (or G_R or G_C), define its *kth degree statistics* inductively as follows:

$$\begin{aligned} \mathcal{D}_0(v) &= \deg(v), \\ \mathcal{D}_{k+1}(v) &= [\mathcal{D}_k(u) : u \in N(v)], \text{ for } k > 0. \end{aligned}$$

Here $N(v)$ is the set of vertices adjacent to v and $\deg(v) = |N(v)|$ is the degree of the vertex v. Note that as G_R and G_C are isomorphic to G, $\mathcal{D}_k(v)$ can be reconstructed from $\mathcal{R}$, $\mathcal{C}$, and *either* the index *or* the value of the row or column v: indeed, if we know that v is the ith row of M, $\mathcal{D}_0(v)$ is obtained by counting the number of columns $c \in \mathcal{C}$ which have a 1 in the ith coordinate. Counting the number of 1s in these columns then gives the 1st degree statistics, and so on. Similarly, knowing the value of v gives us the number of columns which have a 1 in this row, i.e. $\mathcal{D}_0(v)$, as well as their index. A symmetric procedure to the above then computes the 1st and higher degree statistics: for each such column c with index j, count the number of rows which have a 1 in coordinate j, and so on.

In particular, we observe that the index of a row value r, say, can be correctly identified if for all rows $r' \neq r$ there is some k such that $\mathcal{D}_k(r) \neq \mathcal{D}_k(r')$.

We first show that for large p, w.h.p. even $\mathcal{D}_1(v)$ is enough to uniquely identify all rows and columns.

Lemma 5.9. *There exists a $C > 0$ such that for $\frac{C}{n}\log^2 n \leqslant p \leqslant \frac{1}{2}$ and any two distinct rows r and r' of M,*

$$\mathbb{P}\big(\mathcal{D}_1(r) = \mathcal{D}_1(r')\big) = o(n^{-2}).$$

In particular, w.h.p. M is strongly reconstructible.

Proof. Let $r, r' \in \mathcal{R}$. If $\mathcal{D}_1(r) = \mathcal{D}_1(r')$ then the multiset $\mathcal{M}_1$ of degrees of vertices (columns) in $N(r) \setminus N(r')$ is the same as the multiset $\mathcal{M}_2$ of degrees of vertices in $N(r') \setminus N(r)$. The multiset of degrees of vertices in $N(r) \cap N(r')$ contributes equally to $\mathcal{D}_1(r)$ and $\mathcal{D}_1(r')$, so we can ignore these. First, we bound the probability of these multisets being equal in terms of $d = |N(r) \setminus N(r')| = |\mathcal{M}_1|$.

The multisets must have the same size in order to be equal, so $\mathbb{P}(\mathcal{M}_1 = \mathcal{M}_2)$ is bounded by the same probability conditioned on $|\mathcal{M}_2| = d$. If we condition this as well as on $N(r)$ and $N(r')$ (i.e., the entries in rows r and r'), then $\mathcal{M}_1$ and $\mathcal{M}_2$ are i.i.d. random multisets of size d, each of which consists of i.i.d. $\mathrm{Bin}(n-2, p) + 1$ random variables. The $+1$ is because we have exactly one of r or r' counted in the degrees. The elements of the multisets are independent as they depend only on distinct columns, and similarly the multisets themselves are independent as they depend only on disjoint rectangles $(\mathcal{I} \setminus \{r, r'\}) \times (N(r) \setminus N(r'))$ and $(\mathcal{I} \setminus \{r, r'\}) \times (N(r') \setminus N(r))$ respectively, of unconditioned entries in M.

Now if $X \sim \mathrm{Bin}(n-2, p)$ then $p_0 := \max_x \mathbb{P}(X = x) = \Theta(1/\sqrt{np})$, which we may assume is at most $1/(4\log n)$ if C is large enough. Hence, by Lemma 5.8,

$$\mathbb{P}(\mathcal{M}_1 = \mathcal{M}_2) = O(\sqrt{d}) \exp\big(-\tfrac{1}{2p_0}\log(2\pi p_0 d)\big) = O\big(n^{1-2\log(2\pi p_0 d)}\big) = o(n^{-2}),$$

provided $p_0 d > 2$, say. But unconditionally $d \sim \mathrm{Bin}(n, p(1-p))$ stochastically dominates $\mathrm{Bin}(n, p/2)$ and thus by Lemma 1.3,

$$\mathbb{P}(d < np/4) \leqslant \exp(-np/16) = o(n^{-2}).$$

But $d \geqslant np/4$ implies $p_0 d = \Theta(\sqrt{np}) > 2$ for large n. Hence unconditionally,

$$\mathbb{P}\big(\mathcal{D}_1(r) = \mathcal{D}_1(r')\big) = \mathbb{P}\big(\mathcal{M}_1 = \mathcal{M}_2\big) = o(n^{-2}).$$

The union bound now shows that the probability that there are two rows with the same value of $\mathcal{D}_1(r)$ is $o(1)$, and the same also holds for columns. Since the $\mathcal{D}_1(v)$ can be determined either from the values or the indices of the rows or columns v, w.h.p. each row or column value can be associated with a unique index, and so M is strongly reconstructible. □

For smaller values of p we will need to consider $\mathcal{D}_2(v)$. Before we do so, it will be convenient to prove some results about the typical structure of M when $p = O(\frac{1}{n}(\log n)^2)$.

Lemma 5.10. *Let $p = O(\frac{1}{n}\log^2 n)$. Then with high probability, every pair of distinct rows r and r' satisfies $|N(r) \cap N(r')| \leqslant 2$.*

Proof. We have $|N(r) \cap N(r')| \sim \mathrm{Bin}(n, p^2)$, so

$$\mathbb{P}\big(|N(r) \cap N(r')| \geqslant 3\big) \leqslant \tbinom{n}{3}p^6 = o(n^{-2}).$$

A union bound now shows that w.h.p. there is no pair of rows r, r' for which $|N(r) \cap N(r')| \geqslant 3$. □

The next fact roughly states that the 2-balls in our bipartite graph are essentially cycle-free. We will use this to argue that the degree distributions of leaves in these balls are essentially independent.

Lemma 5.11. *For $p = O(\frac{1}{n}\log^2 n)$ w.h.p. there does not exist a pair r, r' of rows for which there are more than two rows $s \neq r, r'$ with at least two neighbours in $N(r) \cup N(r')$.*

Proof. Consider a pair of rows r and r'. If rows $s_1, s_2, s_3 \neq r, r'$ all have at least two neighbours in $N(r) \cup N(r')$, then choose two such neighbours for each of these rows to form a subset $C \subseteq N(r) \cup N(r')$ of columns of size $k := |C| \leqslant 6$ with at least 6 edges from $\{s_1, s_2, s_3\}$ to C. Adding r and r', it follows that there is a set of 5 rows and k columns with $k + 6$ entries equal to 1. The expected number of these configurations is $O(n^{k+5}p^{k+6}) = o(1)$. Summing over $k \leqslant 6$, we see that with high probability there is no configuration of this type, and so no choice of r, r' and the s_i. □

A standard Chernoff bound argument shows the following bound on the maximum number of 1s in any row or column.

Lemma 5.12. *Fix $\delta > 0$. Then there exists a constant $K = K(\delta)$ such that for $p \geqslant \frac{\delta}{n}\log n$, w.h.p., no row or column has degree more than Knp in G.*

Proof. Set $\varepsilon = K - 1$ in Lemma 1.3 and note that $\frac{\varepsilon^2 np}{2+\varepsilon} \geqslant \frac{(K-1)^2\delta}{K+1}\log n \geqslant 2\log n$ for sufficiently large K. Thus the probability that a fixed row or column has degree more than Knp is at most $e^{-2\log n} = 1/n^2$. The result follows from the union bound over the $2n$ rows and columns. □

Call a row or column *heavy* if it has at least $\frac{1}{2}np$ ones, and *light* otherwise.

Lemma 5.13. *Fix $\delta > 0$ and assume $p \geqslant \frac{\delta}{n} \log n$. Then w.h.p. there are at most $n^{1-\delta/9} = o(n)$ light rows. Moreover, w.h.p. each row r is adjacent to at most $K' = K'(\delta)$ light columns.*

Proof. By Lemma 1.3, each row is light with probability at most $e^{-np/8} \leqslant n^{-\delta/8}$, so by Markov there are w.h.p. at most $n^{1-\delta/9}$ light rows.

Now given a column c and a row index $i_0 \in \mathcal{I}$, Lemma 1.3 gives

$$\mathbb{P}\Big(\sum_{i \neq i_0} c_i < \tfrac{1}{2}np - 1\Big) \leqslant \mathbb{P}\Big(\sum_{i \neq i_0} c_i < \tfrac{1}{2}(n-1)p\Big) \leqslant e^{-(n-1)p/8},$$

as $\sum_{i \neq i_0} c_i \sim \mathrm{Bin}(n-1, p)$. Therefore, given a row r, we have that, conditioned on $N(r)$, the probability that a fixed column $c \in N(r)$ is light is at most $e^{-(n-1)p/8}$. Hence, as distinct columns are independent,

$$\mathbb{P}(\text{There are at least } K' \text{ light columns in } N(r)) \leqslant \binom{|N(r)|}{K'} \big(e^{-(n-1)p/8}\big)^{K'}.$$

Now w.h.p. every row r satisfies $|N(r)| \leqslant Knp$ by Lemma 5.12, so for $K' := \lceil 9/\delta \rceil$, say, we have

$$\binom{|N(r)|}{K'} \big(e^{-(n-1)p/8}\big)^{K'} \leqslant e^{K'(\log(Knp) - (n-1)p/8)} = o(n^{-1}).$$

Taking a union bound over the rows of M finishes the argument. □

Thus, it is enough to reconstruct only heavy rows and columns in order to reconstruct $(1 - o(1))$ of the matrix.

Lemma 5.14. *Fix $C, \delta > 0$ and suppose $\frac{\delta}{n} \log n \leqslant p = p(n) \leqslant \frac{C}{n} \log^2 n$. Then w.h.p. for any pair of distinct rows r and r' with r heavy, $\mathcal{D}_2(r) \neq \mathcal{D}_2(r')$. In particular, w.h.p. all heavy row and column values can be matched with their indices.*

Proof. Since p is in the regime where Lemmas 5.10–5.13 apply and their conclusions each hold w.h.p., we may assume for the rest of the proof that these conclusions all simultaneously hold, prior to any conditioning.[1]

Fix two row positions r, r' and assume without loss of generality that $\deg(r) = \deg(r')$ as otherwise $\mathcal{D}_2(r) \neq \mathcal{D}_2(r')$. Reveal the entries in r and r', then reveal all entries in columns of $N(r) \cup N(r')$. Finally, reveal all entries in rows of $N(N(r'))$, so that the full information determining $\mathcal{D}_2(r')$ is revealed. Write $\mathcal{R}$ and $\mathcal{C}$ for the sets of indices of rows and columns revealed so far respectively, namely $\mathcal{R} = \{r, r'\} \cup \{i : m_{ij} = 1 \text{ for } j \text{ such that } r'_j = 1\}$ and $\mathcal{C} = \{j : r_j = 1 \text{ or } r'_j = 1\}$.

[1] Note that if L is the event 'the conclusion of the lemmas hold' and E is any event, $\mathbb{P}(E) = \mathbb{P}(E|L)\mathbb{P}(L) + \mathbb{P}(E|\neg L)\mathbb{P}(\neg L) \to \mathbb{P}(E|L)$ as $n \to \infty$ since L holds w.h.p.. In particular, independence of events is preserved by conditioning on L.

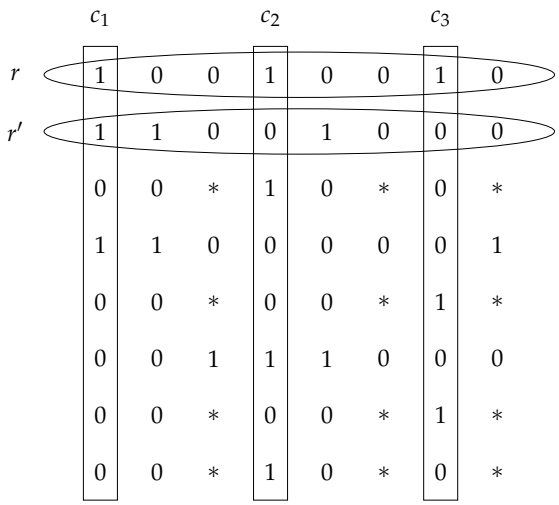


Figure 5.2. The matrix M to be reconstructed, where $*$ represents an entry yet to be revealed. c_1, c_2 are not admissible due to conditions (i) and (ii) respectively, c_3 is admissible. In our regime and for large values of n, almost all columns in $N(r)$ will be admissible.

For the second degree statistics to coincide, there needs to be a matching $\sigma\colon N(r) \to N(r')$ such that $\mathcal{D}_1(c) = \mathcal{D}_1(\sigma(c))$ for all columns $c \in N(r)$. We take a union bound over all $d!$ such matchings where $d = |N(r)| = |N(r')|$.

$$\begin{aligned}\mathbb{P}\big(\mathcal{D}_2(r) = \mathcal{D}_2(r')\big) &\leqslant \sum_{\sigma} \mathbb{P}\big(\forall c \in N(r),\ \mathcal{D}_1(c) = \mathcal{D}_1(\sigma(c))\big) \\ &\leqslant d! \cdot \max_{\sigma} \mathbb{P}\big(\forall c \in N(r),\ \mathcal{D}_1(c) = \mathcal{D}_1(\sigma(c))\big).\end{aligned}$$

It now suffices to bound the last probability for any matching σ.

We first show that essentially all degree distributions $\mathcal{D}_1(c)$ are independent from each other and crucially from the information revealed so far. The probability of each $\mathcal{D}_1(c)$ hitting its prescribed degree distribution is bounded by Lemma 5.8. Taking a product over the appropriate columns c will yield the desired bound.

Say that a column $c \in N(r)$ is *admissible* if none of the following occurs:

(i) $c \in N(r')$;

(ii) there is a $c' \in N(r) \cup N(r')$ with $c' \neq c$ such that $(N(c') \cap N(c)) \setminus \{r, r'\} \neq \varnothing$;

(iii) $\deg(c) < \frac{1}{2}np$.

In particular, if c is admissible and $r'' \in N(c) \setminus \{r\}$, then $r'' \in \mathcal{I} \setminus \mathcal{R}$ so its entries in columns $\mathcal{J} \setminus \mathcal{C}$ have yet to be revealed. See Figure 5.2 for an example, and note that admissibility of $c \in N(r)$ depends only on entries which have already been revealed.

By Lemma 5.10, $N(r)$ and $N(r')$ share at most two elements, so at most two columns c fail because of condition (i). By Lemma 5.11 there are at most two rows $s \neq r, r'$ whose neighbourhood intersects $N(r) \cup N(r')$ in at least two columns. Since these intersections have size at most 4 each, there are at most 8 columns $c \in N(r) \cup N(r')$ for which there is a $c' \in N(r) \cup N(r')$, $c' \neq c$, with $N(c) \cap N(c')$ containing such a row s. In other words, there are at most 8 columns which fail the condition (ii). Lastly, by Lemma 5.13, only constantly many $c \in N(r)$ fail condition (iii).

As r is assumed heavy, for large n there are at least $\frac{1}{3}np$ admissible vertices in $N(r)$.

Notice that if a column c is admissible, then, after omitting r, all rows in its neighbourhood $N(c) \setminus \{r\}$ have had exactly one non-zero entry revealed so far. In particular, the degree of each such row only depends on its $N = n - |N(r) \cup N(r')|$ entries in $\mathcal{J} \setminus \mathcal{C}$ which are yet to be revealed, so there are i.i.d. random variables $X_2, \dots, X_{\deg(c)} \sim \mathrm{Bin}(N, p)$ such that $\mathcal{D}_1(c) = [\deg(r), X_2 + 1, \dots, X_{\deg(c)} + 1]$. Here, the first term set to $\deg(r)$ corresponds to row r being entirely revealed, and the $+1$ terms in other entries correspond to the entire column c being revealed (and therefore exactly one non-zero entry of each row adjacent to c is revealed so far).

In particular, we can apply Lemma 5.8 to bound the probability that these X_i produce a specified multiset. Additionally, the second admissibility condition ensures that neighbourhoods of distinct admissible columns are disjoint. Namely, the degrees of the neighbours of distinct admissible columns c, c' depend on disjoint subsets of unconditioned entries and are therefore independent, so $\mathbb{P}(\mathcal{D}_1(c) = \mathcal{D}_1(c')) = \mathbb{P}(\mathcal{D}_1(c) = \mathcal{M})$ for a random multiset $\mathcal{M}$ of $\deg(c')$ possible values of X_i.

Now for any admissible c and choosing any fixed choice of $\mathcal{D}_1(\sigma(c))$ we have that

$$\mathbb{P}(\mathcal{D}_1(c) = \mathcal{D}_1(\sigma(c))) = O(\deg(c)^{1/2})(2\pi p_0 \deg(c))^{-1/p_0},$$

where $p_0 = \Theta(1/\sqrt{Np})$. However, $|N(r) \cup N(r')| \leqslant 2Knp$, so $p_0 = \Theta(1/\sqrt{np})$ and by assumption $\deg(c) \geqslant \frac{1}{2}np$. Hence,

$$\mathbb{P}(\mathcal{D}_1(c) = \mathcal{D}_1(\sigma(c))) \leqslant \exp\big(-\Theta(\sqrt{np}\log(np))\big).$$

Now, if $\mathcal{D}_1(c) = \mathcal{D}_1(\sigma(c))$ holds for every $c \in N(r)$ then certainly it holds for every admissible c, so,

$$\begin{aligned}\mathbb{P}\big(\forall c \in N(r),\ \mathcal{D}_1(c) = \mathcal{D}_1(\sigma(c))\big) &\leqslant \mathbb{P}\big(\forall c \text{ admissible},\ \mathcal{D}_1(c) = \mathcal{D}_1(\sigma(c))\big) \\ &= \prod_{c \text{ admissible}} \mathbb{P}\big(\mathcal{D}_1(c) = \mathcal{D}_1(\sigma(c))\big) \\ &= \exp\big(-\Omega((np)^{3/2}\log(np))\big),\end{aligned}$$

where we used that there are at least $np/3$ admissible columns. Finally, note that $d! \leqslant (Knp)! = \exp\big(O(np\log(np))\big)$, so that

$$\begin{aligned}\mathbb{P}\big(\mathcal{D}_2(r) = \mathcal{D}_2(r')\big) &\leqslant \exp\big(-\Omega((np)^{3/2}\log(np)) + O(np\log(np))\big)\\ &\leqslant \exp\big(-\Omega(\log^{3/2} n)\big) = o(n^{-2}).\end{aligned}$$

Taking a union bound over all choices of r and r' then gives the result. □

5.3 Full reconstruction

In this section we show that once all but $o(n)$ of the rows and columns have been reconstructed, then with high probability the remaining entries of M can be deduced unless there is some simple obstruction. Recall that we say an entry in M is an isolated 1 if the entry is a 1 but all other entries in the same row or column are zeros.

Lemma 5.15. *Fix $\varepsilon, C > 0$ and suppose $(\frac{1}{3}+\varepsilon)\frac{1}{n}\log n \leqslant p \leqslant \frac{C}{n}\log^2 n$. Then w.h.p. every row and column of M can be determined except possibly for rows and columns where there is an isolated* 1*.*

Proof. We can assume, by Lemma 5.14, that we have already placed all heavy rows and columns in their correct positions. Write $\mathcal{X}$ for the set of rows and $\mathcal{Y}$ for the set of columns that have been placed.

We note that we can place any row r which has a unique neighbourhood $N(r)\cap\mathcal{Y}$ in the already placed columns as we can determine this set from either the value or the index of r. If $N(r)\cap\mathcal{Y}$ is not unique but $|N(r)|$ is unique among rows r' with $N(r')\cap\mathcal{Y} = N(r)\cap\mathcal{Y}$ then we can still place r, as r is distinguished by its degree (and we can determine the degrees of rows in any position from $\mathcal{C}$). Moreover, we may assume that in the $k\times\ell$ submatrix A formed from the unplaced rows and columns, every row and every column contains a 1. Indeed, if A contained a zero row, corresponding to the row r of M, say, then $N(r) = N(r)\cap\mathcal{Y}$. But for any row r' with $N(r')\cap\mathcal{Y} = N(r)\cap\mathcal{Y}$, either $N(r') = N(r)$, in which case the row values of r and r' are identical, or $|N(r')| > |N(r)|$ and these rows can be distinguished by their degrees. Thus the position of any row value r is uniquely identified up to permutation of equal rows.

Now suppose the row r is still unplaced but $N(r)\cap\mathcal{Y}\neq\varnothing$. Then there must be another row r' with $|N(r)| = |N(r')|$ and $N(r)\cap\mathcal{Y} = N(r')\cap\mathcal{Y}$, but $N(r)\neq N(r')$. As $|N(r)| = |N(r')|$ there are distinct columns c, c' with $c\in N(r)\setminus N(r')$ and $c'\in N(r')\setminus N(r)$. Clearly $c, c'\notin\mathcal{Y}$ so are also unplaced. We may also assume that $N(c)\cap\mathcal{X} = N(c')\cap\mathcal{X}$ as otherwise r and r' could be distinguished by the multisets $[N(c)\cap\mathcal{X} : c\in N(r)\setminus$

$N(r')]$ and $[N(c) \cap \mathcal{X} : c \in N(r') \setminus N(r)]$ (as these multisets would then be disjoint), both of which can be identified from either the values or positions of r and r'.

We now bound the number of 4-tuples (r, r', c, c') which could satisfy these conditions. More specifically we count the number of 4-tuples (r, r', c, c') satisfying the following slightly weaker conditions:

(i) $c \in N(r) \setminus N(r')$ and $c' \in N(r') \setminus N(r)$;
(ii) $N(r) \cap N(r') \neq \varnothing$;
(iii) $N(r) \cap C = N(r') \cap C$ where $C = \{c'' \neq c, c' : |N(c'') \setminus \{r, r'\}| \geqslant \frac{1}{2}np\}$;
(iv) $N(c) \cap R = N(c') \cap R$ where $R = \{r'' \neq r, r' : |N(r'') \cap C| \geqslant \frac{1}{2}np\}$.

Note that all columns in C and all rows in R are heavy, so assumed already placed. Fixing (r, r', c, c'), we note that by a similar calculation as in Lemma 5.13, any column $c'' \neq c, c'$ lies in C with probability at least $1 - n^{-1/25}$ independently of (i) and (ii). Indeed, even conditioned on $N(r)$ and $N(r')$, for any $c'' \neq c, c'$,

$$\mathbb{P}(c'' \notin C) = \mathbb{P}\big(\operatorname{Bin}(n-2, p) < \tfrac{1}{2}np\big) \leqslant e^{-(1-o(1))np/8} \leqslant n^{-1/25},$$

independently for each c''. Let E be the event that $|C| < n - 2 - 2n^{24/25} = (1 - o(1))n$. Then as the number of $c'' \neq c, c'$ not in C is stochastically dominated by a $\operatorname{Bin}(n - 2, n^{-1/25})$ random variable, by Lemma 1.3,

$$\mathbb{P}(E) \leqslant e^{-(1/3)(n-2)n^{-1/25}} = n^{-\omega(1)}.$$

Now conditioned on $N(r)$, $N(r')$ and C, and assuming E does not hold, any row $r'' \neq r, r'$ lies in R with probability at least $1 - n^{-1/25}$. Indeed,

$$\mathbb{P}(r'' \notin R) \leqslant \mathbb{P}\big(\operatorname{Bin}(|C|, p) < \tfrac{1}{2}np\big) \leqslant e^{-(1-o(1))np/8} \leqslant n^{-1/25},$$

as an entry being 1 in row r'' is positively correlated with the condition that its column is in C.

Now, given (r, r', c, c'), the probability that (i) holds is $p^2(1-p)^2$. Conditioned on this, (ii) holds with probability at most np^2. Conditioned on (i) and (ii), (iii) holds with probability at most

$$\big((1-p)^2 + p^2 + 2p(1-p)n^{-1/25}\big)^{n-2} = e^{-2pn+o(1)} \leqslant n^{-2/3-2\varepsilon+o(1)}.$$

Conditioned on this the probability that (iv) holds but E does not occur is then at most

$$\big((1-p)^2 + p^2 + 2p(1-p)n^{-1/25}\big)^{n-2} = e^{-2pn+o(1)} \leqslant n^{-2/3-2\varepsilon+o(1)}.$$

Thus the expected number of such 4-tuples is at most

$$n^4 \cdot \left(\mathbb{P}(E) + p^2(1-p)^2 \cdot np^2 \cdot n^{-2/3-2\varepsilon+o(1)} \cdot n^{-2/3-2\varepsilon+o(1)}\right) = n^{-1/3-4\varepsilon+o(1)} = o(1).$$

Hence we may assume $N(r) \cap \mathcal{Y} = \varnothing$ for all unplaced rows r and similarly $N(c) \cap \mathcal{X} = \varnothing$ for all unplaced columns c.

Now consider the graph G restricted to the unplaced rows and columns. The above argument shows that we can assume this forms a union of components in G of total cardinality at most $o(n)$, where isolated vertices correspond to zero rows or zero columns and isolated edges correspond to isolated 1s in M. Thus is is enough to show that G contains no components with between 3 and $o(n)$ vertices. We count the expected number of such components by counting the number of possible choices of spanning trees for such components. We get that the expected number of these components is then at most

$$\begin{aligned}\sum_{k=3}^{o(n)} \binom{2n}{k} p^{k-1} k^{k-2} (1-p)^{k(n-k)} &\leqslant \sum_{k=3}^{o(n)} \left(2ne \cdot e^{-p(n-k)}\right)^k p^{k-1} \\ &\leqslant \sum_{k=3}^{o(n)} n^{1-(1/3+\varepsilon+o(1))k} = o(1). \qquad \square\end{aligned}$$

Remark 5.16. For $p < \frac{1}{3n} \log n$ another obstruction to reconstructibility appears, namely pairs of rows (or columns) each of which contains two 1s which themselves are the unique 1 in their column (or row). In graph terms these consist of at least two components that are isomorphic to a path on 3 vertices (with the central vertices both in the same bipartite class). In general more complex tree components on k vertices appear for $p < \frac{1}{kn} \log n$, and if two isomorphic copies of a tree T exist (with the isomorphism respecting the bipartition of G) then the matrix M fails to be reconstructible as we can interchange the vertices in one bipartite class of T with their counterparts in another copy of T without affecting the multisets of rows and columns. This does however affect the matrix M for $k \geqslant 2$. For $k = 1$ pairs of isolated vertices in the same class correspond to pairs of zero rows or zero columns which is the main obstacle to *strong* reconstructibility, but do not prevent reconstructing M as no edges are changed when they are swapped. It should be noted that isolated tree components are not the only obstacle to reconstructibility. For example, for $p < \frac{1}{4n} \log n$, we can have paths on 5 vertices with only the middle vertex possibly joined to other vertices. This corresponds to say two rows r, r' and three columns c, c', c'' with $N(c) = \{r\}$, $N(c') = \{r'\}$, $N(r) = \{c, c''\}$, $N(r') = \{c', c''\}$. In this case r and r' can be swapped giving a different matrix with the same multisets of rows and columns.

Proof of Theorem 5.1. If $(\frac{1}{3}+\varepsilon)\frac{1}{n}\log n \leqslant p \leqslant \frac{1}{2}$ then by Lemma 5.9 or Lemma 5.15 we can w.h.p. reconstruct M up to rows and columns with isolated 1s. If there is at most one isolated 1 then clearly we can reconstruct the whole of M. If there are two or more isolated 1s then we can't reconstruct M as permuting the rows containing these isolated 1s will give a different matrix with the same row and column multisets. Hence the result follows from Lemma 5.7, with the explicit constant in part *(c)* being as in Lemma 5.7.

If $p < (\frac{1}{3}+\varepsilon)\frac{1}{n}\log n$ then again by Lemma 5.7 we have that w.h.p. there are either two isolated 1s or the matrix has fewer than two 1s in total. □

5.4 A fast algorithm for reconstruction

Lemma 5.17. *There exists an algorithm that w.h.p. either identifies all rows and columns, or finds a pair of isolated* 1*s and takes* $O(n^2)$ *time.*

Proof. We will restrict our attention to the case when $p = O(\frac{1}{n}\log^2 n)$ as for larger p the result follows from [13].

Constructing a list of all neighbours of every row and column index and every row and column value in G_R and G_C takes $O(n^2)$ time as we have to scan every vector in $\mathcal{R}\cup\mathcal{C}$. As w.h.p. there are only $O(\log^2 n)$ neighbours of any index or value, constructing $\mathcal{D}_2(v)$ for every v then takes only $O(n\,\mathrm{polylog}\,n)$ time. Sorting and finding matches between indices and values then again takes $O(n\,\mathrm{polylog}\,n)$ time. Identifying any isolated 1s also takes $O(n\,\mathrm{polylog}\,n)$ time. So unless $p \geqslant (\frac{1}{3}+\varepsilon)\frac{1}{n}\log n$ by Lemma 5.7 we will w.h.p. have terminated with either a pair of isolated 1s which demonstrates non-reconstructibility or a reconstructed matrix with at most one non-zero entry. We may thus assume $p \geqslant (\frac{1}{3}+\varepsilon)\frac{1}{n}\log n$ from now on.

The final part of the algorithm relies on calculating the multisets $[N(c)\cap\mathcal{X} : c \in N(r)]$ for rows r and similarly for columns and identifying rows or columns that are uniquely determined. Again, calculating these multisets takes only $O(n\,\mathrm{polylog}\,n)$ time. If this fails to identify all rows and columns then as in the proof of Lemma 5.15 w.h.p. the remaining rows and columns all contain isolated 1s, which we can easily check.

Thus overall the algorithm takes $O(n^2)$ time (with most of the time taken up in the initial scanning of rows and columns to find their neighbours) and correctly identifies M w.h.p. or finds a pair of isolated 1s showing that reconstruction is impossible. □

5.5 Conclusion

We have shown that there is a sharp threshold for reconstructibility at $p \sim \frac{1}{2n} \log n$, and a sharp threshold for strong reconstructibility at $p \sim \frac{1}{n} \log n$. These results both assume that we know *all* the rows and columns in $\mathcal{R}$ and $\mathcal{C}$. But what if we are missing some of the rows and columns? For example, suppose we are given $\mathcal{R}$, but only a (random) subset of cn elements from $\mathcal{C}$: for what range of c and p can we reconstruct M with high probability?

It would also be very interesting to investigate the problem when there are errors in our data. For example, suppose every entry in each row from $\mathcal{R}$ is given incorrectly with probability q: when can we reconstruct M with high probability? This seems to be interesting even when $p = 1/2$ and q is a small constant.

A similar question arises when we have missing data for both rows and columns: when can we reconstruct almost all of M? And, in a different direction, what happens if the data is corrupted adversarially?

Chapter 6

Reconstructing CAT(0) cube complexes from their boundary metric

This chapter is based on joint work with Haslegrave, Tan and Scott [64].

The reconstruction of higher-dimensional structures from lower-dimensional information has been an important area of research for many years. For example, the question of whether a Riemannian manifold with boundary is determined by its spectrum was popularised in a famous article of Mark Kac [77]; and there is a huge body of research on reconstructing discrete objects from their projections [66].

A particularly natural question of this type is whether the internal structure of an object can be determined from distances between boundary points. In Riemannian geometry, the notion of reconstruction from a distance function on the boundary of a geometric object is well-established in the realm of boundary rigidity questions. Broadly, a Riemannian manifold (M, g) is said to be *boundary rigid* if its associated metric d_g (which is defined on any two points, including the interior) is determined up to isometry by its boundary distance function given by the restriction $d_g|_{\partial M \times \partial M}$. In 1981, Michel [91] conjectured that every simple compact Riemannian manifold with boundary is boundary rigid. The 2-dimensional case was verified by Pestov and Uhlmann [99]. In higher dimensions, however, the conjecture is wide open and has only been verified for a few classes [21, 29].

There has been less work on analogous questions for discrete structures. Haslegrave, answering a question of Benjamini [19], proved the following result in two dimensions.

Theorem 6.1 (Haslegrave [63])**.** *Let Q be a planar quadrangulation with a simple closed boundary such that all internal vertices have degree at least* 4*. Then the distances between*

boundary vertices determine Q up to isomorphism.

Here, the distances are taken to be in the graph metric and the condition on the boundary of Q can be restated by saying that Q is (isomorphic to) a planar quadrangulation of a disc. Moreover, the degree condition is necessary for reconstruction; for instance, one can always 'hide' a square inside another by choosing an existing face of a planar disc quadrangulation, adding a new neighbour inside the face for each of the four vertices bounding the face, and joining them up so as to split the face into five new faces, creating a different quadrangulation with the same boundary distances.

Theorem 6.1 can be viewed as a discrete analogue of the 2-dimensional boundary rigidity result of Pestov and Uhlmann [99]. The discrete case should be more approachable than the continuous one in general. Indeed, this can be seen for instance in the fact that much stronger restrictions on the boundary are required in the latter, and that the proof of Theorem 1 even provides an efficient algorithm for reconstructing Q from the given data which is not expected in the continuous case.

In this chapter, we look at generalising Theorem 6.1 to higher dimensions where the natural counterpart for a quadrangulation is a cube complex – just as a quadrangulation can be formed by gluing Euclidean squares (or 2-dimensional cubes) along edges, a k-dimensional cube complex is informally a complex formed by gluing together cubes of dimension at most k along subcubes. This leads to the following question.

Question 6.2. *Under what conditions is a finite k-dimensional cube complex $X \hookrightarrow \mathbb{R}^k$ reconstructible up to combinatorial type from its boundary distances?*

Question 6.2 requires not only determining the full 1-dimensional structure from boundary distances alone, but also reconstructing the higher-dimensional structure from the 1-dimensional. This is not really true of Theorem 6.1: while the goal is to reconstruct a 2-dimensional complex, any polyhedral graph has a unique embedding in the sphere [120], and hence a unique embedding in the plane with a designated outer face, so the second step is immediate in this case.

We provide an answer to Question 6.2 for 3-dimensional complexes, with a well-studied condition which directly generalises the one stated in Theorem 6.1. We require complexes to satisfy the CAT(0) property, which entails both a global topological condition (simply connectedness) and a local negative curvature condition (Gromov's link condition, which states that the link of every vertex is a flag complex). Intuitively, the latter condition is natural as it allows us to fill in boundaries of cubes in the 1-skeleton with actual cubes, as required for the second step mentioned above. Analogously to CAT(0) spaces, CAT(0) cube complexes form a large, popular class of complexes possessing

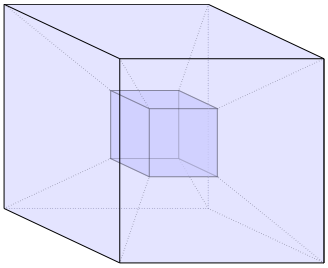


Figure 6.1. A cube hidden within another. Links of vertices of the 'hidden' cube are not flag: four cubes meet at each such vertex, but the complex has no cells of dimension 4.

useful convexity properties (see Section 6.3). This makes them a natural choice of setting for boundary rigidity problems. Our main theorem is the following.

Theorem 6.3. *Suppose that X is a finite CAT(0) cube complex admitting an embedding in $\mathbb{R}^3$, with a labelling of vertices in ∂X. Let D be the matrix of pairwise distances between vertices of ∂X with respect to the graph metric on the 1-skeleton of X. Then the combinatorial type of X is reconstructible from D.*

Like in the 2D case, our proof provides an efficient algorithm for reconstructing the combinatorial type of X, and the labelling of boundary vertices is preserved throughout this process.

Both the simply connectedness and flag conditions are used essentially in numerous places throughout the proof. It is also possible to see directly that they cannot be omitted from the statement. If we do not require links to be flag, then one could 'hide' a cube inside another as in Figure 6.1: geodesics between vertices on the outer cube are unaffected by the presence of the inner cube. This is a 3-dimensional analogue of 'hiding' a square within another in a quadrangulation of the disc, as described earlier. Another example is given by taking a $3 \times 3 \times 3$ block of cubes and considering the cube complexes formed by removing the top two cubes in the centre column and by removing the top and bottom cubes in the centre column (depicted in Figure 6.2). These two cube complexes do not satisfy the flag condition at any vertex of the (possibly missing) middle cube, and it is easily seen that they have the same boundary distances since all vertices are on the boundary and both complexes have the same edges. Similarly, we can see that contractibility, which implies simply connectedness, is necessary as it would be impossible to differentiate between a single square with or without a face only from the boundary distances.

CAT(0) cube complexes are ubiquitous in modern geometric group theory. While the importance of the condition in the context of reconstruction may not be immediately clear, it does in fact directly generalise earlier conditions. To see this, we note that

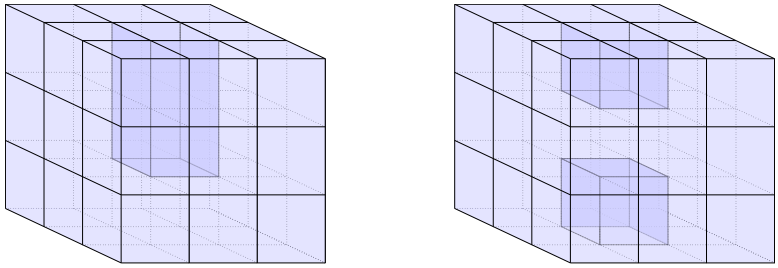


Figure 6.2. Two pure cube complexes with the same boundary distances. The dark shaded cubes are removed.

Gromov's link condition reduces to the degree condition of Theorem 6.1 for cube complexes of top dimension at most 2, and local negative curvature is also one of the key assumptions used by Besson, Courtois and Gallot in the continuous setting [21]. Furthermore, the fact that 2-dimensional cube complexes in Theorem 6.1 are contractible is also captured in the CAT(0) condition, which follows from the Cartan–Hadamard theorem. In fact CAT(0) complexes have the even stronger property of collapsibility (see [3]).

Question 6.2 asks whether one can recover combinatorial information of a cube complex X from some partial combinatorial information, namely the distance (in the graph of the entire complex) between any two vertices on the boundary. In order for this question to be well-defined, we first need to make precise the notion of *boundary*. In the present chapter we mostly work with the natural extrinsic notion of *geometric boundary*: given an embedding $X \hookrightarrow \mathbb{R}^k$, we define ∂X to be the topological (induced from the Euclidean metric) boundary of X. This notion implicitly depends on the dimension k of the space that we embed X into. Indeed, if X has no cells of dimension at least k then it is its own boundary in any embedding $X \hookrightarrow \mathbb{R}^k$, so reconstructing X from boundary information is trivial. Hence, we always consider k-dimensional cube complexes embedded in k-dimensional Euclidean space.

A convenient observation is that for finite cube complexes the geometric boundary is independent of the embedding we choose, so long as one exists; indeed, for finite cube complexes with an embedding in $\mathbb{R}^k$ the geometric boundary admits an intrinsic description. In Section 6.1, we will discuss a combinatorial notion of boundary for cube complexes that provides such a description, which we then use in Section 6.8. Since this combinatorial boundary does not require an embedding, it suggests a natural generalisation of Question 6.2: can we reconstruct CAT(0) cube complexes which do not necessarily admit embeddings in a Euclidean space of their top dimension? We give a positive answer for cube complexes of top dimension at most 2, providing in particular a strengthening of Theorem 6.1.

Theorem 6.4. *Let X be a CAT(0) cube complex of top dimension at most 2 with finitely many cells and D its matrix of pairwise distances between vertices on the combinatorial boundary of X. Then, the combinatorial type of X is reconstructible from D.*

Extensions

Our results prompt the natural question of whether the CAT(0) property is also necessary and sufficient for reconstructing higher dimensional cube complexes. In [64], we conjectured that this is the case, and moreover, in light of Theorem 6.4, that it is not necessary to have an embedding in a Euclidean space of the same dimension as the complex.

Conjecture 6.5. *For any $k \geqslant 3$, any finite CAT(0) cube complex can be reconstructed up to combinatorial type from its boundary distances.*

After the appearance of the results from this chapter as a preprint, Chalopin and Chepoi [33] verified Conjecture 6.5 using an argument based on properties of median graphs.

Another direction is to consider the situation in the world of simplicial complexes.

Question 6.6. *Under what conditions is a finite k-dimensional simplicial complex reconstructible up to combinatorial type from its boundary distances?*

The case of 2-dimensional complexes embeddable in $\mathbb{R}^2$ has been dealt with in [63], where it is shown that all internal vertices having degree at least 6 is a sufficient condition.

In recent work in that direction, Blufstein, Chalopin and Chepoi [22] further extended Conjecture 6.5 to systolic and Helly complexes. Systolic complexes are simplicial complexes whose 1-skeleta are bridged graphs, namely graphs where all isometric cycles have length three. These complexes are often thought of as simplicial analogues of CAT(0) cube complexes, so this result gives a partial answer to Question 6.6. Helly graphs are graphs in which balls satisfy the Helly property, namely that any set of balls whose pairwise intersections are non-empty has non-empty intersection; they are discrete analogues of so-called hyperconvex spaces.

Organisation

This work makes significant use of notions and tools from algebraic topology and from the theory of CAT(0) cube complexes. Since there is a good deal of terminology involved, we postpone technical discussions and first introduce notation and necessary

theory in Section 6.1. With this background in hand, we give a skeletal version of the proof of Theorem 6.3 in Section 6.2. This provides a roadmap for Section 6.4 through Section 6.7, which are devoted to the different aspects of the main proof, while Section 6.3 collects and proves basic technical lemmas needed to make our arguments rigorous. The reader may wish to skip this section at first and use it as a reference. In Section 6.8, we provide a brief sketch of the proof of Theorem 6.4.

6.1 Definitions

In this section, we define the key objects and terminology that we will be working with. This will then allow us to give a broad outline of the proof of Theorem 6.3.

We start with some standard topological notions. For simplicity of exposition, we will restrict some definitions to the cases that we require, although they may exist in much greater generality. We refer to [65] for a detailed account of the concepts from point-set and algebraic topology, and [121] for specifics on cube complexes.

A *CW complex* X is an object that is obtained by taking an infinite union of its k-skeleta, where X^k denotes the *k-skeleton* of X for each integer $k \geqslant 0$, that is, all cells of X of dimension at most k. These are defined inductively by starting with a set X^0 of 0-cells, and for each k constructing the $k+1$ skeleton $X^{k+1} = X^k \bigcup \cup_{\phi_\alpha} D_\alpha$ by gluing each $(k+1)$-dimensional cell D_α onto X^k via a gluing map $\phi_\alpha \colon \partial D_\alpha \to X^k$. We say that a finite CW complex X is *k-dimensional* if cells of top dimension in X are of dimension k.

All CW complexes we consider are *regular*, meaning that their gluing maps are homeomorphisms. Our notation and descriptions below are chosen to reflect the fact that we will require a mix of combinatorial and geometric properties of the objects in question.

We write I for the unit interval $[0,1]$. Let S^k and B^k be the unit sphere and ball in $\mathbb{R}^k$ respectively, i.e.

$$S^k = \{(x_1, x_2, \ldots, x_k) \in \mathbb{R}^k : x_1^2 + \ldots + x_k^2 = 1\},$$
$$B^k = \{(x_1, x_2, \ldots, x_k) \in \mathbb{R}^k : x_1^2 + \ldots + x_k^2 \leqslant 1\}.$$

Geometric boundary

Let X be a regular CW complex whose gluing maps are isometries. Given an embedding $X \hookrightarrow \mathbb{R}^n$, the *(geometric) boundary* of X, denoted ∂X, is the set of points in X for which every neighbourhood intersects both X and $\mathbb{R}^n \setminus X$.[1] The *interior* of X is then

[1] Note that this is an instance of the standard definition of topological boundary.

$\operatorname{int} X := X \setminus \partial X$. Explicitly, this is the set of points $p \in X$ such that the ball $B_\varepsilon(p)$ is contained in X for some sufficiently small $\varepsilon > 0$, where $B_\varepsilon(p)$ denotes the *ball of radius ε centered at p*.

If finiteness of the complex is not assumed, the geometric boundary may depend on the chosen embedding. For example, let Z be the 2-dimensional complex corresponding to the graph with vertex set $\mathbb{N} \cup \{0\}$ and edge set $\{(0,k),(k-1,k) : k \in \mathbb{N}\}$, and whose 2-cells have vertex sets of the form $\{0,k,k+1\}$. One embedding of Z into $D \subseteq \mathbb{C}$ can be obtained by gluing 2-cells to the sectors bounded by consecutive vectors (viewed as edges) from $\{e^{i\pi(1+\sum_{j=0}^{k} 2^{-j})} : k \in \mathbb{N}\}$. The geometric boundary from this embedding is the preimage of the unit circle, so in particular does not contain the (preimage of the) edge $e^{i\pi}$. Yet one can embed Z into $\mathbb{C}$ so that the preimage of this edge is on the geometric boundary by gluing 2-cells to the sectors bounded by consecutive vectors from $\{e^{i\pi\sum_{j=0}^{k} 2^{-j}} : k \in \mathbb{N}\}$ (see Figure 6.3).

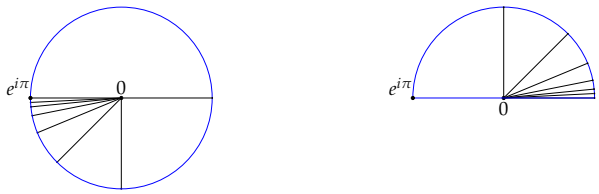


Figure 6.3. Two embeddings of the same (infinite) complex with differing geometric boundary (in blue).

For a regular CW complex X of maximum dimension k whose gluing maps are isometries, we define its *combinatorial boundary* to be the downward closure of the cells of dimension less than k in at most one cell of dimension k. In general, the geometric and combinatorial notions of boundary are different even for complexes admitting embeddings in $\mathbb{R}^k$, as the combinatorial boundary is independent of any embedding. For finite complexes however, these two notions coincide, so long as there exists an embedding in $\mathbb{R}^k$.

Proposition 6.7. *Let X be a regular k-dimensional CW complex whose gluing maps are isometries and which admits an embedding $X \hookrightarrow \mathbb{R}^k$. If X has a finite number of cells, then its combinatorial and geometric boundaries are the same.*

Proof. Let X be a finite regular CW complex of dimension k and fix an embedding in $\mathbb{R}^k$. Suppose that a cell S of dimension less than k is contained in at most one cell of dimension k. Suppose that a point x inside S (that is, in S but not in any lower dimension cell) is not in the geometric boundary. Since the complex is finite, there is

some minimum distance between x and the union of all k-cells not containing S. Take a ball of radius smaller than this. Now any point in the ball that does not intersect the $(k-1)$-skeleton must be in some k-cell. Furthermore, there must be at least two such cells involved, since if there is a unique such cell T then the whole ball is in (the closure of) T, and so, since gluing maps are isometries, x is in the interior of T, contradicting the choice of x. So all of the interior of S is on the geometric boundary. Since the geometric boundary is closed, it also contains all cells of S, so the geometric boundary contains the combinatorial boundary.

Now consider a point x that lies inside the combinatorial boundary. Suppose S is a $(k-1)$-cell contained in two k-cells, and let x be a point in the interior of S. Then a sufficiently small ball around x meets no cells other than these three. By passing to a smaller ball if necessary we can assume that it is divided into two parts by S, either of which contains interior points of the larger cells. Thus each part of the ball is contained in one of the two cells, and x is not on the geometric boundary.

Suppose $x \in S$ is in the geometric boundary, where S is a cell of dimension at most $k-2$ and is the inclusion minimal cell containing x. If S is not in the combinatorial boundary, then every $(k-1)$-cell containing S, of which there is at least one, is in two k-cells. Take a ball around x that is sufficiently small to avoid any $(k-1)$-cell not containing S. This ball contains a point y in the interior of some k-cell (since x lies in the closure of such a cell), and a point z outside the complex (since x is in the geometric boundary). Now the ball is path-connected, even if we remove the $(k-2)$-skeleton from it. Thus there is a path from y to z, which avoids the $(k-2)$-skeleton and must contain a point on the geometric boundary. By the previous paragraph, no point in the interior of a $(k-1)$-cell containing S is in the geometric boundary, a contradiction. □

In particular, this result applies to the cube complexes considered throughout this chapter. Henceforth, all CW complexes are assumed to have a finite number of cells.

Simplicial complexes

An *n-simplex* is an n-dimensional object formed by taking the convex hull of $n+1$ linearly independent vectors. Every n-simplex is homeomorphic to a *standard n-simplex* $\Delta^n := \{(x_0, \ldots, x_n) \subseteq \mathbb{R}^{n+1} : \sum_i x_i = 1 \text{ and } x_i \geqslant 0 \text{ for all } i\}$, which is spanned by the unit vectors along each coordinate axis. We say that $x_0, \ldots, x_n$ span the simplex Δ^n. Low-dimensional simplices are familiar objects: we will call 0-simplices *vertices*, 1-simplices *edges*, 2-simplices *triangles* and 3-simplices *tetrahedra*.

Recall that a *simplicial complex* S is a CW complex whose cells form a collection of simplices such that

- for every simplex in S, all of its simplicial faces are also in S, and
- the intersection of any two simplices in S is a simplicial face of both of them.

The *dimension* of a simplicial complex is the dimension of its top-dimensional simplices. We say that a simplicial complex S is *flag* if whenever there is a collection of k pairwise adjacent vertices (that is, joined by edges), then those k vertices span a $(k-1)$-simplex in S. Informally, this means that there is a k-simplex everywhere there should be one according to the graph of vertices and edges in the complex. Finally, note that the boundary of an n-dimensional simplicial complex has a natural structure as an $(n-1)$-dimensional simplicial complex.

Cube complexes

We now turn to *cube complexes*, which are CW complexes whose n-cells are n-cubes and gluing maps are combinatorial isometries, meaning that faces are glued to faces isometrically. Let the *standard n-cube* be the set of points $Q^n := I^n \subseteq \mathbb{R}^n$ (the standard 0-cube is a single point). Note that each n-cube is endowed with a natural internal coordinate system. By restricting any k of the coordinates to 0 or 1, we obtain an $(n-k)$-cube on the boundary of our n-cube which we call a *cubical face*. In general, an n-cube is any set in $\mathbb{R}^n$ that is homeomorphic to the standard n-cube. The *dimension* of a cube complex is defined analogously to that for simplicial complexes. An n-dimensional cube complex is *pure* if every k-cube with $k < n$ is contained in an n-cube. Since we mainly work with 3-dimensional cube complexes, in the later sections of this chapter we will use the terms *vertices*, *edges*, *faces*, and *cubes* (with no specified dimension) to mean 0-cubes, 1-cubes, 2-cubes and 3-cubes respectively. As such, we will refer to the 1-skeleton X^1 as the *graph* of X, with the corresponding graph theoretic terminology. In particular, unless otherwise stated, a *path* γ in X is a graph path in its graph X^1 and its *length* $|\gamma|$ is its number of edges, and a *walk* is a path where vertices may be repeated. When γ consists of vertices $v_0, \dots, v_k$, in this order, we sometimes use the notation $v_0 \cdots v_k$ for γ.

For $k, \ell \geqslant 0$ we say that an ℓ-cube in a k-dimensional cube complex is *free* if it is not contained in any k-cube, this being one possible structure in a cube complex that is not pure. We will mainly use this term to refer to free faces in 3-dimensional cube complexes, i.e. 2-cubes not contained in 3-cubes. A *cubulation of the ball* is a cube complex that is homeomorphic to B^3. In Section 6.6 we will see a construction which

requires fixing an embedding $X \hookrightarrow \mathbb{R}^3$; as such, when introducing a cube complex X we use this notation to indicate that we have fixed a specific embedding of X. Say that a vertex is the *corner of a cube (face)* if all cells containing it are contained in a unique cube (face).

In this chapter we are interested in the *graph metric* on the 1-skeleton of cube complexes, meaning the length of shortest paths between vertices. When it is unambiguous to do so we will talk about geodesics on cube complexes to mean geodesics on their 1-skeleton. As such, if X is a cube complex, the *1-distance graph induced by the boundary* $G(X)$ is the subgraph of X^1 induced by the vertices ∂X^0. Here and throughout the chapter we abuse notation and write ∂X^i for $(\partial X)^i$. This contains ∂X^1 as a subgraph, which may be a proper subgraph since edges not in ∂X can have endpoints on the boundary. Hence, we can read $G(X)$ off the restriction $D_{\partial X}$ of the distance matrix to the boundary, but not necessarily ∂X^1.

A map $f\colon X \to Y$ between CW complexes X and Y is said to be *combinatorial* if its restriction to the boundary of any cell of X is injective, and if it maps each cell to a cell of same dimension. Two cube complexes X and Y are said to have the same *combinatorial type* if there are bijections $f_i\colon X^i \to Y^i$ for each dimension i such that any two cells σ, τ of X are incident if and only if $f(\sigma)$ is incident to $f(\tau)$ in Y. In this chapter, we will mostly be interested in cube complexes up to combinatorial type, meaning that we consider them to be *distinct* when their combinatorial types differ.

There is an important construction which allows us to encode local structural information from a cube complex via an auxiliary simplicial complex. Given a cube complex X and a vertex $v \in X$, the *link* of v, denoted $\mathrm{link}(v)$, is the simplicial complex where:

- the vertices of $\mathrm{link}(v)$ are in bijection with edges containing v,
- for $n \geqslant 2$, there is an $(n-1)$-simplex with vertices $e_1, \ldots, e_n$ in $\mathrm{link}(v)$ if and only if there is an n-cell C in X containing v where $e_1, \ldots, e_n$ are the edges of C that contain v.

Intuitively, the simplices in $\mathrm{link}(v)$ correspond to 'corners' of cells in X that contain v. A useful alternative perspective, assuming that X is finite and embedded in Euclidean space, is that $\mathrm{link}(v)$ is the intersection of the sphere $S_\varepsilon(v)$ with X for sufficiently small $\varepsilon > 0$. This has a natural simplicial structure.

CAT(0) Cube complexes

At last, we arrive at the key property that we need for reconstruction. A cube complex X is *CAT(0)* if it is simply connected and $\mathrm{link}(v)$ is flag for every $v \in X^0$. The latter

part of this definition is really a condition requiring that the complex has nonpositive curvature, and in fact directly generalises the degree condition in Theorem 6.1. In a disc quadrangulation, the link of each boundary vertex is a path, while the link of an internal vertex is a cycle. Thus a disc quadrangulation is CAT(0) if and only if cycles in links have length at least 4, i.e. if and only if each internal vertex has degree at least 4. However, in three dimensions there is no corresponding equivalence: being CAT(0) implies that every internal vertex has degree at least 6, but the flag condition may fail at an internal vertex even if it has high degree, and it may also fail at a boundary vertex. The structure of links of vertices in CAT(0) cube complexes will be crucial for us. We continue this discussion in Section 6.3.

While there is a great deal of rich theory surrounding CAT(0) cube complexes – especially concerning their applications in geometric group theory, we will only need basic combinatorial considerations for our purposes and will refrain from delving deeper in the existing theory. Two important objects of study will be *immersed hyperplanes* and *disc diagrams*.

Hyperplanes and disc diagrams

For $n \geqslant 1$, a *midcube* in an n-cube $C = I^n$ is a codimension 1 cube M with an embedding $M = I^{n-1} \times \{1/2\}$ in C. As such, C has precisely n distinct midcubes, and the intersection of each midcube with a face of C of codimension at least 2 is again a midcube of that face. Moreover, two midcubes of distinct cubes in a cube complex X intersect in a combinatorial manner, meaning that the natural gluing map between the midcubes is combinatorial. In this way, midcubes of cubes of dimension at least 1 form connected components which we refer to as *hyperplanes*. We emphasize the distinction between a hyperplane H as a standalone cube complex and its natural embedding $H \to X$ as midcubes by referring to the latter as an *immersed hyperplane*. Then, the *(cubical) neighbourhood* of $H \to X$, written $N(H)$ or $H \times I$ (sometimes called the *carrier* of H), is the union of cubes in X containing it. We say that two immersed hyperplanes *cross* (in a complex X) if they contain two midcubes of some cube of any dimension (in this complex X); when a hyperplane crosses itself we say that it *self-crosses*.

A *disc diagram* is a locally injective combinatorial map $D \to X$, where D is a quadrangulation of the disc. We now define some pathological substructures in disc diagrams $D \to X$. In this setting, hyperplanes are 1-dimensional cube complexes, i.e. graphs. A *nonogon* is the neighbourhood of an immersed hyperplane which is a cycle. Bigons can be defined from two non self-intersecting immersed hyperplanes crossing each other at least twice in D: we define a *bigon* to be the cubical neighbourhood of two paths in

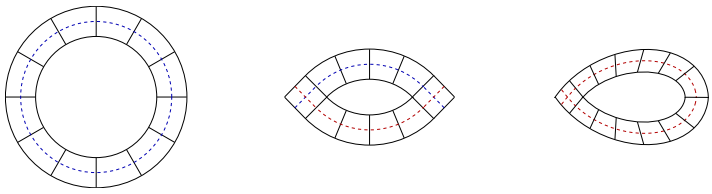


Figure 6.4. From left to right: a nonogon, a bigon and a self-intersecting hyperplane. The relevant immersed hyperplanes are represented in dashed lines.

such hyperplanes crossing each other exactly twice which are inclusion minimal with this property. Notice that these definitions also apply for CAT(0) cube complexes of dimension at most 2.

Both cells and simplices are specified by their vertices, so we will refer to a particular cell or simplex within a complex by a set of vertices. As for general CW complexes, if X is a simplicial complex (or cube complex), we write X^i for the *i-skeleton* of X which is the union of all k-simplices (k-cubes) for $0 \leqslant k \leqslant i$.

In this setting, a *minimal disc diagram* is understood to be, for a fixed cycle γ in X^1, a disc diagram $D \hookrightarrow X$ whose boundary is mapped to γ, chosen so that it minimises the number of faces, edges and vertices. Hyperplanes and minimal disc diagrams are particularly well-behaved in CAT(0) cube complexes, a statement which we make precise in Section 6.3.

6.2 Proof overview and discussion

This section provides, in a skeletal form, the proof of our main theorem. The main purpose is to provide a break-down of the proof into the components that span the remaining sections of this chapter, as well as discuss the necessity of certain approaches.

To prove sufficiency of the conditions in Theorem 6.3, we proceed by induction beginning with complexes with at most one edge. At each step, we aim to reduce the size of our cube complex: we have four processes that each use a particular substructure within the complex to define one or more smaller subcomplexes on which the induction hypothesis can be applied. This entails a number of verifications, namely that:

1) each substructure can be recognised from the boundary distance data that we start with,
2) all resulting subcomplexes still satisfy the CAT(0) property,
3) the boundary distances in all resulting subcomplexes can be recovered, and

4) if we are not in the base case, then at least one of the substructures exists in our finite CAT(0) cube complex so that a reduction can be performed.

Our four chosen substructures are cut-vertices, corners of faces, vertices of degree 3 that are not in any cube, and rows of cubes on the boundary. For the first three of these structures, there is a natural way to reduce our complex into smaller pieces and the corresponding verifications are relatively straightforward. These are detailed in Section 6.4, and allow us to proceed with the assumption that our complex X does not contain any of these three structures in which case we call X 'clean'.

The main work in our proof lies in handling rows of cubes on the boundary – essentially maximal stacks of cubes with one side on the boundary (the precise definition is given in Section 6.5). These structures are a natural choice for induction arguments in CAT(0) cube complexes because of their well-behaved hyperplanes. Here they work nicely in that they can easily be read off the boundary distance matrix, and their removal (for several natural definitions of removal) leaves a subcomplex where the flag condition is preserved at each vertex and their hyperplanes allow us to recover boundary distances to newly created boundary vertices. These properties are proved in Section 6.5. Within our proof, rows of cubes are key to making bulletpoint 4) above true. The intuition for this comes from the controlled case when X is a finite CAT(0) cube complex that is homeomorphic to a ball, where a simple Euler characteristic argument can be used to show that there must exist a row of cubes on the boundary of X.

Unfortunately, the property of being homeomorphic to a ball is not necessarily preserved when removing rows of cubes. See Figure 6.5 for instance: removing the central cube from three cubes glued together to form an 'L' shape. One could hope to reduce the resulting complexes by 'splitting' then appropriately into subcomplexes homeomorphic to balls, but this approach is complicated by the fact that CAT(0) cube complexes may possess 'essential' lower dimensional features, in the sense that removing these yields complexes with non-trivial homotopy. See for example Figure 6.6.

To overcome this, we use a removal process where we leave the 'back wall' of the row of cubes intact, thus ensuring contractibility. However this forces us to adapt our arguments for CAT(0) complexes which may not have fixed Euler characteristic. Our approach is to form an auxiliary complex by 'thickening' the complex X. This procedure, which is the topic of Section 6.6, produces a cube complex $\mathbb{X}$ homeomorphic to the ball B^3 that contains X. Roughly speaking, this is achieved by taking X together with a cubical shell around X. This shell is constructed in such a way that there is a correspondence between the graphs of $\partial\mathbb{X}$ and ∂X (Lemma 6.24). In particular,

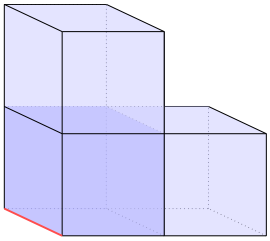


Figure 6.5. CAT(0) cube complex where removing a row of cubes gives a complex which is not homeomorphic to a ball: the dark blue cube is a row of cubes on the boundary (its red edge is a path of length 1 with endpoints and internal vertices of boundary degrees 3 and 4 respectively), yet removing it leaves two cubes sharing an edge.

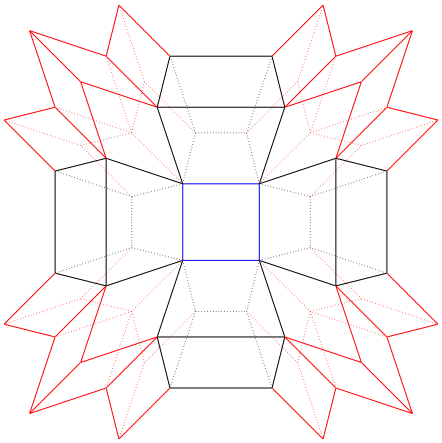


Figure 6.6. CAT(0) cube complex where removing the face bounded by the central blue square gives a complex which is not contractible.

this correspondence allows us to apply the previously mentioned Euler characteristic argument to $\mathbb{X}$ in order to find a row of cubes in $\mathbb{X}$ and then transfer it back to X.

Proof of Theorem 6.3. Let X be a contractible CAT(0) cube complex with an embedding in $\mathbb{R}^3$. We proceed by induction on the number of vertices of ∂X. We may assume without loss of generality that X is clean, as otherwise, by the discussion described in Section 6.4 we can recognise this and perform a reduction to reduce to a complex with fewer vertices on the boundary which satisfies the induction hypothesis.

By Lemma 6.27, the thickening $\mathbb{X}$ of X admits a good configuration $\mathcal{C}$. Under the correspondence described in Lemma 6.24, $\mathcal{C}$ corresponds to a good configuration $\pi(\mathcal{C})$ in X. This good configuration is a pattern in the boundary distance matrix which can be recognised, and by Lemma 6.17, $\pi(\mathcal{C})$ corresponds in turn to a row R of cubes on the boundary of X. Finally, applying the reduction described in Section 6.5, we reduce X to a complex with smaller boundary which satisfies the induction hypothesis. □

6.3 Technical toolbox

Hyperplanes in CAT(0) cube complexes

Let X be a cube complex and $Y \subseteq X$ a subcomplex. Recall that the metric we consider is the graph metric on 1-skeleta. In this setting we say that Y is *convex* if any geodesic with both endpoints in Y is entirely contained in Y.

A consequence of the CAT(0) property in cube complexes is the presence of natural convex subcomplexes, namely neighbourhoods of immersed hyperplanes. This convexity is crucial to recover boundary distances after removing parts of a CAT(0) cube complex, and we use it in the form of the following theorem, based on a paper of Sageev [108].

Theorem 6.8 (Wise [121, Theorem 2.13])**.** *Let X be a $CAT(0)$ cube complex.*

(i) *Each midcube lies in a unique immersed hyperplane.*

(ii) *Hyperplanes are $CAT(0)$ cube complexes.*

(iii) *The cubical neighbourhood $N(H) = H \times [0,1]$ of an immersed hyperplane H is a convex subcomplex.*

(iv) *$X \setminus H$ consists of two connected components.*

Links of vertices

The links of vertices in a complex contain important local information. We will mostly be interested in deducing information about the local structure around vertices from their degree and existing partial information.

Lemma 6.9. *Let X be a finite CAT(0) cube complex with an embedding in $\mathbb{R}^3$ and v a vertex of ∂X. Then:*

(a) *For any subcomplex $Y \subseteq X$ containing v, there is a natural containment map $\mathrm{link}_Y(v) \hookrightarrow \mathrm{link}_X(v)$. In particular, $\mathrm{link}_{\partial X}(v)$ has at least as many components as $\mathrm{link}_X(v)$.*

(b) *If v is contained in a cube C of X, then $\mathrm{link}_{\partial X}(v)$ has at least 3 vertices. If $\mathrm{link}_{\partial X}(v)$ is further a single cycle, then $\mathrm{link}_X(v)$ is homeomorphic to a disk D^2.*

(c) *If $\mathrm{link}_{\partial X}(v)$ does not contain a cycle, then the containment $\mathrm{link}_{\partial X}(v) \subseteq \mathrm{link}_X(v)$ is a bijection.*

(d) *Suppose $\mathrm{link}_{\partial X}(v)$ is a single cycle and $H \subseteq (\mathrm{link}_X(v))^1$ is a subgraph. This inclusion corresponds to an embedding of H in an ε-sphere about v in $\mathbb{R}^3$ which is a planar drawing for H. Suppose that the following hold:*

- *H is a triangulation of* $\mathrm{link}_{\partial X}(v)$*, meaning that* $(\mathrm{link}_{\partial X}(v))^1 \subseteq H$ *and this natural containment is a planar drawing of H such that all vertices of* $\mathrm{link}_{\partial X}(v)$ *lie on the outer face, and every other face is a triangle.*
- *The following diagram, where the maps are the aforementioned natural containments, commutes.*

$$\begin{array}{ccc} \mathrm{link}_{\partial X}(v) & \longrightarrow & H \\ \downarrow & \swarrow & \\ \mathrm{link}_X(v) & & \end{array}$$

Then the above embedding describes an isomorphism between H and $(\mathrm{link}_X(v))^1$.

A useful consequence of the second bullet point is that when X is homeomorphic to the ball B^3, the link of any of its boundary vertices is homeomorphic to a disc D^2.

Proof. (a) Now, the embedding $Y \hookrightarrow X$ gives a natural embedding $\mathrm{link}_Y(v) \hookrightarrow \mathrm{link}_X(v)$ through the identification between k-cells incident to v and $(k-1)$-cells in the links at v. In particular, the natural embedding $\partial X \hookrightarrow X$ guarantees that $\mathrm{link}_{\partial X}(v)$ has at least as many connected components as $\mathrm{link}_X(v)$: a path f (interpreted here as a continuous map $f\colon [0,1] \to \mathrm{link}_{\partial X}(v)$) in $\mathrm{link}_{\partial X}(v)$ extends to a path $f'\colon [0,1] \to \mathrm{link}_X(v)$ by precomposing with the inclusion map $\mathrm{link}_{\partial X}(v) \hookrightarrow \mathrm{link}_X(v)$ as the former is obtained from the latter by removing simplices.

Note that every component of $\mathrm{link}_X(v)$ intersects $\mathrm{link}_{\partial X}(v)$ somewhere as the elements of $\mathrm{link}_{\partial X}(v)$ form the topological boundary of $\mathrm{link}_X(v)$ in the ε-sphere at v. Indeed, any component of $\mathrm{link}_X(v)$ which is not the entire 3-sphere must contain a point x such that all δ-discs centered at x in the ε-sphere contain a point not in $\mathrm{link}_X(v)$, i.e. x is in the topological boundary of $\mathrm{link}_X(v)$.[2]

(b) Consider an ε-sphere around v in X. The connected component of $\mathrm{link}_X(v)$ containing the 2-simplex corresponding to C is 2-dimensional and thus – as $\mathrm{link}_X(v)$ is flag and therefore has no double edges – has at least three vertices on its boundary, as desired.

If $\mathrm{link}_{\partial X}(v)$ is a single cycle, $\mathrm{link}_{\partial X}(v) \subseteq \mathrm{link}_X(v)$ is homeomorphic to S^1 and so separates the ε-sphere into two components homeomorphic to discs D^2. One of these components is $\mathrm{link}_X(v)$, as desired.

(c) If $\mathrm{link}_{\partial X}(v)$ does not contain a cycle, then an ε-sphere about v is not disconnected by removing $\mathrm{link}_{\partial X}(v)$. It follows that $\mathrm{link}_X(v)$ lies entirely on the boundary.

[2]Here, the δ-disc centered at x is with respect to the euclidean metric on the ε-sphere.

(d) First, $\text{link}_{\partial X}(v)$ corresponds to a cycle in an ε-sphere about v, separating this sphere into two parts. Since H is a triangulation and $\text{link}_X(v)$ is flag, $\text{link}_X(v)$ contains at least one 2-simplex and thus (exactly) one of the two parts of the ε-sphere is contained in $\text{link}_X(v)$.

Since links of X are flag and $H \subseteq (\text{link}_X(v))^1$ is a triangulation, each triangle of the planar drawing of H bounds a 2-simplex in $\text{link}_X(v)$. Now, $(\text{link}_{\partial X}(v))^1 \subseteq H$ so $\text{link}_X(v)$ and the 2-simplices bounded by edges of H are both homeomorphic to a disc D^2 with boundary $\text{link}_{\partial X}(v)$, the only difference being that $\text{link}_X(v)$ may contain subdivisions of 2-simplices bounded by edges of H. This in particular gives a planar drawing of H.

Suppose now that H is a proper subcomplex of $(\text{link}_X(v))^1$. Since $(\text{link}_{\partial X}(v))^1 \subseteq H$ and every triangle of H bounds a 2-simplex in $\text{link}_X(v)$, there must be a face F of H that is triangulated in $(\text{link}_X(v))^1$, meaning that there is a single vertex adjacent to every vertex in the triangle. But such a vertex forms a clique of size 4 in the graph of $\text{link}_X(v)$ which, since X is $CAT(0)$, implies that there is a 4-dimensional cell in X, a contradiction. □

Recall that $G(X)$ is the 1-distance subgraph of X induced by the vertices of ∂X. These facts allow us to diagnose structures appearing in X from adjacencies in $G(X)$ and partial information on the structure of X. We elaborate on this in the next lemma.

A *cut-vertex* in X is a vertex $v \in X^0$ such that $X \setminus v$ has at least two non-empty connected components. Recall that a *corner of a cube* in X is a vertex v contained in a unique cube C of X. In particular, $v \in X^0$ and $\deg_X(v) = 3$ and the three faces of C incident to v lie on ∂X.

Lemma 6.10. *Let X be a finite CAT(0) cube complex with an embedding in $\mathbb{R}^3$ and $v \in \partial X^0$.*

(a) *If a vertex v is a cut-vertex of X, then $\text{link}_{\partial X}(v)$ is disconnected.*

(b) *If $\text{link}_{\partial X}(v)$ is a triangle, then v is a corner of a cube.*

(c) *If $\deg_{G(X)}(v) = 4$ and v is in a cube of X, then $\deg_{\partial X}(v) = 4$, which implies that $\text{link}_{\partial X}(v)$ is either a cycle of length 4, a triangle with a pendant edge, or a triangle plus an isolated vertex. Moreover, one of the following holds:*

- *v is incident to exactly one cube and one free face of X;*
- *v is incident to exactly one cube and one edge of X not contained in any face;*
- *v is incident to at least four cubes of X, and the four faces incident to v on ∂X each lie in different cubes; or*

- *v is incident to exactly two cubes in X, each of which contains two of the four faces incident to v on ∂X.*

Proof. (a) If $\mathrm{link}_{\partial X}(v)$ is connected, then by Lemma 6.9 (a) so is $\mathrm{link}_X(v)$. Hence $X \setminus v$ is also connected.

(b) By Lemma 6.9 (d) with $H = \mathrm{link}_{\partial X}(v)$, $\mathrm{link}_{\partial X}(v) \simeq (\mathrm{link}_X(v))^1$ and so $\mathrm{link}_X(v)$ is a single triangle as desired.

(c) Suppose for contradiction that $\deg_{\partial X}(v) \neq 4$. Then, $\deg_{\partial X}(v) \leqslant 3$.

Since v is in a cube of X, by Lemma 6.9 (b) $\mathrm{link}_{\partial X}(v)$ has exactly 3 vertices. In particular, $\mathrm{link}_{\partial X}(v)$ is connected: otherwise each connected component of $\mathrm{link}_{\partial X}(v)$ would have fewer than 3 vertices and hence no cycles, this would lead to a contradiction in view of Lemma 6.9 (c) as $\mathrm{link}_X(v)$ contains at least one 2-simplex. Hence, by Lemma 6.9 (d) v is then the corner of a cube, so $\deg_{G(X)} = 3$, a contradiction.

If $\mathrm{link}_{\partial X}(v)$ does not contain a cycle then Lemma 6.9 (c) contradicts the fact that v lies in a cube. Thus $\mathrm{link}_{\partial X}(v)$ is either a 4-cycle, a 3-cycle with a pendant edge, or a 3-cycle plus an isolated vertex.

In the first case, if two faces incident to v lie in the same cube, then the corresponding edges in $\mathrm{link}_X(v)$ necessarily share an endpoint, and lie in a triangle. Hence, $(\mathrm{link}_X(v))^1$ contains a 4-cycle with an extra edge. By Lemma 6.9 (d), $\mathrm{link}_X(v)$ is a 4-cycle with an extra edge, with each triangle bounding a 2-simplex. Otherwise the four faces of ∂X containing v lie in four different cubes.

In the second case, the triangle bounds a 2-simplex and the final edge corresponds to a free face.

Similarly, in the last case the triangle bounds a 2-simplex and the remaining vertex corresponds to an edge not contained in a face. □

We conclude this technical section with a standard observation that in CAT(0) cube complexes, cycles of length 4 bound a face.

Lemma 6.11. *Let w, x, y, z be vertices forming a square in the graph of X. Then these vertices lie in a face of X.*

Proof. Let C be the square they form. Since X is simply connected, by van Kampen's theorem (see [121, Lemma 3.1]) there is a disc diagram $D \to X$ with $C \simeq \partial D$. Let D be such a disk, chosen to minimise its number of faces. Under such minimality assumptions, D contains in particular no nonogons, no bigons and its hyperplanes do not self-cross (see [121, Lemma 3.2]). Since C has only 4 edges and any hyperplane

crosses the boundary twice, D admits at most two hyperplanes and hence has at most one face: if two distinct faces share an edge, their four midcubes belong to three distinct hyperplanes of D, as they would otherwise force a forbidden structure in D. □

In particular, pairs of vertices with two common neighbours correspond to faces of X. As a consequence, we can easily find the neighbours of a vertex $v \in \partial X$ which lie in the same face: they are precisely those that have a common neighbour in $G(X)$ other than v. Hence, from the 1-skeleton of ∂X we can recover the link in ∂X of every vertex in ∂X.

6.4 Cleaning

Let X be a finite CAT(0) cube complex. Recall that a vertex is the *corner of a cube* if it has degree 3 in X and is contained in a unique cube in X. Analogously, a *corner of a face* is a vertex with degree 2 in X that is contained in a unique face in X. Recall that *cut-vertex* in X is a vertex $v \in X^0$ such that $X \setminus v$ has at least two non-empty connected components. Each such connected component together with v has fewer boundary vertices than X.

In this section, we describe substructures with their recognition and reduction steps in CAT(0) cube complexes. In the order that they will be performed, the cleaning operations are the following:

(1) removing cut-vertices;
(2) removing corners of faces;
(3) removing vertices of degree 3 that are not in a cube.

For (1), the idea of reduction is that if v is a cut-vertex then we will try to apply the induction hypothesis to each connected component of $X \setminus v$ with v added back. If v is one of the features in (2) and (3), then we will apply the induction hypothesis to $X - v$. In order to apply induction, it is important to note that each of the above reduces the number of boundary vertices by at least one. A CAT(0) cube complex with none of the above features is called *clean*.

Note that the order of our cleaning operations is important in the sense that when we go through the steps for a later structure in the list, we sometimes need the assumption that none of the earlier structures are present. Likewise, it is important that we can later assume that our complex is clean to then show that there exists a row of cubes on the boundary.

Removing cutvertices

Recognition. As shown below, a vertex $v \in \partial X^0$ is a cut-vertex in X if and only if it is a cut-vertex in the graph of ∂X^1.

Proposition 6.12. *Let $X \hookrightarrow \mathbb{R}^3$ be a CAT(0) cube complex. Then $X \setminus v$ is connected if and only if $\partial X \setminus v$ is connected.*

We stress that the CAT(0) condition is necessary: consider for instance any cubulation of the space $Y = \{(x, y, z) : x^2 + y^2 + z^2 \leqslant 4 \text{ and } (x-1)^2 + y^2 + z^2 \geqslant 1\}$ consisting of the points between two nested spheres meeting at the single point $p = (2, 0, 0)$. Note that Y is not contractible. Then $Y \setminus p$ is connected while $\partial Y \setminus p$ is disconnected.

Proof. For any cube complex Z and vertex $v \in Z$, note that $Z \setminus v$ is connected if and only if $Z^1 \setminus v$ is connected. Hence, it suffices to show the assertion for 1-skeleta: that $X^1 \setminus v$ is connected if and only if $\partial X^1 \setminus v$ is connected.

If $X \setminus v$ is disconnected, then clearly $\partial X \setminus v$ is disconnected. For the converse, suppose that $\partial X \setminus v$ has at least two connected components. Let A be such a component and write $B := \partial X \setminus (v \cup A)$. Note that since the restriction of an immersed hyperplane of X to ∂X is connected (it is a closed walk) and avoids X^0, it cannot intersect both A and B. At the same time, it must intersect one of the two. Consider two immersed hyperplanes $H, H' \hookrightarrow X$ such that H and H' restricted to the boundary of X are contained in A and B respectively. By Theorem 6.8 the hyperplanes H, H' are CAT(0) cube complexes, and further their top dimension is at most 2. Recall that nonogons in a disc diagram are hyperplanes whose image is a cycle, and that a self-crossing hyperplane is one which contains two midcubes of some face.

Claim 6.13. *The immersed hyperplane H contains no nonogons.*

Proof. Since hyperplanes of CAT(0) cube complexes are themselves CAT(0) cube complexes (Theorem 6.8 (ii)), they are simply connected and in particular contain no nonogons. ■

Claim 6.14. $H \cap H' = \varnothing$.

Proof. Suppose not. Then, since the cubical neighbourhood of $H \cap H'$ in H is a sequence of faces where consecutive ones share edges, there is a disc diagram $D \hookrightarrow H$ which contains $H \cap H'$. As H has no nonogons and is finite, $H \cap H'$ must intersect ∂H non-trivially, contradicting that $\partial X \cap H' \subseteq B$. ■

Let $a \in A$ and $b \in B$ be neighbours of v and suppose for contradiction that there is an (a,b)-path $\gamma \subseteq X^1 \setminus v$. As X is contractible, the closed loop formed by γ together with v bounds a disc diagram $D \hookrightarrow X$. Let $F_1, \dots, F_k$ and $av = e_0, \dots, e_k = bv$ be the clockwise ordering of faces and edges, respectively, incident to v in D. For each $0 \leqslant i \leqslant k$ write H_i for the hyperplane corresponding to the midcube of edge e_i. Then, H_{i-1} crosses H_i in face F_i for each $i = 1, \dots, k$. Since the restrictions to the boundary of H_0 and H_k lie in A and B respectively, this contradicts Claim 6.14 (with $H := H_0$ and $H' := H_k$). □

Reduction. For this step, we use a simple fact from algebraic topology which we prove for completeness. An alternative argument using van Kampen's theorem is also possible.

For a collection of CW complexes (X_i), a *wedge* of the X_is is a complex obtained by choosing a point x_i from each X_i and identifying all chosen x_is. It is usually denoted $\bigvee_i X_i$.

Lemma 6.15. *Let $X = \bigvee_i X_i$ be a wedge of CW complexes X_i with common point x_0. Then if X is contractible, so is each X_i.*

Proof. The map $\pi_n(\bigvee_i X_i) \to \pi_n(\prod_i X_i) \simeq \bigoplus_i \pi_n(X_i)$ induced by inclusion is surjective for each n. By assumption the term on the left hand side is trivial and thus each $\pi_n(X_i)$ is trivial as well for each n. □

Suppose that v is a cut-vertex. Then X can be written as the wedge $\bigvee_i X_i$ of finitely many subcomplexes X_i whose pairwise intersection is $\{v\}$. By Lemma 6.15, each X_i is contractible (in particular, simply connected). Since each connected component of $\mathrm{link}_X(v)$ is a flag complex, so is each $\mathrm{link}_{X_i}(v)$ and thus each X_i is CAT(0).

It is clear that $\partial X_i \subseteq \partial X$ and X can be reconstructed from its subcomplexes X_i. Additionally, if vertices $x, y \in \partial X$ lie in the same subcomplex X_i, then $d_{X_i}(x,y) = d_X(x,y)$ as any x,y-path in X using a vertex not in X_i can be shortened to a path using only vertices in X_i.

Removing corners of faces

Recognition. If X has no cut-vertices, a vertex $v \in \partial X^0$ is the corner of a face precisely if $\deg_{\partial X}(v) = 2$. Indeed, if v is a corner of a face then $\deg_{\partial X}(v) = 2$ by definition. Conversely, if $\deg_{\partial X}(v) = 2$ then by Lemma 6.9 (b), v is not contained in any cube of X and so by Lemma 6.9 (c), $\mathrm{link}_{\partial X}(v) \simeq \mathrm{link}_X(v)$ and therefore is a single edge.

Reduction. Suppose that v is a corner of a face C. Note that C must be free, and hence C is in ∂X. Remove all cells containing v to obtain a proper subcomplex $Y \subseteq X$. As X

deformation retracts onto Y and X is contractible, Y is contractible. For a vertex u not in C, we have $\mathrm{link}_Y(u) = \mathrm{link}_X(u)$ which is therefore a flag complex. For a vertex w in C, $\mathrm{link}_Y(w)$ is obtained from $\mathrm{link}_X(w)$ by removing either a degree 1 vertex together with its incident edge, or an edge not contained in any triangle. In either case, the resulting complex $\mathrm{link}_Y(w)$ is flag and so Y is CAT(0). It is clear that X can be reconstructed from Y. We also have that $\partial Y \subseteq \partial X$ because all vertices of C are in ∂X. Moreover, $\deg_Y(x,y) = \deg_X(x,y)$ for any $x, y \in \partial Y^0$. To see this, let u be the vertex of C not adjacent to v and note that v can be replaced by u in any shortest (x,y)-path in X without changing its length.

Removing vertices of degree 3 not in a cube

Recognition. Suppose X has no cut-vertices nor vertices of degree 2, and $v \in \partial X^0$ has degree 3. Then v is not in a cube of X if and only if two of its neighbours do not have a common neighbour different from v. Indeed, since X has no cut-vertices, it follows from Lemma 6.10 (a) that $\mathrm{link}_{\partial X}(v)$ is connected and is therefore either a path with 2 edges or a triangle. The former case occurs precisely when v is the unique common neighbour of two of its neighbours, and in the latter case, by Lemma 6.10 (b), v is in a cube.

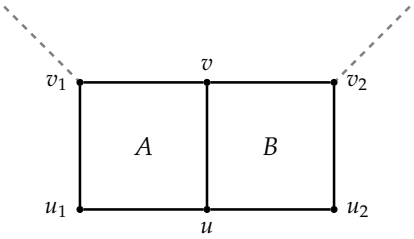


Figure 6.7. Pattern corresponding to a vertex v of degree 3 not in a cube.

Reduction. Suppose X has no cut-vertices nor vertices of degree 2, and that $v \in \partial X^0$ is a vertex of degree 3 not in a cube. This forces the cells containing v to form them pattern depicted in Figure 6.7, where all vertices in the figure are on ∂X. Let $Y \subseteq X$ be the proper subcomplex of X obtained by removing all cells containing v. Using the known pattern, we can reconstruct X from Y.

Since X deformation retracts onto Y via the map collapsing v onto u and the edges v_1v, vv_2 onto v_1u_1u, v_2u_2u respectively, Y is contractible. Moreover, since Y is obtained from X by removing cells of dimension at most 2 not contained in any cubes, all links of vertices in Y are flag complexes. Hence, Y is a CAT(0) cube complex.

Consider vertices $p, q \in \partial Y$; we certainly have $d_Y(p,q) \geqslant d_X(p,q)$. Let H be the unique immersed hyperplane splitting X into parts X_u, X_v with $u \in X_u$, $v \in X_v$. Note that we can recognise the part to which each vertex on ∂X belongs from the known distance matrix. By convexity of the neighbourhood of H, $p \in X_u$ if and only if $d_X(p,u) < d_X(p,v)$. There are three cases to consider to recover distances exactly.

1. If $p, q \in X_u$, then since hyperplanes are convex (Theorem 6.8 (iii)), no shortest p,q-path uses v. Hence, $d_Y(p,q) = d_X(p,q)$.
2. If $p \in X_u, q \in X_v$, we can assume $d_X(p,q) = d_X(p,v) + d_X(v,q)$ as every shortest p,q-path otherwise misses v.

   Since $v, q \in X_v$, either $d_X(v,q) = 1 + d_X(v_1,q)$ or $d_X(v,q) = 1 + d_X(v_2,q)$; assume without loss of generality the latter. As H is the only hyperplane separating u_2 and v_2, we have $d_X(p,u_2) = d_X(p,v_2) - 1 = d_X(p,v)$ and $d_X(u_2,q) = d_X(v_2,q) + 1 = d_X(v,q)$. Thus, there is a shortest (p,q)-path containing u_2. This path does not contain v, since a shortest (p,u_2)-path lies entirely within X_u and v is too far from q to be on a shortest (u_2,q)-path.
3. Finally, suppose that $p, q \in X_v$. To proceed, we claim that every shortest (p,q)-path in X_v is a shortest path in X. This holds since if ρ is a (p,q)-path that intersects X_u in a subpath ρ'_u, the convexity of the tubular neighbourhood of H means that ρ'_u is contained in this neighbourhood. Then we can replace ρ'_u by a projection of the same length ρ'_v in ρ, and this produces a (p,q)-path contained in X_v that is shorter than ρ.

   Now note that v is a cut-vertex in X_v. If p and q are in the same component of $X_v \backslash v$ (this can again be recognised from the distance matrix for ∂X), then the shortest p,q path in X_v avoids v. The preceding claim then implies that $d_Y(p,q) = d_X(p,q) = d_X(p,q)$.

   So suppose that p and q are in different components of $X_v \backslash v$, meaning a shortest (p,q)-path in X uses v. By Theorem 6.8 H is simply connected, so by van Kampen's theorem (see [65, Theorem 1.20]), X_v is also simply connected. Furthermore, since v is not contained in any face in X_v, any (v_1,v_2)-path P avoiding v would form a non-trivial loop with the path $v_1 v v_2$, which is impossible. Hence, any (p,q)-path γ in X^1 avoiding v – i.e. a path in Y^1 – must use at least one vertex outside X_v.

   Let xy be an edge of γ with $x \in X_v$, $y \in X_u$. Then, again by the convexity of the tubular neighbourhood of H, $d_X(p,y) \geqslant d_X(p,x) + 1$ and $d_X(y,q) \geqslant d_X(x,q) + 1$. Hence, $|\gamma| \geqslant d_X(p,y) + d_X(y,q) \geqslant d_X(p,x) + d_X(x,q) + 2$. In particular, taking

γ to be a shortest path allows us to conclude that $d_Y(p,q) \geqslant d_X(p,q)+2$. In fact, by replacing v by u_1uu_2 in γ, we see that $d_Y(p,q) = d_X(p,q)+2$.

6.5 Rows of cubes

A *row of cubes of length k* in X is a tuple of cubes $(C_1,\dots,C_k)$ from X where non-consecutive cubes are disjoint, and for each $i = 2,\dots,k-1$, there are opposite faces F, F' of C_i such that $C_{i-1}\cap C_i = F$ and $C_i \cap C_{i+1} = F'$. If moreover, there is a path $p_0\cdots p_k$ such that $\deg_X(p_0) = \deg_X(p_k) = 3$ and $\deg_X(p_i) = 4$ for $i = 1,\dots,k-1$, and additionally $p_0 \in C_1$, $p_k \in C_k$ and $p_i \in C_{i-1}\cap C_i$ for each $i = 2,\dots,k$, then we say that $(C_1,\dots,C_k)$ is *on the boundary*. In particular, this implies that $p_0,\dots,p_k$ are boundary vertices.

In this section, we describe the recognition and reduction steps for rows of cubes on the boundary. For the former, we will introduce so-called 'row configurations' and 'good row configurations' in Section 6.5. As will follow from Lemma 6.17, when X is a contractible clean CAT(0) cube complex, good row configurations in $G(X)$ correspond to rows of cubes on the boundary. Their existence in clean contractible CAT(0) cube complexes will be discussed in Section 6.6. The latter step is detailed in Section 6.5.

Row configurations

Definition 6.16. A *row configuration of length k* in a graph $H = (V,E)$ is a tuple of labelled vertices

$$(a_i, p_i, b_i \colon i = 0,\dots,k)$$

for some $k \geqslant 1$ such that $\deg_H(p_0) = 3$, $\deg_H(p_k) \neq 4$, $\deg_H(p_i) = 4$ for $i = 1,\dots,k-1$, all edges (referred to as the *edges of $\mathcal{C}$*)

$$\begin{aligned} a_ip_i, b_ip_i, &\qquad i = 0,\dots,k;\\ a_{i-1}a_i, b_{i-1}b_i, p_{i-1}p_i, &\qquad i = 1,\dots,k; \end{aligned}$$

are in E, the p_i are distinct and (a_i), (b_i) are distinct sequences. As it turns out, these sequences do not intersect in row configurations on the boundary of CAT(0) cube complexes. The *oriented* path $p_0p_1\dots p_k$ is called the *spine* of $\mathcal{C}$ and we refer to p_0 and p_k as the *start vertex* (which always has degree 3) and *end vertex* (which always has degree at least 3), respectively. We identify row configurations with the same spine, i.e. those of the form

$$\begin{aligned} &(a_i, p_i, b_i \colon i = 0,\dots,k)\\ &(b_i, p_i, a_i \colon i = 0,\dots,k). \end{aligned}$$

We refer to row configurations with end vertex of degree 3 as *good row configurations* or good configurations for short. When X is a cube complex, we will say 'a (good) row configuration in X' to mean a (good) row configuration in the graph ∂X^1.

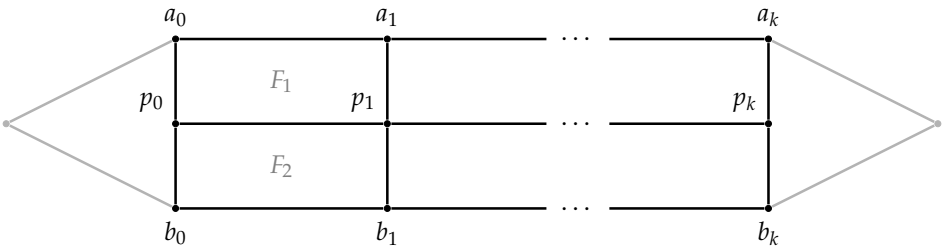


Figure 6.8. Row configuration.

When Y is a clean CAT(0) cube complex, row configurations in $G(Y)$ enjoy multiple useful properties: we show that all edges of a row configuration in $G(Y)$ are in fact edges of $(\partial Y)^1$ and that a row configuration corresponds to a row of cubes.

Lemma 6.17. *Let Y be a clean CAT(0) cube complex. Then a row configuration of length k in $(\partial Y)^1$ or $G(Y)$ corresponds to a row of cubes of length k on the boundary of Y. Namely, given a row configuration $\mathcal{C} := (a_i, p_i, b_i : i = 0, \dots, k)$ in $(\partial Y)^1$ or $G(Y)$, there is a row of cubes $\mathcal{R} := (C_1, \dots, C_k)$ in Y such that the faces bounded by $a_{i-1}, a_i, p_i, p_{i-1}$ and $p_{i-1}, p_i, b_i, b_{i-1}$ are faces of C_i for each $i = 1, \dots, k$, and all edges of $\mathcal{C}$ are in $(\partial Y)^1$. Further, this row of cubes is uniquely determined by k and the vertices p_0, p_1. We refer to $\mathcal{R}$ as the* underlying row of cubes *of $\mathcal{C}$.*

Note that a single row of cubes may be the underlying row of cubes of multiple row configurations in Y (e.g. when Y is a single cube). However, such row configurations are uniquely determined by the first edge of their spine.

Proof. Suppose that $\mathcal{C}$ is a row configuration in $G(Y)$. Since Y is clean p_0 is in a cube, so by Lemma 6.10 (b) and Lemma 6.9 (b), p_0 is the corner of a cube C_1 in Y and moreover edges incident to it in $G(Y)$ are present in $(\partial Y)^1$. In particular, C_1 has faces F_1, F_2 containing vertices $\{a_0, a_1, p_1, p_0\}$ and $\{b_0, b_1, p_1, p_0\}$, respectively. If $a_0 = b_0$, then F_1 and F_2 share three vertices and so, as Y is CAT(0), $F_1 = F_2$ and in particular $a_1 = b_1$. Continuing in this way we see that $(a_i) = (b_i)$, a contradiction. In a similar way, if there is some $0 < i \leqslant k$ with $a_i = b_i$ we can see that $a_0 = b_0$, leading to a contradiction. So without loss of generality, $a_0 \neq b_0$. As p_0 has degree 3 in Y, both F_1 and F_2 are in ∂Y and so a_0a_1, b_0b_1 are edges of $(\partial Y)^1$.

If $\deg_{G(Y)}(p_1) \neq 4$, we are done. Otherwise, $\deg_{G(Y)}(p_1) = 4$ and faces F_1, F_2 both contain p_1 and are faces of the same cube C_1. Furthermore, ∂Y^1 contains the 4-cycles

$p_1a_1a_2p_2$ and $p_1b_1b_2p_2$, which by Lemma 6.11 must each span faces of Y, and so $\mathrm{link}_{\partial Y}(p_1)$ contains a 4-cycle. The only possibility from Lemma 6.10 (c) consistent with this is that there is some cube C_2 such that p_1 is only incident to C_1, C_2 in Y, i.e. C_1, C_2 share a face containing p_1. Moreover, both faces of C_2 incident to p_1 are on ∂Y, and in particular, the edges of these faces are also on ∂Y. Let $\{a_1, a_2, p_2, p_1\}$ and $\{b_1, b_2, p_2, p_1\}$, respectively, be the vertices of the faces of C_2 incident to p_1. Continuing in this way we find the desired row of cubes $(C_1, \dots, C_k)$ with all edges of $\mathcal{C}$ on ∂Y. □

In light of this lemma, there is a one to one correspondence between row configurations in $G(Y)$ and $(\partial Y)^1$: a row configuration in $G(Y)$ is clearly one in $(\partial Y)^1$ since all its edges are in $(\partial Y)^1$, and conversely, given a row configuration $\mathcal{C}$ in $(\partial Y)^1$ labelled as above, we must have $\deg_{\partial Y}(p_i) = \deg_Y(p_i)$ for each i and hence $\mathcal{C}$ is a row configuration in $G(Y)$ as well. Moreover, the underlying row of cubes of a good configuration in $(\partial Y)^1$ is on the boundary, and rows of cubes on the boundary give rise to good configurations in $(\partial Y)^1$.

Removing a row of cubes

Let $\mathcal{C} = (a_i, p_i, b_i : i = 0, \dots, k)$ be a good configuration in ∂X^1 and $(C_1, \dots, C_k)$ its underlying row of cubes as in Lemma 6.17. We define $Y := X - \mathcal{C}$ to be the cube complex obtained from X by removing vertices $p_0, \dots, p_k$ as well as all cubes, faces and edges containing them. For each i, we denote the common neighbour of a_i and b_i by c_i. Note that ∂Y has at most $k-1$ new vertices $c_1, \dots, c_{k-1}$ and thus has at least one fewer vertex than ∂X.

We now check that Y admits an embedding in $\mathbb{R}^3$, is contractible and all links are flag. The first part is clear: an embedding $X \hookrightarrow \mathbb{R}^3$ induces an embedding $Y \hookrightarrow \mathbb{R}^3$. The second part is also straightforward as X deformation retracts onto Y and thus has the same homotopy type. For the third part, observe that for each $i = 1, \dots, k-1$, $\mathrm{link}_Y(c_i)$ is obtained from $\mathrm{link}_X(c_i)$ by removing an edge and the only two triangles containing it. When $i = 0$ or k then $\mathrm{link}_Y(c_i)$ is obtained from $\mathrm{link}_X(c_i)$ by removing an edge as well as the unique triangle containing it. In particular, in either case we have that $\mathrm{link}_Y(c_i)$ is flag, as desired. A similar, simpler argument shows that all $\mathrm{link}_Y(a_i)$s and $\mathrm{link}_Y(b_i)$s are also flag.

It remains to recover the boundary distances. To do so, we will use the following fact, which essentially says that the distance between a vertex p and a pair of adjacent vertices separated by a hyperplane H determines on which side of H the vertex p lies.

In fact, this lemma holds for general CAT(0) cube complexes, with essentially the same proof.

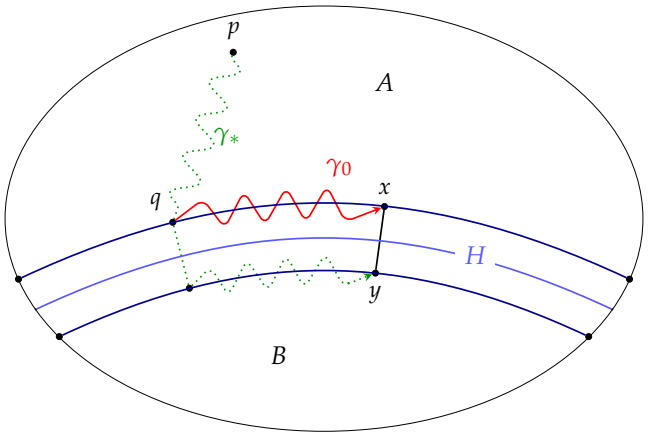


Figure 6.9. Sketch of proof of Lemma 6.18. If the dashed path γ_* is a geodesic, we can project its intersection with the tubular neighbourhood of H to the side of A, giving a subpath γ_0. This produces a (p,x)-path of length at most $d_X(p,y)-1$.

Lemma 6.18. *Let x, y be adjacent vertices of X and $H \hookrightarrow X$ the immersed hyperplane containing the midcube of the edge xy. Let A, B be the connected components of $X \setminus H$, containing x and y respectively. Then, for every $p \in X$,*

$$d_X(p,y) = d_X(p,x) + 1 \Leftrightarrow p \in A.$$

Proof. Suppose first that $p \in A$.

Given a geodesic path from p to x in X^1, we can extend it using the edge xy to form a (p,y)-path of length $d_X(p,x)+1$, and hence $d_X(p,y) \leqslant d_X(p,x)+1$.

For the reverse inequality, consider a geodesic path γ_* from p to y in X^1. Let q be the first intersection point of γ_* with the tubular neighbourhood $N := H \times [0,1]$ of H, where $H \times \{0\} \subseteq A$, $H \times \{1\} \subseteq B$. In particular, if $(h,1) = y$ for $h \in H$ then $(h,0) = x$. The subpath γ of γ_* from q to y has both endpoints in the tubular neighbourhood of H and so, by Theorem 6.8, lies entirely in it.

We show that the (p,y)-path obtained from γ_* by 'pushing' γ to $H \times \{0\}$ uses at least one less edge than γ_*. More precisely, let $\gamma_0 := \{(h,0) : (h,0) \in N\}$ be the projection of γ to $H \times \{0\}$ and define γ' to be γ_* from x to q concatenated with γ' and xy. The length of γ_0 is at most the number of edges of γ which do not cross H, so $|\gamma_0| < |\gamma|$ as γ has at least one edge crossing H. In particular, $|\gamma'| < |\gamma|$ and therefore $d_X(p,x) \leqslant d_X(p,y) - 1$, as desired.

Applying the symmetric argument when $p \in B$ yields the desired equivalence. □

We use the labelling described in the beginning of this section. Additionally, write N for the hyperplane of X containing the midcube of edge $b_i p_i$.

Lemma 6.19. *Given the matrix $D_{\partial X}$ of distances in X between vertices of ∂X, we can deduce the matrix $D_{\partial Y}$ of distances in Y between vertices of ∂Y.*

Proof. First, we claim that for any vertices $x, y \in \partial Y$, we can find a geodesic path avoiding the removed vertices $\{p_i : i = 1, \dots, k\}$. Indeed, given γ a geodesic (x, y)-path in X, by replacing each p_i in γ by c_i we obtain a walk in Y of the same length. In particular, for any $x, y \in \partial Y$ we immediately have $d_X(x, y) = d_Y(x, y)$. These distances are already given by the distance matrix $D_{\partial X}$ for vertices $x, y \in \partial Y \cap \partial X$.

It remains to deduce the distances between boundary vertices and those newly created by the removal process. For each $i = 1, \dots, k-1$, we claim that for any $x \in \partial Y$, the distance to vertex c_i is given by

$$d_Y(c_i, x) = \begin{cases} d_X(x, a_i) + 1 & \text{if } d_X(x, b_i) = d_X(x, p_i) + 1, \\ d_X(x, a_i) - 1 & \text{if } d_X(x, b_i) = d_X(x, p_i) - 1. \end{cases}$$

Let A, B be the two connected components of $X \setminus N$. We apply Lemma 6.18 in X to edges $p_i b_i$ and $a_i c_i$ (whose midcubes are both contained in hyperplane N). On one hand this allows us to determine the component of x as $d_X(x, b_i)$, $d_X(x, p_i)$ are known and

$$\begin{aligned} d_X(x, b_i) = d_X(x, p_i) + 1 &\Leftrightarrow x \in A, \\ d_X(x, b_i) = d_X(x, p_i) - 1 &\Leftrightarrow x \in B. \end{aligned}$$

On the other hand, this information allows us to deduce $d_X(x, c_i) = d_Y(x, c_i)$ as $d_X(x, a_i)$ is known and

$$\begin{aligned} d_X(x, c_i) = d_X(x, a_i) + 1 &\Leftrightarrow x \in A, \\ d_X(x, c_i) = d_X(x, a_i) - 1 &\Leftrightarrow x \in B. \end{aligned}$$

□

6.6 Rows of Cubes II: The Thickening

The purpose of this section is to overcome the difficulty that our complex X is not necessarily homeomorphic to a ball. We describe a method by which to 'thicken' X into a new cube complex, denoted $\mathbb{X}$, that is homeomorphic to the ball, but whose boundary is sufficiently similar to the original that we can use it to deduce the existence of structures on X. Broadly, our thickening process entails gluing a new cube on every face of our cube complex, and then identifying sides of those new cubes to reflect

face incidences in X. This creates a 'shell' of cubes around X which bulks up lower-dimensional free cells and produces a pure contractible cube complex $\mathbb{X}$, the *thickening of* $\mathbb{X}$.

It is good to note that although the construction of the thickening of a cube complex X depends on the choice of embedding $X \hookrightarrow \mathbb{R}^3$, this is not an issue when proving an existence statement for a structure in X. As such, we slightly abuse notation and speak of *the* thickening of X to mean a *fixed instance* of a thickening of X.

Let X be a clean CAT(0) cube complex and $I = [0,1]$ denote the unit interval. Let the *sides* of X refer to faces on ∂X counted with multiplicity, so that free faces are counted twice. For each face $F \in \partial X$, let $\text{sides}(F)$ be the collection of its associated sides, namely a multiset $\{F, F\}$ if F is free and the singleton set $\{F\}$ otherwise. Write $\mathcal{F} := \bigcup_F \text{sides}(F)$ for the multiset of sides over all faces of ∂X.

In the first step of our thickening process, we associate to each side $S \in \mathcal{F}$ a new cube $C_S = I \times S$ with the natural gluing map $\phi_S \colon \{0\} \times S \to S \subseteq X$. We refer to the face $\{1\} \times S \subseteq C_S$ as the *external face* of C_S. As its name suggests, the external face of C_S will be on the boundary of $\mathbb{X}$ (which we verify more formally in Claim 6.21).

We define an intermediate complex $X' = X \sqcup \bigsqcup C_S \Big/ \bigsqcup \phi_S$, where both disjoint unions are taken over $S \in \mathcal{F}$, namely X' is the complex obtained from X by gluing cubes C_S to each side S of ∂X along maps ϕ_S. We write $H_{t,S} \colon S \to C_S \subseteq X'$ for the map $(x,y) \mapsto (t,x,y)$, so that $H_{0,S}(S) = S \subseteq X$ and $H_{1,S}(S)$ is the external face of C_S embedded in X'.

Note that that there exists an embedding of X' in $\mathbb{R}^3$, with $X \hookrightarrow \mathbb{R}^3$ as a subcomplex. Fix such an embedding; for each edge e of ∂X, it induces a cyclic ordering $S_1, \ldots, S_\ell$ of the sides of boundary faces incident to e. For each i, there are two faces of C_{S_i} which are incident to e in X'. One is $\{0\} \times S_i$ and we denote the other S_i^e. See Figure 6.10. For each i, we define gluing maps ϕ_{e,S_i} between S_i^e and S_{i+1}^e whenever the points between S_i^e and S_{i+1}^e in the cyclic ordering around the edge e are not in X, where indices are considered modulo k.

The *thickening* of X, denoted $\mathbb{X}$, is the cube complex $\mathbb{X} = X' \Big/ \bigsqcup \phi_{e,S}$ obtained from X' by gluing cubes associated to consecutive sides in the cyclic ordering around each edge $e \in \partial X^1$ according to gluing maps $\phi_{e,S}$. For each fixed $t \in [0,1]$ we have the following compositions:

$$S \xrightarrow{H_{t,S}} C_S \xrightarrow{\phi_S} X' \xrightarrow{\sqcup \phi_{e,S}} \mathbb{X},$$

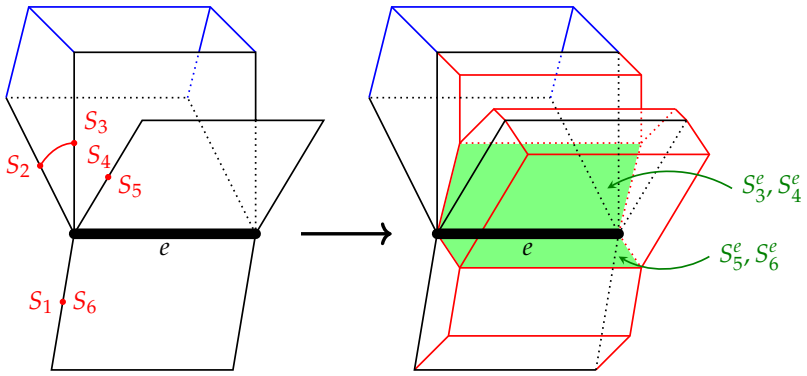


Figure 6.10. Example of gluing faces of cubes C_F together.

where the union in the last map is taken over each edge e of S. The embedding of X' in $\mathbb{R}^3$ allows us for each fixed $t \in [0,1]$ to extend these compositions to a continuous map $H_t \colon L := \bigsqcup S \Big/ \bigsqcup \phi_S \sqcup \bigsqcup \phi_{e,S} \to \mathbb{X}$. Moreover, these maps are continuous with t. Then, $H_0(L) = \partial X$ and $H_1(L) \subseteq \partial \mathbb{X}$. In fact, in Claim 6.20 we will show that $H_1(L) = \partial \mathbb{X}$, so that $\{H_t \colon t \in I\}$ describes a deformation retraction of $\mathbb{X}$ onto X. We define the *projection map* $\pi \colon \partial \mathbb{X} \to \partial X$ as $H_1(x) \mapsto H_0(x)$ for each $x \in L$. Note that π is combinatorial, and that it is well-defined as $H_1 \colon L \to \mathbb{X}$ is an embedding.

Note that each $x \in \partial \mathbb{X}^0$ is adjacent to precisely one vertex $\pi(x)$ of X^0 in $\mathbb{X}$. We sometimes refer to the cubes $\bigsqcup_{S \in \mathcal{F}} C_S$ as the *new cubes* of $\mathbb{X}$.

The following claim is essentially immediate from our construction.

Claim 6.20. *$\mathbb{X}$ is a connected contractible pure 3D cube complex.*

Proof. The maps $\{H_t : t \in I\}$ give a deformation retraction of $\mathbb{X}$ onto X, so contractibility of $\mathbb{X}$ follows from the contractibility of X. Connectedness is also inherited directly from the original complex. The assumptions that X is clean and simply connected mean that the only lower-dimensional cells in X that are not in cubes are the free faces. Since the thickening process ensures that each of these is now in a new cube (two, in fact), we have that $\mathbb{X}$ is pure. ■

Claim 6.21. *The thickened complex $\mathbb{X}$ is a 3-manifold with boundary. Moreover, the boundary of $\mathbb{X}$ is precisely the union of the external faces of all new cubes glued according to restrictions of the maps $\phi_{e,S}$.*

Proof. Let us begin by characterising $\partial \mathbb{X}$. We first note that all external faces must be on the boundary of $\mathbb{X}$, as they are on the boundary of every intermediate cube complex as we go through the construction described. For the other direction, we will rule out the possibility that any faces, edges or vertices not in an external face can be on the

boundary. To this end, note that any (open) face F of X is either in $\mathrm{int}(X)$, or it is in ∂X meaning a new cube is glued onto each side in $\mathrm{sides}(F)$ in the thickening process. In both cases, we see that the interior of F is in the interior of $\mathbb{X}$ from which we conclude that the faces of $\partial\mathbb{X}$ are the external faces. At the same time, for any edge or vertex to be in $\partial\mathbb{X}$, it must be part of a face in $\partial\mathbb{X}$. This follows from the fact that $\mathbb{X}$ and hence its boundary are pure cube complexes (the latter in two dimensions). As we have just seen that all such faces are external, the claim follows.

We now show that $\mathbb{X}$ is a manifold. It is certainly second-countable and Hausdorff, so we just need to verify that it is locally Euclidean. Since X is embedded in Euclidean space, any point in the interior of X certainly has a neighbourhood homeomorphic to the open 3-ball. That leaves us to consider points on the boundary of $\mathbb{X}$. For points p in the interior of a boundary face of $\mathbb{X}$ and $\varepsilon > 0$ sufficiently small, the intersection $B_\varepsilon(p) \cap \mathbb{X}$ is homeomorphic to the half-ball $\{x_1, x_2, x_3 : x_1^2 + x_2^2 + x_3^2 < 1, x_1 \geqslant 0\}$ where points with $x_1 = 0$ map to a disc on the face (this rephrases the fact that an individual cube is homeomorphic to the ball B^3 and hence a 3-manifold with boundary). This is also true for points on the interior of an edge on $\partial\mathbb{X}$, where points with $x_1 = 0$ map to a disc that intersects the two (distinct, by construction of $\mathbb{X}$) boundary faces incident to our edge. By our characterisation of $\partial\mathbb{X}$ in the preceding paragraph, we know that each vertex v on $\partial\mathbb{X}$ is surrounded by external faces of new cubes, and the link of v is a disc by construction. This tells us that the cone on the link, and hence $B_\varepsilon(v) \cap \mathbb{X}$, is homeomorphic to the half-ball whose boundary is the union of a disc on $\partial\mathbb{X}$ and $\mathrm{link}(v)$. ■

It may be interesting to note that the boundary points of $\mathbb{X}$ as a manifold are precisely the points of $\partial\mathbb{X}$.

Corollary 6.22. *$\partial\mathbb{X}$ is homeomorphic to S^2.*

Proof. Since $\mathbb{X}$ is a 3-manifold with boundary, the boundary $\partial\mathbb{X}$ is a surface. Moreover, the contractibility of $\mathbb{X}$ means that $\partial\mathbb{X}$ is a homology 2-sphere. The statement then follows from the classification of surfaces. □

6.7 Existence of rows of cubes on the boundary

In this section, we show the existence of good row configurations in clean CAT(0) cube complexes. In particular, Lemma 6.17 then implies existence of rows of cubes on the boundary.

For this section, let $Y \hookrightarrow \mathbb{R}^3$ be a finite clean CAT(0) cube complex and $\mathbb{Y}$ its thickening.

Row configurations in the thickened complex

Let $v \in \partial\mathbb{Y}^0$, $x := \pi(v)$ and $C_1, \dots, C_k$ be the cyclically ordered cubes incident to the edge vx in $\mathbb{Y}$. Since π is combinatorial, its restriction to D induces a combinatorial map $\pi^* \colon \operatorname{link}_{\partial\mathbb{Y}}(v) \to \operatorname{link}_{\partial Y}(x)$.

Since π^* is combinatorial and only identifies edges corresponding to the two sides of a single face in Y, we have the following fact.

Fact 1. *Under above setup, the image of π^* is a closed walk $W \subseteq \operatorname{link}_{\partial Y}(x)$. This walk is formed by identifying either pairs of edges in $\operatorname{link}_{\partial\mathbb{Y}}(v)$ or vertices of $\operatorname{link}_{\partial\mathbb{Y}}(v)$ together.*

The next lemma essentially states that the projection map π is well-behaved at vertices of degrees 3 and 4: the former come from vertices of degree 3 on ∂Y and under suitable assumptions the latter come from vertices of degree 4. Moreover, all other vertices of $\partial\mathbb{Y}$ have degree at least 4.

Lemma 6.23. *Let $v \in \partial\mathbb{Y}^0$ and $x := \pi(v) \in \partial Y$, i.e. vx is an edge. Then,*

(i) If $\deg_{\partial\mathbb{Y}}(v) = 3$, then $\deg_Y(x) = 3$. Conversely, for every $x \in \partial Y^0$ with $\deg_{\partial Y}(x) = 3$, $\pi^{-1}(x)$ is a singleton vertex and $\deg_{\partial\mathbb{Y}}(\pi^{-1}(v)) = 3$.

(ii) If $\deg_{\partial\mathbb{Y}}(v) \geqslant 4$, then $\deg_{\partial Y}(x) \geqslant 4$.

(iii) Suppose $\deg_{\partial\mathbb{Y}}(v) = 4$ and the faces of $\partial\mathbb{Y}$ incident to it are, in cyclic order, $F_1, \dots, F_4$ with projections $\pi(F_1), \dots, \pi(F_4) \subseteq Y$. Then $\pi(F_1), \dots, \pi(F_4)$ are distinct and if $\pi(F_1)$, $\pi(F_2)$ lie in a common cube C of Y, then $\pi(F_3)$, $\pi(F_4)$ also lie in a common cube D of Y, and $\deg_{\partial Y}(x) = 4$.

Proof. Note that links of vertices in Y do not contain self-loops nor double edges as they are simplicial complexes. In particular, in light of Fact 1, this guarantees that all vertices of $\partial\mathbb{Y}$ have degree at least 3: vertices of degree 1 or 2 would imply self-loops and double edges in boundary links of Y, respectively.

Suppose $\deg_{\partial\mathbb{Y}}(v) = 3$. Then, since $\operatorname{link}_{\partial Y}(\pi(v))$ does not contain self-loops nor double edges, we have that $\pi^*(\operatorname{link}_{\partial\mathbb{Y}}(v))$ is a triangle. Since Y is CAT(0), it follows that this triangle corresponds to a cube $C \subseteq Y$, with faces F_1, F_2, F_3 incident to x. By construction a small ball around x in $\partial\mathbb{Y}$ is covered by $C \cup C_{F_1} \cup C_{F_2} \cup C_{F_3}$, so $\deg_{\mathbb{Y}}(x) = 4$ and $\deg_Y(x) = 3$. The converse follows as Y is clean: vertices of degree 3 on ∂Y are corners of cubes, so the claim follows by construction. This also proves the second bullet point as there are no vertices of degree less than 3 in $\partial\mathbb{Y}$.

Suppose $v \in \partial\mathbb{Y}^0$ satisfies the assumptions of point 3. Then since $\mathrm{link}_{\partial Y}(\pi(v))$ does not contain self-loops nor double edges, we have that $\pi^*(\mathrm{link}_{\partial\mathbb{Y}}(v))$ is a cycle of length 4, so the projected faces $\pi(F_1), \dots, \pi(F_4)$ are distinct. As $\pi(F_1)$ and $\pi(F_2)$ lie in the same cube of Y, by Lemma 6.10 (c), $\pi(v)$ is incident to exactly two cubes in X, has $\deg_{\partial Y}(x) = 4$ and $\pi(F_3)$, $\pi(F_4)$ lie in the same cube, as desired. □

Lemma 6.24. *There is a one to one correspondence between row configurations in $\partial\mathbb{Y}^1$ and ∂Y^1. Namely if $p_0 p_1 \dots p_k$ is the spine of a row configuration in $\partial\mathbb{Y}^1$ then $\pi(p_0)\pi(p_1)\dots\pi(p_k)$ is the spine of a row configuration in ∂Y^1.*

Proof. By Lemma 6.23 (i) and since Y is clean, $\deg_{\partial Y}(\pi(p_0)) = 3$, so $\pi(p_0)$ is the corner of a cube C_1. Now, $\deg_{\partial\mathbb{Y}}(p_1) = 4$ and p_1 has a cyclic ordering of incident faces $F_1, \dots, F_4$. Since π is combinatorial, $\pi(p_0)\pi(p_1)$ is an edge and so $\pi(p_1) \in C_1$.

Hence, without loss of generality we may assume $\pi(F_1), \pi(F_2) \subseteq C_1$ and so by Lemma 6.23 (iii), $\pi(F_3), \pi(F_4)$ both lie in a cube C_2. Moreover, with this labelling of faces, $p_2 \in F_3, F_4$ so $\pi(p_2) \in C_2$. Continuing in this way we find a row of cubes $(C_1, \dots, C_k)$ where $\pi(p_0)$ is a corner of C_0 and $\pi(p_{i-1}), \pi(p_i) \in C_i$ for each $1 \leqslant i \leqslant k$. This implies that $\pi(p_0), \dots, \pi(p_k)$ is the spine of a row configuration in Y, as desired. □

For simplicity, if $\mathcal{C} \subseteq \partial\mathbb{Y}^1$ is a row configuration, we write $\pi(\mathcal{C}) \subseteq \partial Y^1$ for its corresponding projected row configuration in ∂Y.

Existence of good row configurations

Let Y be a clean CAT(0) cube subcomplex of a cubulation of the ball B^3. In this section we show that good row configurations exist in ∂Y^1. By the correspondence between row configurations in ∂Y^1 and those in $G(Y)$ discussed in Section 6.5, this implies that there are good configurations in $G(Y)$. By Lemma 6.24, our argument boils down to showing that good row configurations exist in the thickening $\partial\mathbb{Y}^1$. For this, we use a simple path-counting argument.

Fact 2. *Let $v \in \partial Y^0$ have $\deg_{\partial Y}(v) = 3$. Then, there are at least three row configurations in ∂Y^1 with v as a starting point.*

The above fact is clear from definitions and can be strengthened to exactly three, but this will not be needed for our purposes.

Lemma 6.25. *Let $x \in \partial\mathbb{Y}^0$ be a vertex with $\deg_{\partial\mathbb{Y}}(x) \geqslant 5$. Then there are at most $2\lfloor \frac{d}{3} \rfloor$ row configurations which end at x.*

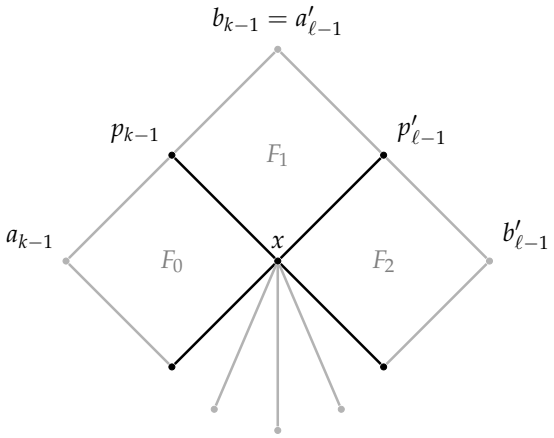


Figure 6.11. Two row configurations sharing endpoint x. In fact, F_0 and F_2 must share an edge.

Note that the condition on the degree is necessary: one can easily find examples of Y with vertices of degree 3 having three row configurations ending at them in the thickening, e.g. when Y a single cube.

Proof. Let

$$\mathcal{C} = (a_i, p_i, b_i : i = 0, \dots, k)$$
$$\mathcal{C}' = (a'_i, p'_i, b'_i : i = 0, \dots, \ell)$$

be two row configurations in $\partial \mathbb{Y}^1$ with a common endpoint $x := p_k = p'_\ell$, as in Figure 6.11.

If $p_{k-1} = p'_{\ell-1}$, then $\mathcal{C}$ and $\mathcal{C}'$ share the last edge of their spine, and by a simple inductive argument it follows that the two spines are the same, meaning $\mathcal{C} = \mathcal{C}'$. Suppose now that $p_{k-1} \neq p'_{\ell-1}$, yet they both lie in face F_1 with vertices $\{p_{k-1}, x, p'_{\ell-1}, y\}$ for $y = b_{k-1} = a'_{\ell-1}$, labelled as in Figure 6.11. In particular, F_0, F_1 and F_2 are pairwise distinct faces of ∂Y.

By Lemma 6.24, the row configurations $\mathcal{C}$, $\mathcal{C}'$ in $\partial \mathbb{Y}^1$ correspond to row configurations $\pi(\mathcal{C}), \pi(\mathcal{C}')$ in ∂Y^1 respectively, which share the face $\pi(F_1)$. Let $(C_1, \dots, C_k)$, $(C'_1, \dots, C'_\ell)$ be the rows of cubes corresponding to $\pi(\mathcal{C})$, $\pi(\mathcal{C}')$, respectively. Then, $C_k = C'_\ell$ as they share a face on ∂Y, Let $\{\pi(p_{k-1}), \pi(p'_{\ell-1}), z\}$ be the three neighbours of $\pi(x)$ in C_k.

Claim 6.26. *There is no row configuration in ∂Y^1 whose spine ends with edge $\pi(x)z$.*

Proof. Since $\pi(F_0), \pi(F_1), \pi(F_2) \subseteq \partial Y$ lie in a common cube C_k of Y, the edges corresponding to them in $\mathrm{link}_{\partial Y}(\pi(x))$ form a triangle T. We first show that there is a face F in ∂Y distinct from $\pi(F_0), \pi(F_1), \pi(F_2)$ which is incident to $\pi(x)z$.

By Lemma 6.23 (ii), $\deg_{\partial Y}(\pi(x)) \geqslant 4$ as $\deg_{\partial \mathbb{Y}}(x) \geqslant 5$ and so there is an edge $\pi(x)v$ in Y such that $v \notin \{\pi(p_{k-1}), \pi(p'_{\ell-1}), z\}$. Since Y is clean, in particular $\pi(x)$ is not a cut vertex. Fix a path $\gamma \subseteq Y^1 \setminus \pi(x)$ from v to z. Notice that the edges $\pi(x)\pi(p_{k-1})$ and $\pi(x)\pi(p'_{\ell-1})$ are only incident to faces of C_k in Y as they are on the spine of row configurations. Hence, in a disc diagram $D \to Y$ with $\partial D = \gamma$, since $\pi(x)z \subseteq C_k$ and $\pi(x)v \not\subseteq C_k$ the face F incident to $\pi(x)z$ in D is not on C_k, as claimed.

This suffices to prove the claim, since if $\pi(x)z$ was the last edge of the spine of a row configuration, then it would only be incident in Y to the two faces of the last cube in the corresponding row of cubes. ■

Indeed, assuming the claim, no row configuration of $\partial \mathbb{Y}^1$ can have spine ending with edge $z'x$, for $z' \in \pi(z)$, as this would lead to a contradiction in light of Lemma 6.24.

In what follows, we restrict our considerations to row configurations in $\partial \mathbb{Y}^1$ which end in x, and faces of $\partial \mathbb{Y}$ incident to x. Let k be the total number of row configurations and write $k = t + 2r$ where t counts *single* row configurations – those sharing no face (incident to x) with other row configurations, and r counts the *pairs* of row configurations as above. Then, each single row configuration forbids two faces while each pair forbids three. In conclusion,

$$k = t + 2r \leqslant \left\lfloor \frac{d - 3r}{2} \right\rfloor + 2r \leqslant 2 \left\lfloor \frac{d}{3} \right\rfloor,$$

as $3r \leqslant d$, concluding the proof. □

Lemma 6.27. *There exists a good configuration in* $\mathbb{Y}$.

Proof. Let n_k denote the number of vertices on $\partial \mathbb{Y}$ with degree k in $\partial \mathbb{Y}$, and $|E|$, $|F|$ denote the number of edges, faces respectively of $\partial \mathbb{Y}$. Then, by Euler's formula and Corollary 6.22,

$$2 = \sum_{k \geqslant 3} n_k - |E| + |F|.$$

Note that every edge is in exactly two faces, and every face is bounded by exactly four edges. Hence $2|F| = |E|$ and so,

$$\begin{aligned} 2 &= \sum_{k\geqslant 3} n_k - \frac{1}{2}|E| \\ n_3 &= 8 + \sum_{k\geqslant 5}(k-4)n_k, \end{aligned}$$

where the last line follows from using that $\sum k \cdot n_k = 2|E|$, multiplying both sides by 4 and rearranging.

We may assume that there exists a vertex of degree at least 5 on ∂X as otherwise a good configuration exists immediately. Let u, v be vertices of $\partial\mathbb{Y}$ with $\deg_{\partial\mathbb{Y}}(v) = 3$ and $\deg_{\partial\mathbb{Y}}(u) = k \geqslant 5$. By Fact 2, v is the starting point of three row configurations, and Lemma 6.25 implies u is the ending point of at most $2\lfloor k/3 \rfloor$ row configurations. Hence, there are at least $3n_3$ row configurations, of which at most $\sum_{k\geqslant 5} 2\lfloor \frac{k}{3} \rfloor n_k$ end in a vertex of degree at least 5. Write N for the number of row configurations with starting and ending point of degree 3. Using the above formula for n_3 we thus have

$$\begin{aligned} N &\geqslant 3n_3 - \sum_{k\geqslant 5} 2\left\lfloor \frac{k}{3} \right\rfloor n_k \\ &= 24 + \sum_{k\geqslant 5} \underbrace{\left(3k - 2\left\lfloor \frac{k}{3} \right\rfloor - 12\right)}_{\geqslant 0 \text{ for } k\geqslant 5} n_k \\ &\geqslant 24. \end{aligned}$$

□

6.8 Finite CAT(0) cube complexes of dimensions 1 and 2

In this section, we sketch the proof of the following strengthening of Theorem 6.1.

Theorem 6.4. *Let X be a CAT(0) cube complex of top dimension at most 2 with finitely many cells and D its matrix of pairwise distances between vertices on the combinatorial boundary of X. Then, the combinatorial type of X is reconstructible from D.*

The arguments we use are very similar to those used in the proof of Theorem 6.3, with much less technical complications.

Proof sketch of Theorem 6.4. If $k = 1$, then X is a tree and its boundary are its leaves, so reconstruction can be achieved by a simple inductive argument. Indeed, the distance between the neighbour of a leaf v and other leaves is one less than the distance between other leaves and v, and this neighbour is again a leaf after removing v precisely if it is not on any geodesic between pairs of leaves.

If $k = 2$, we may first assume that there are no vertices of degree at most 1, as we may then repeatedly remove such vertices. In particular, without loss of generality, every vertex of X is contained in at least one face. Similarly we may assume that X has no cut vertices.

Consider an inclusion maximal disc diagram $D \hookrightarrow X$. In particular, since X has no 3-cubes and is CAT(0), D is a quadrangulation with all internal degrees at least 4, so [63, Lemma 3.4] implies that there is a vertex v of degree 2 on ∂D. Since X has no cut vertices and D is maximal, v is contained in a unique face F of X. Let u be the vertex not adjacent to v in this face. Then, by using Lemma 6.18 similarly as in the argument for the 3-dimensional setting, we can recover distances between boundary vertices and u in the new complex.

Lastly, we show that vertices on ∂X with degree 2 in X can be recognised. Let v be a vertex on ∂X with $d_{\partial X}(v) = 2$ and e, f be the two edges on ∂X incident to v. Write H_e and H_f to denote the hyperplanes in X dual to e, f respectively. We will show that v has degree 2 in X if and only if H_e and H_f cross in X.

Suppose first that H_e and H_f cross in X. Now, by [121, Lemma 3.6] e and f are contained in a same face F of X. In particular, since e and f are both on the boundary, v is only incident to F and therefore has degree 2 in X. Conversely, if $d_X(v) = 2$, H_e and H_f clearly cross in the face incident to v.

We now explain how this condition can be recognised from the boundary distances. Both H_e and H_f split X into two connected components. Let $\partial X^0 = A_e \cup B_e$, $A_f \cup B_f$ be the resulting partitions of the vertices of ∂X. These can be identified from the boundary distances by using Lemma 6.18. If H_e and H_f do not cross in X, then it must be that, without loss of generality, $A_e \subseteq A_f$ or $B_e \subseteq A_f$. Hence, checking that H_e and H_f cross in D amounts to checking that that $A_e \not\subseteq A_f$ and $B_e \not\subseteq A_f$. □

Chapter 7

Small families of partially shattering permutations

This chapter is based on joint work with Girão and Michel [58].

In this chapter, we think of a permutation ρ of a set X to be a total order of the elements of X. We denote this order by $<_\rho$. Note that if $Y \subseteq X$, then ρ induces a permutation on Y. If ρ is a permutation of $A \cup B$, we write $A <_\rho B$ if for all $a \in A$ and $b \in B$ we have $a <_\rho b$.

A family $\mathcal{P}$ of permutations of $[n] = \{1, \dots, n\}$ *shatters* a set $X \subseteq [n]$ if the permutations of $\mathcal{P}$ induce every possible permutation on the elements of X. Shattering families of permutations were first studied by Spencer [113] who asked the following question.

> What is the smallest family of permutations of $[n]$ that shatters all subsets of a fixed size k?

Spencer [113] showed that such families have size $\Theta(\log n)$, with subsequent work improving the constant of the lower bound [73, 54, 101].

A natural refinement of this problem is to consider partial shattering. For $t \geqslant 1$, we say that a family $\mathcal{P}$ *t-shatters* a set X if $\mathcal{P}$ induces at least t distinct permutations on X. Let $f_k(n, t)$ be the minimum number of permutations of $[n]$ that t-shatter all subsets of size k. From above we know that $f_k(n, k!) = \Theta(\log n)$, and monotone permutations can be used to prove that $f_k(n, t) = t$ for $t \leqslant 2$. Moreover, an argument of Hajnal and Spencer [113] shows that $f_k(n, t) = \Theta(\log \log n)$ for $3 \leqslant t \leqslant k$.[1] Therefore, the asymptotic behaviour of $f_k(n, t)$ falls into at least three distinct regimes.

[1] Spencer [113] considered the slightly different problem of determining the minimum number of permutations such that for every subset $X \subseteq [n]$ of size k and every element $x \in X$, there is a permutation where x is the largest element of X. Johnson and Wickes [75] observed that this transfers to $f_k(n, t)$.

In 1996, Füredi [54] asked whether these might be the only possible regimes, even in a much more general version of partial shattering. Let $\mathcal{S}$ be a family of sets of permutations of $[k]$. Then, a family $\mathcal{P}$ of permutations of $[n]$ is *$\mathcal{S}$-mixing* if for every subset $X \subseteq [n]$ of size k, the set of permutations that $\mathcal{P}$ induces on X is a member of $\mathcal{S}$. Moreover, $\mathcal{S}$ is *monotone* if for all $S \in \mathcal{S}$ and $S \subseteq T$ we have $T \in \mathcal{S}$. If $\mathcal{S}$ is the family of sets with at least t permutations of $[k]$, then $\mathcal{S}$ is monotone and $\mathcal{S}$-mixing families are exactly those families that t-shatter all subsets of size k.

Even in this more general $\mathcal{S}$-mixing framework, the minimum size of an $\mathcal{S}$-mixing family in all previously known cases was in one of the three regimes $\Theta(1)$, $\Theta(\log\log n)$, and $\Theta(\log n)$. This prompted Füredi [54] to ask the following question.

Question 7.1 (Füredi, 1996)**.** *If $\mathcal{S}$ is a monotone family of sets of permutations of $[k]$, is the minimum size of an $\mathcal{S}$-mixing family either $\Theta(1)$, $\Theta(\log\log n)$, or $\Theta(\log n)$?*

Johnson and Wickes [75] recently made progress on this question for $f_k(n,t)$. They showed that $f_k(n,t) = \Theta(\log n)$ for $t > 2(k-1)!$. Together with the previously known asymptotics, this yields the following partial classification.

Theorem 7.2 (Johnson, Wickes, 2023)**.** *Let $k \geqslant 3$ be an integer. Then,*

$$f_k(n,t) = \begin{cases} t & \text{for } t \leqslant 2, \\ \Theta(\log\log n) & \text{for } 3 \leqslant t \leqslant k, \\ \Theta(\log n) & \text{for } 2(k-1)! < t \leqslant k!. \end{cases}$$

Moreover, Johnson and Wickes [75] settled the case $k = 3$ completely by additionally proving that $f_3(n,4) = \Theta(\log\log n)$. Given these results, they reiterated Füredi's question and asked specifically whether $f_k(n,t)$ must always fall into one of the three regimes $\Theta(1)$, $\Theta(\log\log n)$, and $\Theta(\log n)$ [17, 75].

We answer the questions of Füredi and of Johnson and Wickes negatively. For $k \geqslant 4$, we show that a fourth regime exists with $f_k(n,t) = \Theta(\sqrt{\log n})$. More generally, we improve the partial classification of the asymptotic behaviour of $f_k(n,t)$ as follows.

Theorem 7.3. *Let $k \geqslant 4$ be an integer. Then,*

$$f_k(n,t) = \begin{cases} t & \text{for } t \leqslant 2, \\ \Theta(\log\log n) & \text{for } 3 \leqslant t \leqslant 2^{\lceil \log_2 k \rceil}, \\ \Theta(\sqrt{\log n}) & \text{for } 2^{\lceil \log_2 k \rceil} < t \leqslant \min\{2k, 2^{\lceil \log_2 k \rceil} + 4\}, \\ \Theta(\log n) & \text{for } 2^{k-1} < t \leqslant k!. \end{cases}$$

For $k = 4$, this settles the asymptotic behaviour of $f_4(n,t)$ completely as all values of t are covered. However, if k is large, there is still an exponential gap between the regime $\Theta(\sqrt{\log n})$ and $\Theta(\log n)$.

Theorem 7.3 is based on a series of new lower and upper bounds on $f_k(n,t)$. First, we show that the lower bound $f_k(n,t) = \Omega(\log n)$ holds for a wider range of values of t.

Theorem 7.4. *Let $k \geqslant 3$ and $t > 2^{k-1}$. Then, $f_k(n,t) = \Omega(\log n)$.*

The main observation for this result is that for any small family $\mathcal{P}$ of permutations of $[n]$, we can construct two large subsets $A, B \subseteq [n]$ that are ordered in the following sense: for each permutation of $\mathcal{P}$ either all elements of A are smaller than all elements of B, or all elements of B are smaller than all elements of A. By recursively constructing ordered sets in A, we find a subset of size k that is only 2^{k-1}-shattered by $\mathcal{P}$.

For smaller values of t, we provide a new lower bound of the form $\Omega(\sqrt{\log n})$.

Theorem 7.5. *Let $k \geqslant 3$ and $t > 2^{\lceil \log_2 k \rceil}$. Then, $f_k(n,t) = \Omega(\sqrt{\log n})$.*

The proof of this result is inspired by the proof of Theorem 7.4. However, instead of only constructing ordered sets in A, we recursively construct ordered sets both in A and in B. We then use a Ramsey-theoretic argument about vertex-coloured binary trees to find a subset of size k that is only $2^{\lceil \log k \rceil}$-shattered.

We note that the lower bound on t in Theorem 7.5 cannot be replaced by anything smaller. Indeed, a careful analysis of the construction of Hajnal and Spencer [113] shows that $f_k(n,t) = \Theta(\log \log n)$ for $3 \leqslant t \leqslant 2^{\lceil \log_2 k \rceil}$. We provide an equivalent construction which proves this and which serves as a motivating example for what follows.

Theorem 7.6. *Let $k \geqslant 3$ and $t \leqslant 2^{\lceil \log_2 k \rceil}$. Then, $f_k(n,t) = \mathcal{O}(\log \log n)$.*

Our construction identifies $[n]$ with $[2]^d$ for $d = \log_2 n$. Then, we consider lexicographic permutations of $[2]^d$ which are permutations where the order of $x, y \in [2]^d$ only depends on the values x_i and y_i for the first position i with $x_i \neq y_i$. We show that an appropriate choice of these permutations ensures that all subsets of size k are $2^{\lceil \log_2 k \rceil}$-shattered.

Finally, for $k \geqslant 4$, we show that there exist $t > 2^{\lceil \log_2 k \rceil}$ with $f_k(n,t) = \mathcal{O}(\sqrt{\log n})$. Together with Theorem 7.5, this establishes the existence of a fourth regime for $f_k(n,t)$.

Theorem 7.7. *Let $k \geqslant 4$ and $t \leqslant \min\{2k, 2^{\lceil \log_2 k \rceil} + 4\}$. Then, $f_k(n,t) = \mathcal{O}(\sqrt{\log n})$.*

This result is proved similar to Theorem 7.6, but we identify $[n]$ with $[2^d]^d$ for $d = \sqrt{\log_2 n}$. Then, most subsets of size k are $2k$-shattered by lexicographic permutations,

and the remaining subsets have a very specific structure. We exploit this structure and add a few more permutations which ensure that all subsets of size k are t-shattered.

We remark that partial shattering with permutations is quite different to partial shattering with sets. A family $\mathcal{F}$ of subsets of $[n]$ *shatters* a set $X \subseteq [n]$ if for every subset $Y \subseteq X$ there exists $F \in \mathcal{F}$ with $F \cap X = Y$, and $\mathcal{F}$ *t-shatters* X if for at least t distinct subsets $Y \subseteq X$ there exists $F \in \mathcal{F}$ with $F \cap X = Y$. The study of shattering families of sets dates back to the seminal works of Sauer [109], Shelah [112], and Vapnik and Chervonenkis [119].

As for permutations, Kleitman and Spencer [80] showed that the minimum number of subsets of $[n]$ that shatter all subsets of size k is $\Theta(\log n)$. However, in the case of partial shattering, the family $\{\varnothing, [n]\}$ 2-shatters all subsets of size k, and every family that 3-shatters all subsets of size k already needs $\Omega(\log n)$ sets.[2] Therefore, partial shattering with sets only has the two regimes $\Theta(1)$ and $\Theta(\log n)$.

The rest of the chapter is structured as follows. In Section 7.1 we discuss an initial approach which led us to the arguments used in the proof of Theorems 7.4 and 7.5. In Section 7.2 we prove the lower bounds of Theorems 7.4 and 7.5. Afterwards, in Section 7.3, we prove the upper bounds of Theorems 7.6 and 7.7. We finish with some open problems in Section 7.4.

7.1 A first approach

An interesting idea that was developped by Johnson and Wickes in [75] to prove lower bounds on $f_k(n,t)$ was to consider how a given pair x, y is 'split' by other elements under $\mathcal{P}$. Using this idea, they showed that $f_k(n,t) = \Theta(\log n)$ when $2(k-1)! < t < k$. In this section, we sketch an approach inspired by this which lay the groundwork for proofs in Section 7.2.

Say that z *dissociates* x, y under a family $\mathcal{P}$ of permutations if there is a $\rho \in \mathcal{P}$ such that z falls between x and y in the ordering induced by ρ, namely $x <_\rho z <_\rho y$.

We can interpret this 'dissociation by z' as assigning to each $\rho \in \mathcal{P}$ a function $g_\rho \colon [n] \setminus z \to \{0,1\}$ where $g_\rho(x) = 0$ precisely when x appears before z under ρ. With this in mind, it is not hard to see that a set $S \subseteq [n] \setminus z$ whose pairs are all dissociated by z has an encoding of its elements as distinct binary strings with $\mathcal{P}$ bits. Such a set therefore satisfies $2^{|\mathcal{P}|} \geqslant |S|$. We further note that the absence of pairs dissociated by z in a set

[2] If $n \geqslant k \cdot 2^{|\mathcal{F}|}$, one of the intersections $\bigcap_{F \in \mathcal{F}} A_F$ with $A_F \in \{F, [n] \setminus F\}$ must have size at least k, and any subset $X \subseteq \bigcap_{F \in \mathcal{F}} A_F$ of size k satisfies $F \cap X \in \{\varnothing, X\}$ for all $F \in \mathcal{F}$.

$S \subseteq [n] \setminus z$ imposes a certain rigidity: under any $\rho \in \mathcal{P}$, either all elements of S precede z or they all succeed it.

These two observations beg for a Ramsey-type argument. For this sketch, we use the following folklore fact that can be proven similarly to the Erdős–Szekeres theorem.

Fact 3. *Any red-blue colouring of the edges of a clique on $(a-1)(b-1)+1$ vertices has either a blue clique K_a or a red path P_b.*

Pick $z \in X_0 = [n]$, colour a pair x, y from the remaining elements blue if it is dissociated by z under $\mathcal{P}$, red otherwise. If there is a blue $K_{\sqrt{n}}$ then $\mathcal{P}$ has size at least order $\log n$ and we are done. Otherwise, by Fact 3 there is a set X_1 of vertices which induce a graph containing a red $P_{\sqrt{n}}$. Note that if w, x and x, y are not dissociated by z then w, y is not dissociated by z either. This implies that the vertices of X_1 in fact induce a red clique, meaning that none of its pairs are dissociated by z. Repeat with X_1.

Continuing in this manner k times, at each step we have $|X_i| \geqslant n^{2^{-k}}$ so a large blue clique at any step implies that $\mathcal{P}$ has order $\log n$. Otherwise, we obtain a sequence of k elements $S = \{z_1, \dots, z_k\}$ where no z_i dissociates z_j, z_h for $j > h > i$. In particular, as for each z_i and $\rho \in \mathcal{P}$ there are only two choices for the position of all subsequent elements z_j relative to z_i, the family $\mathcal{P}$ induces at most 2^{k-1} different orders on S. This shows $f_k(n,t) = \Omega(\log n)$ for $t > 2^{k-1}$.

The proofs presented in the next section are refinements of this argument which, with some additional work and ideas, allow us to extend it to the $2^{\lceil \log_2 k \rceil} < t \leqslant 2^{k-1}$ regime.[3]

7.2 Lower bounds

To prove lower bounds for $f_k(n,t)$, we first need to define the concept of ordered sets. Let $\mathcal{P}$ be a family of permutations of a set X. Then, we say that a pair of disjoint subsets $A, B \subseteq X$ is *$\mathcal{P}$-ordered* if for each permutation $\rho \in \mathcal{P}$ either $A <_\rho B$ or $B <_\rho A$. The following result shows that any set contains large $\mathcal{P}$-ordered subsets.

Lemma 7.8. *Let X be a set and let $\mathcal{P}$ be a family of m permutations of X. Then there exists a $\mathcal{P}$-ordered pair $A, B \subseteq X$ with $\min\{|A|, |B|\} \geqslant \lfloor |X|/2^{m+1} \rfloor$.*

Proof. Let $\mathcal{P} = \{\rho_1, \rho_2, \dots, \rho_m\}$ and define $\mathcal{P}_i = \{\rho_1, \rho_2, \dots, \rho_i\}$. We inductively construct a $\mathcal{P}_i$-ordered pair $A_i, B_i \subseteq X$ with $|A_i| = |B_i| \geqslant \lfloor |X|/2^{i+1} \rfloor$. For $i = 0$, we choose $A_0, B_0 \subseteq X$ as two disjoint subsets of size $\lfloor |X|/2 \rfloor$.

[3] We moreover note that the dependency on k in a bound resulting from the above argument is exponential, as opposed to the linear dependency of the bound proved in Section 7.2.

Now suppose we have constructed A_i and B_i. Let $\ell = |A_i| = |B_i| \geqslant \lfloor |X|/2^{i+1} \rfloor$ and consider the subset L of the ℓ largest elements of $A_i \cup B_i$ under $<_{\rho_{i+1}}$. Note that either $|L \cap A_i| \geqslant \ell/2$ or $|L \cap B_i| \geqslant \ell/2$. Without loss of generality, assume that $|L \cap A_i| \geqslant \ell/2$.

Set $A_{i+1} = L \cap A_i$ and $B_{i+1} = B_i \setminus L$. Since $b <_{\rho_{i+1}} a$ for all $a \in A_{i+1}$ and $b \in B_{i+1}$, the pair (A_{i+1}, B_{i+1}) is $\mathcal{P}_{i+1}$-ordered. We also have $|A_{i+1}| \geqslant \ell/2 \geqslant \lfloor |X|/2^{i+2} \rfloor$ and $|B_{i+1}| = \ell - |L \cap B_i| = \ell - (\ell - |L \cap A_i|) = |A_{i+1}|$, as required.

For $i = m$, the pair (A_m, B_m) is $\mathcal{P}$-ordered with $\min\{|A_m|, |B_m|\} \geqslant \lfloor |X|/2^{m+1} \rfloor$. □

Using $\mathcal{P}$-ordered pairs, we can prove Theorem 7.4. For this, we construct a sequence of nested $\mathcal{P}$-ordered pairs $A_i, B_i \subseteq A_{i-1}$ for $i \in [k]$, where $A_0 = [n]$. Then, if we pick an element $x_i \in B_i$ for each $i \in [k]$, the fact that all pairs are $\mathcal{P}$-ordered implies that $\{x_1, \ldots, x_k\}$ is only 2^{k-1}-shattered by $\mathcal{P}$, as required.

Proof of Theorem 7.4. Let $t > 2^{k-1}$. We claim that $f_k(n,t) > (\log_2 n)/k - 1$. Indeed, suppose that $\mathcal{P}$ is a family of m permutations of $[n]$ with $m \leqslant (\log_2 n)/k - 1$, and so $n \geqslant 2^{k(m+1)}$. We show that there is a subset of size k that is not t-shattered by $\mathcal{P}$.

To do so, we inductively construct a sequence of $\mathcal{P}$-ordered pairs as follows. Start with $A_0 = [n]$. Then, for $i = 1, \ldots, k$, apply Lemma 7.8 to A_{i-1} to obtain a $\mathcal{P}$-ordered pair $A_i, B_i \subseteq A_{i-1}$ with $\min\{|A_i|, |B_i|\} \geqslant \lfloor |A_{i-1}|/2^{m+1} \rfloor$. Note that this implies that $\min\{|A_i|, |B_i|\} \geqslant \left\lfloor n/2^{i(m+1)} \right\rfloor \geqslant \left\lfloor n/2^{k(m+1)} \right\rfloor \geqslant 1$, and so A_i and B_i are non-empty.

For each $i \in [k]$, pick $x_i \in B_i$. We claim that $\mathcal{P}$ does not t-shatter the set $\{x_1, \ldots, x_k\}$. Indeed, let $i \in [k-1]$. Note that $x_j \in A_i$ for all $j > i$. Since (A_i, B_i) is $\mathcal{P}$-ordered, it follows that for each permutation $\rho \in \mathcal{P}$ we either have $x_i <_\rho \{x_{i+1}, \ldots, x_k\}$ or $\{x_{i+1}, \ldots, x_k\} <_\rho x_i$. This provides two choices for the position of x_i relative to $\{x_{i+1}, \ldots, x_k\}$. Moreover, given such a choice for each $i \in [k-1]$, this uniquely determines the permutation that ρ induces on $\{x_1, \ldots, x_k\}$. Hence, $\mathcal{P}$ induces at most 2^{k-1} permutations on $\{x_1, \ldots, x_k\}$, and therefore does not t-shatter $\{x_1, \ldots, x_k\}$. □

To prove Theorem 7.5, we will construct a complete binary tree of nested $\mathcal{P}$-ordered pairs. Namely, for each $\mathcal{P}$-ordered pair (A, B) that we obtain, we recursively construct new $\mathcal{P}$-ordered pairs in each of A and B.

This is useful for the following reason. Suppose that $k = 4$, (A, B) is a $\mathcal{P}$-ordered pair, and two $\mathcal{P}$-ordered pairs $A', B' \subseteq A$ and $A'', B'' \subseteq B$ in the tree are *synchronised* in the sense that for all $\rho \in \mathcal{P}$ we have $A' <_\rho B'$ if and only if $A'' <_\rho B''$. Pick $x_1 \in A'$, $x_2 \in B'$, $x_3 \in A''$, and $x_4 \in B''$. Then, since (A, B) is $\mathcal{P}$-ordered, we have $\{x_1, x_2\} <_\rho \{x_3, x_4\}$ or $\{x_3, x_4\} <_\rho \{x_1, x_2\}$ for all $\rho \in \mathcal{P}$, and since (A', B') and (A'', B'') are synchronised,

we have $x_1 <_\rho x_2$ if and only if $x_3 <_\rho x_4$. This implies that $\mathcal{P}$ induces at most 4 permutations on $\{x_1, x_2, x_3, x_4\}$ which is less than $2^{k-1} = 8$.

To find synchronised $\mathcal{P}$-ordered pairs, we use a Ramsey-theoretic argument about vertex-coloured binary trees. For that, we need to introduce some terminology. A *binary tree* T rooted at a vertex r is a tree where r has degree 0 or 2 and every other vertex has degree 1 or 3. A *leaf* of T is a vertex of degree at most 1. The *layer* of T at height h is the set $N_h(r) \subseteq V(T)$ of those vertices at graph distance exactly h from r. For a vertex $v \in N_h(r)$, its children are its neighbours in $N_{h+1}(r)$, and if $v \neq r$ then its parent is its neighbour in $N_{h-1}(r)$. Clearly, $|N_h(r)| \leqslant 2^h$. If there exists some h such that $|N_h(r)| = 2^h$ and $|N_{h+1}(r)| = 0$, then T is a *complete binary tree* of height h.

Let S be a second binary tree with root q. For us, a *subdivision* of S in T is an injective map $\varphi\colon V(S) \to V(T)$ such that for all vertices $v \in V(S) \setminus \{q\}$ with parent p, $\varphi(v)$ is contained in the subtree of T rooted at $\varphi(p)$. If the vertices of T are coloured (not necessarily properly), we say that the layers of the subdivision are *monochromatic* if $\varphi(N_h(q))$ is monochromatic for all h. We often identify $V(S)$ with $\varphi(V(S))$.

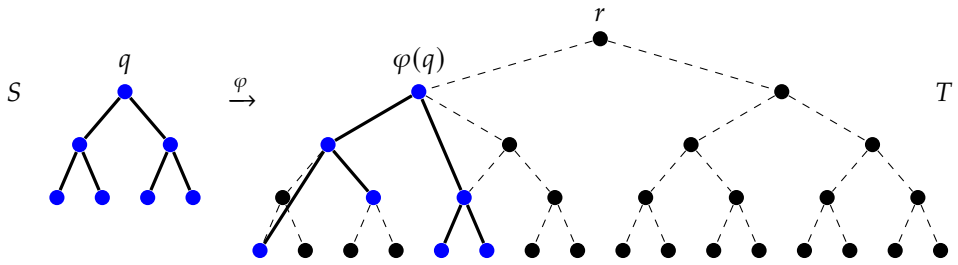


Figure 7.1. A subdivision of S inside T. Edges of T are dashed for better visibility.

In our proof, T will be a complete binary tree whose vertices correspond to nested $\mathcal{P}$-ordered pairs, and we colour the vertices of T in such a way that monochromatic layers correspond to collections of synchronised $\mathcal{P}$-ordered pairs. Then, if we find a subdivision of a complete binary tree of height $\lceil \log_2 k \rceil$ in T with monochromatic layers, we can pick one element from each of its leaves to obtain a set of size k that is only $2^{\lceil \log_2 k \rceil}$-shattered by $\mathcal{P}$, as required.

To do this, we need to find such subdivisions in large vertex-coloured complete binary trees. For all integers $c \geqslant 1$ and $h \geqslant 0$, let $g(c, h)$ be the smallest integer such that every complete binary tree of height $g(c, h)$ whose vertices are coloured with c colours contains a subdivision of a complete binary tree of height h with monochromatic layers. We prove the following upper bound on $g(c, h)$.

Lemma 7.9. *For all integers $c \geqslant 1$ and $h \geqslant 0$ we have $g(c, h) \leqslant h \left\lceil \log_2(c^{h-1} + 1) \right\rceil$.*

Proof. We proceed by induction on h. If $h = 0$, a complete binary tree of height h is a single vertex, and so $g(c, 0) = 0$.

If $h \geqslant 1$, let $d = \lceil \log_2(c^{h-1} + 1) \rceil$ and let T be a complete binary tree with root r and height $d + g(c, h-1)$ that is coloured with c colours. For each vertex $v \in N_d(r)$, the subtree of T rooted at v has height $g(c, h-1)$ and must therefore contain a subdivision S_v of a complete binary tree of height $h-1$ with monochromatic layers. Since such subdivisions admit at most c^{h-1} valid colourings and $|N_d(r)| = 2^d \geqslant c^{h-1} + 1$, there must be two subdivisions S_u and S_v whose colourings coincide. Then, the subdivision formed by S_u, S_v, and the lowest common ancestor of u and v is a subdivision of a complete binary tree of height h with monochromatic layers.

Therefore, we get that

$$g(c,h) \leqslant d + g(c,h-1) \leqslant d + (h-1)\left\lceil \log_2(c^{h-2}+1)\right\rceil \leqslant h\left\lceil \log_2(c^{h-1}+1)\right\rceil. \qquad \square$$

Using the strategy outlined above, we now prove our lower bound for $t > 2^{\lceil \log_2 k \rceil}$.

Proof of Theorem 7.5. Let $h = \lceil \log_2 k \rceil$ and $t > 2^h$. We claim $f_k(n,t) > \sqrt{\log_2 n}/h - 1$. Indeed, suppose that $\mathcal{P}$ is a family of m permutations of $[n]$ with $m \leqslant \sqrt{\log_2 n}/h - 1$, and so using Lemma 7.9 we get $\log_2 n \geqslant h^2(m+1)^2 \geqslant g(2^m, h)(m+1)$. We show that there is a subset of size k that is not t-shattered by $\mathcal{P}$.

Consider a complete binary tree T of height $g(2^m, h)$. We associate to each vertex v of T a set $X_v \subseteq [n]$ and a $\mathcal{P}$-ordered pair $A_v, B_v \subseteq X_v$ as follows. For the root r of T, let $X_r = [n]$ and let $A_r, B_r \subseteq X_r$ be a $\mathcal{P}$-ordered pair given by Lemma 7.8. Then, for a vertex v of T with $\mathcal{P}$-ordered pair (A_v, B_v) and children v_1 and v_2, let $X_{v_1} = A_v$ and $X_{v_2} = B_v$, and let $A_{v_1}, B_{v_1} \subseteq X_{v_1}$ and $A_{v_2}, B_{v_2} \subseteq X_{v_2}$ be $\mathcal{P}$-ordered pairs given by Lemma 7.8. This implies that $|X_w| \geqslant \left\lfloor n/2^{g(2^m,h)(m+1)} \right\rfloor \geqslant 1$ for every $w \in V(T)$, and so X_w is non-empty.

Colour each vertex $v \in V(T)$ with a colour c_v as follows. For $\rho \in \mathcal{P}$, let

$$c_v(\rho) = \begin{cases} 1 & \text{if } A_v <_\rho B_v, \\ 0 & \text{if } B_v <_\rho A_v. \end{cases}$$

Since (A_v, B_v) is $\mathcal{P}$-ordered, this is well-defined. This colours each vertex of T with one of 2^m colours. By Lemma 7.9, T admits a subdivision of a complete binary tree S of height h with monochromatic layers. Let q be the root of S.

For each leaf ℓ of S, pick $y_\ell \in X_\ell$, and let $Y = \{y_\ell : \ell \in N_h^S(q)\}$ be the set of all of these elements. Note that $|Y| = 2^h \geqslant k$. We claim that Y is not t-shattered. Indeed, consider

the colouring that a fixed permutation $\rho \in \mathcal{P}$ induces on the first h layers $N^S_{<h}(q)$ of S by assigning $c_v(\rho)$ to each $v \in N^S_{<h}(q)$. Since each layer of S is monochromatic, this colours $N^S_{<h}(q)$ with one of at most 2^h distinct colourings (across all $\rho \in \mathcal{P}$).

Now, for any two distinct leaves $\ell, \ell' \in N^S_h(q)$, the lowest common ancestor $v \in N^S_{<h}(q)$ of ℓ and ℓ' satisfies $y_\ell \in A_v$ and $y_{\ell'} \in B_v$ (or vice versa). Then, if $c_v(\rho) = 1$, we have $A_v <_\rho B_v$ and so $y_\ell <_\rho y_{\ell'}$, and if $c_v(\rho) = 0$ we have $B_v <_\rho A_v$ and so $y_{\ell'} <_\rho y_\ell$. Thus, the colours $c_v(\rho)$ for the vertices $v \in N^S_{<h}(q)$ uniquely determine the permutation that ρ induces on Y. Since there are at most 2^h distinct colourings of $N^S_{<h}(q)$, it follows that $\mathcal{P}$ induces at most 2^h permutations on Y, and therefore does not t-shatter Y. □

7.3 Upper bounds

To prove our upper bounds on $f_k(n,t)$, we use lexicographic permutations. Identify $[n]$ with a subset of $[b]^d$ where $b \geqslant 2$ and $d \geqslant \lceil \log_b n \rceil$ and define a *lex-permutation* ρ of $[b]^d$ to be a permutation that is constructed as follows. For each $i \in [d]$, pick a permutation ρ_i of $[b]$. Then, for distinct $x, y \in [b]^d$, set $x <_\rho y$ if and only if $x_i <_{\rho_i} y_i$ where $i = i(x,y) = \min\{i \in [d] : x_i \neq y_i\}$. That is, the relative order of x and y in ρ is entirely determined by the first position where x and y differ. It is easy to check that $<_\rho$ is a total order, and so ρ is a permutation.

One can motivate this definition by rephrasing the construction due to Hajnal and Spencer [113] that shows $f_k(n,t) = O(\log\log n)$ for $3 \leqslant t \leqslant k$ under the language of lexicographic permutations. Starting with $[m]$, they construct a maximal so-called k-scrambling family $\mathcal{S}$ for $[m]$ with a specified ordering of its elements and take their ground set $\mathcal{G}$ to be the power set $2^{\mathcal{S}}$ of $\mathcal{S}$. Then, each element $x \in [m]$ induces an ordering $<_x$ of $\mathcal{G}$ as follows. For $G_1, G_2 \in \mathcal{G}$, let S be the minimal element in the symmetric difference $G_1 \triangle G_2$ and set $G_1 <_x G_2$ if $x \in S$ and $G_2 <_x G_1$ otherwise. Crucially, they define the k-scrambling property to be that for any choice $S_1, \ldots, S_k$ of k elements of S and $A \subseteq [k]$, there is some $x \in [m]$ such that $x \in S_i$ precisely when $i \in A$. We can view $\mathcal{G} = 2^{\mathcal{S}}$ as a set of binary strings of length $|\mathcal{S}|$ where each coordinate corresponds to whether it contains a specific element $S \in \mathcal{S}$, and the constructed family of orderings $\{<_x : x \in [m]\}$ as a family of lex-permutations on $\mathcal{G}$ with the k-lex-shattering property defined below.[4]

[4]Although it may seem that our construction was directly motivated by that of Hajnal and Spencer, we first drew our inspiration from a construction of Eichhorn and Mubayi [50] and only noticed similarities a posteriori. We expand on this in Section 7.4.

We say that a family of permutations $\mathcal{P}$ of $[b]^d$ is *k-lex-shattering* if it is a family of lex-permutations such that the following holds. Let $\mathcal{I} \subseteq [d]$ be a subset of size at most k, and for each $i \in \mathcal{I}$ fix a permutation σ_i of a subset $Y_i \subseteq [b]$ of size at most k. Then, there exists a permutation $\rho \in \mathcal{P}$ such that ρ_i induces σ_i on Y_i for all $i \in \mathcal{I}$. We will later show that k-lex-shattering families $2^{\lceil \log_2 k \rceil}$-shatter every subset of size k. However, we first have to show that small k-lex-shattering families exist.

Lemma 7.10. *For $k \geqslant 1$, there exists a k-lex-shattering family of $[b]^d$ with size $\mathcal{O}(\log(bd))$.*

Proof. By Theorem 7.2, there is a family $\mathcal{Q}$ of permutations of $[bd]$ with size $\mathcal{O}(\log(bd))$ that shatters every subset of size k^2. For each $\tau \in \mathcal{Q}$, construct a lex-permutation ρ^τ such that ρ_i^τ is the permutation that τ induces on $(i-1)b + [b] = \{(i-1)b+1, \dots, ib\}$. We claim that $\mathcal{P} = \{\rho^\tau : \tau \in \mathcal{Q}\}$ is k-lex-shattering.

Indeed, let $\mathcal{I} \subseteq [d]$ have size at most k, and for every $i \in \mathcal{I}$ let σ_i be a permutation of a subset $Y_i \subseteq [b]$ of size at most k. Since $\mathcal{Q}$ shatters the set $\bigcup_{i \in \mathcal{I}}((i-1)b + Y_i)$, there exists a permutation $\tau \in \mathcal{Q}$ which induces σ_i on $(i-1)b + Y_i$ for all $i \in \mathcal{I}$. This implies that ρ_i^τ induces σ_i on Y_i for all $i \in \mathcal{I}$. □

To determine the permutations that a k-lex-shattering family $\mathcal{P}$ induces on a subset X, consider the lexicographic structure of X. Let $i(X) = \inf\{i(x,y) : x, y \in X\}$ be the first position where two elements of X differ, and let $S(X) = \{x_{i(X)} : x \in X\}$ denote the *slice* of X at that position if $i(X) < \infty$. This partitions X into the subsets $X_s = \{x \in X : x_{i(X)} = s\}$ for $s \in S(X)$. Note that the sets X_s are pairwise $\mathcal{P}$-ordered and that their relative order in a permutation $\rho \in \mathcal{P}$ is determined by the permutation that $\rho_{i(X)}$ induces on $S(X)$. We will denote this permutation by $\rho(X)$.

If we partition a set X_s in the same way, this yields a partition of X_s into $\mathcal{P}$-ordered sets whose relative order is determined by $\rho(X_s)$. Continuing recursively like this, it follows that the order that ρ induces on X is determined by the permutations $\rho(Y)$ for $Y \subseteq X$. This is captured by the following lemma.

Lemma 7.11. *Let ρ be a lex-permutation of $[b]^d$ and let $X \subseteq [b]^d$. Then, the permutation that ρ induces on X is uniquely determined by the set of permutations $\rho(Y)$ for $Y \subseteq X$.*

Proof. Let $x, y \in X$ and $i = i(x,y)$. Then, $Y = \{x,y\} \subseteq X$ and $\rho(Y)$ is the permutation that ρ_i induces on $S(Y) = \{x_i, y_i\}$. This implies that $x <_\rho y$ if and only if $x_i <_{\rho_i} y_i$ which in turn is equivalent to $x_i <_{\rho(Y)} y_i$. So, the permutation that ρ induces on X is determined by the set of permutations $\rho(Y)$ for $Y \subseteq X$.

Conversely, suppose that ρ and τ are lex-permutations with $\rho(Y) \neq \tau(Y)$ for some $Y \subseteq X$. Let $i = i(Y)$. Since $\rho(Y) \neq \tau(Y)$, there must exist $x, y \in Y$ with $i(x,y) = i$ such that $x_i <_{\rho(Y)} y_i$ and $y_i <_{\tau(Y)} x_i$ (or vice versa). As above, this is equivalent to $x <_\rho y$ and $y <_\tau x$, and so ρ and τ induce different permutations on X. □

Next, we deduce a lower bound on the number of permutations that $\mathcal{P}$ induces on X. Let $\mathcal{I}(X) = \{i(Y) : Y \subseteq X \text{ and } i(Y) < \infty\}$. We rely on the following oberservation.

Lemma 7.12. *Let $\mathcal{P}$ be a k-lex-shattering family of permutations of $[b]^d$ and let $X \subseteq [b]^d$ be a subset of size k. For each $i \in \mathcal{I}(X)$, let $Y_i \subseteq X$ be such that $i(Y_i) = i$. Then, X is $(\prod_{i \in \mathcal{I}(X)} |S(Y_i)|!)$-shattered by $\mathcal{P}$.*

Proof. We claim that $|\mathcal{I}(X)| < k$. Indeed, write $<$ for the standard lexicographic order on $[b]^d$, so $x < y$ if and only if $x_i < y_i$ where $i = i(x,y)$. Note that if $x, y, z \in [b]^d$ satisfy $x < y < z$, then $i(x,z) = \min\{i(x,y), i(y,z)\}$. Therefore, if we write $X = \{x_1, \dots, x_k\}$ with $x_1 < \dots < x_k$, this implies that $\mathcal{I}(X) = \{i(x_\ell, x_{\ell+1}) : \ell \in [k-1]\}$, and so $|\mathcal{I}(X)| \leqslant k-1$ as claimed.

Since $\mathcal{P}$ is k-lex-shattering, it follows that if σ_i is a permutation of $S(Y_i)$ for every $i \in \mathcal{I}(X)$, then there is a permutation $\rho \in \mathcal{P}$ with $\rho(Y_i) = \sigma_i$ for all $i \in \mathcal{I}(X)$. By Lemma 7.11, each choice of the permutations σ_i induces a different permutation on X, and so X is $(\prod_{i \in \mathcal{I}(X)} |S(Y_i)|!)$-shattered by $\mathcal{P}$. □

To maximise this product, consider again the decomposition of X into its lexicographic structure from above. Suppose that we only pick the largest set X_s whenever we recursively continue the decomposition, and let the sets obtained during this process be the sets Y_i. Then, we will show that $\prod_{i \in \mathcal{I}} |S(Y_i)|! \geqslant 2^{\lceil \log_2 k \rceil}$, and so every set of size k is $2^{\lceil \log_2 k \rceil}$-shattered by $\mathcal{P}$. This suffices to prove Theorem 7.6.

For Theorem 7.7, we will additionally infer some structural information about X whenever X is not $2k$-shattered by $\mathcal{P}$. Later, this information will allow us to add more permutations to $\mathcal{P}$ which ensure that X is t-shattered for some $t > 2^{\lceil \log_2 k \rceil}$. Overall, we obtain the following lemma.

Lemma 7.13. *Let $\mathcal{P}$ be a k-lex-shattering family of permutations of $[b]^d$, and let $h = \lceil \log_2 k \rceil$. Then, $\mathcal{P}$ is a family that 2^h-shatters all subsets of $[b]^d$ of size k. Moreover, if $X \subseteq [b]^d$ is a subset of size k that $\mathcal{P}$ does not $2k$-shatter, then the set $\mathcal{I}(X)$ has size h, all $Y \subseteq X$ with $i(Y) < \infty$ satisfy $|S(Y)| = 2$, and all $Y, Z \subseteq X$ with $i(Y) = i(Z) < \infty$ satisfy $S(Y) = S(Z)$.*

Proof. Let $X \subseteq [b]^d$ be a subset of size k and let $\mathcal{I} = \mathcal{I}(X)$.

Claim 7.14. *There exist sets $Y_i \subseteq X$ with $i(Y_i) = i$ for each $i \in \mathcal{I}$ such that $\prod_{i\in\mathcal{I}} |S(Y_i)| \geqslant k$.*

Proof. We prove by induction on j that for all $0 \leqslant j \leqslant d$ there exist sets $Y_i \subseteq X$ with $i(Y_i) = i$ for each $i \in \mathcal{I} \cap [j]$ and a subset $Z \subseteq X$ which satisfy $i(Z) > j$ and $|Z| \cdot \prod_{i\in\mathcal{I}\cap[j]} |S(Y_i)| \geqslant k$. Applying this with $j = d$ proves the claim.

For $j = 0$, let $Z = X$. Now suppose we have constructed such sets for $j-1$. If $i(Z) > j$, the same sets work for j, where we choose any $Y_j \subseteq X$ with $i(Y_j) = j$ if $j \in \mathcal{I}$.

Otherwise, $i(Z) = j$. Let $Y_j = Z$, and for every $s \in S(Z)$ let $Z_s = \{z \in Z : z_j = s\}$. Pick $s \in S(Z)$ such that $|Z_s| \geqslant |Z|/|S(Z)|$. Note that $i(Z_s) > j$ and $|Z_s| \cdot \prod_{i\in\mathcal{I}\cap[j]} |S(Y_i)| \geqslant (|Z|/|S(Z)|) \cdot \prod_{i\in\mathcal{I}\cap[j-1]} |S(Y_i)| \cdot |S(Z)| \geqslant k$. So we have constructed the sets for j. ■

For all $i \in \mathcal{I}$, let $Y_i \subseteq X$ be a set with $i(Y_i) = i$ that maximises $|S(Y_i)|$. By Claim 7.14, $\prod_{i\in\mathcal{I}} |S(Y_i)| \geqslant k$, and by Lemma 7.12 $\mathcal{P}$ induces at least $\prod_{i\in\mathcal{I}} |S(Y_i)|!$ permutations on X. If $|S(Y_i)| \geqslant 3$ for some $i \in \mathcal{I}$, then $\prod_{i\in\mathcal{I}} |S(Y_i)|! \geqslant 2\prod_{i\in\mathcal{I}} |S(Y_i)| \geqslant 2k \geqslant 2^h$. Otherwise, $|S(Y_i)| = 2$ for all $i \in \mathcal{I}$. Then, since $2^{|\mathcal{I}|} = \prod_{i\in\mathcal{I}} |S(Y_i)| \geqslant k$, we have $|\mathcal{I}| \geqslant h$ and so $\prod_{i\in\mathcal{I}} |S(Y_i)|! = 2^{|\mathcal{I}|} \geqslant 2^h$. This shows that X is 2^h-shattered.

Moreover, if X is not $2k$-shattered, these arguments imply that $|S(Y_i)| = 2$ for all $i \in \mathcal{I}$ and $|\mathcal{I}| = h$. Suppose that for some $j \in \mathcal{I}$ there exists a set $Z \subseteq X$ with $i(Z) = j$ and $S(Y_j) \neq S(Z)$. Then, by the maximality of Y_j, $|S(Z)| = 2$ and so $|S(Y_j) \cap S(Z)| \leqslant 1$. In particular, for any pair of permutations σ of $S(Y_j)$ and τ of $S(Z)$, there is a permutation σ_j of $S(Y_j) \cup S(Z)$ that induces σ on $S(Y_j)$ and τ on $S(Z)$. By the same argument as in the proof of Lemma 7.12, it follows that $\mathcal{P}$ induces at least $|S(Z)|! \cdot \prod_{i\in\mathcal{I}} |S(Y_i)|! \geqslant 2k$ permutations on X. This contradicts the fact that X is not $2k$-shattered. □

With $b = 2$ and $d = \lceil \log_2 n \rceil$, Theorem 7.6 is an immediate consequence of Lemmas 7.10 and 7.13. We remark that while the description of our construction differs significantly from that of Spencer [113], these constructions are essentially equivalent.

To prove Theorem 7.7, start with a k-lex-shattering family $\mathcal{P}$. By Lemma 7.13, most subsets of size k are $2k$-shattered by $\mathcal{P}$, and even if a subset X is not $2k$-shattered, it is nevertheless $2^{\lceil \log_2 k \rceil}$-shattered and it has a very specific structure. We exploit this structure by adding new permutations to $\mathcal{P}$ which induce at least four additional permutation on X and thereby ensure that X is $(2^{\lceil \log_2 k \rceil} + 4)$-shattered.

To do so, we add a constant number of new permutations for each position $i \in [d]$. Let $<$ and $>$ denote the standard and reverse permutations of $[b]$.[5] Then, for all $\sigma, \tau \in \{<, >\}$, let $\pi_{i,\sigma,\tau}$ be the permutation of $[b]^d$ with $x <_{\pi_{i,\sigma,\tau}} y$ if and only if either $x_i <_\sigma y_i$, or

[5] That is, for $x, y \in [b]$ set $x <_< y$ if and only if $x < y$, and $x <_> y$ if and only if $x > y$.

$x_i = y_i$ and $x_{i(x,y)} <_\tau y_{i(x,y)}$. That is, $\pi_{i,\sigma,\tau}$ first sorts according to position i and only afterwards behaves like a lex-permutation. Define

$$\mathcal{Q}_i = \{\pi_{i,\sigma,\tau} : \sigma, \tau \in \{<, >\}\}.$$

If X is not $2k$-shattered, we show that for an appropriate $i \in [d]$ the permutations of $\mathcal{Q}_i$ induce four additional permutations on X, as required.

Proof of Theorem 7.7. Let $b = 2^d$, $d = \lceil \sqrt{\log_2 n} \rceil$, and $h = \lceil \log_2 k \rceil$. By Lemma 7.10, there exists a k-lex-shattering family $\mathcal{P}$ of $[b]^d$ with size $\mathcal{O}(\sqrt{\log n})$. Define $\mathcal{R} = \mathcal{P} \cup \bigcup_{i \in [d]} \mathcal{Q}_i$ and note that $\mathcal{R}$ has size $\mathcal{O}(\sqrt{\log n})$.

We claim that every subset $X \subseteq [b]^d$ of size k is $\min\{2k, 2^h + 4\}$-shattered by $\mathcal{R}$. Indeed, suppose that X is not $2k$-shattered by $\mathcal{P}$. Then, by Lemma 7.13, we know that the set $\mathcal{I}(X)$ has size h, all $Y \subseteq X$ with $i(Y) < \infty$ satisfy $|S(Y)| = 2$, and all $Y, Z \subseteq X$ with $i(Y) = i(Z) < \infty$ satisfy $S(Y) = S(Z)$.

Write $X = \{x_1, \ldots, x_k\}$ with $x_1 < \ldots < x_k$. Note that $i(x_\ell, x_{\ell+1}) \in \mathcal{I}(X)$ for all $\ell \in [k-1]$. Since $k \geqslant 4$, one can check that $k - 1 > h = |\mathcal{I}(X)|$. So there must exist $1 \leqslant \ell < m < k$ with $i(x_\ell, x_{\ell+1}) = i(x_m, x_{m+1})$. Let $i = i(x_\ell, x_{\ell+1})$ and $j = i(x_{\ell+1}, x_m)$.

If $j \geqslant i$, then $i(x_{\ell+1}, x_{m+1}) = i$ and $i(x_\ell, x_{m+1}) = i$. So, $Y = \{x_\ell, x_{\ell+1}, x_{m+1}\} \subseteq X$ satisfies $|S(Y)| = 3$. This contradicts the fact that all $Y \subseteq X$ with $i(Y) < \infty$ satisfy $|S(Y)| = 2$. Therefore, $j < i$. In particular, $(x_\ell)_j = (x_{\ell+1})_j < (x_m)_j = (x_{m+1})_j$, and so $(\{x_\ell, x_{\ell+1}\}, \{x_m, x_{m+1}\})$ is $\mathcal{P}$-ordered.

Note that $i(\{x_\ell, x_{\ell+1}\}) = i = i(\{x_m, x_{m+1}\})$. Since all $Y, Z \subseteq X$ with $i(Y) = i(Z) < \infty$ satisfy $S(Y) = S(Z)$, we must therefore have $(x_\ell)_i = (x_m)_i < (x_{\ell+1})_i = (x_{m+1})_i$, and so $(\{x_\ell, x_m\}, \{x_{\ell+1}, x_{m+1}\})$ is $\mathcal{Q}_i$-ordered.

In particular, $\mathcal{P}$ and $\mathcal{Q}_i$ induce different permutations on X. Moreover, note that $\{x_\ell, x_{\ell+1}, x_m, x_{m+1}\}$ is 4-shattered by $\mathcal{Q}_i$. Since Lemma 7.13 implies that X is 2^h-shattered by $\mathcal{P}$, it follows that X is $(2^h + 4)$-shattered by $\mathcal{R}$. □

7.4 Open problems and discussion

In this chapter, we have shown that at least four regimes exist for the asymptotic behaviour of $f_k(n, t)$ when $k \geqslant 4$, and we narrowed the range of values of t for which the asymptotic behaviour is unknown. The main open problem is to determine the asymptotic behaviour of $f_k(n, t)$ for $k \geqslant 5$ and $\min\{2k, 2^{\lceil \log_2 k \rceil} + 4\} < t \leqslant 2^{k-1}$. The following remains possible.

Question 7.15. *For all integers k and t, is the asymptotic behaviour of $f_k(n,t)$ either $\Theta(1)$, $\Theta(\log\log n)$, $\Theta(\sqrt{\log n})$, or $\Theta(\log n)$?*

We conjecture that Theorem 7.4 determines the entire range of values of t that satisfy $f_k(n,t) = \Theta(\log n)$.

Conjecture 7.16. *Let $k \geqslant 3$ and $t \leqslant 2^{k-1}$. Then, $f_k(n,t) = o(\log n)$.*

Since we did not try to optimise the upper bound on t in Theorem 7.7, we believe that constructions similar to ours could prove that $f_k(n,t) = o(\log n)$ for a larger range of values of t. However, they seem to be far away from reaching $t = 2^{k-1}$.

In this context it seems important to mention the following inspiration for our approach that comes from a problem of Erdős and Gyárfás [52]. If $\mathcal{P}$ is a family of permutations of $[n]$, assign a colour $c_{x,y}$ to every pair $x, y \in [n]$ as follows. For $\rho \in \mathcal{P}$, let

$$c_{x,y}(\rho) = \begin{cases} 1 & \text{if } x <_\rho y, \\ 0 & \text{otherwise.} \end{cases}$$

This colouring uses $2^{|\mathcal{P}|}$ colours, and if a subset $X \subseteq [n]$ of size k spans at most ℓ colours, then $\mathcal{P}$ induces at most 2^ℓ permutations on X. We implicitely used this approach to prove Theorems 7.4 and 7.5. Unfortunately, the converse does not hold: X can simultaneously span at least ℓ colours and fail to be 2^ℓ-shattered by $\mathcal{P}$.

Nevertheless, Eichhorn and Mubayi [50] gave a colouring of K_n with $2^{\Theta(\sqrt{\log n})}$ colours such that all subsets of size k span at least $\ell = 2\lceil \log_2 k \rceil - 2$ colours. This inspired our construction for Theorem 7.7, even if our construction does not 2^ℓ-shatter all subsets of size k. Note that Conlon, Fox, Lee, and Sudakov [35] gave a colouring of K_n with $2^{o(\log n)}$ colours such that all subsets of size k span at least $k-1$ colours. We hope that these colourings could inspire further progress towards better upper bounds on $f_k(n,t)$.

Bibliography

[1] Kartick Adhikari and Sukrit Chakraborty. Shotgun assembly of Linial-Meshulam model. *arXiv preprint arXiv:2209.10942*, 2022.

[2] Kartick Adhikari and Sukrit Chakraborty. Shotgun assembly of random geometric graphs. *arXiv preprint arXiv:2202.02968*, 2022.

[3] Karim Adiprasito and Bruno Benedetti. Collapsibility of CAT(0) spaces. *Geometriae Dedicata*, 206(1):181–199, 2020.

[4] Noga Alon, Guoli Ding, Bogdan Oporowski, and Dirk Vertigan. Partitioning into graphs with only small components. *Journal of Combinatorial Theory, Series B*, 87(2):231–243, 2003.

[5] Noga Alon, Guoli Ding, Bogdan Oporowski, and Dirk Vertigan. Partitioning into graphs with only small components. *J. Combin. Theory Ser. B*, 87(2):231–243, 2003.

[6] Omer Angel, Itai Benjamini, and Nizan Horesh. An isoperimetric inequality for planar triangulations. *Discrete & Computational Geometry*, 59:802–809, 2018.

[7] Kenneth I. Appel and Wolfgang Haken. *Every planar map is four colorable*, volume 98. American Mathematical Soc., 1989.

[8] Richard Arratia, Daniela Martin, Gesine Reinert, and Michael S Waterman. Poisson process approximation for sequence repeats, and sequencing by hybridization. *Journal of Computational Biology*, 3(3):425–463, 1996.

[9] K. J. Asciak, M. A. Francalanza, J. Lauri, and W. Myrvold. A survey of some open questions in reconstruction numbers. *Ars Combinatoria*, 97:443–456, 2010.

[10] Armen S. Asratian, Carl Johan Casselgren, and Petros A. Petrosyan. Decomposing graphs into interval colorable subgraphs and no-wait multi-stage schedules. *Discrete Applied Mathematics*, 2022.

[11] Armen S. Asratian and Raffi R. Kamalian. Interval colorings of edges of a multigraph (in Russian). *Applied Mathematics 5, Yerevan State University*, pages 25–34, 1987. For an English translation see *arXiv:1401.8079*.

[12] Armen S. Asratian and Raffi R. Kamalian. Investigation on interval edge-colorings of graphs. *Journal of Combinatorial Theory, Series B*, 62(1):34–43, 1994.

[13] Caelan Atamanchuk, Luc Devroye, and Massimo Vicenzo. An algorithm to recover shredded random matrices. *SIAM Journal on Discrete Mathematics*, 38(3):2509–2529, 2024.

[14] Maria Axenovich, António Girão, Lawrence Hollom, Julien Portier, Emil Powierski, Michael Savery, Youri Tamitegama, and Leo Versteegen. A note on interval colourings of graphs. *European Journal of Combinatorics*, 120:103956, 2024.

[15] Maria Axenovich, António Girão, Emil Powierski, Michael Savery, and Youri Tamitegama. A note on the interval colouring thickness of graphs. *arXiv preprint arXiv:2303.04782*, 2023.

[16] Maria Axenovich and Michael Zheng. Interval colorings of graphs—coordinated and unstable no-wait schedules. *Journal of Graph Theory*, 104(4):757–768, 2023.

[17] Rahil Baber, Natalie Behague, Asier Calbet, David Ellis, Joshua Erde, Ron Gray, Maria-Romina Ivan, Barnabás Janzer, Robert Johnson, Luka Milićević, et al. A collection of open problems in celebration of Imre Leader's 60th birthday. *arXiv preprint arXiv:2310.18163*, 2023.

[18] Paul Balister, Gal Kronenberg, Alex Scott, and Youri Tamitegama. Reconstruction of shredded random matrices. *arXiv preprint arXiv:2401.05058*, 2024.

[19] I. Benjamini. Personal communication, 2020.

[20] Umberto Bertelè and Francesco Brioschi. *Nonserial dynamic programming*. Academic Press, Inc., 1972.

[21] G. Besson, G. Courtois, and S. Gallot. Entropies et rigidités des espaces localement symétriques de courbure strictment négative. *Geometric and Functional Analysis*, 5:731–799, 1995.

[22] Martín Blufstein, Jérémie Chalopin, and Victor Chepoi. Boundary rigidity of systolic and Helly complexes. *arXiv preprint arXiv:2506.13219*, 2025.

[23] Béla Bollobás. Almost every graph has reconstruction number three. *Journal of Graph Theory*, 14(1):1–4, 1990.

[24] Béla Bollobás. *Random graphs*, volume 73 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, second edition, 2001.

[25] Béla Bollobás and Imre Leader. Edge-isoperimetric inequalities in the grid. *Combinatorica*, 11(4):299–314, 1991.

[26] J. A. Bondy. A graph reconstructor's manual. In *Surveys in combinatorics, 1991 (Guildford, 1991)*, volume 166 of *London Math. Soc. Lecture Note Ser.*, pages 221–252. Cambridge Univ. Press, Cambridge, 1991.

[27] J. A. Bondy and R. L. Hemminger. Graph reconstruction—a survey. *Journal of Graph Theory*, 1(3):227–268, 1977.

[28] Marta Borowiecka-Olszewska, Ewa Drgas-Burchardt, Nahid Yelene Javier-Nol, and Rita Zuazua. Consecutive colouring of oriented graphs. *Results in Mathematics*, 76:200, 2021.

[29] Dmitri Burago and Sergei Ivanov. Boundary rigidity and filling volume minimality of metrics close to a flat one. *Annals of Mathematics*, 171:1183–1211, 2010.

[30] Rutger Campbell, J. Pascal Gollin, Kevin Hendrey, Thomas Lesgourgues, Bojan Mohar, Youri Tamitegama, Jane Tan, and David R Wood. Clustered colouring of graph products. *arXiv preprint arXiv:2407.21360*, 2024.

[31] Johannes Carmesin, Reinhard Diestel, Matthias Hamann, and Fabian Hundertmark. k-blocks: a connectivity invariant for graphs. *SIAM Journal on Discrete Mathematics*, 28(4):1876–1891, 2014.

[32] Carl Johan Casselgren and Bjarne Toft. On interval edge colorings of biregular bipartite graphs with small vertex degrees. *Journal of Graph theory*, 80(2):83–97, 2015.

[33] Jérémie Chalopin and Victor Chepoi. Boundary rigidity of CAT(0) cube complexes. *Journal of Combinatorial Theory, Series B*, 169:352–366, 2024.

[34] Julia Chuzhoy and Zihan Tan. Towards tight (er) bounds for the excluded grid theorem. *Journal of Combinatorial Theory, Series B*, 146:219–265, 2021.

[35] David Conlon, Jacob Fox, Choongbum Lee, and Benny Sudakov. The Erdős–Gyárfás problem on generalized Ramsey numbers. *Proceedings of the London Mathematical Society*, 110(1):1–18, 2015.

[36] Alice M. Dean, Joan P. Hutchinson, and Edward R. Scheinerman. On the thickness and arboricity of a graph. *Journal of Combinatorial Theory, Series B*, 52(1):147–151, 1991.

[37] Michelle Delcourt and Luke Postle. Reducing linear Hadwiger's conjecture to coloring small graphs. *Journal of the American Mathematical Society*, 38(2):481–507, 2025.

[38] Guoli Ding and Bogdan Oporowski. Some results on tree decomposition of graphs. *J. Graph Theory*, 20(4):481–499, 1995.

[39] Guoli Ding and Bogdan Oporowski. On tree-partitions of graphs. *Discrete Mathematics*, 149(1-3):45–58, 1996.

[40] Guoli Ding, Bogdan Oporowski, Daniel P. Sanders, and Dirk Vertigan. Partitioning into two graphs with only small components. *Preprint*, 1996.

[41] Jian Ding, Yiyang Jiang, and Heng Ma. Shotgun threshold for sparse Erdős–Rényi graphs. *IEEE Transactions on Information Theory*, 69(11):7373–7391, 2023.

[42] Marc Distel, Kevin Hendrey, Nikolai Karol, David R. Wood, and Jung Hon Yip. Treewidth 2 in the planar graph product structure theorem. *Discrete Math. Theoret. Comput. Sci.*, 27:2:#8, 2025.

[43] Marc Distel, Robert Hickingbotham, Tony Huynh, and David R. Wood. Improved product structure for graphs on surfaces. *Discrete Math. Theor. Comput. Sci.*, 24(2):#6, 2022.

[44] Vida Dujmović, Gwenaël Joret, Piotr Micek, Pat Morin, Torsten Ueckerdt, and David R Wood. Planar graphs have bounded queue-number. *Journal of the ACM (JACM)*, 67(4):1–38, 2020.

[45] Vida Dujmović, Pat Morin, and David R. Wood. Graph product structure for non-minor-closed classes. *J. Combin. Theory Ser. B*, 162:34–67, 2023.

[46] Zdeněk Dvořák and Sergey Norin. Treewidth of graphs with balanced separations. *Journal of Combinatorial Theory, Series B*, 137:137–144, 2019.

[47] Zdeněk Dvořák and David R. Wood. Product structure of graph classes with strongly sublinear separators. *arXiv preprint arXiv:2208.10074*, 2022.

[48] Martin Dyer, Alan Frieze, and Stephen Suen. The probability of unique solutions of sequencing by hybridization. *Journal of Computational Biology*, 1(2):105–110, 1994.

[49] Katherine Edwards, Dong Yeap Kang, Jaehoon Kim, Sang-il Oum, and Paul Seymour. A relative of Hadwiger's conjecture. *SIAM Journal on Discrete Mathematics*, 29(4):2385–2388, 2015.

[50] Dennis Eichhorn and Dhruv Mubayi. Note-edge-coloring cliques with many colors on subcliques. *Combinatorica*, 20(3):441–444, 2000.

[51] Paul Erdős and Siemion Fajtlowicz. On the conjecture of Hajós. *Combinatorica*, 1(2):141–143, 1981.

[52] Paul Erdős and András Gyárfás. A variant of the classical Ramsey problem. *Combinatorica*, 17:459–467, 1997.

[53] Louis Esperet and David R. Wood. Colouring strong products. *European J. Combin.*, page 103847, 2023.

[54] Zoltán Füredi. Scrambling permutations and entropy of hypergraphs. *Random Structures & Algorithms*, 8(2):97–104, 1996.

[55] Julia Gaudio and Elchanan Mossel. Shotgun assembly of Erdős-Rényi random graphs. *Electronic Communications in Probability*, 27:Paper No. 5, 14, 2022.

[56] Krzysztof Giaro and Marek Kubale. Consecutive edge-colorings of complete and incomplete Cartesian products of graphs. *Congressus Numerantium*, 128:143–149, 1997.

[57] António Girão and Zach Hunter. Induced subdivisions in $K_{s,s}$-free graphs with polynomial average degree. *International Mathematics Research Notices*, 2025(4):rnaf025, 2025.

[58] António Girão, Lukas Michel, and Youri Tamitegama. Small families of partially shattering permutations. *arXiv preprint arXiv:2407.05773*, 2024.

[59] György Hajós. Uber eine konstruktion nicht n-farbbarer graphen. *Martin-Luther-Univ, Halle-Wittenberg*, 1961.

[60] Arsen Hambardzumyan and Levon Muradyan. On interval edge-colorings of planar graphs. *arXiv preprint arXiv:2303.11466*, 2023.

[61] H. M. Hansen. Scheduling with minimum waiting periods (in Danish). Master's thesis, Odense University, Odense, Denmark, 1992.

[62] Denis Hanson, C. O. M. Loten, and Bjarne Toft. On interval colourings of bi-regular bipartite graphs. *Ars Combinatoria*, 50:23–32, 1998.

[63] John Haslegrave. Determining triangulations and quadrangulations by boundary distances. *Journal of Combinatorial Theory, Series B*, 163:233–255, 2023.

[64] John Haslegrave, Alex Scott, Youri Tamitegama, and Jane Tan. Boundary rigidity of 3D CAT(0) cube complexes. *European Journal of Combinatorics*, 124:104077, 2025.

[65] Allen Hatcher. *Algebraic Topology*. Cambridge University Press, 2002.

[66] Gabor T. Herman and Attila Kuba. *Discrete Tomography: Foundations, Algorithms, and Applications*. Birkhäuser, Boston, MA, 1999.

[67] Robert Hickingbotham and David R. Wood. Shallow minors, graph products and beyond-planar graphs. *SIAM J. Discrete Math.*, 38(1):1057–1089, 2024.

[68] Robert Hickingbotham and David R. Wood. Structural properties of graph products. *J. Graph Theory*, 109(2):107–136, 2025.

[69] Lawrence Hollom, Julien Portier, and Leo Versteegen. On interval colourings of graphs. *arXiv preprint arXiv:2303.05505*, 2023.

[70] Han Huang and Konstantin Tikhomirov. Shotgun assembly of unlabeled Erdős–Rényi graphs. *Probability Theory and Related Fields*, pages 1–50, 2025.

[71] Freddie Illingworth, Alex Scott, and David R. Wood. Product structure of graphs with an excluded minor. *Trans. Amer. Math. Soc. Ser. B*, 11:1233–1248, 2024.

[72] Freddie Illingworth and David Wood. Dominating K_t-models. *arXiv:2405.14299* preprint, 2024.

[73] Yoshiyasu Ishigami. An extremal problem of d permutations containing every permutation of every t elements. *Discrete Mathematics*, 159(1-3):279–283, 1996.

[74] Tommy R. Jensen and Bjarne Toft. *Graph Coloring Problems*. John Wiley & Sons, 1995.

[75] J. Robert Johnson and Belinda Wickes. Shattering k-sets with permutations. *Order*, pages 1–18, 2023.

[76] Tom Johnston, Gal Kronenberg, Alexander Roberts, and Alex Scott. Shotgun assembly of random graphs. *Probability Theory and Related Fields*, pages 1–40, 2025.

[77] Mark Kac. Can one hear the shape of a drum? *American Mathematical Monthly*, 73(4):1–23, 1966.

[78] Paul Joseph Kelly. *On isometric transformations*. PhD thesis, University of Wisconsin, 1942.

[79] Jon Kleinberg, Rajeev Motwani, Prabhakar Raghavan, and Suresh Venkatasubramanian. Storage management for evolving databases. In *Proceedings 38th Annual Symposium on Foundations of Computer Science*, pages 353–362. IEEE, 1997.

[80] Daniel J. Kleitman and Joel Spencer. Families of k-independent sets. *Discrete Mathematics*, 6(3):255–262, 1973.

[81] A. V. Kostochka. On $K_{s,t}$ minors in (s+t)-chromatic graphs. *Journal of Graph Theory*, 65(4):343–350, 2010.

[82] Alexandr V. Kostochka. The minimum Hadwiger number for graphs with a given mean degree of vertices. *Metody Diskret. Analiz.*, (38):37–58, 1982.

[83] Alexandr V. Kostochka. Lower bound of the Hadwiger number of graphs by their average degree. *Combinatorica*, 4(4):307–316, 1984.

[84] Alexandr V. Kostochka. $K_{s,t}$ minors in (s+t)-chromatic graphs, II. *J. Graph Theory*, 75(4):377–386, 2014.

[85] Josef Lauri and Raffaele Scapellato. *Topics in graph automorphisms and reconstruction*, volume 432 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, Cambridge, second edition, 2016.

[86] Nathan Linial, Jiří Matoušek, Or Sheffet, and Gábor Tardos. Graph colouring with no large monochromatic components. *Combin. Probab. Comput.*, 17(4):577–589, 2008.

[87] R. J. Lipton and Robert Endre Tarjan. A planar separator theorem. *SIAM Journal of Applied Mathematics*, 36(2):177–189, 1979.

[88] Chun-Hung Liu, Sergey Norin, and David R. Wood. Product structure and tree decompositions. *arXiv preprint arXiv:2410.20333*, 2024.

[89] Jiří Matoušek and Jaroslav Nešetřil. *Invitation to discrete mathematics*. Clarendon Press, 1998.

[90] Jiří Matoušek and Aleš Přívětivý. Large monochromatic components in two-colored grids. *SIAM J. Discrete Math.*, 22(1):295–311, 2008.

[91] R. Michel. Sur la rigidité imposée par la longueur des géodésiques. *Inventiones Mathematicae*, 65:71–83, 1981.

[92] Michael Mitzenmacher and Eli Upfal. *Probability and computing*. Cambridge University Press, Cambridge, second edition, 2017. Randomization and probabilistic techniques in algorithms and data analysis.

[93] Elchanan Mossel and Nathan Ross. Shotgun assembly of labeled graphs. *IEEE Transactions on Network Science and Engineering*, 6(2):145–157, 2019.

[94] Elchanan Mossel and Nike Sun. Shotgun assembly of random regular graphs. *arXiv preprint arXiv:1512.08473*, 2015.

[95] Abolfazl S. Motahari, Guy Bresler, and David N. C. Tse. Information theory of DNA shotgun sequencing. *Institute of Electrical and Electronics Engineers. Transactions on Information Theory*, 59(10):6273–6289, 2013.

[96] Vladimir Müller. Probabilistic reconstruction from subgraphs. *Commentationes Mathematicae Universitatis Carolinae*, 17(4):709–719, 1976.

[97] Sergey Norin, Luke Postle, and Zi-Xia Song. Breaking the degeneracy barrier for coloring graphs with no K_t minor. *Advances in Mathematics*, 422:109020, 2023.

[98] Sergey Norin and Zi-Xia Song. Breaking the degeneracy barrier for coloring graphs with no K_t minor. *arXiv preprint arXiv:1910.09378v1*, 2019.

[99] L. Pestov and G. Uhlmann. Two dimensional compact simple Riemannian manifolds are boundary distance rigid. *Annals of Mathematics*, 161:1093–1110, 2005.

[100] Petros A. Petrosyan and Hrant H. Khachatrian. Interval non-edge-colorable bipartite graphs and multigraphs. *Journal of Graph Theory*, 76(3):200–216, 2014.

[101] Jaikumar Radhakrishnan. A note on scrambling permutations. *Random Structures & Algorithms*, 22(4):435–439, 2003.

[102] Bruce Reed and Paul Seymour. Fractional colouring and Hadwiger's conjecture. *Journal of combinatorial theory, Series B*, 74(2):147–152, 1998.

[103] Herbert Robbins. A remark on Stirling's formula. *American Mathematical Monthly*, 62(1):26–29, 1955.

[104] Neil Robertson, Paul Seymour, and Robin Thomas. Hadwiger's conjecture for K_6-free graphs. *Combinatorica*, 13(3):279–361, 1993.

[105] Neil Robertson and Paul D. Seymour. Graph minors. II. algorithmic aspects of tree-width. *Journal of algorithms*, 7(3):309–322, 1986.

[106] Neil Robertson and Paul D. Seymour. Graph minors. V. excluding a planar graph. *Journal of Combinatorial Theory, Series B*, 41(1):92–114, 1986.

[107] Vojtech Rödl and Beata Wysocka. Note on regular subgraphs. *Journal of Graph Theory*, 24(2):139–154, 1997.

[108] Michah Sageev. Ends of group pairs and non-positively curved cube complexes. *Proceedings of the London Mathematical Society*, 3(3):585–617, 1995.

[109] Norbert Sauer. On the density of families of sets. *Journal of Combinatorial Theory, Series A*, 13(1):145–147, 1972.

[110] Alex Scott and Paul Seymour. A survey of χ-boundedness. *Journal of Graph Theory*, 95(3):473–504, 2020.

[111] S. V. Sevastianov. Interval colorability of the edges of a bipartite graph (in Russian). *Metody Diskretnogo Analiza*, 50:61–72, 1990. For an English translation see https://www.math.kit.edu/iag6/~axenovich/seite/publications/media/sevastianov-translation.pdf.

[112] Saharon Shelah. A combinatorial problem; stability and order for models and theories in infinitary languages. *Pacific Journal of Mathematics*, 41(1):247–261, 1972.

[113] Joel H. Spencer. Minimal scrambling sets of simple orders. *Acta Mathematica Academiae Scientiarum Hungaricae*, 22:349–353, 1972.

[114] Michael Stiebitz, Diego Scheide, Bjarne Toft, and Lene M. Favrholdt. *Graph Edge Coloring: Vizing's Theorem and Goldberg's Conjecture*. John Wiley & Sons, 2012.

[115] Andrew Thomason. An extremal function for contractions of graphs. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 95, pages 261–265. Cambridge University Press, 1984.

[116] Carsten Thomassen. Some remarks on Hajós' conjecture. *Journal of Combinatorial Theory, Series B*, 93(1):95–105, 2005.

[117] S. M. Ulam. *A collection of mathematical problems*. Interscience Tracts in Pure and Applied Mathematics, no. 8. Interscience Publishers, New York-London, 1960.

[118] Jan Van Den Heuvel and David R. Wood. Improper colourings inspired by Hadwiger's conjecture. *Journal of the London Mathematical Society*, 98(1):129–148, 2018.

[119] Vladimir N. Vapnik and A. Ya Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. In *Measures of complexity: festschrift for Alexey Chervonenkis*, pages 11–30. Springer, 2015.

[120] H. Whitney. 2-isomorphic graphs. *American Journal of Mathematics*, 55:245–254, 1933.

[121] Daniel T. Wise. *From Riches to Raags: 3-Manifolds, Right-Angled Artin Groups, and Cubical Geometry*. American Mathematical Society, 2012.

[122] Michael Zheng. On interval colorings of graphs. Bachelor's thesis, Karlsruhe Institute of Technology, Karlsruhe, Germany, 2022.