On Treewidth and Graph Minors

Daniel John Harvey

Submitted in total fulfilment of the requirements of the degree of
Doctor of Philosophy
February 2014

Department of Mathematics and Statistics
The University of Melbourne

Produced on archival quality paper

Abstract

Both treewidth and the Hadwiger number are key graph parameters in structural and algorithmic graph theory, especially in the theory of graph minors. For example, treewidth demarcates the two major cases of the Robertson and Seymour proof of Wagner's Conjecture. Also, the Hadwiger number is the key measure of the structural complexity of a graph. In this thesis, we shall investigate these parameters on some interesting classes of graphs.

The *treewidth* of a graph defines, in some sense, how "tree-like" the graph is. Treewidth is a key parameter in the algorithmic field of fixed-parameter tractability. In particular, on classes of bounded treewidth, certain NP-Hard problems can be solved in polynomial time. In structural graph theory, treewidth is of key interest due to its part in the stronger form of Robertson and Seymour's Graph Minor Structure Theorem. A key fact is that the treewidth of a graph is *tied* to the size of its largest grid minor. In fact, treewidth is tied to a large number of other graph structural parameters, which this thesis thoroughly investigates. In doing so, some of the tying functions between these results are improved. This thesis also determines exactly the treewidth of the line graph of a complete graph. This is a critical example in a recent paper of Marx, and improves on a recent result by Grohe and Marx. By extending the techniques used, we also determine the treewidth of the line graph of a complete multipartite graph, up to lower order terms in general, and exactly whenever the complete multipartite graph is regular. This generalises a result by Lucena. We also determine a lower bound on the treewidth of any line graph; this result is similar to a question about the Hadwiger number of line graphs posed by Seymour, which was recently proven by DeVos *et al.*. Finally, we prove a result on the treewidth of the Kneser graph; in doing so we also prove a generalisation of the famous Erdős-Ko-Rado Theorem.

The *Hadwiger number* of a graph is the size of its largest complete minor. One of the most important conjectures in modern mathematics is Hadwiger's Conjecture, which conjectures that the Hadwiger number of a graph is at least its chromatic number. A related question is determining what lower bound on the average degree is required to ensure the existence of a K_t-minor (or, more generally, an H-minor for any graph H). The K_t-minor case has been thoroughly studied, and independently answered by Kostochka and Thomason. In this thesis we answer a slightly different question and present an algorithm for finding, in $O(n)$ time, an H-minor forced by high average degree. Finally, this thesis determines a weakening of Hadwiger's Conjecture on the class of circular arc graphs, an interesting generalisation of the class of interval graphs, and in the process of doing so, proves some useful results about linkages in interval graphs.

Declaration

This is to certify that:

(i) the thesis comprises only my original work towards the PhD except where indicated in the Preface,

(ii) due acknowledgement has been made in the text to all other material used, and

(iii) the thesis is less that 100,000 words in length, exclusive of tables, maps, bibliographies and appendices.

Preface

No work presented in this thesis has been submitted for any other kind of qualification, and no work in this thesis was carried out prior to PhD enrolment.

Other than in Chapter 1 and Chapter 2, all results presented are new. (Some of the results presented in Chapter 2, an investigation into parameters tied to treewidth, are improvements on previously presented results. These results are noted in the text.)

Chapters 2,3,4,5,6 and 8 were carried out in collaboration with my supervisor, David Wood. Chapter 7 is the result of joint work with Vida Dujmović, Gwenaël Joret, Bruce Reed and David Wood. An important example in Chapter 5 was provided by Bruce Reed.

The results of Chapter 7 extend results which have been recently published in the SIAM Journal on Discrete Mathematics [26]. The results in Chapter 6 have recently appeared in the Electronic Journal of Combinatorics [44]. The results in Chapter 3 have been accepted by the Journal of Graph Theory [45]. (An earlier version of this paper, including the results of Chapter 4, is available on the arXiv [42].) Chapter 2 is also available on the arXiv [43].

In all cases I, Daniel Harvey, was the primary and corresponding author.

Acknowledgements

I'd like to acknowledge the following individuals for their assistance during my PhD candidature. Firstly, I'd like to thank my supervisor David Wood for the enormous amount of assistance provided throughout my time as a PhD student.

For their guidance, I'd like to thank the members of my Advisory Committee—Sanming Zhou, Peter Forrester and Graham Farr.

I'd like to thank my collaborators Vida Dujmović, Gwenaël Joret and Bruce Reed for their work on the paper which would become Chapter 7.

Further thanks to my officemates and fellow students Michael Payne, Guangjun Xu and Ricky Rotheram for their advice and support.

I also wish to acknowledge the several reading and seminar groups that ran during my candidature. Firstly, the Graph Theory Reading Group, and its current leader Arun Mani. Secondly, the Discrete Structures and Algorithms seminar group. Finally, the Theoretical Research in Computer Science (TRICS) group, run by Tony Wirth and (previously) Kerri Morgan.

Thanks to Alex Scott for pointing out references [36, 37, 99] in Chapter 6. Thanks also to Jacob Fox for helpful conversations with regards to Chapter 2.

A final note of thanks to my parents David and Jan Harvey for their emotional support during the last four years.

Table of Contents

List of Figures

Chapter 1

Introduction and Literature Review

1.1 Graph Minors

The numbered Theorems 1.1 to 1.13 form the core results presented in this thesis. In this chapter we present the statements of these theorems and provide the necessary background and context; the remaining chapters contain thorough proofs of these results.

A *graph* G is a set of elements $V(G)$, called *vertices*, together with a set $E(G)$ of unordered pairs of vertices, called *edges*. This is the definition of a *simple graph*; for our purposes, all graphs are simple unless otherwise noted. All of the basic definitions and notation in this thesis are as in the standard textbook by Diestel [22], unless otherwise noted.

A graph H is a *minor* of a graph G if a graph isomorphic to H can be constructed from G by repeated applications of *vertex deletion*, *edge deletion* and *edge contraction*. Vertex deletion and edge deletion are self-evident. Edge contraction is defined as follows: given an edge $vw \in E(G)$, replace vertices v and w with a single new vertex x adjacent to all vertices initially adjacent to v or w. If H is a minor of G, we say that G contains an H-minor. The study of graph minors is of key importance in modern graph theory. We now discuss some of these key results.

An important early result is the Kuratowski-Wagner Theorem [66, 113]. A graph is *planar* if it can be drawn in the plane so that no two edges cross.

Theorem (Kuratowski-Wagner Theorem [113]). *A graph G is planar if and only if it contains neither K_5 nor $K_{3,3}$ as a minor.*

We say a class of graphs $\mathfrak{G}$ is *minor-closed* if $H \in \mathfrak{G}$ whenever $G \in \mathfrak{G}$ and H is a minor of G. The class of planar graphs is minor-closed; clearly vertex and edge deletion do not turn a planar graph non-planar, and edge contraction does not create any crossing edges since it is possible to move two vertices together by "shortening" the edge until they merge together. The Kuratowski-Wagner Theorem gives a finite *forbidden minor characterisation* of the planar graphs—that is, a finite set of graphs such that every graph that is non-planar must contain a minor in that set, and every graph that is planar contains no minor in that set. While it is easy to find an infinite set of graphs with this property (simply take all graphs not in $\mathfrak{G}$), it is not obvious that a finite set must exist.

A similar result holds for every non-trivial minor-closed class.

Theorem (Graph Minor Theorem [88])**.** *Let $\mathfrak{G}$ be a minor-closed class of graphs, other than the class of all graphs. Then $\mathfrak{G}$ has a finite forbidden minor characterisation.*

This is one of the key theorems in modern structural graph theory. It was originally referred to as Wagner's Conjecture, before being proven by Robertson and Seymour in their twenty-three paper sequence "Graph Minors" [88]. (Wagner, however, stated that he had never conjectured such a result [22].) The Graph Minor Theorem gives rise to a $O(n^3)$ time algorithm for determining whether or not a given n-vertex graph G is in a fixed minor-closed class $\mathfrak{G}$ [93]. Specifically, this algorithm determines whether G contains a fixed graph H as a minor in $O(n^3)$ time; given a forbidden minor characterisation for $\mathfrak{G}$, it is sufficient to check whether or not each graph is a minor of G. However, this algorithm requires that the finite forbidden minor characterisation be known for $\mathfrak{G}$, which is not the case for most classes. It also depends on the number of graphs in this characterisation. This is finite and does not depend on n (which is why the total algorithm is still only $O(n^3)$ time; there is one iteration of the loop for each graph in the characterisation), but it is often enormous (see [13], for example).

At the heart of Robertson and Seymour's proof is the Graph Minor Structure Theorem [91, 94]. There are several different versions of the Graph Minor Structure Theorem (see Kawarabayashi and Mohar [53] for an overview), but essentially it shows that a minor-closed class either has bounded treewidth (which we define in Section 1.2), or if the class has unbounded treewidth, then any graph in it can be constructed with a restricted series of operations. A simplified version of this theorem follows:

Theorem (Robertson and Seymour [94])**.** *Let H be a non-planar graph and let $\mathfrak{G}$ be the family of graphs with no H-minor, and denote by $\Sigma_1, \dots, \Sigma_s$ all the connected surfaces (up to homeomorphism) in which H cannot be drawn. Then every graph in $\mathfrak{G}$ can be constructed by clique-sums from those that can "almost" be drawn in some Σ_i.*

Note the following facts. Firstly, we shall deal with the "missing" case when H is planar in Section 1.2. Secondly, by "almost", we mean that if G can "almost" be drawn in Σ_i, then G is embeddable in Σ_i except for a small number of *apex vertices* and a small number of *vortices.* Apex vertices are allowed to be adjacent to any other vertex without their incident edges "counting" with respect to the embedding. Vortices are subgraphs of G with bounded pathwidth (again defined in Section 1.2), which meet Σ_i without intersecting each other or other vertices except in a very restricted way. By "a small number", we mean depending only on H, and not on G or $|G|$ itself. Thirdly, a graph G is a *clique sum* of graphs G_1 and G_2 if G can be constructed by choosing cliques of equal size in G_1 and G_2, identifying them, and then possibly deleting some edges from the identified clique. Finally, note that this version of the Graph Minor Structure Theorem was insufficient for Robertson and Seymour to use to obtain their result, as they explain in [94]. The stronger version of the Graph Minor Structure Theorem uses the concept of *tangles.* We discuss tangles in Section 2.6.

A key parameter in the Graph Minor Structure Theorem, and the piece of the puzzle most interesting to us, is the parameter known as *treewidth.*

1.2 Treewidth

Let G be a graph. A *tree decomposition* of G is a pair $(T, (B_x \subseteq V(G))_{x \in V(T)})$ consisting of

- a tree T, and
- a collection of *bags* B_x containing vertices of G, indexed by the nodes of T,

such that:

- for all $v \in V(G)$, the set $\{x \in V(T) : v \in B_x\}$ induces a non-empty subtree of T, and
- for all $vw \in E(G)$, there is some bag B_x containing both v and w.

The *width* of a tree decomposition is the size of the largest bag in the tree decomposition, minus 1. The *treewidth* of a graph, denoted by $\mathsf{tw}(G)$, is the minimum width over all tree decompositions of G. To illustrate these concepts, we provide an example tree decomposition in Figure 1.1. Often, for the sake of simplicity, we will refer to a tree decomposition simply as T, leaving the set of bags implied whenever this is unambiguous. For similar reasons, often we say that bags X and Y are adjacent (or we refer to an edge XY), instead of the more accurate statement that the nodes of T indexing X and Y are adjacent. Usually, we refer to the vertices of G as vertices, but refer to the vertices of T as *nodes*, as in the above definition. This is also done to avoid confusion.

Define the *pathwidth* of a graph G, denoted $\mathsf{pw}(G)$, to be the minimum width of a tree decomposition that has a path as the underlying tree. Such a tree decomposition is called a *path decomposition*. Pathwidth was initially defined by Robertson and Seymour [89]. It follows that $\mathsf{pw}(G) \geq \mathsf{tw}(G)$, for all graphs G.

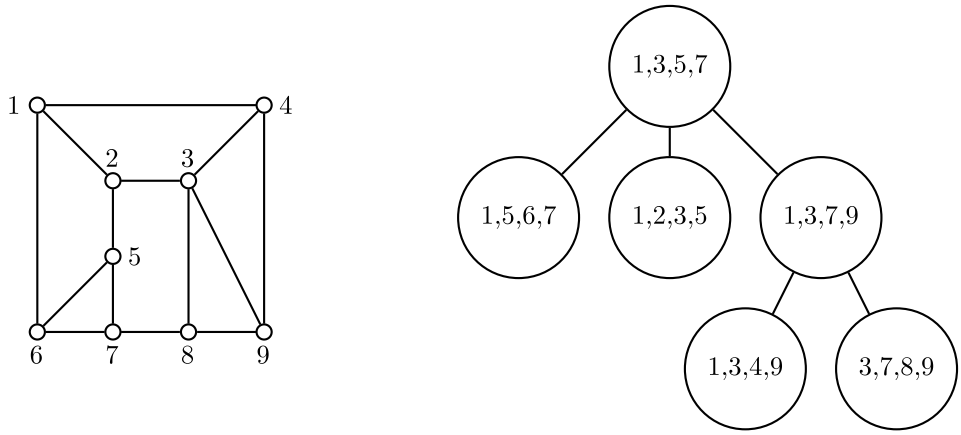

Figure 1.1: A graph G with 9 vertices, and a minimum width tree decomposition of G.

Treewidth was initially defined by Halin [41], who defined it in terms of S-functions on graphs. An S-function is a graph parameter $f(G)$ that behaves similarly to the Hadwiger number (which we define in Section 1.3). Specifically, an S-function is a function defined on all graphs that is non-increasing when taking a minor, has a value of 0 for the graph with no vertices, increases by 1 when adding a new vertex adjacent to all others in a graph, and is at most $\max\{f(G_1), f(G_2)\}$ when G is a clique sum of G_1 and G_2.

The modern definition of treewidth was provided by Robertson and Seymour [90]. Intuitively, a graph with low treewidth is simple and treelike—note that a tree itself has treewidth 1. (In fact, ensuring this fact is the reason for the minus 1 in the definition of width.) On the other hand, a complete graph K_n has treewidth $n-1$. This is a consequence of the more general result that $\mathsf{tw}(G) \geq \omega(G) - 1$, where $\omega(G)$ is the order of the largest clique (that is, complete subgraph) of the graph G. This follows from the fact that a set of subtrees of a tree satisfies the *Helly property*: for any set of pairwise intersecting subtrees of a tree T, there exists some vertex of T at which all the subtrees intersect. To see this, root the tree at some point, and root each subtree at the point of minimum distance to the root of T. Then all subtrees must intersect the root of a subtree which is at furthest distance from the root of T. As a result of this, given any clique of G, simply consider the subtrees of tree decomposition T induced by $\{x \in V(T) : v \in B_x\}$ for each vertex v in the clique. These trees pairwise intersect, so there is some bag containing all vertices of the clique and $\mathsf{tw}(G) \geq \omega(G) - 1$.

For a given constant c, the class of graphs with treewidth at most c is minor-closed. This follows from the fact that treewidth does not increase when taking a minor. For any acceptable graph minor operation performed on a graph G, there is a corresponding modification that can be performed on a tree decomposition of G; for a vertex deletion, simply remove that vertex from all bags, and for edge contraction, replace the endpoints of the edge with the new vertex in all bags in which either of the endpoints appeared. (Do nothing for edge deletion.) In fact, these graphs with bounded treewidth are a good example of a "well-behaved" minor-closed class. Treewidth is of major interest in the field of algorithm design, particularly in the field of fixed-parameter tractability. Many NP-Hard problems can be solved on graphs with bounded treewidth in polynomial time [6].

When adding a new dominating vertex to a graph G, essentially the only modification available to the tree decomposition is to add that vertex to every bag. (For a special kind of tree decomposition, which we call a *normalised tree decomposition*, adding the new vertex to every bag is exactly the only possible operation. There is always a normalised tree decomposition of G with width $\mathsf{tw}(G)$; we discuss this result in Lemma 2.2.) This increases $\mathsf{tw}(G)$ by 1.

Finally, note that if G is a clique sum of G_1 and G_2 such that C_1 and C_2 are the identified cliques in G_1, G_2 respectively, it is possible to "paste" two tree decompositions (one for G_1, one for G_2) together at bags containing C_1 and C_2 respectively. This proves that $\mathsf{tw}(G) + 1$ is an S-function. (We must take $\mathsf{tw}(G) + 1$ in order to ensure the correct value when G is a vertex-less graph.) In fact, the S-functions form a complete lattice, where the unit element, or top, is $\mathsf{tw}(G) + 1$ [41]. This makes an interesting contrast to the Hadwiger number, which we shall note later. While Halin identified $\mathsf{tw}(G) + 1$ as the top of the lattice of S-functions, Halin defined the function in terms of chordal graphs. We explore the connection between treewidth and chordal graphs in Section 2.4.

As part of their proof of the Graph Minor Structure Theorem, Robertson and Seymour used treewidth to prove what is known as the Grid Minor Theorem.

Theorem (Grid Minor Theorem [91])**.** *If G contains no H-minor where H is the $r \times r$-grid, then $\mathsf{tw}(G) \leq g(r)$, where $\mathsf{tw}(G)$ is the treewidth of G and $g(r)$ is a function that only depends on r.*

A similar result to the above follows whenever H is any planar graph—this follows from the above since any planar graph is a minor of a sufficiently large grid. This deals with the "missing" planar case in the Graph Minor Structure Theorem that was mentioned previously—when H is planar, the class of graphs with no H-minor has bounded treewidth, and the Graph Minor Theorem is (relatively) easy to prove for such graphs. Given that an $r \times r$-grid has treewidth r (a result we provide a proof of in Section 2.10), it also follows that the treewidth and the size of the largest grid minor of G are *tied*[†]. Two graph parameters α and β are *tied* if there exists a function f such that, for every graph G, $\alpha(G) \leq f(\beta(G))$ and $\beta(G) \leq f(\alpha(G))$. If f is a polynomial, then we say α and β are *polynomially tied*. So if G contains an $r \times r$-grid minor but no $(r+1) \times (r+1)$-grid minor, then $r \leq \mathsf{tw}(G) \leq g(r+1)$. Many other graph parameters are also tied to treewidth. In Chapter 2, we provide a thorough investigation of a large number of graph parameters that are polynomially tied to treewidth. The existence of these polynomial ties is well known, except the existence of a *polynomial* tie between treewidth and maximum order of a grid minor; this result was recently announced by Chekuri and Chuzhoy [14]. Chapter 2 includes a few improvements over these known results.

Theorem 1.1. *The following graph parameters are polynomially tied:*

- *treewidth,*
- *bramble number,*
- *minimum integer k such that G is a spanning subgraph of a k-tree,*
- *minimum integer k such that G is a spanning subgraph of a chordal graph with no $(k+2)$-clique,*
- *separation number,*
- *branchwidth,*
- *tangle number,*
- *lexicographic tree product number,*
- *Cartesian tree product number,*
- *linkedness,*
- *well-linked number,*
- *maximum order of a grid minor,*
- *maximum order of a grid-like-minor,*
- *Hadwiger number of the Cartesian product $G \square K_2$ (viewed as a function of G),*
- *fractional Hadwiger number,*
- *r-integral Hadwiger number for each $r \geq 2$.*

The relationships between treewidth and these other graph parameters are of use when solving problems in both algorithmic and structural graph theory. For example, the relationship between treewidth and graph separators (via the separation number seen above) is of key interest when considering the algorithmic aspects of treewidth. A separator

[†]Occasionally, other authors use the term *comparable* [30].

is a small set of vertices whose deletion "separates" the remaining vertices into components which are at most half the total size (or thereabouts). A separator is of particular use when dynamic programming is used to solve graph problems; sometimes, it is possible to delete the separator, recursively solve the problem on the remaining components, and then combine the solutions on the components to obtain a solution for the original graph. The tie between treewidth and separation number is also of key importance to our work in Chapter 6. The parameters linkedness and well-linked number were used in more recent proofs of the Grid Minor Theorem—we discuss this in more detail in the appropriate sections. Finally, recall the concept of tangles is of fundamental importance to the most powerful version of the Graph Minor Structure Theorem.

Theorem 1.1 proves several general results with respect to treewidth. It is also worth considering more specific treewidth results with regards to specific classes of graphs. We consider both the line graphs and the Kneser graphs.

Given a graph G, the *line graph* $L(G)$ is the graph with vertex set $E(G)$ and with an edge between two vertices of $L(G)$ if the corresponding edges in G share an endpoint. Line graphs have many interesting properties; for example, the neighbourhood of any vertex in a line graph can be covered by two cliques. (Any graph with this property is called a *quasi-line graph*.) Both the line graphs and the quasi-line graphs satisfy Hadwiger's Conjecture; see Reed and Seymour [84] and Chudnovsky and Ovetsky Fradkin [15] respectively. (We will discuss Hadwiger's Conjecture in more detail in Section 1.3.) In recent papers by Marx [77] and Grohe and Marx [39], the treewidth of the line graph of a complete graph is a critical example. For a graph G, let $G \cdot K_q$ denote the *lexicographic product* of G with K_q, that is, the graph created by replacing each vertex of G with a clique of size q and replacing each edge between two vertices with all of the edges between the two new cliques. Marx [77] shows that if $\mathsf{tw}(G) \geq k$, then $G \cdot K_p$ contains $L(K_k) \cdot K_q$ as a minor (for appropriate choices of p and q, depending on k and $|V(G)|$). Then Grohe and Marx [39] show that $\mathsf{tw}(L(K_n)) \geq \frac{\sqrt{2}-1}{4}n^2 + O(n)$.

In Chapter 3, we determine exactly the treewidth of $L(K_n)$. In doing so, we prove that the optimal tree decomposition is also a path decomposition. Hence we prove the following theorem.

Theorem 1.2.

$$\mathsf{tw}(L(K_n)) = \mathsf{pw}(L(K_n)) = \begin{cases} (\frac{n-1}{2})(\frac{n-1}{2}) + n - 2 & \text{, } \textit{if } n \textit{ is odd} \\ (\frac{n-2}{2})(\frac{n}{2}) + n - 2 & \text{, } \textit{if } n \textit{ is even.} \end{cases}$$

In Chapter 4, we extend our techniques from Chapter 3 to determine, up to lower order terms, the treewidth and pathwidth of $L(K_{n_1,\dots,n_k})$. (Here $K_{n_1,\dots,n_k}$ is a complete multipartite graph, the graph with vertex set partitioned into parts of size $n_1, \dots, n_k$ and edges between any two vertices in different parts.)

Theorem 1.3. *If $k \geq 2$ and $n = |V(K_{n_1,\dots,n_k})|$, then*

$$\frac{1}{2}\left(\sum_{1 \leq i < j \leq k} n_i n_j\right) - n(k-1) + \frac{3}{4}k(k-1) - 1$$
$$\leq \mathsf{tw}(K_{n_1,\dots,n_k}) \leq \mathsf{pw}(K_{n_1,\dots,n_k}) \leq$$
$$\frac{1}{2}\left(\sum_{1 \leq i < j \leq k} n_i n_j\right) + \frac{1}{2}n(k+5) + \frac{1}{4}k(k-1) - 4.$$

In the case where the complete multipartite graph is regular (that is, $n_1 = \cdots = n_k$), we determine the treewidth and pathwidth exactly.

Theorem 1.4. *If $k \geq 2$ and $n_1 = n_2 = \cdots = n_k = c \geq 1$, then*

$$\mathsf{tw}(L(K_{n_1,\dots,n_k})) = \mathsf{pw}(L(K_{n_1,\dots,n_k})) = \begin{cases} \frac{c^2k^2}{4} - \frac{c^2k}{4} + \frac{ck}{2} - \frac{c}{2} + \frac{k}{4} - \frac{5}{4}\,, & \text{if } k \text{ odd and } c \text{ odd} \\ \frac{c^2k^2}{4} - \frac{c^2k}{4} + \frac{ck}{2} - \frac{c}{2} - 1 & , \text{if } c \text{ even} \\ \frac{c^2k^2}{4} - \frac{c^2k}{4} + \frac{ck}{2} - \frac{c}{2} + \frac{k}{4} - \frac{3}{2}\,, & \text{if } k \text{ even and } c \text{ odd.} \end{cases}$$

Previously, Lucena [72] determined the treewidth of the Cartesian product $K_n \square K_n$, which is isomorphic to $L(K_{n,n})$. Theorem 1.4 generalises this result.

Finally, in Chapter 5, we prove a more general lower bound on the treewidth and pathwidth of any line graph.

Theorem 1.5. *For every graph G with minimum degree $\delta(G)$,*

$$\mathsf{tw}(L(G)) \geq \frac{2}{9}\delta(G)^2 - 1.$$

Theorem 1.6. *For every graph G with minimum degree $\delta(G)$,*

$$\mathsf{pw}(L(G)) \geq \frac{1}{4}\delta(G)^2 - 1.$$

This result is tight up to lower order terms.

It is well known that $\mathsf{tw}(G) \geq \delta(G)$ for every graph G; see Section 2.2 for a proof. So Theorems 1.5 and 1.6 are strengthenings of this result for line graphs.

We also consider the Kneser graphs. Let $[n] := \{1, \dots, n\}$. For any set $S \subseteq [n]$, a subset of S of size k is called a *k-set*, or occasionally a k-set in S. Let $\binom{S}{k}$ denote the set of all k-sets in S. We say two sets *intersect* when they have non-empty intersection. The

Kneser graph $\mathsf{Kneser}(n,k)$ is the graph with vertex set $\binom{[n]}{k}$, such that two vertices are adjacent if they are disjoint.

Kneser graphs were first investigated by Kneser [56]. The chromatic number of $\mathsf{Kneser}(n,k)$ was shown to be $n-2k+2$ by Lovász [71], as Kneser originally conjectured. This was an important proof due to the development of the topological methods involved. Many other proofs of this result have been found, for example consider [118], which gives a more combinatorial version. The Kneser graph is also of interest with regards to fractional chromatic number [97]. When $k=2$ and $n=5$, $\mathsf{Kneser}(n,k)$ is the famous Petersen graph, and so in some sense the Kneser graphs form a generalisation of the Petersen graph.

In Chapter 6, we determine exactly the treewidth of the Kneser graph, when n is sufficiently large with respect to k, and also when $k=2$.

Theorem 1.7. *Let $G = \mathsf{Kneser}(n,k)$ with $n \geq 4k^2 - 4k + 3$ and $k \geq 3$. Then*

$$\mathsf{tw}(G) = \binom{n-1}{k} - 1.$$

Theorem 1.8. *Let $G = \mathsf{Kneser}(n,2)$. Then*

$$\mathsf{tw}(G) = \begin{cases} 0 & \text{if } n \leq 3 \\ 1 & \text{if } n = 4 \\ 4 & \text{if } n = 5 \\ \binom{n-1}{2} - 1 & \text{if } n \geq 6. \end{cases}$$

We also provide a weaker bound on the treewidth when n is smaller, that is, a weaker result for a weaker assumption.

Theorem 1.9. *Let $G = \mathsf{Kneser}(n,k)$ with $n \geq \frac{1}{2}(\sqrt{5k^2-12k+8}+3k+2)$ and $k \geq 3$. Then*

$$\binom{n-1}{k} - \binom{n-1}{k-1} - 1 \leq \mathsf{tw}(G) \leq \binom{n-1}{k} - 1.$$

Note that since $k \geq 3$, Theorem 1.9 holds when $n \geq 3k-1$.

Define an *independent set* of a graph G to be a set of pairwise non-adjacent vertices. Then define the *independence number* $\alpha(G)$ to be the size of the largest independent set of G.

In proving Theorems 1.7, 1.8 and 1.9, we also prove some results generalising the famous Erdős-Ko-Rado Theorem [28, 50], which states:

Theorem (Erdős-Ko-Rado Theorem [28]). *Let $G = \mathsf{Kneser}(n,k)$ for some $n \geq 2k$. Then*

$$\alpha(G) = \binom{n-1}{k-1}.$$

If $n \geq 2k+1$ and $\mathcal{A}$ is an independent set such that $|\mathcal{A}| = \binom{n-1}{k-1}$, then $\mathcal{A} = \{v | i \in v\}$ for a fixed element $i \in [n]$.

Note that normally this theorem is phrased in terms of intersecting k-sets in $[n]$ rather than in terms of independent sets in the Kneser graph. It can also be stated in terms of the maximum order of a clique in the complement of $\mathsf{Kneser}(n, k)$. Recent work has been done extending this theorem by weakening the requirement that all the k-sets intersect—allowing a certain amount of intersection between the k-sets, for example [36, 37, 99]. (In our terms, this means considering an induced subgraph of the Kneser graph with bounded maximum degree.) We instead determine the following generalisation.

Theorem 1.10. *Say $c \in [\frac{2}{3}, 1)$ and $n \geq \max\{4k^2 - 4k + 3, \frac{1}{1-c}(k^2 - 1) + 2\}$. If H is a complete multipartite subgraph of the complement of $\mathsf{Kneser}(n, k)$ such that no colour class contains more than $c|H|$ vertices, then $|H| \leq \binom{n-1}{k-1}$.*

Theorem 1.10 is stated in terms of the complement of $\mathsf{Kneser}(n, k)$, but instead of determining an upper bound on the clique size, we determine an upper bound on the order of the largest complete multipartite subgraph (under some restrictions).

1.3 Hadwiger's Conjecture

A *k-colouring* (sometimes called a *k-vertex-colouring*) of a graph G is a function that assigns one of k colours to each vertex of G such that no two adjacent vertices are assigned the same colour. A graph with a k-colouring is called *k-colourable.* The *chromatic number* of G, denoted $\chi(G)$, is the smallest integer k such that G is k-colourable. Colouring is a very well established field of research [46]. A well-known theorem in this field is the Four Colour Theorem.

Theorem (Four Colour Theorem [2])**.** *If G is a planar graph, then $\chi(G) \leq 4$.*

The *Hadwiger number* of a graph G, denoted $\mathsf{had}(G)$, is the size of the largest complete minor of G. The Hadwiger number, like treewidth, is an S-function as defined by Halin [41], and forms the zero or bottom of the lattice of S-functions. Given that the class of graphs with treewidth at most c is minor-closed (as proved in Section 1.2) and that $\mathsf{tw}(K_n) = n-1$, it follows that $\mathsf{had}(G) \leq \mathsf{tw}(G)+1$ for every graph G. By the Kuratowski-Wagner Theorem, a planar graph G has $\mathsf{had}(G) \leq 4$. The obvious extension is to ask if all graphs with $\mathsf{had}(G) \leq 4$ have a 4-colouring. This is true, and follows from a result of Wagner [112] (actually proved before the Four Colour Theorem) which classifies all graphs with $\mathsf{had}(G) \leq 4$ in the following way:

Theorem (Wagner's Theorem [112]). *A graph G contains no K_5-minor if and only if G is constructable from planar graphs and the Wagner graph V_8 (the 8 vertex cycle with opposite vertices also adjacent) by repeated clique sums on cliques of at most 3 vertices.*

The above theorem shows that it is sufficient to 4-colour the planar graphs to prove all graphs with $\mathsf{had}(G) \leq 4$ have a 4-colouring—the graph V_8 can be coloured with 3 colours and clique sums of 4-colourable graphs are also 4-colourable. The similarities between this theorem of Wagner and the Graph Minor Structure Theorem should be noted—in fact, the Graph Minor Structure Theorem can be seen as a generalisation of Wagner's Theorem to general minor-closed classes in the same way as the Graph Minor Theorem is an extension of the Kuratowski-Wagner Theorem.

Hadwiger's Conjecture [40] is a generalisation of this result.

Conjecture (Hadwiger's Conjecture [40]). *For every graph G, $\mathsf{had}(G) \geq \chi(G)$.*

Hadwiger's Conjecture is one of the most important conjectures in modern graph theory; see Toft [107] for a survey. Beyond the Four Colour Theorem, Hadwiger's Conjecture has been proven if $\mathsf{had}(G) \leq 5$ [95]. (Note that for $\mathsf{had}(G) < 4$, Hadwiger's Conjecture is comparatively easy to prove.)

Given a graph G define the *degeneracy* of G, denoted $\mathsf{degen}(G)$, to be the smallest integer k such that every subgraph of G contains a vertex of degree at most k. By use of the following greedy algorithm, it is possible to see that $\chi(G) \leq \mathsf{degen}(G) + 1$. For a graph G with $\mathsf{degen}(G) = k$, choose a vertex v of degree at most k. Since every subgraph of $G - v$ is also a subgraph of G, it follows that $\mathsf{degen}(G - v) \leq k$, and so by induction, $G - v$ is $(k + 1)$-colourable. Since v has at most k neighbours, this can be extended to a $(k + 1)$-colouring of G. Given that the minimum degree is bounded above by the treewidth, and that treewidth of a subgraph of G is at most $\mathsf{tw}(G)$, this also shows that $\chi(G) \leq \mathsf{tw}(G) + 1$.

In order to prove Hadwiger's Conjecture, it would be sufficient to prove that every graph with average degree at least $t - 1$ contains a K_t-minor. While this is true when $t \leq 3$, it is not true in general. It is true, however, that there exists a function $f(t) > t - 1$ such that every graph with average degree $d(G) \geq f(t)$ contains a K_t-minor. Mader [73] initially proved this result, and showed that $f(t) = 2^{t-2}$ was enough to force the existence of a K_t-minor in G. This bound was later improved by Mader [74] to $16t \log_2(t)$, and then to $\Theta(t\sqrt{\log t})$ by Thomason [104] and Kostochka [57, 58]. This result is best possible, since certain random graphs achieve this bound. Thomason [105] determined the asymptotic constant for this bound.

While determining the average degree required to force a K_t-minor is of key interest due to its relation to Hadwiger's Conjecture, it is also worth considering the average degree required to force an H-minor, for an arbitrary graph H. Define

$$g(H) := \inf\{D : \text{every graph } G \text{ with average degree } d(G) \geq D \text{ contains an } H\text{-minor}\}.$$

Myers and Thomason [79] essentially determined $g(H)$ for dense graphs H. The result of Chudnovsky *et al.* [16] proves $g(K_{2,t}) = t + 1$. Kostochka and Prince [59, 60] determined asymptotically (in s) exact bounds on $g(K_{s,t})$. Recently, Reed and Wood [87] determined upper bounds on $g(H)$ for sparse graphs H.

On the other hand, given that a graph with high average degree does contain a K_t-minor, it is also worth considering how such a minor can be found efficiently. This was first posed by Reed and Wood [85], who showed that in $O(n + m)$ time it is possible to find a K_t-minor in an n-vertex m-edge graph as long as $d(G) \geq 2^{t-2}$. The author, in joint work with Dujmović, Joret, Reed and Wood [26] improved this, showing that only $d(G) \geq (2 + \epsilon)g(K_t)$ was required, for all $\epsilon > 0$, and for sufficiently large t as a function of ϵ. In Chapter 7, we give the result of [26], extended to quickly find an H-minor in a dense graph G for any fixed graph H, rather than just a K_t-minor.

Theorem 1.11. *For every fixed t-vertex graph H, there exists a $O(n)$ time algorithm that, given an n-vertex graph G with $d(G) \geq 2(g(H) + t)$, finds an H-minor in G.*

Reed and Wood [85] used their algorithm as a subroutine for finding separators in H-minor free graphs (also see Wulff-Nilsen [116] for a related separator result). This result has subsequently been used in other algorithms for H-minor free graphs, in particular, shortest path algorithms by Tazari and Müller-Hannemann [102] and Wulff-Nilsen [115], and a maximum matching algorithm by Yuster and Zwick [117]. The algorithm presented in Chapter 7 speeds up all these results, in terms of the dependence on H. (Note the original version presented in [26] also provides this speed up.)

Finally, recall that Robertson and Seymour [93] describe a $O(n^3)$ time algorithm that tests whether a given n-vertex graph contains a fixed graph H as a minor. The time complexity was improved to $O(n^2)$ by Kawarabayashi *et al.* [55]. Kawarabayashi and Reed have announced a $O(n \log n)$ time algorithm for this problem. The algorithm described by Theorem 1.11 is weaker than these results in the sense that it only works on graphs of high average degree, but it is stronger in the sense that it is faster.

The algorithmic extension of the average degree result is of key interest, but obviously Hadwiger's Conjecture itself remains unsolved. Given the difficulties in proving the entire conjecture, one direction of interest is to consider Hadwiger's Conjecture for interesting classes of graphs. For example, Hadwiger's Conjecture is true for line graphs [84] and quasi-line graphs [15], as mentioned previously. As such, consider the following definition.

An *intersection graph* G is a graph where the vertex set is a collection of sets, and any two vertices are adjacent if their corresponding sets intersect. There are many interesting classes of intersection graphs. The previously mentioned Kneser graphs are complements of the intersection graphs on $\binom{[n]}{k}$. The *interval graphs* are the intersection graphs where the underlying sets are intervals on the real line. We say the *maximum load* of an interval graph $\mathcal{L}(G)$ is the maximum number of intervals at any point of the line. The clique size $\omega(G)$ of an interval graph is equal to its maximum load $\mathcal{L}(G)$, which is also equal to $\mathsf{pw}(G) + 1$. For every graph G with $\mathsf{pw}(G) = k$, G is a spanning subgraph of an interval graph G' with $\mathsf{pw}(G') = \mathcal{L}(G') - 1 = \omega(G') - 1 = k$. It is also possible to greedily colour any interval graph with $\mathcal{L}(G)$ colours; simply traverse the intervals left to right and colour an interval v with a colour not being used by any neighbour of v when reaching its left endpoint. At that point only at most $\mathcal{L}(G)-1$ of its neighbours are coloured; those intervals also at the left endpoint of v. (Because the class is closed under taking induced subgraphs, this means every interval graph is a *perfect graph*; that is, every induced subgraph H of an interval graph G has $\chi(H) = \omega(H)$.) So Hadwiger's Conjecture holds trivially for the interval graphs. Hence, we consider a class of graphs that is slightly more complex.

A *circular arc graph* is an intersection graph where the underlying sets are arcs on a circle. (Note that *arcs* will always refer to arcs on the circle, not directed edges as is sometimes the case elsewhere.) The circular arc graphs are a generalisation of the interval graphs—any interval graph is also a circular arc graph. Define $\mathcal{L}(G)$ for a circular arc graph analogously to the interval graph case. (Note here that there may be multiple possible representations of a graph G as a circular arc graph, and that these representations may have different maximum load. For example, it is possible to represent K_3 either with $\mathcal{L}(K_3) = 3$ or $\mathcal{L}(K_3) = 2$. In that sense, $\mathcal{L}(G)$ is not well defined. However, if we fix a collection of arcs and consider the graph which arises from that collection, then $\mathcal{L}(G)$ is well defined.) Let $\beta(G)$ denote the *cover number* of a circular arc graph G, the minimum number of arcs required to completely cover the circle. (If no set of arcs of G completely covers the circle, then say $\beta(G) = \infty$.)

A circular arc graph is not necessarily perfect; for example, any odd cycle is a circular arc graph. However, the circular arc graphs are χ-bounded. That is, the chromatic number of a circular arc graph is bounded from above by a function of its maximum clique size; this is equivalent to saying that χ and ω are tied for circular arc graphs. (Such a result does not hold for general graphs. The *girth* of a graph G is the size of its smallest cycle, and if this is at least 4, then $\omega(G) \leq 2$. A famous result of Erdős [27] states that there exist graphs with arbitrarily large girth (that is, $\omega(G) = 2$) and arbitrarily large chromatic number.)

Specifically, for a circular arc graph G, $\chi(G) \leq \frac{3}{2}\omega(G)$, which was shown by Karapetjan [48]. This extends a result of Tucker [109], who showed that if $\beta(G) > 3$, then $\chi(G) \leq \frac{3}{2}\mathcal{L}(G) = \frac{3}{2}\omega(G)$. The bound on $\chi(G)$ in terms of $\mathcal{L}(G)$ was improved by Valencia-Pabon [110] to $\chi(G) \leq \lceil (\frac{\beta(G)-1}{\beta(G)-2})\mathcal{L}(G) \rceil$.

Circular arc graphs are also interesting due to Hajós' Conjecture, a strengthening of Hadwiger's Conjecture.

Conjecture (Hajós' Conjecture)**.** *For every graph G, if $\chi(G) = k$ then G contains a subdivision of K_k as a subgraph.*

The graph G contains a subdivision of K_k if K_k can be constructed from G by vertex deletion, edge deletion and the *smoothing operation*. The *smoothing operation* allows us to remove a degree 2 vertex and replace it with an edge between its two neighbours. This is the exact reverse of the subdivision operation, hence the name. This can also be seen as a very specific form of edge contraction; an edge can be contracted only if one endpoint of the edge has degree 2. (Sometimes, if G contains a subdivision of a graph H as a subgraph, we say G contains H as a *topological minor*.) Since the acceptable operations are weaker, Hajós' Conjecture is stronger than Hadwiger's Conjecture. Also, since the acceptable operations are weaker, Hajós Conjecture is false. This was shown by Catlin [12], giving the following counterexample. Consider the lexicographic product of $C_{2n+1} \cdot K_k$: when $n \geq 2$ then the largest complete graph contained as a subdivision has $2k + 1$ vertices but the graph requires $2k + \lceil \frac{k}{n} \rceil$ colours. This shows that Hajós' Conjecture fails for $\chi(G) \geq 8$. By taking an appropriate induced subgraph of this counterexample, it can be shown that Hajós' Conjecture also fails for $\chi(G) = 7$. The conjecture holds trivially for $\chi(G) \leq 3$ and Dirac [24] showed it holds for $\chi(G) = 4$. The conjecture remains open for $\chi(G) \in \{5, 6\}$. More information on Hajós conjecture can be found in Thomassen [106].

Interestingly, this counterexample given by Catlin is a circular arc graph.

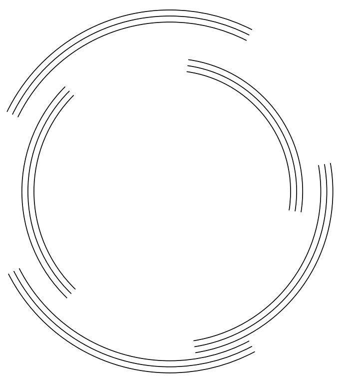

Figure 1.2: Catlin's counterexample to Hajós' Conjecture, when $n = 2$ and $k = 3$, viewed as a circular arc graph.

As a result, the circular arc graphs are an interesting class to consider with respect to Hadwiger's Conjecture; if Hadwiger's Conjecture is false, any counterexample is also a counterexample to Hajós' Conjecture, as Hajós' Conjecture is strictly stronger. Say a circular arc graph is *proper* if no arc on the circle covers another. Hadwiger's Conjecture has been proven for proper circular arc graphs [4]. Also note that a proper circular arc graph is a quasi-line graph (mentioned in Section 1.2); for any vertex v, all neighbours are either at the left end or the right end of the arc of v, and as such form two cliques.

If $\beta(G) > 3$, then $\mathcal{L}(G) = \omega(G)$, since any clique must be a set of vertices at a single point on the circle. This subclass of the circular arc graphs is of interest itself, and these graphs are sometimes called *normal Helly circular arc graphs.* Much work has been done attempting to find a complete list of minimal forbidden induced subgraphs for the class of normal Helly circular arc graphs; this work was completed recently by Grippo and Safe [38]. See also Lin *et al.* [69] for an in-depth study of normal Helly circular arc graphs. The associated matrices of normal Helly circular arc graphs have many useful properties, see Curtis *et al.* [18], Gavril [34] and Tucker [108], for example. However, these properties are not used in the work of this thesis.

In Chapter 8, we prove the following.

Theorem 1.12. *For a normal Helly circular arc graph G,* $\mathsf{had}(G) \geq \chi(G) - 1$.

This is a weakening of Hadwiger's Conjecture for circular arc graphs. In order to prove Theorem 1.12, we attempt to construct a complete minor in a vertex-minimum counterexample G by starting with a maximum size clique (of which all the vertices are at a point) and attempting to build a set of paths around the rest of the circle to obtain the $\mathsf{had}(G) - \omega(G)$ extra vertices the minor requires. If we are able to do this, then G is not a counterexample. If we are unable to do this, we use the information obtained when building the paths to instead recolour G with less than $\chi(G)$ colours, which contradicts the definition of $\chi(G)$. Thus either way Theorem 1.12 is proven.

In proving Theorem 1.12, we also developed some useful results about interval graphs. Say $2k$ distinct vertices $s_1, \ldots, s_k, t_1, \ldots, t_k$ can be *linked* if there exists a set of k pairwise vertex disjoint paths $P_1, \ldots, P_k$ such that P_i starts at s_i and ends at t_i. The paths $P_1, \ldots, P_k$ are called a *linkage*. For a graph G, if $|V(G)| \geq 2k$ and if any $2j$ distinct vertices (where $j \leq k$) can be linked, then we say the graph G is *k-linked*. We call the vertices $s_1, \ldots, s_k$ *sources* and $t_1, \ldots, t_k$ *targets*.

(Note that unfortunately the terminology here is very similar to the terminology of Section 2.8. Do note that here we refer to a *graph* being k-linked, whereas in Section 2.8 it is a *set of vertices* which is k-linked. We shall endeavour to ensure that these terms are not confused.)

Finding linkages (sometimes referred to as the *disjoint paths problem*) is a well studied problem in graph theory—consider the survey by Schrijver [98]. Robertson and Seymour [93] considered this problem in one of their "Graph Minors" papers. As they note, finding a linkage is closely related to determining if a graph G contains a subdivision of a graph H; find disjoint paths for each edge of H and then use the smoothing operation. (Note that some interpretations of a linkage allow sources and targets to be non-distinct and only require the paths to be internally vertex disjoint. We prefer the version of Diestel [22] due to the restrictions we require in Theorem 1.13.)

Clearly, if a graph is k-linked then it is also k-connected. A graph is k-connected if there is a set of k internally vertex disjoint paths between any two sets $\{s_1, \ldots, s_k\}$ and $\{t_1, \ldots, t_k\}$, however there is no guarantee that the paths will "start" and "end" at the correct vertices. As a concrete example, a cycle is 2-connected but not 2-linked; if we choose the vertices s_1, s_2, t_1, t_2 such the vertices appear in that order if we take a clockwise traversal of the cycle, then s_1, s_2, t_1, t_2 cannot be linked.

However, there is a function $f(k)$ such that if a graph is $f(k)$-connected then it is also k-linked. (That is, these graph parameters are tied.) An initial proof was given by both Jung [47] and Larman and Mani [67], however $f(k)$ was exponential. The function $f(k)$ was improved by Robertson and Seymour [93] and then by Bollobás and Thomason [10], who proved $f(k)$ was linear. Recent results by Kawarabayashi *et al.* [54] and Thomas and Wollan [103] proved that every $10k$-connected graph is k-linked. Kawarabayashi [51] used the result of Thomas and Wollan [103] to determine a lower bound on the connectivity of a counterexample to Hadwiger's Conjecture.

The function $f(k)$ can be improved when considering chordal graphs.

Lemma (Böhme *et al.* [9])**.** *If G is a $(2k-1)$-connected chordal graph, then G is k-linked.*

This result is tight, and it is also tight for interval graphs, which are a subclass of the chordal graphs. (See Chapter 9 for a proof of this result.)

In Chapter 9, our major result is an improvement of the Böhme *et al.* result for interval graphs, under slightly restricted circumstances.

Theorem 1.13. *Let G be a $\lceil \frac{3k}{2} \rceil$-connected interval graph, and let $s_1, \ldots, s_k, t_1, \ldots, t_k$ be $2k$ pairwise distinct vertices, such that no source s_i and no target t_j are adjacent, and such that s_i is left of t_i for all i. Then $s_1, \ldots, s_k, t_1, \ldots, t_k$ can be linked.*

Theorem 1.13 is tight, which we also prove. We also provide some interesting results about linkages in powers of paths, and about Hadwiger's number for powers of cycles.

1.4 A Unifying Example

Finally, we provide a unifying example which ties together many of the major concepts of the thesis, including treewidth, pathwidth, the Hadwiger number, circular arc graphs and interval graphs. The k^{th}*-power of a cycle* C_n^k is the graph formed by taking the n-vertex cycle and adding edges between any two vertices at distance at most k. This graph is a circular arc graph, as follows. Label n points on the cycle $0, \dots, n-1$ in clockwise order, and then place the arcs $\{[i, i+k(\text{mod } n)] : 0 \leq i \leq n-1\}$. (Note that this means C_n^k is also a proper circular arc graph.) With this representation, $\mathcal{L}(C_n^k) = k+1$. In fact, the k^{th}-power of a cycle is arguably the simplest n-vertex circular arc graph with maximum load $k+1$.

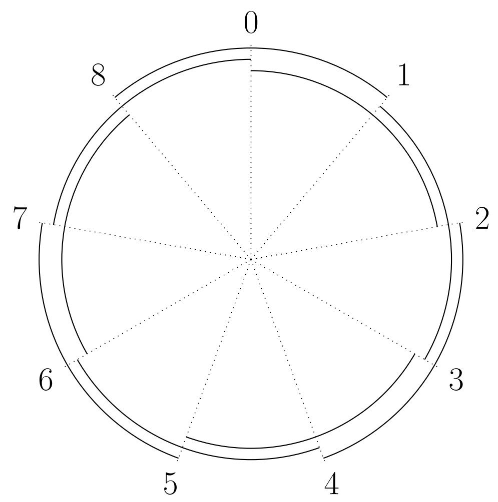

Figure 1.3: C_9^2 represented as a circular arc graph.

Li and Liu [68] previously proved that Hadwiger's Conjecture holds for the power of a cycle. We now prove an upper bound on $\mathsf{had}(C_n^k)$, which we denote by G.

First, choose a point on the circle of load k (for example, a point between any two of the labelled points $0, \dots, n-1$) and delete all vertices at that point. Denote this graph H. (The graph H is actually something called the power of a path, which we discuss in detail in Chapter 9.) The graph H is an interval graph—simply "unroll" it, and treat the point of deletion as both $+\infty$ and $-\infty$. (This idea will be used substantially in Chapter 8.) Recall $\mathcal{L}(H) = \mathsf{pw}(H) + 1$ for an interval graph, and $\mathcal{L}(H) \leq k+1$ since vertex deletions do not increase the maximum load. Thus $\mathsf{tw}(H) \leq \mathsf{pw}(H) \leq k$. We construct a path decomposition for G from a width k path decomposition of H by placing the k vertices of $G - H$ into every bag. Thus $\mathsf{tw}(G) \leq \mathsf{pw}(G) \leq 2k$. Since $\mathsf{had}(G) \leq \mathsf{tw}(G) + 1$, it follows

$\mathsf{had}(C_n^k) \leq 2k + 1$.

Now, we consider a lower bound on $\mathsf{had}(C_n^k)$ in certain circumstances. Assume that $n \equiv 1 \mod k$ and that $n \geq 2k + 1$. Label the vertices $0, \ldots, n-1$ so that vertex i corresponds to the arc $[i, i+k (\mathrm{mod}\ n)]$. If we contract the edges $\{1, k+1\}, \ldots, \{k, 2k+1\}$, the resultant graph created by these contractions is C_{n-k}^k. Since $n - k \equiv 1 \mod k$, by repeating this process it follows C_{2k+1}^k is a minor of C_n^k when $n \equiv 1 \mod k$ and $n \geq 2k+1$. However, C_{2k+1}^k is the complete graph on $2k + 1$ vertices. Thus if $n \equiv 1 \mod k$ and $n \geq 2k + 1$, then $\mathsf{had}(C_n^k) \geq 2k + 1$. In Chapter 9 we provide a slightly weaker lower bound independent of the modulus of n. These bounds together show that in some cases, $\mathsf{had}(C_n^k) = 2k + 1$.

The proof of the upper bound shows how, on occasion, we are able to use treewidth to determine an upper bound on the Hadwiger number. (However, this does not always work well—as previously mentioned, an $r \times r$-grid has treewidth r but Hadwiger number at most 4 due to planarity.)

Part I

Treewidth

Chapter 2

Parameters Tied to Treewidth

2.1 Introduction

A *graph parameter* is a real-valued function α defined on all graphs such that $\alpha(G_1) = \alpha(G_2)$ whenever G_1 and G_2 are isomorphic. Two graph parameters α and β are *tied* if there exists a function f such that for every graph G,

$$\alpha(G) \leq f(\beta(G)) \text{ and } \beta(G) \leq f(\alpha(G)).$$

Moreover, say that α and β are *polynomially tied* if f is a polynomial.

In this chapter, we draw on results in the literature and prove the following theorem.

Theorem 1.1. *The following graph parameters are polynomially tied:*

- *treewidth,*
- *bramble number,*
- *minimum integer k such that G is a spanning subgraph of a k-tree,*
- *minimum integer k such that G is a spanning subgraph of a chordal graph with no $(k+2)$-clique,*
- *separation number,*
- *branchwidth,*
- *tangle number,*
- *lexicographic tree product number,*
- *Cartesian tree product number,*
- *linkedness,*
- *well-linked number,*
- *maximum order of a grid minor,*
- *maximum order of a grid-like-minor,*

- *Hadwiger number of the Cartesian product $G\Box K_2$ (viewed as a function of G),*
- *fractional Hadwiger number,*
- *r-integral Hadwiger number for each $r \geq 2$.*

Fox [30] states (without proof) a theorem similar to Theorem 1.1 with the parameters treewidth, bramble number, separation number, maximum order of a grid minor, fractional Hadwiger number, and r-integral Hadwiger number for each $r \geq 2$. This result of Fox was motivation for the research in this chapter.

We investigate the parameters in Theorem 1.1, showing where these parameters have been useful, and provide proofs that each parameter is tied to treewidth (except in a few cases). In a number of cases we improve known bounds, provide simpler proofs, and show that the inequalities presented are tight. For the sake of completeness, we include a few well-known proofs. The following graph is a key example.

Say n, k are positive integers. Let $\psi_{n,k}$ be the graph with vertex set $A \cup B$, where A is a clique on n vertices, B is an independent set on kn vertices, and $A \cap B = \emptyset$, such that each vertex of A is adjacent to exactly $k(n-1)$ vertices of B and each vertex of B is adjacent to exactly $n-1$ vertices of A. (Note it always possible to add edges in this fashion; pair up each vertex in A with k vertices in B such that all pairs are disjoint, and then add all edges from A to B except those between paired vertices.)

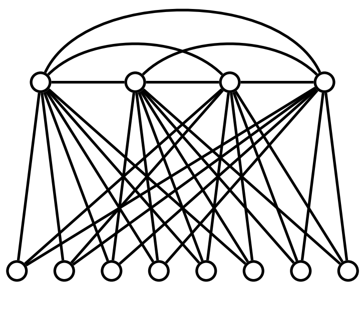

Figure 2.1: The graph $\psi_{4,2}$.

The following result will be useful when proving the tightness of several bounds.

Lemma 2.1. *If $n, k \geq 1$, then $tw(\psi_{n,k}) = n - 1$.*

Proof. Construct the following tree decomposition. Let T be a star with kn leaves, and let each vertex of B correspond to a unique leaf node. In the bag indexed by the centre node, place all the vertices of A. In the bag indexed by a leaf corresponding to $v \in B$, place $\{v\} \cup N(v)$. Since B is an independent set, this is a valid tree decomposition. (Note it is the tree decomposition described by Lemma 6.5.) Since $|N(v)| = n - 1$ and $|A| = n$, the width of this tree decomposition is $n - 1$. Given that the treewidth is at least the minimum degree, which is also $n-1$ (consider a vertex in B), our statement is proven. $\square$

2.2 Basics

The definition of treewidth was given in Chapter 1. Say a tree decomposition is *normalised* if each bag has the same size and $|X - Y| = |Y - X| = 1$ whenever XY is an edge. The following result is well-known.

Lemma 2.2. *If a graph G has a tree decomposition of width k, then G has a normalised tree decomposition of width k.*

Proof. Let T be a tree decomposition of G with width k. Thus T contains a bag of size $k + 1$. If some bag of T does not contain $k + 1$ vertices, then since T is connected, there exist adjacent bags X and Y such that $|X| = k + 1$ and $|Y| < k + 1$. Then $X - Y$ is non-empty; take some vertex of $X - Y$ and add it to Y. This increases $|Y|$, so repeat this process until all bags have size $k + 1$.

Now, consider an edge XY. Since $|X| = |Y|$, it follows $|X - Y| = |Y - X|$. If $|X - Y| > 1$, then let $v \in X - Y$ and $u \in Y - X$. Subdivide the edge XY of T and call the new bag Z. Let $Z = (X - \{v\}) \cup \{u\}$. Now $|X - Z| = 1$ and $|Y - Z| = |Y - X| - 1$, so repeat this process until $|X - Y| = |Y - X| \leq 1$ for each pair of adjacent bags. Finally, if XY is an edge and $|X - Y| = 0$, then contract the edge XY, and let the bag at the contracted node be X. Repeat this process so that if X and Y are a pair of adjacent bags, then $|X - Y| = |Y - X| = 1$. All of these operations preserve tree decomposition properties and width. Hence this modified T is our desired normalised tree decomposition. $\square$

As a result of Lemma 2.2, it follows that $\mathsf{tw}(G) \geq \delta(G)$, where $\delta(G)$ is the minimum degree of G. Consider a leaf bag, which must contain a vertex $v \in V(G)$ that is no other bag. Hence this leaf bag must contain the entire neighbourhood of v, and as such the bag contains at least $\delta(G) + 1$ vertices.

Given a graph H, an *H-model* of G is a set of pairwise vertex-disjoint connected subgraphs of G, each called a *branch set*, indexed by the vertices of H, such that if $vw \in E(H)$, then there exists an edge between the branch sets indexed by v and w. If G contains an H-model, then repeatedly contract the edges inside each branch set and delete extra vertices and edges to obtain a copy of H. Thus if G contains an H-model, then H is a minor of G. Similarly, if H is a minor of G, "uncontract" each vertex in the minor to obtain an H-model of G. Models are helpful when dealing with questions relating to minors, since they describe how the H-minor "sits" in G.

2.3 Brambles

Two subgraphs A and B of a graph G *touch* if $V(A) \cap V(B) \neq \emptyset$, or some edge of G has one endpoint in A and the other endpoint in B. A *bramble* in G is a set of connected subgraphs of G that pairwise touch. A set S of vertices in G is a *hitting set* of a bramble $\mathcal{B}$ if S intersects every element of $\mathcal{B}$. The *order* of $\mathcal{B}$ is the minimum size of a hitting set. (We often refer to such a hitting set as a *minimum hitting set.*) The *bramble number* of G is the maximum order of a bramble in G. Brambles were first defined by Seymour and Thomas [101], where they were called *screens of thickness* k. Seymour and Thomas proved the following result.

Theorem 2.3 (Treewidth Duality Theorem (Seymour and Thomas [101]))**.** *For every graph* G*,*

$$\mathsf{tw}(G) = \mathsf{bn}(G) - 1.$$

Proof. Here, we present a short proof showing one direction of this result. The other (more difficult) direction can be found in [101]; see Bellenbaum and Diestel [5] for a shorter proof. Let β be a bramble in G of maximum order, and let T be the underlying tree in a tree decomposition of G. For a subgraph $A \in \beta$, let T_A be the subgraph of T induced by the nodes of T whose bags contain vertices of A. Since A is connected, T_A is also connected. Similarly, if $A, B \in \beta$, then since these subgraphs touch, there is a node of T in both T_A and T_B. So the set of subtrees $\{T_A : A \in \beta\}$ pairwise intersect. By the Helly Property of trees, there is some node x that is in all such T_A. The bag indexed by x contains a vertex from each $A \in \beta$, so it is a hitting set of β. Hence that bag has order at least $\mathsf{bn}(G)$, and so $\mathsf{tw}(G) \geq \mathsf{bn}(G) - 1$. □

Note that Theorem 2.3 means that the bramble number is equal to the size of the largest bag in a minimum width tree decomposition.

Brambles are useful for proving a lower bound on the treewidth of a graph. Consider the following: given a valid tree decomposition T for a graph G, then $\mathsf{tw}(G)$ is at most the width of T. Brambles provide the equivalent functionality for the lower bound—given a valid bramble of a graph G, it follows that the bramble number is at least the order of that bramble, giving us a lower bound on the treewidth. (For examples of this, see Bodlaender *et al.* [8], Lucena [72] and Lemma 2.23.)

By considering the definition of an H-model in Section 2.2, it should be clear that a K_t-model forms a bramble with a minimum hitting set of order t. As a result, $\mathsf{had}(G) \leq \mathsf{bn}(G)$ for every graph G. (This is essentially an alternate proof that $\mathsf{had}(G) \leq \mathsf{tw}(G) + 1$.)

2.4 k-Trees and Chordal Graphs

In certain applications, such as graph drawing [20, 25] or graph colouring [1, 65], it often suffices to consider only the edge-maximal graphs of a given family to obtain a result. The language of k-trees and chordal graphs provides an elegant description of the edge-maximal graphs with treewidth at most k.

A vertex v in a graph G is *k-simplicial* if it has degree k and its neighbours induce a clique. A graph G is a *k-tree* if either:

- $G = K_{k+1}$, or
- G contains a k-simplicial vertex v and $G - v$ is also a k-tree.

Note that there is some discrepancy over this definition; certain authors use K_k in the base case. This means that K_k is a k-tree, but creates no other changes. It is well known that k-trees have a strong tie to treewidth; see Lemma 2.4.

A graph is *chordal* if it contains no induced cycle of length at least four. That is, every cycle that is not a triangle contains a chord. Gavril [35] showed that the chordal graphs are exactly the intersection graphs of subtrees of a tree T. So construct a tree decomposition with underlying tree T as follows. Think of each $v \in V(G)$ as a subtree of T; place v in the bags indexed by the nodes of that subtree. It can easily be seen that this is a tree decomposition of G in which every bag is a clique (that is, every possible edge exists), since should two vertices share a bag, then their subtrees intersect and the vertices are adjacent. It also follows that the graph arising from a tree decomposition with all possible edges (that is, two vertices are adjacent if and only if they share a bag) is a chordal graph. Chordal graphs are therefore interesting by being the edge-maximal graphs for a fixed tree-width. The initial definition of $\mathsf{tw}(G)$ by Halin [41] is that $\mathsf{tw}(G) + 1$ is equal to the minimum chromatic number of any chordal graph which contains G. This is identical to the second equality below, given that chordal graphs are perfect.

Lemma 2.4. *For every graph G,*

$$\begin{aligned}\mathsf{tw}(G) &= \min\{k : G \textit{ is a spanning subgraph of a } k\textit{-tree }\}.\\ &= \min\{k : G \textit{ is a spanning subgraph of a chordal graph with no } (k+2)\textit{-clique }\}.\end{aligned}$$

Proof. For simplicity, let $a(G) = \min\{k : G$ is a spanning subgraph of a k-tree$\}$ and $b(G) = \min\{k : G$ is a spanning subgraph of a chordal graph with no $(k+2)$-clique$\}$.

First, we show that $b(G) \leq a(G)$. Fulkerson and Gross [32] showed that a graph H is chordal if and only if it has a *perfect elimination ordering*; that is, an ordering of the vertex set such that for each $v \in V(H)$, v and all vertices adjacent to v which are after

v in the ordering form a clique. If H is an $a(G)$-tree such that G is a spanning subgraph of H, then there is a simple perfect elimination ordering for H. (Repeatedly delete the $a(G)$-simplicial vertices to obtain $K_{a(G)+1}$, and consider the order of deletion.) So H is chordal. It is clear that each v has only $a(G)$ neighbours after it in this ordering, so H contains no $(a(G)+2)$-clique. (For any clique, consider the first vertex of the clique in the ordering, and note at most $a(G)$ other vertices are in the clique.) Thus $b(G) \leq a(G)$.

Second, we show that $a(G) \leq \mathsf{tw}(G)$. Assume for the sake of a contradiction that G is a vertex-minimal counterexample, and say G has treewidth k. It is easy to see $a(G) \leq \mathsf{tw}(G)$ when G is complete, so assume otherwise. Let T be a tree decomposition of G with minimum width. By Lemma 2.2, assume T is normalised. Note, since G is not complete, T contains more than one bag. Let G' be the graph created by taking G and adding all edges vw, where v and w share some bag of T. So G is a spanning subgraph of G' and T is a tree decomposition of G' as well as G. By the normalisation, there is a vertex $v \in V(G')$ such that v appears in a leaf bag B of T and nowhere else. Hence v has exactly k neighbours in G', which form a clique since they are all in B. Since it is smaller than the minimal counterexample, $a(G'-v) \leq \mathsf{tw}(G'-v) \leq k$. Since $G'-v$ contains a $(k+1)$-clique (consider a bag of T other than B), it follows $a(G'-v) \geq k$. Thus $a(G'-v) = k$, and $G'-v$ is a spanning subtree of a k-tree H. Since v is k-simplicial in G', it follows G' (and thus G) is a spanning subgraph of a k-tree, which contradicts our assumption.

Finally, we show that $\mathsf{tw}(G) \leq b(G)$. The graph G is a spanning subgraph of chordal graph H with no $(b(G)+2)$-clique. There is a tree decomposition of H where every bag is a clique; this means it has width at most $b(G)$. This tree decomposition is also a tree decomposition for G, so $\mathsf{tw}(G) \leq b(G)$.

Hence, it follows that $b(G) \leq a(G) \leq \mathsf{tw}(G) \leq b(G)$, which is sufficient to prove our desired result. $\square$

2.5 Separators

For a graph G, a set $S \subseteq V(G)$, and some $c \in [\frac{1}{2}, 1)$, a (k, S, c)-*separator* is a set $X \subseteq V(G)$ with $|X| \leq k$, such that no component of $G - X$ contains more than $c|S - X|$ vertices of S. Note that a (k, S, c)-separator is also a (k, S, c')-separator for all $c' \geq c$. Define the *separation number* $\mathsf{sep}_c(G)$ to be the minimum integer k such that there is a (k, S, c)-separator for all $S \subseteq V(G)$. We also consider the following variant: a $(k, S, c)^*$-*separator* is a set $X \subseteq V(G)$ with $|X| \leq k$ such that no component of $G - X$ contains more than $c|S|$ vertices of $S - X$. Define $\mathsf{sep}^*_c(G)$ analogously to $\mathsf{sep}_c(G)$, but with respect to these variant separators. It follows from the definition that $\mathsf{sep}^*_c(G) \le \mathsf{sep}_c(G)$.

Separators can be seen as a generalisation of the ideas presented in the famous planar separator theorem [70], which essentially states that a planar graph G with n vertices contains a $(O(\sqrt{n}), V(G), \frac{2}{3})^*$-separator. Unfortunately, the precise definition of a separator and the separation number is inconsistent across the literature. The above definition is an attempt to unify all existing definitions. Robertson and Seymour [90] gave the first lower bound on $\mathsf{tw}(G)$ in terms of separators, though they do not use the term, nor do they give an explicit definition of separation number. This definition is equivalent to our standard definition but with c fixed at $\frac{1}{2}$. Grohe and Marx [39], give the above variant definition, with c fixed at $\frac{1}{2}$, and instead call it a *balanced separator*. Reed [82] defines separators using our standard definition, with $c = \frac{2}{3}$. Bodlaender [7] defines "type-1" and "type-2" separators (see below for an explanation), which have variable proportion (i.e. allow for different values of c), but are not defined on sets other than $V(G)$. Sometimes [7, 30, 39] instead of considering components in $G - X$, separators are defined as partitioning the vertex set of $G - X$ into exactly two parts A and B, such that no edge has an endpoint in both parts and $|A \cap S|, |B \cap S| \le c|S|$. (In fact, Bodlaender [7] uses both this definition and the standard "components of $G - X$" definition as the difference between type-1 and type-2 separators.) As long as $c \ge \frac{2}{3}$, this is equivalent to considering the components, since Lemma 2.5 and Corollary 2.6 allow partitioning of the components into parts A and B. However, for lower values of c this no longer holds, for example, if $c = \frac{1}{2}$, it is possible that each component contains exactly $\frac{1}{3}$ of the vertices of S, so there is no acceptable partition into A and B. As a result, $c = \frac{2}{3}$ and $c = \frac{1}{2}$ are the most "natural" choices for c.

Fortunately, $\mathsf{sep}_c(G), \mathsf{sep}^*_c(G), \mathsf{sep}_{c'}(G)$ and $\mathsf{sep}^*_{c'}(G)$ are all tied for all $c, c' \in [\frac{1}{2}, 1)$.[†]

Robertson and Seymour [90] proved that

$$\mathsf{sep}_{\frac{1}{2}}(G) \le \mathsf{tw}(G) + 1.$$

(Of course, they did not use our notation.) Robertson and Seymour [90, 93] also proved that

$$\mathsf{tw}(G) + 1 \le 4\,\mathsf{sep}_{\frac{2}{3}}(G) - 2. \tag{2.1}$$

(Reed [81, 82] gives a more accessible proof of this upper bound.) Flum and Grohe [29] proved that

$$\mathsf{tw}(G) \le 3\,\mathsf{sep}^*_{\frac{1}{2}}(G) - 2. \tag{2.2}$$

[†]Fox [30] defines a separator to be a set $X \subseteq V(G)$ that partitions $V(G)$ into $X \cup A \cup B$ with no $A - B$ edge and $|A|, |B| \le \frac{2}{3}|V(G)|$. Fox then defines the separation number to be the minimum integer k such that each subgraph of G contains a separator of size k. However, we will not consider this definition here.

We now provide a series of lemmas to prove slightly stronger results. Specifically, Lemma 2.8 replaces the multiplicative constant "4" by "3" in (2.1), and the multiplicative constant "3" by "2" in (2.2).

First, we prove a useful well-known lemma for dealing with components of a graph.

Lemma 2.5. *For every graph G and for all sets $X, S \subseteq V(G)$ such that no component of $G - X$ contains more than half of the vertices of $S - X$, it is possible to partition the components of $G - X$ into at most three parts such that no part contains more than half the vertices of $S - X$.*

Proof. If $G-X$ contains at most three components, the claim follows immediately. Hence assume $G - X$ contains at least four components. Initially, let each part simply contain a single component. Merge parts as long as the merge does not cause the new part to contain more than half the vertices of $S-X$. Now if two parts contain more than $\frac{1}{4}$ of the vertices of $S - X$ each, then all other parts (of which there must be at least two) contain, in total, less than $\frac{1}{2}$ of the vertices of $S-X$. Then merge all other parts together, leaving the partition with exactly three parts. Alternatively only one part (at most) contains more than $\frac{1}{4}$ of the vertices of $S - X$. So at least three parts contain at most $\frac{1}{4}$ of the vertices of $S - X$, and so merge two of them. This lowers the number of parts in the partition. As long as there are four or more parts, one of these operations can be performed, so repeat until at most three parts remain. $\square$

Corollary 2.6. *For every graph G and for all sets $X, S \subseteq V(G)$ such that no component of $G-X$ contains more than two-thirds of the vertices of $S-X$, it is possible to partition the components of $G - X$ into at most two parts such that no part contains more than two-thirds the vertices of $S - X$.*

This corollary follows by a very similar argument to Lemma 2.5.

The following argument is similar to that provided in [90].

Lemma 2.7 (Robertson and Seymour [90])**.** *For every graph G and for all $c \in [\frac{1}{2}, 1)$,*

$$\mathsf{sep}_c(G) \leq \mathsf{tw}(G) + 1.$$

Proof. Fix $S \subseteq V(G)$ and let $k := \mathsf{tw}(G) + 1$. It is sufficient to construct a $(k, S, \frac{1}{2})$-separator for G. The graph G has a normalised tree decomposition T with maximum bag size k, by Lemma 2.2. Consider a pair of adjacent bags X, Y. Let T_X and T_Y be the subtrees of $T - XY$ containing bags X and Y respectively. Let $U_X \subseteq V(G)$ be the set of vertices only appearing in bags of T_X, and U_Y the set of vertices only appearing in bags of T_Y. Then $U_X, X \cap Y, U_Y$ is a partition of $V(G)$ such that no edge has an endpoint in

U_X and U_Y. Each component of $G-(X\cap Y)$ is contained entirely within U_X or U_Y. Say $Q\subseteq V(G)$ is *large* if $|Q\cap S|>\frac{1}{2}|S-(X\cap Y)|$.

If neither U_X or U_Y is large, then no component of $G-(X\cap Y)$ is large. Hence $X\cap Y$ is a $(|X\cap Y|,S,\frac{1}{2})$-separator. Since $|X\cap Y|\leq|Y|\leq k$, this is sufficient.

Alternatively, for all edges $XY\in E(T)$, exactly one of U_X and U_Y is large. (If both sets are large, then $|S-(X\cap Y)|=|U_X\cap S|+|U_Y\cap S|>|S-(X\cap Y)|$, which is a contradiction.) Orient the edge $XY\in E(T)$ towards X if U_X is large, or towards Y if U_Y is large.

Now there must be a bag B with outdegree 0. If B is a $(|B|,S,\frac{1}{2})$-separator, then since $|B|=k$, the result is achieved. Otherwise, exactly one component C of $G-B$ is large. The vertices of C only appear in the bags of a single subtree of $T-B$. Label that subtree as T', and let A denote the bag of T' adjacent to B. Recall there is a partition $V(G)$ into $U_A, A\cap B, U_B$ where $|U_B\cap S|>\frac{1}{2}|S-(A\cap B)|$, since the edge AB is oriented towards B. Hence $|U_A\cap S|<\frac{1}{2}|S-(A\cap B)|$. Also note the vertices of $G-B$ that only appear in the bags of T' are exactly the vertices of U_A. Hence $C\subseteq U_A$, and $|U_A\cap S|>\frac{1}{2}|S-B|$.

So $\frac{1}{2}|S-B|<|U_A\cap S|<\frac{1}{2}|S-(A\cap B)|$. By our normalisation, $|A\cap B|=|B|-1$. So $|S-B|\geq|S-(A\cap B)|-1$. Thus $|S-(A\cap B)|-1<2|U_A\cap S|<|S-(A\cap B)|$, which is a contradiction since $|S-(A\cap B)|-1$, $2|U_A\cap S|$ and $|S-(A\cap B)|$ are all integers. □

Now we provide a proof of the upper bound.

Lemma 2.8. *For every graph G, for all $c\in[\frac{1}{2},1)$,*

$$\mathsf{bn}(G)\leq\frac{1}{1-c}\mathsf{sep}_c^*(G).$$

Proof. Say β is an optimal bramble of G with a minimum hitting set H. That is, $|H|=\mathsf{bn}(G)$. For the sake of a contradiction, assume that $(1-c)\mathsf{bn}(G)>\mathsf{sep}_c^*(G)$. So there is a $(\mathsf{sep}_c^*(G),H,c)^*$-separator X. If X is a hitting set for β then $\mathsf{bn}(G)\leq|X|\leq sep_c^*(G)<(1-c)\mathsf{bn}(G)$, which is a contradiction. So X is not a hitting set for β. Thus some bramble element of β is entirely within a component of $G-X$. Only one such component can contain bramble elements. Call this component C. Then we can hit every bramble element of β with the vertices of X or the vertices of H inside C, that is, $X\cup(H\cap V(C))$ is a hitting set. Since X is a $(\mathsf{sep}_c^*(G),H,c)^*$-separator, $|H\cap V(C)|\leq c|H|$. Thus $|X\cup(H\cap V(C))|=|X|+|H\cap V(C)|\leq|X|+c|H|\leq\mathsf{sep}_c^*(G)+c|H|<(1-c)|H|+c|H|=|H|$. Thus $X\cup(H\cap V(C))$ is a hitting set smaller than the minimum hitting set, a contradiction. □

Hence, from the above it follows that for $c\in[\frac{1}{2},1)$,

$$\mathsf{sep}_c^*(G)\leq\mathsf{sep}_c(G)\leq\mathsf{tw}(G)+1=\mathsf{bn}(G)\leq\frac{1}{1-c}\mathsf{sep}_c^*(G)\leq\frac{1}{1-c}\mathsf{sep}_c(G).$$

Each of the above inequalities is tight. In particular, the second and third inequalities are tight for K_n. We now show that the first and fourth inequalities are tight.

Lemma 2.9. *For a given $c \in [\frac{1}{2}, 1)$, if n, k are integers such that $k > \frac{c}{1-c} + 1$ and $n \geq \frac{k-1}{1-c}$, then $\mathsf{sep}^*_c(\psi_{n,k}) = \mathsf{sep}_c(\psi_{n,k}) = n$.*

Proof. Let $G := \psi_{n,k}$. It follows from the definition that $\mathsf{sep}^*_c(G) \leq \mathsf{sep}_c(G)$. Hence it is sufficient to show that $\mathsf{sep}^*_c(G) \geq n$ and that $\mathsf{sep}_c(G) \leq n$. We prove these facts in order.

To prove our first inequality, we shall show that if X is a $(|X|, V(G), c)^*$-separator, then $|X| \geq n$. Suppose for the sake of a contradiction that $|X| \leq n - 1$. If $|X| \leq n - 2$, then $G - X$ contains at least two vertices of A, and so $G - X$ is connected. Alternatively, if $|X| = n - 1$, then either at least two vertices of A remain and $G - X$ is connected, or $X \subset A$. Thus, there are two cases to consider: firstly, when $G - X$ is connected, and secondly, when $X \subset A$ and $|X| = n - 1$. In the first case, the only component of $G - X$ contains at least $(k+1)n - (n-1) = kn+1$ vertices. Since X is a $(|X|, V(G), c)^*$-separator, this component can contain at most $c((k + 1)n)$ vertices. So $kn + 1 \leq c(kn + n)$. Thus $(1 - c)kn \leq cn - 1 < cn$, and $k < \frac{c}{1-c}$, which contradicts our assumption on k. In the second case, $G - X$ contains one component that contains a single vertex of A and all but k vertices of B. Label this component C and note it contains $nk - k + 1$ vertices. Every other component is an isolated vertex. Since X is a $(|X|, V(G), c)^*$-separator, $|C| \leq c(kn + n)$. So $(n - 1)k + 1 \leq c(kn + n)$, and thus $n(k - ck - c) \leq k - 1$. Since $k > \frac{c}{1-c} + 1$, it follows that $n(1 - c) < k - 1$, which contradicts our lower bound on n. Thus X is not a $(|X|, V(G), c)^*$-separator when $|X| \leq n - 1$, and $\mathsf{sep}^*_c(G) \geq n$.

Secondly, it suffices to show that for any $S \subseteq V(G)$ there exists a set X such that $|X| \leq n$ and no component in $G - X$ contains more than $\frac{1}{2}|S - X|$ vertices of $S - X$. Now if $|S| \leq n$, then simply set $X = S$. If $|S \cap B| \geq 2$, then set $X = A$. Then each component C of $G - X$ is an isolated vertex, and $|C \cap S| \leq 1 = \frac{1}{2} \cdot 2 \leq \frac{1}{2}|S - X|$. Finally, the only other possibility is that $|S| = n + 1$ and $|S \cap B| = 1$. In that case, label the vertex of $S \cap B$ as v and let $X = N(v)$. So $|X| = n - 1$, $X \subset A$, and since all of A must be in S, only two vertices of S are in $G - X$, that is, v and the one vertex of A not adjacent to v. Since these vertices are in different components, no component of $G - X$ is too large with respect to S, and X is the required $(|X|, S, c)$-separator. Hence $\mathsf{sep}_c(G) \leq n$. □

2.6 Branchwidth and Tangles

A *branch decomposition* of a graph G is a pair (T, θ) where T is a tree with each node having degree 3 or 1, and θ is a bijective mapping from the edges of G to the leaves of T. A vertex x of G is *across* an edge e of T if there are edges xy and xz of G mapped to leaves in different subtrees of $T - e$. The *order* of an edge e of T is the number of vertices of G across e. The *width* of a branch decomposition is the maximum order of an edge. Finally, the *branchwidth* $\mathsf{bw}(G)$ of a graph G is the minimum width over all branch decompositions of G. Note that if $|E(G)| \leq 1$, there are no branch decompositions of G, in which case we define $\mathsf{bw}(G) = 0$. Robertson and Seymour [92] first defined branchwidth, where it was defined more generally for hypergraphs; here we just consider the case of simple graphs.

Tangles were first defined by Robertson and Seymour [92]. Their definition is in terms of sets of separations of graphs. (Note, importantly, that a *separation* is not the same as a *separator* as defined in Section 2.5.) We omit their definition and instead present the following, initially given by Reed [82].

A set τ of connected subgraphs of a graph G is a *tangle* if for all sets of three subgraphs $A, B, C \in \tau$, there exists either a vertex v of G in $V(A \cap B \cap C)$, or an edge e of G such that each of A, B and C contain at least one endpoint of e. Clearly a tangle is also a bramble—this is the main advantage of this definition. The *order* of a tangle is equal to its order when viewed as a bramble. The *tangle number* $\mathsf{tn}(G)$ is the maximum order of a tangle in G.

When defined with respect to hypergraphs, treewidth and tangle number are tied to the maximum of branchwidth and the size of the largest edge. So for simple graphs, there are a few exceptional cases when $\mathsf{bw}(G) < 2$, which we shall deal with briefly. If G is connected and $\mathsf{bw}(G) \leq 1$, then G contains at most one vertex with degree greater than 1 (that is, G is a star), and $\mathsf{bn}(G) = \mathsf{tn}(G) \leq 2$. Henceforth, assume $\mathsf{bw}(G) \geq 2$.

Robertson and Seymour [92] prove the following relation between tangle number and branchwidth; we omit the proof. Instead we show that $\mathsf{tn}(G), \mathsf{bw}(G), \mathsf{bn}(G)$ and $\mathsf{tw}(G)$ are all tied by small constant factors.

Theorem 2.10 (Robertson and Seymour [92])**.** *For a graph G, if $\mathsf{bw}(G) \geq 2$, then*

$$\mathsf{bw}(G) = \mathsf{tn}(G).$$

Robertson and Seymour [92] proved that $\mathsf{bn}(G) \leq \frac{3}{2}\mathsf{tn}(G)$. Reed [82] provided a short proof that $\mathsf{bn}(G) \leq 3\,\mathsf{tn}(G)$. Here, we modify Reed's proof to show that $\mathsf{bn}(G) \leq 2\,\mathsf{tn}(G)$.

Lemma 2.11. *For every graph G,*

$$\mathsf{tn}(G) \leq \mathsf{bn}(G) \leq 2\,\mathsf{tn}(G).$$

Proof. Since every tangle is also a bramble, $\mathsf{tn}(G) \leq \mathsf{bn}(G)$.

To prove that $\mathsf{bn}(G) \leq 2\,\mathsf{tn}(G)$, let $k := \mathsf{bn}(G)$, and say β is a bramble of G of order k. Consider a set $S \subseteq V(G)$ with $|S| < k$. If two components of $G - S$ entirely contain a bramble element of β, then those two bramble elements do not touch. Alternatively, if no component of $G - S$ entirely contains a bramble element, then all bramble elements use a vertex in S, and S is a hitting set of smaller order than the minimum hitting set. Thus exactly one component S' of $G - S$ entirely contains a bramble element of β. Clearly, $V(S') \cap S = \emptyset$.

Define $\tau := \{S' : S \subseteq V(G), |S| < \frac{k}{2}\}$. To prove that τ is a tangle, let T_1, T_2, T_3 be three elements of τ. Say $T_i = S'_i$ for each i. Since $|S_1 \cup S_2| < k$, some bramble element B_1 of β does not intersect $S_1 \cup S_2$. Similarly, some bramble element B_2 does not intersect $S_2 \cup S_3$. Since B_1 does not intersect S_1, it is entirely within one component of $G - S_1$, that is, $B_1 \subseteq T_1$. Similarly, $B_1 \subseteq T_2$ and $B_2 \subseteq T_2 \cap T_3$. Since $B_1, B_2 \in \beta$, they either share a vertex v, or there is an edge e with one endpoint in B_1 and the other in B_2. In the first case, $v \in V(T_1 \cap T_2 \cap T_3)$. In the second case, one endpoint of e is in $T_1 \cap T_2$, the other in $T_2 \cap T_3$. It follows that τ is a tangle. The order of τ is at least $\frac{k}{2}$, since a set X of size less than $\frac{k}{2}$ has a defined $X' \in \tau$, and so X does not intersect all subgraphs of τ. Then $\mathsf{tn}(G) \geq \frac{k}{2}$. □

We now provide a proof for a direct relationship between branchwidth and treewidth. Note again these proofs are modified versions of those in [92].

Lemma 2.12 (Robertson and Seymour [92]). *For a graph G, if $\mathsf{bw}(G) \geq 2$ then*

$$\mathsf{bw}(G) \leq \mathsf{tw}(G) + 1 \leq \frac{3}{2}\,\mathsf{bw}(G).$$

Proof. We prove the second inequality first. Assume no vertex is isolated. Let $k := \mathsf{bw}(G)$, and let (T, θ) be a branch decomposition of order k. We construct a tree decomposition with T as the underlying tree, and where B_x will denote the bag indexed by each node x of T. A node x in T has degree 3 or 1. If x has degree 1, then let B_x contain the two endpoints of $e = \theta^{-1}(x)$. If x has degree 3, then let B_x be the set of vertices that are across at least one edge incident to x. We now show that this is a tree decomposition. Every vertex appears at least once in the tree decomposition. Also, for every edge $vw \in E(G)$, the bag of the leaf node $\theta(vw)$ contains both v and w. If we consider vertex $v \in V(G)$ incident with vw and vu, then v is across every edge in T on the path from $\theta(vw)$ to $\theta(vu)$. Thus, v is in every bag indexed by a node on that path. Such a path exists for all neighbours w, u of v. It follows that the subtree of nodes indexing bags containing v form a subtree of T. Thus $(T, (B_x)_{x \in V(T)})$ is a tree decomposition of G. A bag indexed by a leaf node has size 2. If x is not a leaf, then B_x contains the vertices that are across at least one edge incident to x. Suppose v is across exactly one such edge e. Then there exists $\theta(vw)$ and $\theta(vu)$ in different subtrees of $T - e$. Without loss of generality, $\theta(vw)$ is in the subtree containing x. But then the path from x to $\theta(vw)$ uses one of the other two edges incident to x. Hence if v is in B_x then v is across at least two edges incident to x. If the sets of vertices across the three edges incident to x are A, B and C respectively, then $|A|+|B|+|C| \geq 2|B_x|$. But $|A|+|B|+|C| \leq 3k$. Therefore, regardless of whether x is a leaf, $|B_x| \leq \max\{2, \frac{3}{2}k\} = \frac{3}{2}k$ (since $k \geq 2$). Therefore $\mathsf{tw}(G)+1 \leq \frac{3}{2}k$.

Now we prove the first inequality. Let $k := \mathsf{tw}(G)+1$. Hence there exists a tree decomposition $(T, (B_x)_{x \in V(T)})$ with maximum bag size k; choose this tree decomposition such that T is node-minimal, and such that the subtree induced by $\{x \in V(T) : v \in B_x\}$ is also node-minimal for each $v \in V(G)$. If $k < 2$, then G contains no edge, and $\mathsf{bw}(G) = 0$. Now assume $k \geq 2$ and $E(G) \neq \emptyset$. Since the first inequality is trivial when G is complete, we assume otherwise, and thus T is not a single node.

Note the following facts: if x is a node of T with degree 2, then there exists some pair of adjacent vertices v, w such that B_x is the only bag containing v and w. (Otherwise, T would violate the minimality properties.) Similarly, if x is a leaf node, then there exists some $v \in V(G)$ such that B_x is the only bag containing v. The bag B_x also contains the neighbours of v, but nothing else.

Now, for every edge $vw \in E(G)$, choose some bag B_x containing v and w. Unless x is a leaf with $B_x = \{v, w\}$, add to T a new node y adjacent to x, such that $B_y = \{v, w\}$. Clearly $(T, (B_x)_{x \in V(T)})$ is still a tree decomposition of the same width. From our above facts, every leaf node is either newly constructed or was already of the form $B_x = \{v, w\}$. Also, every node that previously had degree 2 now has higher degree. A node that was previously a leaf either remains a leaf, or now has degree at least 3. So no node of the new T has degree 2.

If a node x has degree greater than 3, then delete the edges from x to two of its neighbours (denoted y, z), and add to T a new node s adjacent to x, y and z. Let $B_s := B_x \cap (B_y \cup B_z)$. Clearly this is still a tree decomposition of the same width. Now the degree of x has been reduced by 1, and the new node has degree 3. Repeat this process until all nodes have either degree 3 or 1.

Since each leaf bag contains exactly the endpoints of an edge (and no edge has both endpoints in more than one leaf), there is a bijective mapping θ that takes $vw \in E(G)$ to the leaf node containing v and w. Together with T, this gives a branch decomposition of G. If $xy \in E(T)$, then all edges of G across xy are in $B_x \cap B_y$. So the order of this branch decomposition is at most k. Thus $\mathsf{bw}(G) \leq \mathsf{tw}(G)+1$.

(Note that our minimality properties would imply that $|B_x \cap B_y| < k$, however converting the tree to ensure that all nodes have degree 3 or 1 does not necessarily maintain this.) □

Robertson and Seymour [92] showed the bounds in Lemma 2.12 are tight. The upper bound on $\mathsf{tw}(G)$ in Lemma 2.12 is tight for K_n when n is divisible by 3, since $\mathsf{tw}(K_n) = n-1$ and $\mathsf{bw}(K_n) = \mathsf{tn}(K_n) = \frac{2}{3}n$. The lower bound on $\mathsf{tw}(G)$ is tight when $n \geq 4$ and G is the graph $K_{n,n}$ minus a perfect matching. In this case $\mathsf{tw}(G) + 1 = \mathsf{bw}(G) = \mathsf{tn}(G) = n$.

2.7 Tree Products

For a tree T, let $T \cdot K_k$ denote the lexicographic product of T with K_k. That is, $T \cdot K_k$ is the graph created by taking T and replacing each vertex with a clique of k vertices, and replacing each edge with all possible edges between the two new cliques. The *lexicographic tree product number* of G, denoted $\mathsf{ltp}(G)$, is the minimum integer k such that G is a minor of the graph $T \cdot K_k$ for some tree T.

Lemma 2.13. *For every graph G,*

$$\mathsf{ltp}(G) - 1 \leq \mathsf{tw}(G) \leq 2\,\mathsf{ltp}(G) - 1.$$

Proof. We prove the first inequality. Let $\mathsf{tw}(G) + 1 = k$, and let T be the underlying tree of a tree decomposition of G with maximum bag size k. It is sufficient to show G is a minor of $T \cdot K_k$. For each vertex v of G, let the branch set R_v contain a single vertex of each k vertex clique of $T \cdot K_k$ that corresponds to a bag of T containing v. It is possible to ensure that no vertex in placed in more than one branch set since each clique contains k vertices and each bag contains at most k vertices. Each of these branch set is connected due to the properties of a tree decomposition and the structure of $T \cdot K_k$. Similarly, for each edge vw of G there is an edge between R_v and R_w. So G is a minor of $T \cdot K_k$.

Now we prove the second inequality. We first prove $\mathsf{tw}(T \cdot K_k) \leq 2k - 1$ for every tree T, as follows. Take the tree T and subdivide each edge; this will be the underlying tree of the tree decomposition. Then place each vertex of $T \cdot K_k$ in the bag indexed by the corresponding node, and in the bags indexed by the neighbours of the corresponding node (that is, the subdivided nodes from the incident edges). It is clear that this is a valid tree decomposition. A bag indexed by a subdivided edge node contains $2k$ vertices (all the vertices in both neighbouring bags, which contain at most k vertices each). Since these are the largest bags, $\mathsf{ltp}(T \cdot K_k) \leq 2k - 1$. If $\mathsf{ltp}(G) = k$, then G is a minor of some $T \cdot K_k$. Thus $\mathsf{tw}(G) \leq \mathsf{tw}(T \cdot K_k) \leq 2\,k - 1 = 2\,\mathsf{ltp}(G) - 1$, as required. □

If T is a tree, let $T^{(k)}$ denote the Cartesian product of T with K_k. That is, the graph with vertex set $\{(x, i) : x \in T, i \in \{1, \ldots, k\}\}$ and with an edge between (x, i) and (y, j) when $x = y$, or when $xy \in E(T)$ and $i = j$. Then define the *Cartesian tree product number* of G, $\mathsf{ctp}(G)$, to be the minimum integer k such that G is a minor of $T^{(k)}$. The parameter $\mathsf{ctp}(G)$ was first defined by van der Holst [111] and Colin de Verdière [17], however they did not use that name or notation. (Instead, they called it *largeur d'arborescence*, and denoted it by $\mathsf{la}(G)$.) They also proved the following result. We provide a different proof.

Lemma 2.14 (Colin de Verdière [17], van der Holst [111])**.** *For every graph G,*

$$\mathsf{ctp}(G) - 1 \leq \mathsf{tw}(G) \leq \mathsf{ctp}(G).$$

Proof. Let $k := \mathsf{tw}(G)$. By Lemma 2.4, G is the spanning subgraph of a chordal graph G' that contains a $(k+1)$-clique but no $(k+2)$-clique. Let $(T, (B_x \subseteq V(G))_{x \in V(T)})$ be a minimum width tree decomposition of G'. This has width k and is also a tree decomposition of G. To prove the first inequality, it is sufficient to show that G is a minor of $T^{(k+1)}$. Let c be a $(k+1)$-colouring of G'. (It is well known that chordal graphs are perfect.) For each $v \in V(G)$, define the connected subgraph R_v of $T^{(k+1)}$ such that $R_v := \{(x, c(v)) : v \in B_x\}$. If $(x, i) \in V(R_v) \cap V(R_w)$ then both v and w are in B_x and $c(v) = c(w) = i$. But if v and w share a bag then $vw \in E(G')$, which contradicts the vertex colouring c. So the subgraphs R_v are pairwise disjoint, for all $v \in V(G)$. If $vw \in E(G)$, then v and w share a bag B_x. Hence there is an edge $(x, c(v))(x, c(w))$ between the subgraphs R_v and R_w. Hence the R_v subgraphs form a G-model of $T^{(k+1)}$.

Now we prove the second inequality. Let $k := \mathsf{ctp}(G)$, and choose tree T such that G is a minor of $T^{(k)}$. Since $\mathsf{tw}(G) \leq \mathsf{tw}(T^{(k)})$, it is sufficient to show that $\mathsf{tw}(T^{(k)}) \leq k$. Let T' be the tree T with each edge subdivided k times. Label the vertices created by subdividing $xy \in E(T)$ as $xy(1), \dots, xy(k)$, such that $xy(1)$ is adjacent to x and $xy(k)$ is adjacent to y. Construct $(T', (B_x \subseteq V(G))_{x \in V(T')})$ as follows. For a vertex $x \in T$, let $B_x = \{(x, i) | i \in \{1, \dots, k\}\}$. For a subdivision vertex $xy(j)$, let $B_{xy(j)} = \{(x, i), (y, i') | 1 \leq i' \leq j \leq i \leq k\}$. This is a valid tree decomposition with maximum bag size $k+1$. Hence $\mathsf{tw}(T^{(k)}) \leq k$ as required. □

We now show the first inequalities in Lemmas 2.13 and 2.14 are tight.

Lemma 2.15. *If $n \geq 3$, then $\mathsf{ltp}(\psi_{n,1}) = \mathsf{ctp}(\psi_{n,1}) = \mathsf{tw}(\psi_{n,1}) + 1 = n$.*

Proof. Let $G := \psi_{n,1}$. Since $T^{(k)} \subseteq T \cdot K_k$, it follows that $\mathsf{ltp}(G) \leq \mathsf{ctp}(G)$. Also, by Lemma 2.13 and Lemma 2.14, $\mathsf{ltp}(G) \leq \mathsf{tw}(G) + 1$ and $\mathsf{ctp}(G) \leq \mathsf{tw}(G) + 1$. Hence it is sufficient to show that $\mathsf{tw}(G) + 1 \leq n$, and that $\mathsf{ltp}(G) \geq n$. The first inequality follows from Lemma 2.1.

Now we show that $\mathsf{ltp}(G) \geq n$. Suppose for the sake of a contradiction that $\mathsf{ltp}(G) \leq n - 1$. So there is some tree T such that G is a minor of $T \cdot K_{n-1}$. We can assume that

T is a node-minimal such tree. Label the vertices of A by $1, \dots, n$ and the vertices of B by $1', \dots, n'$, such that each i, i' pair is non-adjacent. For a node $x \in T$, let C_x be the corresponding $(n-1)$ vertex clique in $T \cdot K_{n-1}$. For each $v \in V(G)$, let X_v be the branch set corresponding to v in $T \cdot K_{n-1}$. Say two branch sets *touch* if there is an edge between them.

Pick a leaf node x of T and let y be the parent of x. By node minimality, there at least one vertex $v \in V(G)$ such that $X_v \subseteq C_x$. We claim there is exactly one, and that it must be a vertex of B. Suppose for the sake of a contradiction that there exists vertex v such that $X_v \subseteq C_x$ and $v \in A$. Now X_v can touch at most $(n-1)-1+(n-1) = 2n-3$ other branch sets. But the degree of v is $(n-1)+(n-1) = 2n-2$. Hence if $X_v \subseteq C_x$ then $v \in B$. Now if there are two vertices v, w such that $X_v, X_w \subseteq C_x$, then note that $v, w \in B$ and so $A = N(w) \cup N(v)$. But then the branch sets of all vertices of A intersect $C_x \cup C_y$, and so some vertex of A has its branch set entirely inside C_x, since $|C_y| = n-1$. However, this contradicts our previous result, and this completes the proof of the claim.

Without loss of generality, say that $X_{1'} \subseteq C_x$. Now the vertex $1'$ has neighbourhood $2, \dots, n$, and all of those branch sets must intersect C_y, since they are not entirely inside C_x but must touch $X_{1'}$. Since $|C_y| = n-1$, these are the only branch sets intersecting C_y. Now X_1 is entirely inside exactly one component of $(T \cdot K_{n-1}) - C_y$; let z the node of T such that z is adjacent to y and C_z is inside the component containing X_1. Since 1 is adjacent to $2', \dots, n'$, it follows $X_{2'}, \dots, X_{n'}$ are also entirely inside this component. By node minimality we can assume that one of the branch sets $X_2, \dots, X_n$ does not intersect C_z, without loss of generality it is X_2. Since $X_1, X_{3'}, \dots, X_{n'}$ must touch X_2, it follows $X_1, X_{3'}, \dots, X_{n'}$ intersect C_z. Since $|C_z| = n-1$, these are the only branch sets intersecting C_z. However, consider the branch set $X_{2'}$. It is entirely in the component of $(T \cdot K_{n-1}) - C_y$ containing C_z, but also must touch X_3. However, X_3 does not intersect this component at all (since it does not intersect C_z), and $X_{2'}$ also does not intersect C_z. Thus $X_{2'}$ and X_3 cannot touch. This gives the desired contradiction. □

Also see Markov and Shi [76] for a similar result. The second inequalities in Lemmas 2.13 and 2.14 are tight for K_n (for Lemma 2.13, ensure that n is even).

2.8 Linkedness

Reed [82] introduced the following definition. For a positive integer k, a set S of vertices in a graph G is *k-linked* if for every set $X \subseteq V(G)$ such that $|X| < k$ there is a component of $G - X$ that contains more than half of the vertices in S. The *linkedness* of G, denoted by $\mathsf{link}(G)$, is the maximum integer k for which G contains a k-linked set. Linkedness is used by Reed [82] in his proof of the Grid Minor Theorem.

Lemma 2.16 (Reed [82])**.** *For every graph G,*

$$\mathsf{link}(G) \leq \mathsf{bn}(G) \leq 2\,\mathsf{link}(G).$$

Proof. We first prove that $\mathsf{link}(G) \leq \mathsf{bn}(G)$. Let $k := \mathsf{link}(G)$, and let S be a k-linked set of G. So for every set $X \subseteq V(G)$ such that $|X| < k$, there exists some component of $G - X$ that contains more than half of the vertices of S. Let C_X denote this component, and then let $\beta = \{C_X | X \subseteq V(G), |X| < k\}$. Clearly each element of β is connected, and any two elements touch since they each contain more than half the vertices of S. Thus β is a bramble. Let H be a hitting set of β. If $|H| < k$, then there exists some $C_H \in \beta$, but $H \cap C_H = \emptyset$, and H is not a hitting set. Thus $|H| \geq k$ and so $\mathsf{bn}(G) \geq k$, as required.

Now we prove that $\mathsf{bn}(G) \leq 2\,\mathsf{link}(G)$. Assume for the sake of a contradiction that $\mathsf{bn}(G) > 2\,\mathsf{link}(G)$. Let $k := \mathsf{link}(G)$, so G is not $(k+1)$-linked. Let H be a minimum hitting set for a bramble β of G of largest order. Since H is not $(k+1)$-linked, there exists a set X of order at most k such that no component of $G - X$ contains more than half of the vertices in H. Note that at most one component of $G - X$ can entirely contain a bramble element of β (otherwise two bramble elements do not touch). If no component of $G - X$ entirely contains a bramble element of β, then X is a hitting set for β of order $|X| \leq k < \frac{1}{2}\mathsf{bn}(G)$, which contradicts the order of the minimum hitting set. Finally, if a component of $G - X$ entirely contains some bramble element of β, then let $H' \subset H$ be the set of vertices of H in that component. Now H' intersects all of the bramble elements contained in the component (since those bramble elements do not intersect any other vertices of H), and X intersects all remaining bramble elements, as in the previous case. Thus, $H' \cup X$ is a hitting set for β. However, $|X| \leq k < \frac{1}{2}\mathsf{bn}(G)$, and by the choice of X, $|H'| \leq \frac{1}{2}|H| = \frac{1}{2}\mathsf{bn}(G)$. So $|H' \cup X| = |H'| + |X| < \mathsf{bn}(G)$, again contradicting the order of the minimum hitting set. □

When n is even, $\mathsf{link}(K_n) = \frac{n}{2}$, so the second inequality in Lemma 2.16 is tight. We now show that the first inequality is tight.

Lemma 2.17. *If $k \geq 2$ and $n \geq 3$, then $\mathsf{link}(\psi_{n,k}) = \mathsf{bn}(\psi_{n,k}) = n$.*

Proof. Let $G := \psi_{n,k}$. Then $\mathsf{bn}(G) = n$ by Lemma 2.1 and $\mathsf{link}(G) \leq \mathsf{bn}(G)$ by Lemma 2.16. Hence it is sufficient to show that $\mathsf{link}(G) \geq n$. To do this, we shall show that $V(G)$ is an n-linked set; that is, if X is a set of vertices and $|X| < n$, then $G - X$ must contain a component containing more than half of $V(G)$.

If $G - X$ contains at least two vertices of A, then it is connected. Since $|X| \leq n - 1$, if $G - X$ is not connected, then $|X| = n - 1$ and $X \subset A$. In the first case, $|G - X| \geq$

$(k+1)n-n+1=kn+1$ and $|G|=(k+1)n=kn+n$. Since $k\geq 1$, $G-X$ contains more than half of $V(G)$, as required. In the second case, there exists a component containing a single vertex of A and all $kn-k$ of its neighbours in B. Since $k\geq 2$ and $n\geq 3$, it follows that $kn-k+1>\frac{1}{2}(kn+n)$. Hence the large component of $G-X$ is large enough. Thus $V(G)$ is an n-linked set, as required. □

2.9 Well-linked and k-Connected Sets

For a graph G, a set $S\subseteq V(G)$ is *well-linked* if for every pair $A,B\subseteq S$ such that $|A|=|B|$, there exists a set of $|A|$ vertex-disjoint paths from A to B. If we can ensure these vertex-disjoint paths also have no internal vertices in S, then S is *externally-well-linked*. The notion of a well-linked set was first defined by Reed [82], while a similar definition was used by Robertson *et al.* [96]. Reed also described externally-well-linked sets in the same paper (but did not define it explicitly) and stated but did not prove that S is well-linked iff S is externally-well-linked. We provide a proof below. The *well-linked number* of G, denoted $\mathsf{wl}(G)$, is the size of the largest well-linked set in G.

Lemma 2.18 (Reed [82])**.** *S is well-linked iff S is externally-well-linked.*

Proof. It should be clear that if S is externally-well-linked that S is well-linked, so we prove the forward direction. Let $S\subseteq V(G)$ be well-linked. It is sufficient to show that for all $A,B\subseteq S$ with $|A|=|B|$ there are $|A|$ vertex-disjoint paths from A to B that are internally disjoint from S. Define $C:=S-(A\cup B)$ and $A':=A\cup C$ and $B':=B\cup C$. Now $S=A'\cup B'$. Since S is well-linked and $|A'|=|B'|$, there are $|A'|$ vertex-disjoint paths between A' and B'. Each such path uses exactly one vertex from A' and one vertex from B'. Thus, if $v\in C\subseteq A\cap B$, then the path containing v must simply be the singleton path $\{v\}$. Thus this path set contains a set of singleton paths for each vertex of C and, more importantly, a set of paths starting in $A'-C=A$ and ending at $B'-C=B$. Since every vertex of S is in either A' or B', and each path starts at a vertex in A' and ends at one in B', no internal vertex of these paths is in S. This is the required set of disjoint paths from A to B that are internally disjoint from S. □

Reed [82] proved that $\mathsf{bn}(G)\leq\mathsf{wl}(G)\leq 4\,\mathsf{bn}(G)$. We now provide Reed's proof of the first inequality.

Lemma 2.19 (Reed [82])**.** *For every graph G,*

$$\mathsf{bn}(G)\leq\mathsf{wl}(G).$$

Proof. Assume for the sake of a contradiction that $\mathsf{wl}(G) < \mathsf{bn}(G)$. Let β be a bramble of largest order, and H a minimal hitting set of β. Thus H is not well-linked (since $|H| = \mathsf{bn}(G) > \mathsf{wl}(G)$). Choose $A, B \subseteq H$ such that $|A| = |B|$ but there are not $|A|$ vertex-disjoint paths from A to B. By Menger's Theorem, there exists a set of vertices C with $|C| < |A|$ such that after deleting C, there is no A-B path in G. Now consider a bramble element of β. If two components of $G - C$ entirely contain bramble elements, then those bramble elements cannot touch. Thus, it follows that at most one component of $G - C$ entirely contains some bramble element. Label this component C'; if no such component exists label C' arbitrarily. Since C' does not contain vertices from both A and B, without loss of generality we assume $A \cap C' = \emptyset$. Thus all bramble elements entirely within C' are hit by vertices of $H - A$, and all others are hit by C. So $(H - A) \cup C$ is a hitting set for β, but $|(H - A) \cup C| = |H| - |A| + |C| < |H|$, contradicting the minimality of H. Hence $\mathsf{bn}(G) \leq \mathsf{wl}(G)$. □

We now modify the proof of Reed's second inequality to give the following stronger result.

Lemma 2.20. *For every graph G,*

$$\mathsf{wl}(G) \leq 3\,\mathsf{link}(G) \leq 3\,\mathsf{bn}(G).$$

Proof. We show that $\mathsf{wl}(G) \leq 3\,\mathsf{link}(G)$. For the sake of a contradiction, say $3\,\mathsf{link}(G) < \mathsf{wl}(G)$. Define $k := \frac{1}{3}\mathsf{wl}(G)$. Let S be the largest well-linked set. That is, $|S| = \mathsf{wl}(G)$. By Lemma 2.18 S is externally-well-linked. The set S is not $\lceil k \rceil$-linked since $\mathsf{link}(G) < \lceil k \rceil$. Hence there exists a set $X \subseteq V(G)$ with $|X| < \lceil k \rceil$ such that $G - X$ contains no component containing more than $\frac{1}{2}|S|$ vertices of S. Since $|X|$ is an integer, $|X| < k$. Let $a := |X \cap S|$.

Using an argument similar to Lemma 2.5, the components of $G - X$ can be partitioned into two or three parts, each with at most $\frac{1}{2}|S|$ vertices of S. Some part contains at least a third of the vertices of $S - X$. Let A be the set of vertices in S contained in that part, and let B be the set of vertices in S in the other parts of $G - X$. Now $\frac{1}{2}|S| \geq |A| \geq \frac{1}{3}|S - X| = \frac{1}{3}(|S| - a)$, and so $|B| \geq |S| - |S \cap X| - |A| \geq |S| - a - \frac{1}{2}|S|$. Remove vertices arbitrarily from the largest of A and B until these sets have the same order. Hence $|A| = |B|$ and $|A| \geq \min\{\frac{1}{3}(|S| - a), \frac{1}{2}|S| - a\}$. Since $A, B \subseteq S$ and S is externally-well-linked, there are $|A|$ vertex-disjoint paths from A to B with no internal vertices in S. Since A and B are in different components of $G - X$, these paths must use vertices of X, but more specifically, vertices of $X - S$. Thus there are at most $|X - S|$ such paths. Thus $|A| \leq |X - S| < k - a$.

Either $\frac{1}{3}(|S| - a) \leq |A| < k - a$ or $\frac{1}{2}|S| - a \leq |A| < k - a$, so $|S| < 3k$. However, $|S| = \mathsf{wl}(G) = 3k$, which is a contradiction.

The final inequality follows from Lemma 2.16. □

Lemma 2.19 is tight since $\mathsf{bn}(K_n) = \mathsf{wl}(K_n) = n$. We do not know if Lemma 2.20 is tight, but $\mathsf{wl}(G) \leq 2\,\mathsf{bn}(G) - 2$ would be best possible since $\mathsf{bn}(K_{2n,n}) = n + 1$ and $\mathsf{wl}(K_{2n,n}) = 2n$ (the larger part is the largest well-linked set).

Diestel *et al.* [23] defined the following: $S \subseteq V(G)$ is *k-connected* in G if $|S| \geq k$ and for all subsets $A, B \subseteq S$ with $|A| = |B| \leq k$, there are $|A|$ vertex-disjoint paths from A to B. If we can ensure these vertex-disjoint paths have no internal vertex or edge in $G[S]$, then S is *externally k-connected*. This notion was used in [23] to prove a short version of the grid minor theorem.

Note the obvious connection to well-linked sets: X is well-linked iff X is $|X|$-connected. Also note that Diestel [22], in his treatment of the grid minor theorem, provides a slightly different formulation of externally k-connected sets, which only requires vertex-disjoint paths between A and B when they are disjoint subsets of S. These definitions are equivalent, which can be proven using a similar argument as in Lemma 2.18. Diestel [22] also does not use the concept of k-connected sets, just externally k-connected sets.

Diestel *et al.* [23] prove the following, but due to its similarity between k-connected sets and well-linked sets, we omit the proof.

Lemma 2.21 (Diestel *et al.* [23])**.** *If G has $\mathsf{tw}(G) < k$ then G contains $(k+1)$-connected set of size $\geq 3k$. If G contains no externally $(k+1)$-connected set of size $\geq 3k$, then $\mathsf{tw}(G) < 4k$.*

2.10 Grid Minors

As previously mentioned in Chapter 1, a key part of the Graph Minor Structure Theorem is as follows: given a fixed planar graph H, there exists some integer r_H such that every graph with no H-minor has treewidth at most r_H. This cannot be generalised to when H is non-planar, since there exist planar graphs, the grids, with unbounded treewidth. (By virtue of being planar, the grids do not contain a non-planar H as a minor.) In fact, since every planar graph is the minor of some grid, it is sufficient to just consider the grids, which leads to the Grid Minor Theorem:

Theorem 2.22 (Robertson and Seymour [91])**.** *For each integer k there is a minimum integer $f(k)$ such that every graph with treewidth at least $f(k)$ contains the $k \times k$ grid as a minor.*

All of our previous sections have provided parameters with linear ties to treewidth. However, the order of the largest grid minor is not linearly tied to treewidth. The initial bound on $f(k)$ by Robertson and Seymour [91] was an iterated exponential tower. Later, Robertson *et al.* [96] improved this to $f(k) \leq 20^{2k^5}$. They also note, by use of a probabilistic argument, that $f(k) \geq \Omega(k^2 \log k)$. Diestel *et al.* [23] obtained an upper bound of $2^{5k^5 \log k}$, which is actually slightly worse than the bound provided by Robertson, Seymour and Thomas, but with a more succinct proof. Kawarabayashi and Kobayashi [52] proved that $f(k) \leq 2^{O(k^2 \log k)}$, and Seymour and Leaf [100] proved that $f(k) \leq 2^{O(k \log k)}$. A recently announced result of Chekuri and Chuzhoy [14] gave a polynomial bound of $f(k) \leq O(k^{228})$. Together with the following well-known lower bound, this would imply that treewidth and the order of the largest grid-minor are polynomially tied.

Lemma 2.23. *If G contains a $k \times k$ grid minor, then* $\mathsf{tw}(G) \geq k$.

Proof. If H is a minor of G then $\mathsf{tw}(H) \leq \mathsf{tw}(G)$. Thus it suffices to prove that the $k \times k$ grid H has treewidth at least k, which is implied if $\mathsf{bn}(H) \geq k + 1$. Consider H drawn in the plane. For a subgraph S of H, define a *top vertex* of S in the obvious way. (Note it is not necessarily unique.) Similarly define *bottom vertex*, *left vertex* and *right vertex*. Let subgraph H' of H be the top-left $(k - 1) \times (k - 1)$ grid in H. A *cross* is a subgraph containing exactly one row and column from H', and no vertices outside H'. Let X denote the bottom row of H, and Y the right column without its bottom vertex. Let $\beta := \{X, Y, \text{ all crosses}\}$. A pair of crosses intersect in two places. There is an edge from a bottom vertex of a cross to X and a right vertex of a cross to Y. There is also an edge from the right vertex of X to the bottom vertex of Y. Hence β is a bramble. If Z is a hitting set for β, it must contain $k - 1$ vertices of $V(H')$, for otherwise a row and column are not hit, and so a cross is not hit. The set Z must also contain two other vertices to hit X and Y. So $|Z| \geq k + 1$, as required. □

2.11 Grid-like Minors

A *grid-like-minor of order t* of a graph G is a set of paths $\mathcal{P}$ in G with a bipartite intersection graph that contains a K_t-minor. Note that if the intersection graph of $\mathcal{P}$ is partitioned A and B, then we can think of the set of paths A as being the “rows” of the “grid”, and the set B being the “columns”. Also note that an actual $k \times k$ grid gives rise to a set $\mathcal{P}$ with an intersection graph $K_{k,k}$ and as such contains a complete minor of order k+1. Let $\mathsf{glm}(G)$ be the maximum order of a grid-like-minor of G. Grid-like-minors were first defined by Reed and Wood [86] as a weakening of a grid minor; see Section 2.10. As a result of this weakening, it is easier to tie $\mathsf{glm}(G)$ to $\mathsf{tw}(G)$. This notion has also been applied to prove computational intractability results in monadic second order logic; see Ganian *et al.* [33], Kreutzer [61] and Kreutzer and Tazari [62, 63].

The following definitions were introduced by Fox [30].[†] Given a graph G, consider a bramble β together with a function w which assigns a weight to each subgraph in β, such that for any vertex v, the sum of the weights of the bramble elements containing v is at most 1. Let $h(\beta, w) = \sum_{X \in \beta} w(X)$. The *fractional Hadwiger number* of G, denoted $\mathsf{had}_f(G)$, is the maximum of $h(\beta, w)$ over all β, w where the weights assigned by w are non-negative real numbers. For a positive integer r, the *r-integral Hadwiger number* of G, denoted $\mathsf{had}_r(G)$, is the maximum of $h(\beta, w)$ over all β, w where the weights assigned by w are integer multiples of $\frac{1}{r}$. It is clear that $\mathsf{had}_f(G) \geq \mathsf{had}_r(G)$ for every G and positive integer r. As an example, the branch sets of a $K_{\mathsf{had}(G)}$-minor form a bramble, and we set the weight of each branch set to be 1. Thus $\mathsf{had}_f(G) \geq \mathsf{had}_r(G) \geq \mathsf{had}(G)$ for all positive integers r.

The graph $G \square K_2$ (that is, the Cartesian product of G with K_2) consists of two disjoint copies of G with an edge between corresponding vertices in the two copies. Label the vertices of K_2 as 1 and 2, so a vertex of $G \square K_2$ has the form (v, i) where $v \in V(G)$ and $i \in \{1, 2\}$. The following proof is due to Reed and Wood [86].

Lemma 2.24 (Reed and Wood [86])**.** *For every graph G,*

$$\mathsf{glm}(G) \leq \mathsf{had}(G \square K_2).$$

Proof. Let $t := \mathsf{glm}(G)$. It suffices to show there exists a K_t-model in $G \square K_2$. If S is a subgraph of G, define (S, i) to be the subgraph of $G \square K_2$ induced by $\{(v, i) | v \in S\}$. Let H be the intersection graph of a set of paths $\mathcal{P}$ with bipartition A, B, such that H contains a K_t-minor. For each $P \in \mathcal{P}$, let $P' := (P, i)$ where $i = 1$ if $P \in A$, and $i = 2$ if $P \in B$.

If $PQ \in E(H)$, then without loss of generality $P \in A$ and $Q \in B$, and there exists a vertex v such that $v \in V(P) \cap V(Q)$. Then the edge $(v, 1)(v, 2) \in E(G \square K_2)$ has one endpoint in P' and the other in Q'. So $P' \cup Q'$ is connected.

Let $X_1, \dots, X_t$ be the branch sets of a K_t-model in H. Define $X'_i := \bigcup_{P \in X_i} P'$. Now each X'_i is connected. It is sufficient to show, for $i \neq j$, that $V(X'_i \cap X'_j) = \emptyset$ and there exists an edge of $G \square K_2$ with one endpoint in X'_i and the other in X'_j. If there exists $v \in V(X'_i \cap X'_j)$ then there exists P' such that $v \in P'$ and $P' \in X'_i \cap X'_j$. But then $P \in X_i \cap X_j$, which is a contradiction. So $V(X'_i \cap X'_j) = \emptyset$. Also, since $X_1, \dots, X_t$ is a K_t-model of H, there exists some $PQ \in E(H)$ such that $P \in X_i$ and $Q \in X_j$. From above, there exists an edge between P' and Q' in $G \square K_2$, which is sufficient. $\square$

Lemma 2.25. *For every graph G and integer $r \geq 2$,*

$$\mathsf{had}(G \square K_2) \leq 3\,\mathsf{had}_r(G),$$

[†]However, Fox also states the definitions were independently introduced by Seymour.

and if r is even then

$$\mathsf{had}(G\square K_2) \leq 2\,\mathsf{had}_r(G).$$

Proof. Let $k := \mathsf{had}(G\square K_2)$, and let $X_1, \dots, X_k$ be the branch sets of the complete minor in $G\square K_2$. Let X'_j be the induced subgraph of G on the vertex set $\{v|(v,i) \in X_j\}$. Now $\beta = \{X'_1, \dots, X'_k\}$ is a bramble in G, since each X'_j is connected, and since each pair of subgraphs X'_a and X'_b either has an edge between them (corresponding to an edge between X_a and X_b in one copy of G) or a vertex in common (corresponding to an edge between X_a and X_b between the copies of G). Let w weight each element of β by $\lfloor \frac{r}{2} \rfloor / r$. If r is even, then this is $\frac{1}{2}$. If r is odd, then $r \geq 3$ and $\lfloor \frac{r}{2} \rfloor / r \geq \frac{1}{3}$. Since every vertex v of G is in at most two bramble elements (one for $(v,1)$ and one for $(v,2)$), the sum of the weights of the bramble elements containing v is at most 1. Then $h(\beta, w) = k\lfloor \frac{r}{2} \rfloor / r \geq k\frac{1}{3}$, and so $\mathsf{had}_r(G) \geq h(\beta, w) \geq k\frac{1}{3}$, as required. However, if r is even then we can improve the lower bound on $h(\beta, w)$ to $k\frac{1}{2}$ and $\mathsf{had}_r(G) \geq k\frac{1}{2}$, as required. $\square$

Lemma 2.26. *For every graph G,*

$$\mathsf{had}_f(G) \leq \mathsf{bn}(G).$$

Proof. Let β, w be the bramble and weight function in G which maximise $h(\beta, w)$. Then let H be a minimum order hitting set for β. It is sufficient to show that $|H| \geq h(\beta, w)$. Let $s = \sum_{v \in H} \sum_{X \in \beta : v \in X} w(X)$. Since $\sum_{X \in \beta : v \in X} w(X) \leq 1$, for all vertices v, it follows $s \leq |H|$. However, since H is a hitting set, s counts the weight of each bramble element at least once, and thus $s \geq h(\beta, w)$. This proves our result. $\square$

Note Lemma 2.26 is tight; consider $G = K_n$.

Wood [114] proved that $\mathsf{had}(G\square K_2) \leq 2\,\mathsf{tw}(G) + 2$ and Reed and Wood [86] proved that $\mathsf{glm}(G) \leq 2\,\mathsf{tw}(G) + 2$. More precisely, Lemmas 2.24, 2.25 and 2.26 imply that

$$\mathsf{glm}(G) \leq \mathsf{had}(G\square K_2) \leq 2\,\mathsf{had}_2(G) \leq 2\,\mathsf{had}_f(G) \leq 2\,\mathsf{bn}(G) = 2\,\mathsf{tw}(G) + 2.$$

Also, for every integer $r \geq 2$,

$$\mathsf{glm}(G) \leq 3\,\mathsf{had}_r(G) \leq 3\,\mathsf{had}_f(G) \leq 3\,\mathsf{bn}(G) = 3\,\mathsf{tw}(G) + 3.$$

Conversely, Reed and Wood [86] proved that

$$\mathsf{tw}(G) \leq c\,\mathsf{glm}(G)^4 \sqrt{\log \mathsf{glm}(G)}$$

for some constant c. Thus tw, glm, $\mathsf{had}(\,\cdot\,\square K_2)$, had_f and had_r (for each $r \geq 2$) are tied by polynomial functions.

2.12 Fractional Open Problems

Given a graph G define a *b-fold colouring* for G to be an assignment of b colours to each vertex of G such that if two vertices are adjacent, they have no colours assigned in common. We can consider this a generalisation of standard graph colouring, which is equivalent when $b = 1$. A graph G is *$a : b$-colourable* when there is a b-fold colouring of G with a colours in total. Then define the *b-fold chromatic number* $\chi_b(G) := \min\{a | G$ is $a : b$-colourable$\}$. So $\chi_1(G) = \chi(G)$. Then, define the *fractional chromatic number* $\chi_f(G) = \lim_{b\to\infty} \frac{\chi_b(G)}{b}$. See Scheinerman and Ullman [97] for an overview of the topic. Reed and Seymour [83] proved that $\chi_f(G) \leq 2\,\mathsf{had}(G)$. Hence there is a linear relationship between the fractional chromatic number and Hadwiger's number. We have

$$\chi_f(G) \leq \chi(G) \quad \text{and} \quad \mathsf{had}(G) \leq \mathsf{had}_f(G) \leq \mathsf{tw}(G) + 1.$$

Recall Hadwiger's Conjecture asserts that $\chi(G) \leq \mathsf{had}(G)$, thus bridging the gap in the above inequalities. Recall that $\chi(G) \leq \mathsf{tw}(G) + 1$. Thus the following two questions provide interesting weakenings of Hadwiger's Conjecture:

Conjecture 2.27. *For every graph G, $\chi(G) \leq \mathsf{had}_f(G)$.*

Conjecture 2.28. *For every graph G, $\chi_f(G) \leq \mathsf{had}_f(G)$.*

Finally, note that the results of Section 2.11 prove that had_3 is bounded by a polynomial function of had_2. It remains an open question whether $\mathsf{had}_3(G) \leq c\,\mathsf{had}_2(G)$ for some constant c.

Chapter 3

Treewidth of the Line Graph of a Complete Graph

3.1 Introduction

The definition of treewidth was given in Chapter 1. Recall that the pathwidth of a graph G, denoted $\mathsf{pw}(G)$, to be the minimum width of a tree decomposition where the underlying tree is a path. (We call such a tree decomposition a *path decomposition*.) It follows from the definition that $\mathsf{pw}(G) \geq \mathsf{tw}(G)$ for all graphs G.

Also recall the *line graph* $L(G)$ of a graph G is the graph with $V(L(G)) = E(G)$, such that two vertices of $L(G)$ are adjacent when the corresponding edges of G are incident at a vertex.

In this chapter, we determine $\mathsf{tw}(L(K_n))$ exactly. As it turns out, the minimum width tree decomposition that we construct is also a path decomposition. Hence we prove the following result.

Theorem 1.2.

$$\mathsf{tw}(L(K_n)) = \mathsf{pw}(L(K_n)) = \begin{cases} (\frac{n-1}{2})(\frac{n-1}{2}) + n - 2 & \text{, } \textit{if } n \textit{ is odd} \\ (\frac{n-2}{2})(\frac{n}{2}) + n - 2 & \text{, } \textit{if } n \textit{ is even.} \end{cases}$$

Note the following conventions: if S is a subgraph of a graph G and $x \in V(G) - V(S)$, then let $S \cup \{x\}$ denote the subgraph of G with vertex set $V(S) \cup \{x\}$ and edge set $E(S) \cup \{xy : y \in S, xy \in E(G)\}$. Similarly, if $u \in V(S)$, let $S - \{u\}$ denote the subgraph with vertex set $V(S) - \{u\}$ and edge set $E(S) - \{uw : w \in S - \{u\}\}$.

3.2 Line-Brambles and the Treewidth Duality Theorem

Recall the definition of a bramble given in Section 2.3. We employ the following standard approach for determining the treewidth and pathwidth of a particular graph G. First construct a bramble of large order, thus proving a lower bound on $\mathsf{tw}(G)$. Then to prove an upper bound, construct a path decomposition of small width. Given that $\mathsf{tw}(G) \leq \mathsf{pw}(G)$, this is sufficient to prove Theorem 1.2.

In order to construct a bramble of the line graph $L(G)$, define the following:

Definition A *line-bramble* $\mathcal{B}$ of G is a collection of connected subgraphs of G satisfying the following properties:

- For all $X \in \mathcal{B}$, $|V(X)| \geq 2$.
- For all $X, Y \in \mathcal{B}$, $V(X) \cap V(Y) \neq \emptyset$.

Define a *hitting set* for a line-bramble $\mathcal{B}$ to be a set of edges $H \subseteq E(G)$ that intersects each $X \in \mathcal{B}$. Then define the *order* of $\mathcal{B}$ to be the size of the minimum size hitting set H of $\mathcal{B}$. (Recall we often refer to such a hitting set simply as a *minimum hitting set* of $\mathcal{B}$.)

Lemma 3.1. *Given a line-bramble $\mathcal{B}$ of a graph G, there is a bramble $\mathcal{B}'$ of $L(G)$ of the same order.*

Proof. Let X be an element of line-bramble $\mathcal{B}$ and let $\mathcal{B}' = \{E(X) \subseteq L(G) | X \in \mathcal{B}\}$ (here we interpret $E(X)$ as an induced subgraph of $L(G)$). Since X is connected and $|V(X)| \geq 2$, the subgraph X contains an edge. So $E(X)$ induces a non-empty connected subgraph of $L(G)$. Consider $E(X)$ and $E(Y)$ in $\mathcal{B}'$. Thus $V(X) \cap V(Y) \neq \emptyset$. Let v be a vertex in $V(X) \cap V(Y)$. Then there exists some $xv \in E(X)$ and $vy \in E(Y)$, and thus in $L(G)$ there is an edge between the vertex xv and the vertex vy. Hence $E(X)$ and $E(Y)$ touch, and so $\mathcal{B}'$ is a bramble of $L(G)$. All that remains is to ensure $\mathcal{B}$ and $\mathcal{B}'$ have the same order. If H is a minimum hitting set for $\mathcal{B}$, then H is also a set of vertices in $L(G)$ that intersects a vertex in each $E(X) \in \mathcal{B}'$. So H is a hitting set for $\mathcal{B}'$ of the same size. Conversely, if H' is a minimum hitting set of $\mathcal{B}'$, then H' is a set of edges in G that contains an edge in each $X \in \mathcal{B}$. So H' is a hitting set for $\mathcal{B}$. Thus the orders of $\mathcal{B}$ and $\mathcal{B}'$ are equal. $\square$

Hence, in order to determine a lower bound on the bramble number $\mathsf{bn}(L(G))$, it is sufficient to construct a line-bramble of G of large order. We will now define a particular line-bramble for any graph G with $|V(G)| \geq 3$.

Definition Given a graph G and a vertex $v \in V(G)$, the *canonical line-bramble for v* of G is the set of connected subgraphs X of G such that either $|V(X)| > \frac{|V(G)|}{2}$, or

$|V(X)| = \frac{|V(G)|}{2}$ and X contains v. Note that if $|V(G)|$ is odd, then no elements of the second type occur.

Lemma 3.2. *For every graph G with $|V(G)| \geq 3$ and for all $v \in V(G)$, the canonical line-bramble for v, denoted $\mathcal{B}$, is a line-bramble of G.*

Proof. By definition, each element of $\mathcal{B}$ is a connected subgraph. Since $|V(G)| \geq 3$, each element of $\mathcal{B}$ contains at least two vertices. All that remains to show is that each pair of subgraphs X,Y in $\mathcal{B}$ intersect in at least one vertex. If $|V(X)| = |V(Y)| = \frac{|V(G)|}{2}$, then X and Y intersect at v. Otherwise, without loss of generality, $|V(X)| > \frac{|V(G)|}{2}$ and $|V(Y)| \geq \frac{|V(G)|}{2}$. If $V(X) \cap V(Y) = \emptyset$, then $|V(X) \cup V(Y)| = |V(X)| + |V(Y)| > |V(G)|$, which is a contradiction. $\square$

Let $v \in V(G)$ be an arbitrary vertex and let H be a minimum hitting set of $\mathcal{B}$, the canonical line-bramble for v. Consider the graph $G - H$. Since H is a set of edges, the set $V(G - H) = V(G)$. Then each component of $G - H$ contains at most $\frac{|V(G)|}{2}$ vertices, otherwise some component of $G - H$ contains an element of $\mathcal{B}$ that does not contain an edge of H. Similarly, if a component contains $\frac{|V(G)|}{2}$ vertices, it cannot contain the vertex v. Thus, our hitting set H must be large enough to separate G into such components. The next lemma follows directly:

Lemma 3.3. *Let G be a graph with $|V(G)| \geq 3$, let v be a vertex of G, and let $\mathcal{B}$ be the canonical line-bramble for v. Then $H \subseteq E(G)$ is a hitting set of $\mathcal{B}$ if and only if every component of $G - H$ contains at most $\frac{|V(G)|}{2}$ vertices, and v is not in a component that contains exactly $\frac{|V(G)|}{2}$ vertices.*

Note the similarity between this characterisation and the *bisection width* of a graph (see Díaz *et al.* [21], for example), which is the minimum number of edges between any $A, B \subset V(G)$ where $A \cap B = \emptyset$ and $|A| = \lfloor \frac{|V(G)|}{2} \rfloor$ and $|B| = \lceil \frac{|V(G)|}{2} \rceil$. (Later we show that most of our components have maximum or almost maximum allowable order.) From the above lemma we prove the following, which is slightly more useful in practice.

Lemma 3.4. *Let G be a graph with $|V(G)| \geq 3$, let v be a vertex of G, let $\mathcal{B}$ be the canonical line-bramble for v, and let H be a hitting set of $\mathcal{B}$. Then $G - H$ contains at least three components.*

Proof. By Lemma 3.3, we have an upper bound on the order of the components of $G - H$. If $G - H$ contains only one component, clearly this component is too large. If $G - H$ contains only two components and n is odd, then one of the components must contain more than $\frac{n}{2}$ vertices. If $G - H$ contains only two components and n is even, it is possible that both components contain exactly $\frac{n}{2}$ vertices, however one of these components must contain v. Thus $G - H$ contains at least three components. $\square$

Given that the components of $G - H$ are what is really important, we can also prove the following lemma.

Lemma 3.5. *Let G be a graph with $|V(G)| \geq 3$, let v be a vertex of G, and let $\mathcal{B}$ be the canonical line-bramble for v. If H is a minimum hitting set for $\mathcal{B}$, then no edge of H has both endpoints in the same component of $G - H$.*

Proof. For the sake of a contradiction assume that both endpoints of an edge $e \in H$ are in the same component of $G - H$. Then consider the set $H - e$. By Lemma 3.3, $H - e$ is a hitting set of $\mathcal{B}$, since the vertex sets of the components of $G - H$ have not changed. But $H - e$ is smaller than the minimum hitting set H, a contradiction. $\square$

3.3 Proof of Result

Let $G := K_n$. When $n \leq 2$, Theorem 1.2 holds trivially, so assume $n \geq 3$. Firstly, we determine a lower bound on the treewidth by considering a canonical line-bramble for v, denoted $\mathcal{B}$. Given that K_n is regular, it suffices to choose a vertex v of K_n arbitrarily.

If H is a minimum hitting set of a canonical line-bramble $\mathcal{B}$, label the components of $G - H$ as $Q_1, \ldots, Q_p$ such that $|V(Q_1)| \geq |V(Q_2)| \geq \cdots \geq |V(Q_p)|$. We refer to this as labelling the components *descendingly*.

Consider a pair of components (Q_i, Q_j) where $i < j$ and the components are labelled descendingly. Call this a *good pair* if one of the following conditions hold:

1. $|V(Q_i)| < \frac{n}{2} - 1$,
2. n is even, $|V(Q_i)| = \frac{n}{2} - 1$, $V(Q_j) \neq \{v\}$, and $v \notin V(Q_i)$.

Lemma 3.6. *Let G be a complete graph with $n \geq 3$ vertices, let v be a vertex of G, let $\mathcal{B}$ be the canonical line-bramble for v, and let H be a minimum hitting set of $\mathcal{B}$. If $Q_1, \ldots, Q_p$ are the components of $G - H$ labelled descendingly, then $Q_1, \ldots, Q_p$ does not contain a good pair.*

Proof. Say (Q_i, Q_j) is a good pair. Let x be a vertex of Q_j, such that if (Q_i, Q_j) is of the second type, then $x \neq v$. Let H' be the set of edges obtained from H by removing the edges from x to Q_i and adding the edges from x to Q_j. We add all these edges by Lemma 3.5. Then the components for $G - H'$ are $Q_1, \ldots, Q_{i-1}, Q_i \cup \{x\}, Q_{i+1}, \ldots, Q_{j-1}, Q_j - \{x\}, Q_{j+1}, \ldots Q_p$. By Lemma 3.3, to ensure H' is a hitting set, it suffices to ensure that $V(Q_i) \cup \{x\}$ is sufficiently small, since all other components are the same as in $G - H$, or smaller. If (Q_i, Q_j) is of the first type, then $|V(Q_i) \cup \{x\}| = |V(Q_i)| + 1 < \frac{n}{2}$. If (Q_i, Q_j) is of the second type, $|V(Q_i) \cup \{x\}| = \frac{n}{2}$, but it does not contain v. Thus, by Lemma 3.3, H' is a hitting set. However, $|H'| = |H| - |V(Q_i)| + |V(Q_j)| - 1 \leq |H| - 1$, which contradicts that H is a minimum hitting set. □

Lemma 3.7. *Let $G, v, \mathcal{B}$ and H be as in Lemma 3.6. Then $G - H$ contains exactly three components.*

Proof. By Lemma 3.4, $G - H$ contains at least three components. Now, assume $G - H$ contains at least four components and label the components of $G - H$ descendingly. We show that these components contain a good pair, contradicting Lemma 3.6.

If n is odd, there is a good pair of the first type when any two components contain less than $\frac{n-1}{2}$ vertices. Thus at least three components have order at least $\frac{n-1}{2}$. Then $|V(G)| \geq 3(\frac{n-1}{2}) + 1 > n$ when $n \geq 2$, which is a contradiction.

If n is even, there is a good pair of the first type when any two components contain less than $\frac{n}{2} - 1$ vertices. Similarly to the previous case, $|V(G)| \geq 3(\frac{n}{2} - 1) + 1 > n$, again a contradiction when $n > 4$. If $n = 4$ then each component is a single vertex. Take Q_i, Q_j to be two of these components, neither of which contain the vertex v. Then (Q_i, Q_j) is a good pair of the second type. Hence $G - H$ does not contain more than three components, and as such it contains exactly three components. □

Lemma 3.8. *Let $G, v, \mathcal{B}$ and H be as in Lemma 3.6, and let the components of $G - H$ be labelled descendingly. If n is odd then $|V(Q_1)| = |V(Q_2)| = \frac{n-1}{2}$ and $|V(Q_3)| = 1$. If n is even then $|V(Q_1)| = \frac{n}{2}, |V(Q_2)| = \frac{n}{2} - 1$ and $|V(Q_3)| = 1$.*

Proof. Lemma 3.7 shows that $G - H$ contains exactly three components. By Lemma 3.6, (Q_2, Q_3) is not a good pair. Hence $|V(Q_1)| \geq |V(Q_2)| \geq \frac{n-1}{2}$ when n is odd, and $|V(Q_1)| \geq |V(Q_2)| \geq \frac{n}{2} - 1$ when n is even, or else there is a good pair of the first type. When n is odd, it follows from Lemma 3.3 that $|V(Q_1)| = |V(Q_2)| = \frac{n-1}{2}$, and so $|V(Q_3)| = 1$. When n is even, however, $\frac{n}{2} - 1 \leq |V(Q_1)|, |V(Q_2)| \leq \frac{n}{2}$. Since Q_3 is not empty, it follows that $|V(Q_3)| = 1$ or 2. If $|V(Q_3)| = 1$, then $|V(Q_1)| = \frac{n}{2}, |V(Q_2)| = \frac{n}{2} - 1$ and $|V(Q_3)| = 1$, as required. Otherwise, $|V(Q_1)|, |V(Q_2)| = \frac{n}{2} - 1$ and $|V(Q_3)| = 2$. But then at least one of Q_1, Q_2 does not contain v, and $V(Q_3) \neq \{v\}$. Thus either (Q_1, Q_3) or (Q_2, Q_3) is a good pair of the second type, contradicting Lemma 3.6. □

Lemma 3.9. *Let $G, v, \mathcal{B}$ and H be as in Lemma 3.6. Then $|H| = (\frac{n-1}{2})(\frac{n-1}{2}) + (n - 1)$ when n is odd, and $|H| = (\frac{n-2}{2})(\frac{n}{2}) + (n - 1)$ when n is even.*

Proof. From Lemma 3.8 we know the order of the components of $G - H$. By Lemma 3.5, H contains exactly every edge between each pair of components, and since G is complete there is an edge for each pair of vertices. From this it is easy to calculate $|H|$. □

Lemma 3.9 and the Treewidth Duality Theorem imply:

Corollary 3.10. *Let G be a complete graph with $n \geq 3$ vertices. Then*

$$\mathsf{pw}(L(K_n)) \geq \mathsf{tw}(L(K_n)) = \mathsf{bn}(L(K_n)) - 1 \geq \begin{cases} (\frac{n-1}{2})(\frac{n-1}{2}) + (n-2) & \text{, if } n \text{ is odd} \\ (\frac{n-2}{2})(\frac{n}{2}) + (n-2) & \text{, if } n \text{ is even.} \end{cases}$$

Now, to obtain an upper bound on $\mathsf{pw}(L(G))$, construct a path decomposition of $L(G)$. First, label the vertices of G by $1, \dots, n$. Let T be an n-node path, also labelled by $1, \dots, n$. The bag A_i for the node labelled i, is defined such that $A_i = \{ij : j \in V(G)\} \cup \{uw : u < i < w\}$. For a given A_i, call the edges of $\{ij : j \in V(G)\}$ *initial edges* and call the edges of $\{uw : u < i < w\}$ *crossover edges.* (Note here these edges of G are really acting as vertices of $L(G)$, but refer to them as edges for simplicity.)

Lemma 3.11. *Let G be a complete graph with $n \geq 3$ vertices. Then $(T, \{A_1, \dots, A_n\})$ is a path decomposition for $L(G)$ of width $(\frac{n-1}{2})(\frac{n-1}{2}) + (n-2)$ if n is odd, or $(\frac{n-2}{2})(\frac{n}{2}) + (n-2)$ if n is even.*

Proof. Each edge uw of G appears in A_u and A_w as an initial edge. Observe that uw is in A_i if and only if $u \leq i \leq w$, so the nodes indexing the bags containing uw form a connected subtree of T. Finally, all of the edges incident at the vertex u appear in A_u, and the same holds for w, so if two edges of G are adjacent in $L(G)$, they share a bag.

Now determine the size of A_i. The bag A_i contains $n-1$ initial edges and $(i-1)(n-i)$ crossover edges. So $|A_i| = (n-1) + (i-1)(n-i)$. This is maximised when $i = \frac{n+1}{2}$ if n is odd, and when $i = \frac{n}{2}$ or $\frac{n+2}{2}$ if n is even. From this it is possible to calculate the largest bag size, and hence the width of T. □

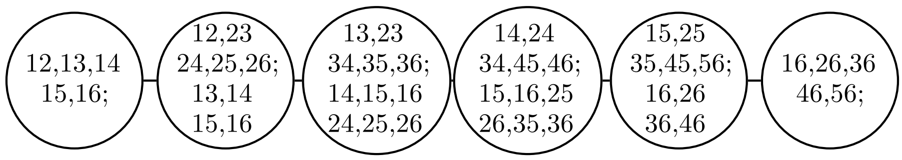

Figure 3.1: The described path decomposition for $L(K_6)$.

Lemma 3.11 gives an upper bound on $\mathsf{pw}(L(K_n))$ and also on $\mathsf{tw}(L(K_n))$. This, combined with the lower bound in Corollary 3.10, completes the proof of Theorem 1.2.

Chapter 4

Treewidth of the Line Graph of a Complete Multipartite Graph

4.1 Introduction

Recall the line graph $L(G)$ of a graph G is the graph with $V(L(G)) = E(G)$, such that two vertices of $L(G)$ are adjacent when the corresponding edges of G are incident at a vertex.

A complete multipartite graph $K_{n_1,n_2,\ldots,n_k}$ is a graph with k colour classes, of order $n_1, \ldots, n_k$ respectively, containing an edge between every pair of differently coloured vertices. A complete graph K_n is a complete multipartite graph with n colour classes each containing a single vertex. By extending the techniques from Chapter 3, we determine bounds on the treewidth of the line graph of a complete multipartite graph. Given that our constructed tree decomposition is once again a path decomposition, we get the following result.

Theorem 1.3. *If $k \geq 2$ and $n = |V(K_{n_1,\ldots,n_k})|$, then*

$$\frac{1}{2}\left(\sum_{1 \leq i < j \leq k} n_i n_j\right) - n(k-1) + \frac{3}{4}k(k-1) - 1$$

$$\leq \mathsf{tw}(K_{n_1,\ldots,n_k}) \leq \mathsf{pw}(K_{n_1,\ldots,n_k}) \leq$$

$$\frac{1}{2}\left(\sum_{1 \leq i < j \leq k} n_i n_j\right) + \frac{1}{2}n(k+5) + \frac{1}{4}k(k-1) - 4.$$

Theorem 1.3 implies that when $n_1 = \cdots = n_k = c$, (that is, when our complete multipartite graph is regular) then $\mathsf{tw}(L(K_{c,\ldots,c})) \approx \frac{k^2c^2}{4}$ (ignoring the lower order terms).

We improve this result, obtaining an exact answer for the treewidth of the line graph of a regular complete multipartite graph.

Theorem 1.4. *If* $k \geq 2$ *and* $n_1 = n_2 = \cdots = n_k = c \geq 1$, *then*

$$\mathsf{tw}(L(K_{n_1,\ldots,n_k})) = \mathsf{pw}(L(K_{n_1,\ldots,n_k})) = \begin{cases} \frac{c^2k^2}{4} - \frac{c^2k}{4} + \frac{ck}{2} - \frac{c}{2} + \frac{k}{4} - \frac{5}{4}\,, & \textit{if } k \textit{ odd and } c \textit{ odd} \\ \frac{c^2k^2}{4} - \frac{c^2k}{4} + \frac{ck}{2} - \frac{c}{2} - 1 & , \textit{if } c \textit{ even} \\ \frac{c^2k^2}{4} - \frac{c^2k}{4} + \frac{ck}{2} - \frac{c}{2} + \frac{k}{4} - \frac{3}{2}\,, & \textit{if } k \textit{ even and } c \textit{ odd.} \end{cases}$$

Note that this implies Theorem 1.2 is a special case of Theorem 1.4.

Recall the following conventions: if S is a subgraph of a graph G and $x \in V(G)-V(S)$, then let $S \cup \{x\}$ denote the subgraph of G with vertex set $V(S) \cup \{x\}$ and edge set $E(S) \cup \{xy : y \in S, xy \in E(G)\}$. Similarly, if $u \in V(S)$, let $S - \{u\}$ denote the subgraph with vertex set $V(S) - \{u\}$ and edge set $E(S) - \{uw : w \in S - \{u\}\}$.

4.2 Line-Brambles of a Complete Multipartite Graph

Recall the definitions and results of Section 3.2, as they form the basis of this section. Let $n := |V(G)| = n_1 + \cdots + n_k$. If $n = k$, then $G = K_n$ and Theorem 1.2 determines $\mathsf{tw}(G)$ exactly, so we may assume $n > k$. As stated in Theorem 1.3, we assume that $k \geq 2$, since if $k = 1$, then $L(G)$ is the null graph. Let X_i be the i^{th} colour class of G, with order n_i. Call X_i *odd* or *even* depending on the parity of $|X_i|$.

As in Chapter 3, consider a canonical line-bramble for v denoted $\mathcal{B}$. However, we shall choose vertex v from a colour class of largest order. Note that such a vertex has minimum degree. Let H be a hitting set of $\mathcal{B}$, and label the components of $G - H$ by $Q_1, \ldots, Q_p$. Denote H and the labelling of its components together as $(H, (Q_1, \ldots, Q_p))$. Choose $(H, (Q_1, \ldots, Q_p))$ such that the following conditions hold, in order of preference:

(0) $|H|$ is minimised,
(1) $|V(Q_1)|$ is maximised,
(2) $|V(Q_2)|$ is maximised,
⋮
(p) $|V(Q_p)|$ is maximised,
(p+1) v is in the component of highest possible index.

By condition (0), H is a minimum hitting set. Note, as a result of this that $|V(Q_1)| \geq |V(Q_2)| \geq \cdots \geq |V(Q_p)|$, otherwise we can keep H and easily find a better choice of labelling. Call a choice of $(H, (Q_1, \ldots, Q_p))$ that satisfies these conditions a *good labelling*.

Consider a pair of components (Q_i, Q_j) where $i < j$ and $Q_1, \dots, Q_p$ is from a good labelling. We call this a *good pair* when for all $x \in Q_j$ there exists $y \in Q_i$ such that xy is an edge, and one of the following holds:

1. $|V(Q_i)| < \frac{n}{2} - 1$,
2. n is even, $|V(Q_i)| = \frac{n}{2} - 1$, $v \notin V(Q_i)$ and $V(Q_j) \cap X_s \neq \{v\}$ for all colour classes X_s.

Note this is very similar to the definition of a good pair in Chapter 3. However, when considering the complete multipartite graph we also need to consider the interaction between the colour classes and the hitting set H, hence the slightly more complex definition above. The following lemma is essentially a more complex version of Lemma 3.6, adjusted for our updated definition of a good pair.

Lemma 4.1. *Let G be a complete multipartite graph $G := K_{n_1,\dots,n_k}$ such that $k \geq 2$ and $n > k$, let v be a vertex of G chosen from a largest colour class, let $\mathcal{B}$ be a canonical line-bramble for v, and let $(H, (Q_1, \dots, Q_p))$ be a good labelling. Then $Q_1, \dots, Q_p$ does not contain a good pair.*

Proof. Assume (Q_i, Q_j) is a good pair. For each X_s that intersects Q_j, let x_s be some vertex of $Q_j \cap X_s$. If (Q_i, Q_j) is of the second type, choose each $x_s \neq v$. Let H_s be the set of edges created by taking H and removing the edges from x_s to Q_i, then adding the edges from x_s to $(Q_j - X_s)$. Thus we have removed $|V(Q_i)| - |V(Q_i) \cap X_s|$ edges and have added $|V(Q_j)| - |V(Q_j) \cap X_s|$ (by Lemma 3.5).

Suppose that $|V(Q_j)| - |V(Q_j) \cap X_s| > |V(Q_i)| - |V(Q_i) \cap X_s|$ for each X_s that intersects Q_j. Then

$$\sum_{s: X_s \cap V(Q_j) \neq \emptyset} |V(Q_j)| - |V(Q_j) \cap X_s| > \sum_{s: X_s \cap V(Q_j) \neq \emptyset} |V(Q_i)| - |V(Q_i) \cap X_s|.$$

However, since we are cycling through all colour classes that intersect Q_j,

$$\sum_{s: X_s \cap V(Q_j) \neq \emptyset} |V(Q_j) \cap X_s| = |V(Q_j)|.$$

If there are r such colour classes, then

$$(r-1)|V(Q_j)| > r|V(Q_i)| - \sum_{s: X_s \cap V(Q_j) \neq \emptyset} |V(Q_i) \cap X_s| \geq (r-1)|V(Q_i)|.$$

This implies $|V(Q_j)| > |V(Q_i)|$, which is a contradiction of condition (i). Hence, for some s, $|V(Q_j)| - |V(Q_j) \cap X_s| \leq |V(Q_i)| - |V(Q_i) \cap X_s|$. Fix such an s.

A component of $G - H_s$ is either one of $Q_1, \dots, Q_{i-1}, Q_{i+1}, \dots, Q_{j-1}, Q_{j+1}, \dots, Q_p$, or $Q_i \cup \{x_s\}$ (which is connected since x_s has a neighbour in Q_i), or strictly contained within Q_j. Since H is a hitting set, to prove H_s is a hitting set it suffices to show that $Q_i \cup \{x_s\}$ is sufficiently small, by Lemma 3.3. If (Q_i, Q_j) is of the first type, then $|V(Q_i) \cup \{x_s\}| = |V(Q_i)| + 1 < \frac{n}{2}$. So $V(Q_i) \cup \{x_s\}$ is sufficiently small. If (Q_i, Q_j) is of the second type, $|V(Q_i) \cup \{x_s\}| = \frac{n}{2}$, but it does not contain v. Thus H_s is a hitting set. However, $|H_s| = |H| - (|V(Q_i)| - |V(Q_i) \cap X_s|) + (|V(Q_j)| - |V(Q_j) \cap X_s|) \leq |H|$. If $|H_s| < |H|$, then condition (0) is contradicted. If $|H_s| = |H|$, since $|V(Q_i) \cup \{x_s\}| > |V(Q_i)|$ and only components of higher index have become smaller, H_s is a better choice of minimum hitting set by condition (i), which is a contradiction. □

If G is a star $K_{1,n-1}$, then $L(G) \cong K_{n-1}$ and $\mathsf{tw}(L(G)) = n - 2$, which satisfies Theorem 1.3. For technical reasons, most of our following results shall assume that G is not a star. (The underlying reason for this is that Theorem 1.3 is a more useful result when the colour classes are close to being the same size, and less useful when they more irregular. The star is the most irregular complete bipartite graph.) If $(H, (Q_1, \dots, Q_p))$ is a good labelling where $p \geq 4$ and $Q_2, \dots, Q_p$ are all singleton sets and contained within one colour class, then say that $(H, (Q_1, \dots, Q_p))$ is a *rare configuration*.

Lemma 4.2. *Let G be a complete multipartite graph $G := K_{n_1,\dots,n_k}$ such that G is not a star and such that $k \geq 2$ and $n > k$. Also let v be a vertex of G chosen from a largest colour class, let $\mathcal{B}$ be a canonical line-bramble for v, and let $(H, (Q_1, \dots, Q_p))$ be a good labelling. Then $(H, (Q_1, \dots, Q_p))$ is not a rare configuration.*

Proof. Assume G is a rare configuration, but G is not a star. Let X_s be the colour class of $Q_2, \dots, Q_p$. Since $p \geq 4$, we may choose $j \in \{2, \dots, p\}$ such that $V(Q_j) \neq \{v\}$.

Suppose that one of the following conditions hold:

- $|V(Q_1)| < \frac{n}{2} - 1$,
- n is even, $|V(Q_1)| = \frac{n}{2} - 1$ and $v \notin V(Q_1)$.

The component Q_1 must contain at least two vertices not in X_s since G is not a star or an independent set (since $k \geq 2$). So for each $x \in V(Q_2) \cup \dots \cup V(Q_p)$, there is some $y \in V(Q_1)$ such that $y \notin X_s$, so the edge xy exists. Then (Q_1, Q_j) is a good pair, which contradicts Lemma 4.1. Thus by Lemma 3.3,

$$|Q_1| = \begin{cases} \frac{n-1}{2} & \text{, if } n \text{ is odd} \\ \frac{n}{2} - 1 & \text{, if } n \text{ is even and } v \in V(Q_1) \\ \frac{n}{2} & \text{, if } n \text{ is even and } v \notin V(Q_1) \end{cases}$$

Since at least two vertices of Q_1 are not in X_s, we may choose $y \in (V(Q_1) - \{v\}) - X_s$. Say $y \in X_t$. We can assume that $v \in V(Q_1)$ or $v \in V(Q_p)$, since if $v \in V(Q_2) \cup \ldots V(Q_{p-1})$, then we can relabel the components $Q_2, \ldots, Q_p$ to obtain a choice of $(H, (Q_1, \ldots, Q_p))$ which is better with regards to the condition (p+1). Thus let $z \in V(Q_2)$, and so $z \neq v$ since $p \geq 4$. Let H' be the set of edges created by taking H and removing the edges from y to $Q_3 \cup \cdots \cup Q_{p-1}$, adding the edges from y to $Q_1 - X_t$, and removing the edges from z to $Q_1 - \{y\}$. Then the three components of $G - H'$ are $Q_1 \cup \{z\} - \{y\}$, $\{y\} \cup Q_3 \cup \cdots \cup Q_{p-1}$, Q_p. The component $Q_1 \cup \{z\} - \{y\}$ is connected since $Q_1 - \{y\}$ contains a vertex not in X_s and $z \in X_s$. Similarly, $\{y\} \cup Q_3 \cup \cdots \cup Q_{p-1}$ is connected since $y \in X_t$ and all vertices of $Q_3 \cup \cdots \cup Q_{p-1}$ are in X_s.

By Lemma 3.3, to show H' is a hitting set, it is sufficient to show that no component of $G - H'$ is too large. Since $|V(Q_1 \cup \{z\} - \{y\})| = |V(Q_1)|$ and $v \neq z$ and H is a hitting set, $Q_1 \cup \{z\} - \{y\}$ is sufficiently small. Similarly Q_p is sufficiently small. However, $|V(\{y\} \cup Q_3 \cup \cdots \cup Q_{p-1})| = p - 2$. Since $|V(Q_1)| + \cdots + |V(Q_p)| = n$, it follows that $p - 2 = n - |V(Q_1)| - 1$. In order to show this is sufficiently small, we need to consider the parity of n, which we consider below. Also note,

$$|H'| = |H| - (p - 3) + (|V(Q_1)| - |V(Q_1) \cap X_t|) - (|V(Q_1)| - 1 - |V(Q_1) \cap X_s|).$$

Since $|V(Q_1) \cap X_t| \geq 1$ and $|V(Q_1) \cap X_s| \leq |V(Q_1)| - 2$, we have $|H'| \leq |H| - (p - 1) + |V(Q_1)| = |H| + 2|V(Q_1)| - n$. This also depends on the parity of n. Now we consider these separate cases to check the order of $\{y\} \cup Q_3 \cup \cdots \cup Q_{p-1}$ and $|H'|$.

Firstly, say n is odd. In this case $|V(Q_1)| = \frac{n-1}{2}$, so then $|V(\{y\} \cup Q_3 \cup \cdots \cup Q_{p-1})| = p - 2 = n - \frac{n-1}{2} - 1 = \frac{n-1}{2}$, and so $\{y\} \cup Q_3 \cup \cdots \cup Q_{p-1}$ is sufficiently small, and H' is a hitting set. Also, $|H'| \leq |H| + 2(\frac{n-1}{2}) - n < |H|$, which contradicts condition (0). Secondly, say n is even and $v \in V(Q_1)$. Then $|V(Q_1)| = \frac{n}{2} - 1$, implying $p - 2 = \frac{n}{2}$, and $|H'| \leq |H| - 2$. This contradicts condition (0). Finally, say n is even and $v \notin V(Q_1)$. Then $|V(Q_1)| = \frac{n}{2}$ and $v \in V(Q_p)$. Then $p - 2 = \frac{n}{2} - 1$, and $|H'| \leq |H| + 2(\frac{n}{2}) - n = |H|$. However, note that the order of the second largest component of $G - H'$ is $p - 2 = \frac{n}{2} - 1$, whereas for $G - H$ the order of the second largest component is 1. Since G is a rare configuration but not a star, $n \geq 5$, since $|V(Q_1)| \geq 2$ and $p \geq 4$, implying $\frac{n}{2} - 1 > 1$. Thus H' is a better choice of minimum hitting set, by condition (2).

Thus, in either case, if G is not a star, but is a rare configuration, then there is a contradiction to one of our conditions on $(H, (Q_1, \ldots, Q_p))$. □

The following lemma is similar to Lemma 3.7. The main difficulty here a rare configuration, for which we proved Lemma 4.2 in order to avoid.

Lemma 4.3. *Let $G, v, \mathcal{B}$ and $(H, (Q_1, \ldots, Q_p))$ be as in Lemma 4.2. Then $G - H$ contains exactly three components.*

Proof. $G - H$ contains at least three components, by Lemma 3.4. Assume for the sake of a contradiction that $G - H$ contains $p > 3$ components. Since $p \geq 4$, if all components but Q_1 are singleton sets in the one colour class, then we have a rare configuration. By Lemma 4.2, this cannot occur. Thus either Q_2 is not a singleton set, or $Q_2, \ldots, Q_p$ are not all in one colour class. Consider a pair (Q_i, Q_j), where $i \in \{1, 2\}$, $i < j$ and if $|V(Q_i)| = 1$ then Q_j and Q_i are not in the same colour class. We can find such a pair for $i = 1$ and for $i = 2$ since this is not a rare configuration. In either case, for all $x \in V(Q_j)$ there exists a $y \in V(Q_i)$ such that xy is an edge, since there is always some $y \in V(Q_i)$ of a different colour class to x. Since (Q_i, Q_j) is not a good pair by Lemma 4.1, we know $|V(Q_i)|$ is too large. In particular, if n is odd, $|V(Q_1)| = |V(Q_2)| = \frac{n-1}{2}$. However, since each component must contain a vertex and $p \geq 4$, the sum of the orders of the components is at least $2(\frac{n-1}{2}) + 2 > n$, which is a contradiction. If n is even and v is in neither Q_1 nor Q_2, then $|V(Q_1)| = |V(Q_2)| = \frac{n}{2}$, which again means the sum of the orders of the components is too large. Finally, if n is even and without loss of generality $v \in V(Q_2)$, then $|V(Q_1)| = \frac{n}{2}$ and $|V(Q_2)| = \frac{n}{2} - 1$, which still gives a contradiction on the orders of the components. Hence $G - H$ contains exactly three components. □

Again, the following lemma is similar to Lemma 3.8.

Lemma 4.4. *Let $G, v, \mathcal{B}$ and $(H, (Q_1, \ldots, Q_p))$ be as in Lemma 4.2. If n is odd, then $p = 3$, $|V(Q_1)| = |V(Q_2)| = \frac{n-1}{2}$ and $|V(Q_3)| = 1$. If n is even, then $p = 3$, $|V(Q_1)| = \frac{n}{2}, |V(Q_2)| = \frac{n}{2} - 1$ and $|V(Q_3)| = 1$.*

Proof. Lemma 4.3 shows that $G - H$ contains exactly three components. Recall that in a good labelling that $|V(Q_1)| \geq |V(Q_2)| \geq |V(Q_3)|$. If $|V(Q_1)| = 1$, then $n = 3$, and since $\frac{n-1}{2} = 1$, then our statement holds in this case. Thus we can assume $n \geq 4$ and $|V(Q_1)| \geq 2$. Hence (Q_1, Q_j) is a good pair for $j > 1$ unless Q_1 is too large. If n is odd, then $|V(Q_1)| = \frac{n-1}{2}$. If $|V(Q_2)| = 1$, $\frac{n-1}{2} + 1 + 1 = n$, implying $n = 3$. So $|V(Q_2)| \geq 2$, and (Q_2, Q_3) is a good pair unless $|V(Q_2)| = \frac{n-1}{2}$, in which case $|V(Q_3)| = 1$.

If n is even and $v \in V(Q_1)$, then $|V(Q_1)| = \frac{n}{2} - 1$. Again, if $|V(Q_2)| = 1$ then $\frac{n}{2} - 1 + 1 + 1 = n$, implying $n = 2$. So $|V(Q_2)| \geq 2$, and (Q_2, Q_3) is a good pair unless $|V(Q_2)| = \frac{n}{2}$, implying $|V(Q_3)| = 1$. (Note here we'd need to relabel the components so they are in descending order of size.) Finally, if n is even and $v \notin V(Q_1)$, then $|V(Q_1)| = \frac{n}{2}$. If $|V(Q_2)| = 1$, then $\frac{n}{2} + 1 + 1 = n$, implying $n = 4$. However, then $|V(Q_3)| = 1$ and our statement holds. If $n \geq 5$, then $|V(Q_2)| \geq \frac{n}{2} - 1$ else (Q_2, Q_3) is a good pair. Since we must have three components, $|V(Q_2)| = \frac{n}{2} - 1$ and $|V(Q_3)| = 1$. Either way, our components have the desired size. $\square$

Lemma 4.5. *Let $G, v, \mathcal{B}$ and $(H, (Q_1, \dots, Q_p))$ be as in Lemma 4.2. If $v \notin Q_3$, then the vertex in Q_3 is in a different colour class to v.*

Proof. By Lemma 4.4, $|V(Q_3)| = 1$. Let x be the vertex in Q_3. Assume for the sake of contradiction that x, v are in colour class X_s. If n is odd then $v \in V(Q_1)$ or $V(Q_2)$, but these components have the same order, by Lemma 4.4. If n is even, $v \in V(Q_2)$, since otherwise v is in a component of order $\frac{n}{2}$, again by Lemma 4.4. So without loss of generality, $v \in V(Q_2)$. Define the hitting set H' as follows: from H, add all the edges from v to $Q_2 - X_s$ (by Lemma 3.5), and then remove the edges from x to $Q_2 - X_s$. Since $xv \notin E(G)$, the components of $G - H'$ are $Q_1, (Q_2 - \{v\}) \cup \{x\}$ and $\{v\}$ (since x, v are in the same colour class, $(Q_2 - \{v\}) \cup \{x\}$ is connected). The orders of the components have not changed, and v has not been placed into a component of order $\frac{n}{2}$, so this is a hitting set by Lemma 3.3. Note $|H'| = |H| + (|V(Q_2)| - |V(Q_2) \cap X_s|) - (|V(Q_2)| - (|V(Q_2) \cap X_s|)) = |H|$, so the order of the hitting set H has not changed. Since v is now in a component of higher index, this contradicts condition (p+1). $\square$

The previous lemmas give a good idea of the structure of the components of $G - H$. When dealing with a complete graph, this was sufficient. However, in the case of a complete multipartite graph, we also need to know how the components of $G - H$ interact with the colour classes of G. As we might expect, in the optimal case, each colour class is essentially split evenly across the two large components Q_1 and Q_2. In order to show this, however, we need to be careful about the parity of n and the parities of $n_1, \dots, n_k$. (Obviously, an odd number of vertices cannot be broken into two equal sized parts.) Recall that we label the colour classes $X_1, \dots, X_n$. For the following section, we assume that G is a complete multipartite graph such that $k \geq 2$ and G is not a star, and as such we have only three components by Lemma 4.3.

Definition Let $X_i^* := X_i \cap (V(Q_1) \cup V(Q_2))$, and say X_i^* is *even* or *odd* depending on the parity of its order.

Definition

- A colour class X_i is called *balanced* if $|V(Q_1) \cap X_i| = |V(Q_2) \cap X_i|$.
- A colour class X_i is Q_1-*skew* if $|V(Q_1) \cap X_i| \geq |V(Q_2) \cap X_i| + 1$. When $|V(Q_1) \cap X_i| = |V(Q_2) \cap X_i| + 1$, we say X_i is *just-Q_1-skew*.
- A colour class X_i is Q_2-*skew* if $|V(Q_1) \cap X_i| + 1 \leq |V(Q_2) \cap X_i|$. When $|V(Q_1) \cap X_i| + 1 = |V(Q_2) \cap X_i|$, we say X_i is *just-Q_2-skew*.

- (X_i, X_j) is called a *skew pair* if X_i is Q_1-skew and X_j is Q_2-skew.

For simplicity, if X_i is Q_1-skew or Q_2-skew, then we say X_i is *skew*. Similarly if X_i is just-Q_1-skew or just-Q_2-skew, then we say X_i is *just-skew*.

We say G is an *exception* if n is even, and there is a colour class X_s such that $|V(Q_1) \cap X_s| = |V(Q_1)| - 1$ and $|V(Q_2) \cap X_s| = |V(Q_2)| - 1$. We define the exception for technical reasons, however we need to avoid the exception in the following lemma only.

Lemma 4.6. *Let G be a complete multipartite graph $G := K_{n_1,\dots,n_k}$ such that $k \geq 2$ and $n > \max\{4, k\}$. Also assume G is neither a star nor an exception. Let v be a vertex of G chosen from a largest colour class, let $\mathcal{B}$ be a canonical line-bramble for v, and let $(H, (Q_1, Q_2, Q_3))$ be a good labelling. If (X_i, Xj) is a skew pair, then both X_i and X_j are just-skew.*

Proof. Since no colour class can be both Q_1-skew and Q_2-skew, $i \neq j$. Since $n \geq 5$, by Lemma 4.4, both Q_1 and Q_2 contain at least two vertices, and thus intersect at least two colour classes.

First, we show that both X_i^* and X_j^* contain a vertex other than v. If $X_i^* = \emptyset$, then X_i is not skew. So now assume $X_i^* \neq \emptyset$. Similarly, $X_j^* \neq \emptyset$. If $X_i^* = \{v\}$, then by Lemma 4.5, $X_i \cap V(Q_3) = \emptyset$, and so $X_i = \{v\}$. But since v is in a largest colour class, every colour class has order one, and as such $k = n$, which contradicts one of our assumptions on n. Thus both X_i^* and X_j^* contain a vertex other than v, and since X_i is Q_1-skew and X_j is Q_2-skew, there are vertices $x \in (V(Q_1) \cap X_i) - \{v\}$ and $y \in (V(Q_2) \cap X_j) - \{v\}$. Then define the hitting set H' as follows: remove the edges from x to $V(Q_2)$ from H, add all the edges from x to $V(Q_1) - X_i$ (these edges are not in H by Lemma 3.5), remove the edges from y to $V(Q_1) - \{x\}$, and add the edges from y to $V(Q_2) \cup \{x\}$. Now $G - H'$ contains components $(Q_1 - \{x\}) \cup \{y\}, (Q_2 - \{y\}) \cup \{x\}$ and Q_3, assuming that $(Q_1 - \{x\}) \cup \{y\}$ and $(Q_2 - \{y\}) \cup \{x\}$ are in fact connected (which we now prove).

If $(Q_1 - \{x\}) \cup \{y\}$ is not connected, then it intersects only one colour class by Lemma 3.5, which must be X_j since $y \in X_j$. Since $x \in X_i$, it follows that $|V(Q_1) \cap X_j| = |V(Q_1)| - 1$. Since X_j is Q_2-skew,

$$|V(Q_1)| = |V(Q_1) \cap X_j| + 1 \leq |V(Q_2) \cap X_j| \leq |V(Q_2)|.$$

Since $|V(Q_1)| \geq |V(Q_2)|$, we have $|V(Q_1)| = |V(Q_2)|$, and each inequality in the above equation is an equality. In particular, $|V(Q_2) \cap X_j| = |V(Q_2)|$, and thus $V(Q_2) \subseteq X_j$. But Q_2 intersects at least two colour classes, which is a contradiction. Thus $(Q_1 - \{x\}) \cup \{y\}$ is a connected component of $G - H'$.

If $(Q_2 - \{y\}) \cup \{x\}$ is not connected, then it intersects only one colour class, which must be X_i since $x \in X_i$. Since $y \in X_j$, it follows that $|V(Q_2) \cap X_i| = |V(Q_2)| - 1$. Since X_i is Q_1-skew,

$$|V(Q_1)| \geq |V(Q_1) \cap X_i| \geq |V(Q_2) \cap X_i| + 1 = |V(Q_2)|.$$

By Lemma 4.4, either $|V(Q_1)| = |V(Q_2)|$ (when n is odd) or $|V(Q_1)| = |V(Q_2)|+1$ (when n is even). If $|V(Q_1) \cap X_i| = |V(Q_1)|$, then $V(Q_1) \subseteq X_i$, contradicting our result that Q_1 intersects at least two colour classes. Otherwise $|V(Q_1) \cap X_i| = |V(Q_1)| - 1$, which can only happen when n is even. In this case, since $|V(Q_1) \cap X_i| = |V(Q_1)| - 1$ and $|V(Q_2) \cap X_i| = |V(Q_2)| - 1$, G is an exception. This contradiction shows that $(Q_2 - \{y\}) \cup \{x\}$ is a connected component of $G - H'$.

Thus $G - H'$ contains components $(Q_1 - \{x\}) \cup \{y\}, (Q_2 - \{y\}) \cup \{x\}$ and Q_3. Hence the orders of the components have not changed. Since the vertex v has not changed components, H' is a legitimate hitting set. But since H is the minimum hitting set by condition (0), $|H'| \geq |H|$. Hence

$$\begin{aligned}|H'| =& |H| - (|V(Q_2)| - |V(Q_2) \cap X_i|) + (|V(Q_1)| - |V(Q_1) \cap X_i|) \\ &- (|V(Q_1)| - 1 - |V(Q_1) \cap X_j|) + (|V(Q_2)| + 1 - |V(Q_2) \cap X_j|) \\ \geq& |H|.\end{aligned}$$

Which implies

$$|V(Q_2) \cap X_i| + |V(Q_1) \cap X_j| \geq |V(Q_1) \cap X_i| + |V(Q_2) \cap X_j| - 2.$$

Since X_i is Q_1-skew and X_j is Q_2-skew,

$$|V(Q_1) \cap X_i| + |V(Q_2) \cap X_j| - 2 \geq |V(Q_2) \cap X_i| + |V(Q_1) \cap X_j| \geq |V(Q_1) \cap X_i| + |V(Q_2) \cap X_j| - 2.$$

This only holds if every inequality is actually an equality. That is, X_i is just-Q_1-skew and X_j is just-Q_2-skew. □

Lemma 4.6 exists to help prove the following far more useful lemma.

Lemma 4.7. *Let G be a complete multipartite graph $G := K_{n_1,\ldots,n_k}$ such that $k \geq 2$, $n > k$ and such that G is not a star. Let v be a vertex of G chosen from a largest colour class, let $\mathcal{B}$ be a canonical line-bramble for v, and let $(H, (Q_1, Q_2, Q_3))$ be a good labelling. If X_i is skew, then X_i is just-skew.*

Proof. Suppose G is not an exception and $n > 4$. If there exists a Q_1-skew colour class X_s and a Q_2-skew colour class X_t, then either (X_s, X_i) or (X_i, X_t) is a skew pair, and by Lemma 4.6, X_i is just-skew, as required.

Alternatively, either no colour class is Q_1-skew or no colour class is Q_2-skew. Suppose, for the sake of contradiction, there is a skew colour class X_j that is not just-skew. In the first case, for all ℓ, $|V(Q_1) \cap X_\ell| \leq |V(Q_2) \cap X_\ell|$, and $|V(Q_1) \cap X_j| + 2 \leq |V(Q_2) \cap X_j|$. Thus

$$\begin{aligned} |V(Q_1)| + 2 =& (\sum_{1 \leq \ell \leq k, \ell \neq j} |V(Q_1) \cap X_\ell|) + |V(Q_1) \cap X_j| + 2 \\ \leq& (\sum_{1 \leq \ell \leq k, \ell \neq j} |V(Q_2) \cap X_\ell|) + |V(Q_2) \cap X_j| = |V(Q_2)|. \end{aligned}$$

This contradicts $|V(Q_1)| \geq |V(Q_2)|$. Similarly, in the second case, $|V(Q_1)| \geq |V(Q_2)| + 2$, which contradicts Lemma 4.4. Thus if $n \geq 5$ and G is not an exception, then our statement holds.

Consider the case when G is an exception. Then $|V(Q_1) \cap X_s| = |V(Q_1)| - 1$ and $|V(Q_2) \cap X_s| = |V(Q_2)| - 1$. Since n is even, by Lemma 4.4, $|V(Q_1)| = |V(Q_2)| + 1$, so X_s is just-skew. There are exactly two other vertices of $Q_1 \cup Q_2$, one in each component, which we label x and y respectively. If x and y are in the same colour class, then that colour class is balanced. Otherwise, x and y are in different colour classes, each of which intersects $Q_1 \cup Q_2$ in one vertex. Such a colour class is just-skew, as required.

Finally, consider the case $n \leq 4$. Then $|V(Q_1) \cup V(Q_2)| \leq 3$. Thus either $|V(Q_1)| = |V(Q_2)| = 1$, or $|V(Q_1)| = 2$ and $|V(Q_2)| = 1$. If X_i is not just-skew, then X_i contains at least two vertices in some component. Thus, the only possibility to consider is when $|V(Q_1) \cap X_i| = 2$. But then Q_1 is not connected, since both vertices are in the same colour class, which contradicts the fact that Q_1 is a connected component.

Thus X_i is just-skew. □

From Lemma 4.7 and Lemma 4.4, we get the following results about $|Q_1 \cap X_i|$ and $|Q_2 \cap X_i|$:

Corollary 4.8. *Let $G, v, \mathcal{B}$ and $(H, (Q_1, Q_2, Q_3))$ be as in Lemma 4.7. If a colour class X_i does not intersect Q_3, then*

- *if X_i is balanced, then $|Q_1 \cap X_i| = |Q_2 \cap X_i| = \frac{n_i}{2}$*
- *if X_i is Q_1-skew, then $|Q_1 \cap X_i| = \frac{n_i+1}{2}$ and $|Q_2 \cap X_i| = \frac{n_i-1}{2}$*
- *if X_i is Q_2-skew, then $|Q_1 \cap X_i| = \frac{n_i-1}{2}$ and $|Q_2 \cap X_i| = \frac{n_i+1}{2}$*

Corollary 4.9. *Let $G, v, \mathcal{B}$ and $(H, (Q_1, Q_2, Q_3))$ be as in Lemma 4.7. If a colour class X_i does intersect Q_3, then $|V(Q_3) \cap X_i| = 1$ and*

- *if X_i is balanced, then $|Q_1 \cap X_i| = |Q_2 \cap X_i| = \frac{n_i-1}{2}$*
- *if X_i is Q_1-skew, then $|Q_1 \cap X_i| = \frac{n_i}{2}$ and $|Q_2 \cap X_i| = \frac{n_i-2}{2}$*

- *if X_i is Q_2-skew, then* $|Q_1 \cap X_i| = \frac{n_i-2}{2}$ *and* $|Q_2 \cap X_i| = \frac{n_i}{2}$

As we might expect, there is approximately the same number of Q_1 and Q_2-skew colour classes, since the extra vertices they contribute to Q_1 and Q_2 balance each other out.

Lemma 4.10. *Let $G, v, \mathcal{B}$ and $(H, (Q_1, Q_2, Q_3))$ be as in Lemma 4.7. If n is odd, then there is an equal number of Q_1-skew and Q_2-skew colour classes. If n is even, then there is one more Q_1-skew colour class than there are Q_2-skew colour classes.*

Proof. Say there are a Q_1-skew colour classes and b Q_2-skew colour classes. By Lemma 4.7, if X_i is Q_1-skew, then $|V(Q_1) \cap X_i| = |V(Q_2) \cap X_i| + 1$, and if X_i is Q_2-skew, then $|V(Q_1) \cap X_i| = |V(Q_2) \cap X_i| - 1$. Thus

$$|V(Q_1)| = \sum_{1 \le i \le k} |V(Q_1) \cap X_i| = (\sum_{1 \le i \le k} |V(Q_2) \cap X_i|) + a - b = |V(Q_2)| + a - b.$$

If n is odd, then by Lemma 4.4, $|V(Q_1)| = |V(Q_2)|$, so $a = b$, as required. When n is even, $|V(Q_1)| = |V(Q_2)| + 1$, so $a = b + 1$. $\square$

From Lemma 4.3, Lemma 4.4, Corollary 4.8 and Corollary 4.9, we get the following result that summarises this section:

Theorem 4.11. *Let G be a complete multipartite graph $G := K_{n_1,\dots,n_k}$ such that $k \ge 2, n > k$ and such that G is not a star. Let v be a vertex of G chosen from a largest colour class, let $\mathcal{B}$ be a canonical line-bramble for v, and let $(H, (Q_1, \dots, Q_p))$ be a good labelling. Then $p = 3$. If n is odd, then $|V(Q_1)| = |V(Q_2)| = \frac{n-1}{2}$ and $|V(Q_3)| = 1$, and if n is even, then $|V(Q_1)| = \frac{n}{2}$, $|V(Q_2)| = \frac{n}{2} - 1$ and $|V(Q_3)| = 1$. For a colour class X_i,*

$$\lceil \frac{n_i - 2}{2} \rceil \le |V(Q_1) \cap X_i|, |V(Q_2) \cap X_i| \le \lfloor \frac{n_i + 1}{2} \rfloor.$$

The bounds on $|V(Q_1) \cap X_i|$ and $|V(Q_2) \cap X_i|$ in the above theorem are a little weak. However, it is difficult to improve Theorem 4.11 in general, given the broad set of possibilities for the sizes of the different colour classes. Given a specific set of colour classes, it would be possible to use the lemmas in this section to give a stronger bound.

Now we can use Theorem 4.11 to determine a lower bound on $\mathsf{tw}(L(G))$.

Theorem 4.12. *Let G be a complete multipartite graph $G := K_{n_1,\dots,n_k}$ where $k \ge 2$. Then* $\mathsf{tw}(L(G)) + 1 = \mathsf{bn}(L(G)) \ge \frac{1}{2}\left(\sum_{1 \le i < j \le k} n_i n_j\right) + \frac{3}{4}k^2 - kn - \frac{3}{4}k + n.$

Proof. First, consider the case when $k \ge 2, n > k$ and G is not a star. Then choose some vertex v in a largest colour class of G, a canonical line-bramble for v denoted $\mathcal{B}$ and a good labelling $(H, (Q_1, \dots, Q_p))$. It is sufficient to determine a lower bound on $|H|$, since

H is a minimum hitting set for $\mathcal{B}$ by condition (0), and since $\mathcal{B}$ forces the existence of a bramble of $L(G)$ of the same order by Lemma 3.1. Using Theorem 4.11, we can determine the structure of H. The set H contains all edges with an endpoint in Q_1 and an endpoint in Q_2; simply count these edges. By Theorem 4.11, $|V(Q_1) \cap X_i|, |V(Q_2) \cap X_i| \geq \lceil \frac{n_i}{2} - 1 \rceil$. Since $n_i, n_j \geq 1$, it follows that $|V(Q_1) \cap X_i||V(Q_2) \cap X_j| \geq \left(\frac{n_i}{2} - 1\right)\left(\frac{n_j}{2} - 1\right) - \frac{1}{4}$. So we count the edges from Q_1 to Q_2 as follows:

$$\begin{aligned}\sum_{i \neq j} |V(Q_1) \cap X_i||V(Q_2) \cap X_j| \geq & \sum_{i \neq j} \left(\frac{n_i}{2} - 1\right)\left(\frac{n_j}{2} - 1\right) - \frac{1}{4} \\ = & \frac{1}{4}\left(\sum_{i \neq j} n_i n_j\right) - (k-1)n + \frac{3}{4}k(k-1) \\ = & \frac{1}{2}\left(\sum_{1 \leq i < j \leq k} n_i n_j\right) + \frac{3}{4}k^2 - kn - \frac{3}{4}k + n.\end{aligned}$$

This gives the required lower bound on $|H|$ in this case.

It remains to check the cases when either $n = k$ or G is a star. When $n = k$, G is simply a complete graph, and our lower bound follows by Theorem 1.2. If G is a star, then $L(G)$ is a complete graph, and the lower bound follows by inspection. □

Using the same techniques as in the proof of Lemma 4.12, we can also determine an upper bound on $|H|$. We do this now. Note when considering the upper bound, we also need to account for the edges from Q_3 into the components Q_1, Q_2, but there are not many of these edges.

Lemma 4.13. *Let $G, v, \mathcal{B}$ and $(H, (Q_1, \ldots, Q_p))$ be as in Theorem 4.11. Then* $|H| \leq \frac{1}{2}\left(\sum_{1 \leq i < j \leq k} n_i n_j\right) + \frac{1}{2}n(k+1) + \frac{1}{4}k(k-1) - 1.$

Finally, our results in this section give a more detailed understanding of H when G is regular. (This ties in with our previous statement about stronger bounds being possible whenever specific information is known about the size of the colour classes.)

Theorem 4.14. *Let G be a complete regular k-partite graph $G := K_{c,\ldots,c}$, such that $k \geq 2$ and $n > k$, let v be a vertex of G chosen from a largest colour class, let $\mathcal{B}$ be a canonical line-bramble for v, and let $(H, (Q_1, \ldots, Q_p))$ be a good labelling. Then $p = 3$. If n is odd, then $|V(Q_1)| = |V(Q_2)| = \frac{n-1}{2}$ and $|V(Q_3)| = 1$ and*

- *for one colour class X_i, we have $|V(Q_1) \cap X_i| = |V(Q_2) \cap X_i| = \frac{c-1}{2}$ and $|V(Q_3) \cap X_i| = 1$,*
- *for $\frac{k-1}{2}$ other colour classes X_i, we have $|V(Q_1) \cap X_i| = \frac{c+1}{2}$ and $|V(Q_2) \cap X_i| = \frac{c-1}{2}$,*

- *for the remaining* $\frac{k-1}{2}$ *colour classes* X_i*, we have* $|V(Q_1) \cap X_i| = \frac{c-1}{2}$ *and* $|V(Q_2) \cap X_i| = \frac{c+1}{2}$*.*

If n *is even, then* $|V(Q_1)| = \frac{n}{2}$*,* $|V(Q_2)| = \frac{n}{2} - 1$ *and* $|V(Q_3)| = 1$*. If* n *is even and* c *is odd, then*

- *for one colour class* X_i*, we have* $|V(Q_1) \cap X_i| = |V(Q_2) \cap X_i| = \frac{c-1}{2}$ *and* $|V(Q_3) \cap X_i| = 1$*,*
- *for* $\frac{k}{2}$ *other colour classes* X_i*, we have* $|V(Q_1) \cap X_i| = \frac{c+1}{2}$ *and* $|V(Q_2) \cap X_i| = \frac{c-1}{2}$*,*
- *for the remaining* $\frac{k}{2} - 1$ *colour classes* X_i*, we have* $|V(Q_1) \cap X_i| = \frac{c-1}{2}$ *and* $|V(Q_2) \cap X_i| = \frac{c+1}{2}$*.*

Finally, if n *is even and* c *is even, then*

- *for one colour class* X_i*, we have* $|V(Q_1) \cap X_i| = \frac{c}{2}$*,* $|V(Q_2) \cap X_i| = \frac{c}{2} - 1$ *and* $|V(Q_3) \cap X_i| = 1$*,*
- *for the other* $k - 1$ *colour classes* X_i*, we have* $|V(Q_1) \cap X_i| = |V(Q_2) \cap X_i| = \frac{c}{2}$*.*

Proof. Since G is regular and $n > k \geq 2$, G is not a star. The statements about the number and order of the components of $G - H$ all follow from Lemma 4.3 and Lemma 4.4. Since $n = ck$, when n is odd, c is odd and k is odd. When n is even, at least one of c and k are even. Then from Corollary 4.8, Corollary 4.9 and Lemma 4.10, the rest of the theorem follows. □

4.3 Path Decompositions

As we did in Chapter 3, we define a path decomposition for $L(G)$ to determine an upper bound on $\mathsf{tw}(G)$ and $\mathsf{pw}(G)$. Let T denote the underlying path. Since T is a path, it makes sense to refer to a bag *left* or *right* of another bag, depending on the relative positions of the corresponding nodes in T. If a bag is to the right of another bag and the nodes which index them are adjacent in T, then we say it is *directly right.* Similarly define *directly left.* In order to construct this path decomposition, we will use a hitting set H from a good labelling, as well as the information we have about Q_1, Q_2 and Q_3. For a vertex u of G, let $\deg_i(u)$ be the number of edges in G incident to u with the other endpoint in the component Q_i.

First, label the vertices of Q_1 by $x_1, \ldots, x_{|V(Q_1)|}$ in some order, which we will specify later. Similarly, label the vertices of Q_2 by $y_1, \ldots, y_{|V(Q_2)|}$, again in an order we will later specify. Finally, by Theorem 4.11, Q_3 contains a single vertex, which we label z.

Then define the following bags:

- $\gamma := H = \{uw \in E(G) : u, w$ are in different components of $G - H\}$,
- for $1 \leq i \leq |V(Q_1)|$,

$$\alpha_i := \{x_\ell u \in E(G) : u \in V(Q_1), 1 \leq \ell \leq i\}$$
$$\cup \{x_j w \in E(G) : w \in V(G) - V(Q_1), i \leq j \leq |V(Q_1)|\},$$

- for $1 \leq i \leq |V(Q_2)|$,

$$\beta_i := \{y_\ell u \in E(G) : u \in V(Q_2), 1 \leq \ell \leq i\}$$
$$\cup \{y_j w \in E(G) : w \in V(G) - V(Q_2), i \leq j \leq |V(Q_2)|\}.$$

Each bag is indexed by a node of T. Left-to-right, the nodes of T index the bags in the following order: $\beta_{|V(Q_2)|}, \ldots, \beta_1, \gamma, \alpha_1, \ldots, \alpha_{|V(Q_1)|}$. Let $\mathcal{X}$ denote the collection of bags. We claim this defines a tree decomposition $(T, \mathcal{X})$ for $L(G)$, independent of our ordering of Q_1 and Q_2.

Lemma 4.15. *Let G be a complete multipartite graph $G := K_{n_1,\ldots,n_k}$ such that $k \geq 2$ and $n > k$. Also say G is not a star. Let v be a vertex of G chosen from a largest colour class, let $\mathcal{B}$ be a canonical line-bramble for v, and let $(H, (Q_1, Q_2, Q_3))$ be a good labelling. Then $(T, \mathcal{X})$ is a tree decomposition of $L(G)$, independent of the ordering used on Q_1 and Q_2.*

Proof. Consider $uw \in E(G)$. We require that the nodes indexing the bags containing uw induce a non-empty connected subpath of T. Firstly, assume that u and w are in different components of $G - H$. If $u = x_i$ and $w = y_j$, then $uw \in \beta_j, \ldots, \beta_1, \gamma, \alpha_1, \ldots, \alpha_i$, meaning uw is in precisely this sequence of bags. If $u = x_i$ and $w = z$, then $uw \in \gamma, \alpha_1, \ldots, \alpha_i$. If $u = y_j$ and $w = z$, then $uw \in \beta_j, \ldots, \beta_1, \gamma$.

Secondly, assume that u and w are in the same component of $G - H$, which is either Q_1 or Q_2, since by Theorem 4.11, $|V(Q_3)| = 1$. If $u, w \in V(Q_1)$, then let $u = x_i$ be the vertex of smaller label. Then $uw \in \alpha_i, \ldots, \alpha_{|V(Q_1)|}$. If $u, w \in V(Q_2)$, then similarly let $u = y_i$ be the vertex of smaller label. Then $uw \in \beta_{|V(Q_2)|}, \ldots, \beta_i$. This shows that the nodes indexing the bags containing uw induce a non-empty connected subpath of T.

All that remains is to show that if two edges are incident at a vertex in G (that is, the edges are adjacent in $L(G)$), then there is a bag of $\mathcal{X}$ containing both of them. Now if the shared vertex of the two edges is $x_i \in V(Q_1)$, then by inspection both edges are in α_i. If the shared vertex is $y_j \in V(Q_2)$, then both edges are in β_j. Finally, if the shared vertex is z, then both edges are in γ. □

Now we determine the width of $(T, \mathcal{X})$, which is one less than the order of the largest bag. To do so, we use a specific labelling of $Q_1 \cup Q_2$. We do this in two different ways, depending on whether or not G is regular.

In our first ordering, label the vertices $x_1, \ldots, x_{|V(Q_1)|}$ in order of non-decreasing size of the colour class containing x_i, and do the same for $y_1, \ldots, y_{|V(Q_2)|}$. We denote this ordering as the *red ordering*.

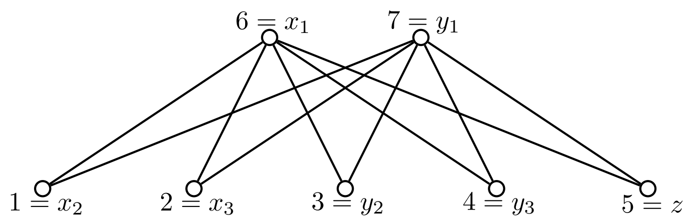

Figure 4.1: A red ordering for $L(K_{5,2})$. Here $v = 5$, and $Q_1 = \{1,2,6\}, Q_2 = \{3,4,7\}, Q_3 = \{5\}$. The lemmas in the previous section actually give (without loss of generality) two different possibilities for Q_1, Q_2, Q_3. The choice shown is the better of the two, which we determine simply by inspecting both.

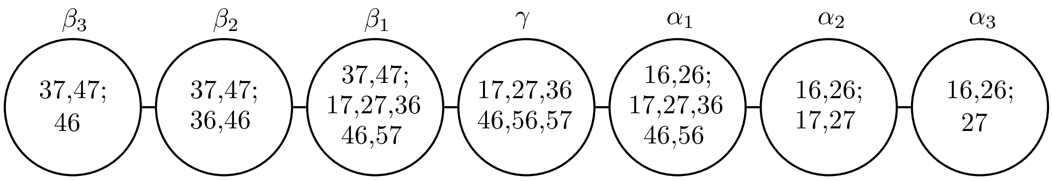

Figure 4.2: A path decomposition for $L(K_{5,2})$ using the red ordering.

Lemma 4.16. *Let $G, v, \mathcal{B}, (H, (Q_1, Q_2, Q_3))$ and $(T, \mathcal{X})$ be as in Lemma 4.15, but assume the ordering on Q_1 and Q_2 is the red ordering. Then $|\alpha_i| \leq |\alpha_1| + n - 2$, for all $1 \leq i \leq |V(Q_1)|$.*

Proof. We will show that $|\alpha_i| \leq |\alpha_{i-1}| + 2$ for all i. This implies that $|\alpha_i| \leq |\alpha_1| + 2(i-1)$. Since $i \leq |V(Q_1)|$ and $|V(Q_1)| \leq \frac{n}{2}$ by Lemma 3.3, this is sufficient.

$$\begin{aligned}
\alpha_i =& \{x_\ell u, x_j w \in E(G) : u \in V(Q_1), w \in V(G) - V(Q_1), 1 \leq \ell \leq i, i \leq j \leq |V(Q_1)|\} \\
=& \{x_\ell u \in E(G) : u \in V(Q_1), 1 \leq \ell \leq i\} \\
& \cup \{x_j w \in E(G) : w \in V(G) - V(Q_1), i \leq j \leq |V(Q_1)|\}.
\end{aligned}$$

This is a disjoint union. Let X_s, X_t be the colour classes such that $x_{i-1} \in X_s$ and $x_i \in X_t$, and note that it is possible $s = t$. Then

$$\begin{aligned}
|\alpha_i| - |\alpha_{i-1}| =& |\{x_\ell u \in E(G) : u \in V(Q_1), 1 \leq \ell \leq i\}| \\
&- |\{x_\ell u \in E(G) : u \in V(Q_1), 1 \leq \ell \leq i-1\}| \\
&+ |\{x_j w \in E(G) : w \in V(G) - V(Q_1), i \leq j \leq |V(Q_1)|\}| \\
&- |\{x_j w \in E(G) : w \in V(G) - V(Q_1), i-1 \leq j \leq |V(Q_1)|\}| \\
\leq& \deg_1(x_i) - |\{x_{i-1} w \in E(G) : w \in V(G) - V(Q_1)\}| \\
=& \deg_1(x_i) - (\deg_G(x_{i-1}) - \deg_1(x_{i-1})) \\
=& \deg_1(x_i) - (n - n_s - \deg_1(x_{i-1})) \\
=& |V(Q_1)| - |V(Q_1 \cap X_t)| - (n - n_s - |V(Q_1)| + |V(Q_1 \cap X_s)|) \\
=& 2|V(Q_1)| + n_s - |V(Q_1 \cap X_t)| - n - |V(Q_1 \cap X_s)|.
\end{aligned}$$

Assume for the sake of contradiction that $|\alpha_i| - |\alpha_{i-1}| > 2$. Then:

$$2|V(Q_1)| + n_s > n + |V(Q_1) \cap X_s| + |V(Q_1) \cap X_t| + 2.$$

By the ordering of the vertices in Q_1, $n_t \geq n_s$. Then by Theorem 4.11,

$$|V(Q_1) \cap X_s| + |V(Q_1) \cap X_t| \geq \frac{n_s - 2}{2} + \frac{n_t - 2}{2} \geq n_s - 2.$$

Hence $2|V(Q_1)| + n_s > n + n_s - 2 + 2$; that is, $2|V(Q_1)| > n$. But $|V(Q_1)| > \frac{n}{2}$ contradicts Lemma 4.4. $\square$

By symmetry we have:

Lemma 4.17. *Let $G, v, \mathcal{B}, (H, (Q_1, Q_2, Q_3))$ and $(T, \mathcal{X})$ be as in Lemma 4.15, but assume the ordering on Q_1 and Q_2 is the red ordering. Then $|\beta_i| \leq |\beta_1| + n - 2$, for all $1 \leq i \leq |V(Q_2)|$.*

Using Lemmas 4.16 and 4.17, we determine the maximum size of a bag in this path decomposition in terms of $|H|$.

Lemma 4.18. *Let $G, v, \mathcal{B}, (H, (Q_1, Q_2, Q_3))$ and $(T, \mathcal{X})$ be as in Lemma 4.15. The maximum bag size of $(T, \mathcal{X})$, using the red ordering, is at most $|H| + 2n - 2$.*

Proof. By Lemma 4.16 and Lemma 4.17, the maximum size of a bag right of γ is at most $|\alpha_1| + n - 2$, and left of γ it is $|\beta_1| + n - 2$. By inspection, the edges in $\alpha_1 - \gamma$ are all adjacent to x_1. Hence there are at most n of them. Thus $|\alpha_1| \leq |\gamma| + n$. Similarly $|\beta_1| \leq |\gamma| + n$. Since $\gamma = H$, this is sufficient. $\square$

Given this, we can determine an upper bound on $\mathsf{tw}(L(G))$.

Theorem 4.19. *Let G be a complete multipartite graph $G := K_{n_1,\dots,n_k}$ where $k \geq 2$. Then*

$$\mathsf{tw}(G) \leq \frac{1}{2}\left(\sum_{1\leq i<j\leq k} n_i n_j\right) + \frac{1}{4}k^2 + \frac{1}{2}kn - \frac{1}{4}k + \frac{5}{2}n - 4.$$

Proof. If $G, v, \mathcal{B}, (H, (Q_1, Q_2, Q_3))$ and $(T, \mathcal{X})$ are as in Lemma 4.15, then we have a tree decomposition of width at most $|H| + 2n - 3$ by Lemma 4.18. (Note since $k \geq 2$, it follows $n \geq 2$ and so $2n - 3$ is positive.) Then our result follows from Lemma 4.13. In the remaining cases, G is either a complete graph or a star, and this result follows by Theorem 1.2 or inspection, respectively. $\square$

Thus Theorem 1.3 follows from Theorem 4.12 and Theorem 4.19.

When G is regular, that is, $n_1 = \cdots = n_k$, we can get a more accurate bound on the treewidth. Define $c := n_1$ to be the size of each colour class. We need a different ordering of the vertices $x_1, \dots, x_{|Q_1|}$ and $y_1, \dots, y_{|Q_2|}$ to obtain our result. In order to do this, we recall the notion of a skew colour class, as defined in Section 4.2, and the associated results. First consider a colour class X_i that does not intersect Q_3. If X_i is balanced, then say every vertex of X_i is *Type 1*. If X_i is Q_1-skew, then each vertex in $Q_1 \cap X_i$ is *Type 1* and each vertex in $Q_2 \cap X_i$ is *Type 2*. If X_i is Q_2-skew, then each vertex in $Q_1 \cap X_i$ is *Type 2* and each vertex in $Q_2 \cap X_i$ is *Type 1*. Finally, each vertex in the remaining colour class (that does intersect Q_3) is *Type 3*. Thus each vertex of $V(G) - z$ is either Type 1, 2 or 3. Label the vertices of Q_1 in order $x_1, \dots, x_{|V(Q_1)|}$ by first labelling Type 1 vertices, then Type 2 vertices, and finally Type 3 vertices. Do the same for $y_1, \dots, y_{|V(Q_2)|}$. We denote this ordering as the *blue ordering*.

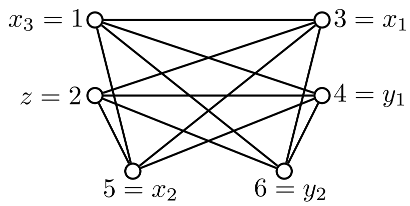

Figure 4.3: A blue ordering for $L(K_{2,2,2})$. Here, $v = 2$ and $Q_1 = \{1, 3, 5\}, Q_2 = \{4, 6\}, Q_3 = \{2\}$. Thus $1, 2$ are Type 3 vertices, and all other vertices are Type 1.

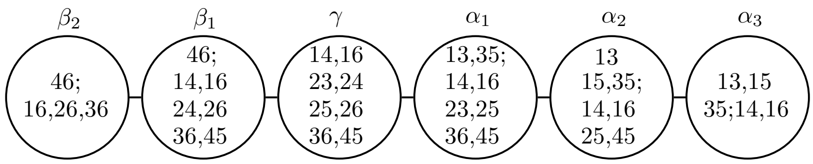

Figure 4.4: A path decomposition for $L(K_{2,2,2})$ using the blue ordering.

Lemma 4.20. *Let G be a complete k-partite graph with $n > k$, let v be a vertex of G, let $\mathcal{B}$ be a canonical line-bramble for v and let $(H, (Q_1, Q_2, Q_3))$ be a good labelling. If $k \geq 3$, then Q_1 contains at least two Type 1 vertices, and Q_2 contains at least one Type 1 vertex. If $k = 2$ and $c \geq 3$, then Q_1 contains at least two Type 1 vertices, and Q_2 contains at least one Type 1 or Type 2 vertex.*

Proof. If X_i is a colour class that does not intersect Q_3, then it intersects both of Q_1 and Q_2—if not, then by Lemma 4.7, $|X_i| = 1$ and G is a complete graph. Since we are trying to find Type 1 and Type 2 vertices, from now on we only consider colour classes that do not intersect Q_3. If $k \geq 5$, then there are at least four colour classes that do not intersect Q_3. From Theorem 4.14, there are either at least two Q_1-skew and Q_2-skew colour classes, or at least four balanced colour classes. Even if each such colour class intersects each of Q_1 and Q_2 only once, there are still enough colour classes of the correct skew to get all our required Type 1 vertices. Similarly, if $k = 4$ and c is odd, then there are two Q_1-skew colour classes and one Q_2-skew colour class, and if $k = 4$ and c is even, there are three balanced colour classes. This is again sufficient.

If $k = 3$, then by Theorem 4.14 again, there are enough Q_2-skew or balanced colour classes to ensure that Q_2 contains at least one Type 1 vertex. However, if n is odd, there is only one Q_1-skew colour class. In this case, c is odd, and so $c \geq 3$. Thus that colour class contains at least two vertices in Q_1. Thus Q_1 contains two Type 1 vertices.

Now assume $k = 2$ and $c \geq 3$. If c is odd, there is one Q_1-skew colour class, again by Theorem 4.14. This colour class contains at least two vertices in Q_1 and one in Q_2, which satisfies our requirement, now that Q_2 only requires a Type 2 vertex. If c is even, then there is one balanced colour class. Since $c \geq 3$ and even, it follows that $c \geq 4$ and each component contains two vertices from this colour class. This is sufficient. □

The following lemma strengthens Lemma 4.16 for the case when G is regular.

Lemma 4.21. *Let G be a complete k-partite graph with $n > k$ and $k \geq 2$, let v be a vertex of G, let $\mathcal{B}$ be a canonical line-bramble for v and let $(H, (Q_1, Q_2, Q_3))$ be a good labelling.*

Let $(T, \mathcal{X})$ be our tree decomposition where Q_1 and Q_2 are ordered by the blue ordering. If $k \geq 3$ or $c \geq 3$, then $|\alpha_1| \geq |\alpha_2| \geq \cdots \geq |\alpha_{|V(Q_1)|}|$.

Proof. We will show that $|\alpha_i| \leq |\alpha_{i-1}|$ for all i. We can write α_i as the disjoint union

$$\begin{aligned}\alpha_i =&\{x_\ell u \in E(G) : u \in V(Q_1), 1 \leq \ell \leq i\}\\ &\cup\{x_j w \in E(G) : w \in V(G) - V(Q_1), i \leq j \leq |V(Q_1)|\}.\end{aligned}$$

Let X_s, X_t be the colour classes such that $x_{i-1} \in X_s$ and $x_i \in X_t$, and note that it is possible that $s = t$. Define $r := |\{x_i x_f \in E(G) : f < i\}|$. Then

$$\begin{aligned}|\alpha_i| - |\alpha_{i-1}| =&|\{x_\ell u \in E(G) : u \in V(Q_1), 1 \leq \ell \leq i\}|\\ &- |\{x_\ell u \in E(G) : u \in V(Q_1), 1 \leq \ell \leq i-1\}|\\ &+ |\{x_j w \in E(G) : w \in V(G) - V(Q_1), i \leq j \leq |V(Q_1)|\}|\\ &- |\{x_j w \in E(G) : w \in V(G) - V(Q_1), i-1 \leq j \leq |V(Q_1)|\}|\\ =&\deg_1(x_i) - r - |\{x_{i-1} w \in E(G) | w \in V(G) - V(Q_1)\}|\\ =&\deg_1(x_i) - r - (\deg_G(x_{i-1}) - \deg_1(x_{i-1}))\\ =&\deg_1(x_i) - r - (n - n_s - \deg_1(x_{i-1}))\\ =&|V(Q_1)| - |V(Q_1 \cap X_t)| - r - (n - n_s - |V(Q_1)| + |V(Q_1 \cap X_s)|)\\ =&2|V(Q_1)| + n_s - r - |V(Q_1 \cap X_t)| - n - |V(Q_1 \cap X_s)|.\end{aligned}$$

Assume for the sake of contradiction that $|\alpha_i| - |\alpha_{i-1}| > 0$. Then:

$$2|V(Q_1)| + n_s > n + r + |V(Q_1) \cap X_s| + |V(Q_1) \cap X_t|.$$

There are two cases to consider. Firstly, say that both x_{i-1} and x_i are Type 1. So X_s and X_t are both balanced or Q_1-skew, and neither intersects Q_3. Since G is regular, $n_t = n_s$. Then by Corollary 4.8, $|V(Q_1) \cap X_s| + |V(Q_1) \cap X_t| \geq \frac{n_s}{2} + \frac{n_t}{2} = n_s$. Hence $2|V(Q_1)| + n_s > n + n_s + r \geq n + n_s$, so $2|V(Q_1)| > n$, which contradicts Lemma 4.4.

Secondly, since we ordered our vertices by non-decreasing type, we can assume x_i does not have Type 1. However, by Lemma 4.20, Q_1 contains at least two Type 1 vertices, x_a and x_b. Note if two vertices of Q_1 are in the same colour class, they have the same type, so we know that x_a and x_b are in a different colour class to x_i. Also, $a, b < i$, thus $r \geq 2$. Since $n_t = n_s$, by Theorem 4.11, $|V(Q_1) \cap X_s| + |V(Q_1) \cap X_t| \geq \frac{n_s-2}{2} + \frac{n_t-2}{2} = n_s - 2$. Hence $2|V(Q_1)| + n_s > n + n_s - 2 + r \geq n + n_s$, so $2|V(Q_1)| > n$, which again contradicts Lemma 4.4. □

We must also consider the equivalent argument for bags to the left of γ, as we did in the general case. However, here the arguments are not quite the same.

Lemma 4.22. *Let $G, v, \mathcal{B}, (H, (Q_1, Q_2, Q_3))$ and $(T, \mathcal{X})$ be as in Lemma 4.21. If $k \geq 3$ or $c \geq 3$, then $|\beta_1| \geq |\beta_2| \geq \cdots \geq |\beta_{|V(Q_2)|}|$.*

Proof. We will show that $|\beta_i| \leq |\beta_{i-1}|$ for all i. We can write β_i as the disjoint union

$$\begin{aligned}\beta_i = &\{y_\ell u \in E(G) : u \in V(Q_2), 1 \leq \ell \leq i\}\\ &\cup \{y_j w \in E(G) : w \in V(G) - V(Q_2), i \leq j \leq |V(Q_2)|\}.\end{aligned}$$

Let X_s, X_t be the colour classes such that $y_{i-1} \in X_s$ and $y_i \in X_t$, and note that it is possible that $s = t$. Define $r := |\{y_i y_f \in E(G) : f < i\}|$. Then

$$\begin{aligned}|\beta_i| - |\beta_{i-1}| = &|\{y_\ell u \in E(G) : u \in V(Q_2), 1 \leq \ell \leq i\}|\\ &- |\{y_\ell u \in E(G) : u \in V(Q_2), 1 \leq \ell \leq i-1\}|\\ &+ |\{y_j w \in E(G) : w \in V(G) - V(Q_2), i \leq j \leq |V(Q_2)|\}|\\ &- |\{y_j w \in E(G) : w \in V(G) - V(Q_2), i-1 \leq j \leq |V(Q_2)|\}|\\ = &\deg_2(y_i) - r - |\{y_{i-1} w \in E(G) | w \in V(G) - V(Q_2)\}|\\ = &\deg_2(y_i) - r - (\deg_G(y_{i-1}) - \deg_2(y_{i-1}))\\ = &\deg_2(y_i) - r - (n - n_s - \deg_2(y_{i-1}))\\ = &|V(Q_2)| - |V(Q_2 \cap X_t)| - r - (n - n_s - |V(Q_2)| + |V(Q_2 \cap X_s)|)\\ = &2|V(Q_2)| + n_s - r - |V(Q_2 \cap X_t)| - n - |V(Q_2 \cap X_s)|.\end{aligned}$$

Assume for the sake of contradiction that $|\beta_i| - |\beta_{i-1}| > 0$. Then:

$$2|V(Q_2)| + n_s > n + r + |V(Q_2) \cap X_s| + |V(Q_2) \cap X_t|.$$

There are two cases to consider. Firstly, say that neither of y_i and y_{i-1} have Type 3. So neither X_s nor X_t intersects Q_3. Since G is regular, $n_t = n_s$. By Corollary 4.8, $|V(Q_2) \cap X_s| + |V(Q_2) \cap X_t| \geq \frac{n_s - 1}{2} + \frac{n_t - 1}{2} = n_s - 1$. Hence $2|V(Q_2)| + n_s > n + r + n_s - 1 \geq n + n_s - 1$, and so $2|V(Q_2)| > n - 1$. However, Theorem 4.14 states that $|V(Q_2)| \leq \frac{n-1}{2}$, so this is a contradiction.

Secondly, y_i has Type 3. By Lemma 4.20, Q_2 contains at least one non-Type 3 vertex; this will be of a different colour class to y_i and have a lower numbered index. Hence $r \geq 1$. By Theorem 4.11, $|V(Q_2) \cap X_s| + |V(Q_2) \cap X_t| \geq \frac{n_s - 2}{2} + \frac{n_t - 2}{2} = n_s - 2$, and hence $2|V(Q_2)| + n_s > n + r + n_s - 2 \geq n + n_s - 1$. Again, this contradictions Theorem 4.14. □

Now we prove the most important fact, that γ is the largest bag.

Lemma 4.23. *Let $G, v, \mathcal{B}, (H, (Q_1, Q_2, Q_3))$ and $(T, \mathcal{X})$ be as in Lemma 4.21. If $k \geq 3$ or $c \geq 3$, then $|\alpha_1| \leq |\gamma|$ and $|\beta_1| \leq |\gamma|$.*

Proof. By inspection, $\alpha_1 = \{x_1u, uw \in E(G) : u \in V(Q_1), w \in V(G) - V(Q_1)\}$. Thus the edges of the form x_1u are the only edges in α_1 not in γ, and the edges between Q_2 and Q_3 (all of which are adjacent to z) are the only edges in γ not in α_1. Thus $|\alpha_1| + \deg_2(z) - \deg_1(x_1) = |\gamma|$. Suppose for the sake of contradiction that $|\alpha_1| > |\gamma|$. Say $x_1 \in X_s$ and $z \in X_t$. By Lemma 4.20, x_1 has Type 1, so $s \neq t$. Substituting $\deg_2(z) = |V(Q_2)| - |V(Q_2) \cap X_t|$ and $\deg_1(x_1) = |V(Q_1)| - |V(Q_1) \cap X_s|$ gives

$$|V(Q_1)| - |V(Q_2)| > |V(Q_1) \cap X_s| - |V(Q_2) \cap X_t|.$$

By Theorem 4.14, $|V(Q_1)| - |V(Q_2)| \leq 1$. Similarly, since X_t intersects Q_3, $|V(Q_2) \cap X_t| = \frac{c-1}{2}$ if c is odd, and $|V(Q_2) \cap X_t| = \frac{c-2}{2}$ if c is even. Since $X_s \cap Q_3 = \emptyset$ and x_1 has Type 1, $|V(Q_1) \cap X_s| \geq \frac{c}{2}$. Hence $|V(Q_1) \cap X_s| - |V(Q_2) \cap X_t| \geq \frac{1}{2}$ if c is odd, or 1 if c is even. However, this value is an integer, so $|V(Q_1) \cap X_s| - |V(Q_2) \cap X_t| \geq 1$, implying $|V(Q_1)| - |V(Q_2)| > 1$, which is a contradiction of Theorem 4.14.

Now we consider $\beta_1 = \{y_1u, uw \in E(G) : u \in V(Q_2), w \in V(G) - V(Q_2)\}$. Suppose for the sake of contradiction that $|\beta_1| > |\gamma|$. Let $y_1 \in X_s$ and $z \in X_t$. By Lemma 4.20, x_1 has Type 1 or Type 2, so $s \neq t$. Performing substitutions as we did in the α_1 case gives

$$|V(Q_2)| - |V(Q_1)| > |V(Q_2) \cap X_s| - |V(Q_1) \cap X_t|.$$

Since X_s does not intersect Q_3 and X_t does, by Theorem 4.14, $|V(Q_2) \cap X_s| \geq \frac{c-1}{2}$ and $|V(Q_1) \cap X_t| = \frac{c-1}{2}$ or $\frac{c}{2}$. Thus $|V(Q_2) \cap X_s| - |V(Q_1) \cap X_t| \geq 0$ or $-\frac{1}{2}$, but since it is an integer, $|V(Q_2) \cap X_s| - |V(Q_1) \cap X_t| \geq 0$, implying $|V(Q_2)| - |V(Q_1)| > 0$, which contradicts Theorem 4.14. □

By Lemmas 4.21, 4.22 and 4.23, γ is the largest bag in all but a few cases. Recall $\gamma = H$. Hence, together with Theorem 4.14, we get the following result.

Theorem 4.24. *If G is a regular k-partite graph such that $n > k$ and $k \geq 2$, and either $k \geq 3$ or $c \geq 3$, then*

$$\mathsf{tw}(L(G)) = |H| - 1.$$

We now accurately determine $|H|$ when G is regular.

We can determine $|H|$ by calculating the number of edges between Q_1 and Q_2, and the number of edges adjacent to $z \in Q_3$. Theorem 4.14 gives us all we require. It follows that:

$$|H| = \begin{cases} \frac{c^2k^2}{4} - \frac{c^2k}{4} + \frac{ck}{2} - \frac{c}{2} + \frac{k}{4} - \frac{1}{4} & \text{, if } ck \text{ odd} \\ \frac{c^2k^2}{4} - \frac{c^2k}{4} + \frac{ck}{2} - \frac{c}{2} & \text{, if } c \text{ even} \\ \frac{c^2k^2}{4} - \frac{c^2k}{4} + \frac{ck}{2} - \frac{c}{2} + \frac{k}{4} - \frac{1}{2} & \text{, if } k \text{ even and } c \text{ odd.} \end{cases}$$

Theorem 1.4 always assumes that $k \geq 2$. If $n > k$ and either $k \geq 3$ or $c \geq 3$, the above results prove the theorem. When $n = k$ (which is also the case when $c = 1$), we determine the treewidth using by Theorem 1.2. The final case is when $k = 2$ and $c = 2$. Here G is a 4-cycle, and thus $L(G)$ is also a 4-cycle. Since $\mathsf{tw}(K_{2,2}) = 2$ satisfies our result by inspection, this proves Theorem 1.4.

Chapter 5

Treewidth of General Line Graphs

5.1 Introduction

Given our results in Chapters 3 and 4, we desire, if possible, to extend our results to the class of general line graphs. Line-brambles, described in Section 3.2, work in the context of an arbitrary graph G, and can always be used to construct lower bounds on $\mathsf{tw}(L(G))$.

Recall that in Section 2.2, we proved the following well-known result.

Lemma 5.1. *Let $\delta(G)$ be the minimum degree of a graph G. Then for every graph G,*

$$\mathsf{tw}(G) \geq \delta(G).$$

In general, $\delta(L(G)) \geq 2\,\delta(G) - 2$, and this is tight when G contains two adjacent vertices of minimum degree. Thus, the following theorems are a significant strengthening of Lemma 5.1 for the class of line graphs.

Theorem 1.5. *For every graph G with minimum degree $\delta(G)$,*

$$\mathsf{tw}(L(G)) \geq \frac{2}{9}\delta(G)^2 - 1.$$

Theorem 1.6. *For every graph G with minimum degree $\delta(G)$,*

$$\mathsf{pw}(L(G)) \geq \frac{1}{4}\delta(G)^2 - 1.$$

This result is tight up to lower order terms.

Ignoring the lower order terms, Theorem 1.6 is tight for the line graph of a complete graph—it is not possible to improve the $\frac{1}{4}$ coefficient in general. Similarly, Theorem 1.5 is close to being tight, since $\frac{1}{4} - \frac{2}{9} = \frac{1}{36}$. If Theorem 1.5 is tight, then any family of graphs that proves this must have differing treewidth and pathwidth. (This precludes any of our results from Chapters 3 and 4 proving that Theorem 1.5 is tight.)

It is well known that a graph G with average degree $d(G)$ contains a subgraph H with $\delta(H) > \frac{1}{2}d(G)$. Given that $L(H)$ is a subgraph of $L(G)$ when H is a subgraph of G, it follows from Theorem 1.5 that

$$\mathsf{tw}(L(G)) \geq \mathsf{tw}(L(H)) \geq \frac{2}{9}\delta(H)^2 - 1 > \frac{1}{18}d(G)^2 - 1. \tag{5.1}$$

It follows equivalently from Theorem 1.6 that

$$\mathsf{pw}(L(G)) > \frac{1}{16}d(G)^2 - 1. \tag{5.2}$$

Note the similarity of (5.1) and (5.2) to a recent conjecture by Paul Seymour. This conjecture was proven by DeVos *et al.* [19], using the theory of immersions.

Theorem 5.2 (DeVos *et al.* [19]). *For every graph G with average degree $d(G)$,*

$$\mathsf{had}(L(G)) \geq c\, d(G)^{\frac{3}{2}}$$

for some constant $c > 0$. The exponent $\frac{3}{2}$ is best possible when G is the complete graph.

Given that $\mathsf{tw}(G) \geq \mathsf{had}(G) - 1$ for all graphs G, Theorem 5.2 implies

$$\mathsf{tw}(L(G)) \geq c\, d(G)^{\frac{3}{2}} - 1.$$

(5.1) and (5.2) replace the exponent of $\frac{3}{2}$ on $d(G)$ with an exponent of 2. Given that the exponent in Theorem 5.2 is best possible, (5.1) and (5.2) cannot be proven via a lower bound on $\mathsf{had}(L(G))$.

5.2 The General Lower Bound

To prove our results, we first prove a few facts about arbitrary tree decompositions of line graphs.

Let $(T, (B_x : x \in V(T)))$ be a tree decomposition. Root the tree T at an arbitrary leaf node r. Then orient every edge down the tree T, away from r. We can ensure that every node in T has outdegree at most 2, as follows. Let u be a node in T with outdegree greater than 2. Pick two out-neighbours of u and label them x and y. Delete the edges ux, uy and create a new node z, and add the edges uz, zx, zy (where the first node is the tail of the edge). In the bag B_z, place all vertices of the set $(B_x \cup B_y) \cap B_u$. This modified tree decomposition is still valid, and maintains the same width. However, z has outdegree exactly 2, and u has outdegree lowered by 1. Repeat this process until no node has outdegree greater than 2. Call such a tree decomposition a *degree-3* tree decomposition. So for each graph G there is a degree-3 tree decomposition of width $\mathsf{tw}(G)$.

Let $(T, (B_x : x \in V(T)))$ be a degree-3 tree decomposition for the line graph $L(G)$. For every edge $vw \in E(G) = V(L(G))$, let S_{vw} be the subtree of T induced by the nodes that index a bag containing vw. A node x of T is a *base node* of a vertex v of G if B_x contains every edge of G that is incident to v.

Lemma 5.3. *Let $(T, (B_x : x \in V(T)))$ be a tree decomposition for a graph $L(G)$ such that each S_{vw} is node-minimal. Then every non-isolated vertex of G has exactly one base node, and S_{vw} is a path between the base nodes of v and w.*

Proof. As previously, refer to vertices of $L(G)$ as edges for simplicity.

All of the edges incident at a vertex v of G form a clique in $L(G)$. By the Helly property, there is a bag containing all of the edges incident to v. Thus for all non-isolated vertices of G there is at least one base node.

Suppose for the sake of a contradiction that some non-isolated vertex v has more than one base node in T. Given that each S_{vw} is connected, it follows that there must be two adjacent base nodes of v in T. Let e be the edge between these two nodes, and label the subtrees of $T - e$ as T_1 and T_2. Without loss of generality, there exists some vertex $w \in N(v)$ such that w has a base node in T_1. Say we remove the edge vw from all bags indexed by nodes in T_2. Given that all edges incident to vw appear in a base node of v or w, both of which exist in T_1, it follows that this is still a tree decomposition. However, $|S_{vw}|$ has decreased (since we have removed vw from the other base node of v), and as such this contradicts the fact that S_{vw} was node-minimal. Hence each non-isolated v has exactly one base node in T.

It remains to show that each S_{vw} is a path. If x is a leaf node of S_{vw} (note that this is not necessarily a leaf of T itself), then by the node-minimality of S_{vw} it must be the case that if we removed vw from B_x, then we would no longer have a valid tree decomposition of $L(G)$. Hence either $S_{vw} = \{x\}$, or there is some edge incident to vw that is present in B_x but not in any other bag indexed by S_{vw}.

If $S_{vw} = \{x\}$, then S_{vw} is a path, as required.

Otherwise, B_x is the only bag containing both vw and some neighbour of vw. Without loss of generality, that neighbour is vz. Since B_x is the only bag containing vw and vz, it must be the only base node for v. Hence, any leaf in S_{vw} is the only base node for v and/or the only base node for w, so it follows S_{vw} is a path, since it may have at most two leaves. □

Everything in the proof of Lemma 5.3 also holds for path decompositions.

Corollary 5.4. *Let $(T, (B_x : x \in V(T)))$ be a path decomposition for a line graph $L(G)$ such that each S_{vw} is node-minimal. Then every non-isolated vertex of G has exactly one base node, and S_{vw} is a path between the base nodes of v and w.*

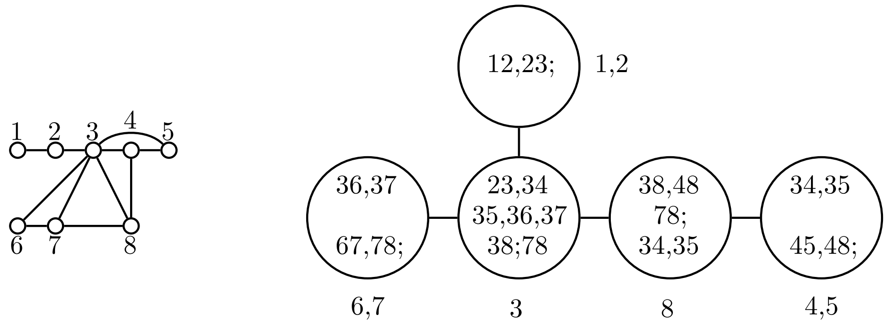

Figure 5.1: An example graph, together with a tree decomposition of its line graph. The tree decomposition has base nodes labelled, unique base nodes for each non-isolated vertex, and each $S_{(v,w)}$ as a path.

What Lemma 5.3 and Corollary 5.4 show is that, in some sense, there is only one way to create a tree (or path) decomposition for a line graph—start with a tree T, place base nodes for each non-isolated vertex (note isolated vertices have no influence on the line graph) and place edges between the base nodes. The difficulty is determining the best tree T and the best method of placing the base nodes. Note that the path decompositions constructed in the proofs of Theorems 1.2, 1.3 and 1.4 all have this structure.

Let $(T, (B_x : x \in V(T)))$ be a degree-3 tree decomposition for $L(G)$ such that each S_{vw} is node-minimal. For any node u of T, there are at most three subtrees in $T - u$, since u has at most one parent (that is, one in-neighbour) and two children (that is, two out-neighbours). Label these subtrees T_0 (for the parent subtree) and T_1, T_2 (for the child subtrees). Should one or more of T_0, T_1 and T_2 not exist, then let that $T_i = \emptyset$.

For a subset $Z \subseteq V(T)$, define $b(Z) = |\{v \in V(G) | v \text{ has its unique base node in } Z\}|$.

Lemma 5.5. *Let $(T, (B_x : x \in V(T)))$ be a degree-3 tree decomposition for a line graph $L(G)$ such that each S_{vw} is node-minimal. There exists a node u of $V(T)$ such that one of the following holds:*

1. $b(\{u\}) \geq \frac{1}{3}\delta(G)$
2. *At least one of* $b(V(T_1)), b(V(T_2))$ *or* $b(V(T_1) \cup V(T_2))$ *are in* $[\frac{1}{3}\delta(G), \frac{2}{3}\delta(G)]$.

Note that T_1, T_2 are defined with respect to $T - u$.

Proof. Traverse T, starting at r, as follows. If we are at a node u, and one of the child subtrees has $b(V(T_i)) > \frac{2}{3}\delta(G)$, then (and only then) traverse to the root of that subtree. Note that we only ever travel in the direction of an oriented edge (and, obviously, a tree is acyclic), so eventually this traversal must halt. Say the traversal halts at node u. Thus $b(V(T_1)), b(V(T_2)) \leq \frac{2}{3}\delta(G)$. If $b(V(T_i)) \geq \frac{1}{3}\delta(G)$ for $i = 1$ or 2, then u satisfies the second outcome. Otherwise $b(V(T_1) \cup V(T_2)) = b(V(T_1)) + b(V(T_2)) < \frac{2}{3}\delta(G)$. If $b(V(T_1) \cup V(T_2)) \geq \frac{1}{3}\delta(G)$, then u still satisfies the second outcome. Hence we assume otherwise.

If $u \neq r$, then the parent of u exists, and we traversed from the parent of u to u. Thus, $b(\{u\}) + b(V(T_1)) + b(V(T_2)) \geq \frac{2}{3}\delta(G)$, by the rules of the traversal. Alternatively, if $u = r$, then note there is no parent subtree T_0, and so $b(\{u\}) + b(V(T_1)) + b(V(T_2)) = |V(G)| \geq \delta(G) \geq \frac{2}{3}\delta(G)$. Hence, in either case $b(\{u\}) > \frac{2}{3}\delta(G) - \frac{1}{3}\delta(G) \geq \frac{1}{3}\delta(G)$, satisfying the first outcome. ☐

Let $(T, (B_x : x \in V(T)))$ be a path decomposition for $L(G)$ such that each S_{vw} is node-minimal. Root T at an endpoint r and orient all edges away from r. For any node u of T, let T_0 denote the parent subtree and T_1 the child subtree. Should either of T_0 and T_1 not exist, then let that $T_i = \emptyset$.

Lemma 5.6. *Let $(T, (B_x : x \in V(T)))$ be a path decomposition for a line graph $L(G)$ such that each S_{vw} is node-minimal, rooted at an endpoint r with all edges oriented away from the root. There exists a node u of $V(T)$ such that $b(\{u\} \cup T_0) \geq \lceil \frac{1}{2}\delta(G) \rceil$, but no node closer to the root has this property. Note that T_0 is defined with respect to $T - u$.*

Proof. If u is the node furthest from r, then $b(\{u\} \cup T_0) = |V(G)| \geq \lceil \frac{1}{2}\delta(G) \rceil$. Thus some node with the desired property holds; it is sufficient to choose the closest such node to r. ☐

Take u to be the node guaranteed by Lemma 5.5. Say u has Type 1 or Type 2 depending on which outcome of the lemma holds (if both hold, choose arbitrarily). If u has Type 1, then let $U \subset V(G)$ be a set of $\lceil \frac{1}{3}\delta(G) \rceil$ vertices with base node u. If u has Type 2, then let U be all the vertices with base node in T_1, T_2 or $T_1 \cup T_2$, depending on which value is in $[\frac{1}{3}\delta(G), \frac{2}{3}\delta(G)]$. If u is the node guaranteed by Lemma 5.6, then let $U \subset V(G)$ be a set of $\lceil \frac{1}{2}\delta(G) \rceil$ vertices with base node in $T_0 \cup \{u\}$, such that all vertices with a base node in T_0 are in U. (The choice of u allows this.)

Lemma 5.7. *Let $(T, (B_x : x \in V(T)))$ be a degree-3 tree decomposition for a line graph $L(G)$ such that each S_{vw} is node-minimal, and let u be the node guaranteed by Lemma 5.5. If $vw \in E(G)$, $v \in U$ and $w \in V(G) - U$, then $vw \in B_u$.*

Proof. If u has Type 1, then every $v \in U$ has base node u, so every edge of G with at least one endpoint in U is in B_u.

Say u has Type 2. Since w is not in U, w has base node not in the same subtree of $T - u$ as v. This is because U contains all vertices of G with a base node in a given subtree (or pair of subtrees). Hence the path from the base node of v to the base node of w must travel through u and so $vw \in B_u$. □

Lemma 5.8. *Let $(T, (B_x : x \in V(T)))$ be a path decomposition for a line graph $L(G)$ such that each S_{vw} is node-minimal, rooted at an endpoint r with all edges oriented away from the root, and let u be the node guaranteed by Lemma 5.6. If $vw \in E(G)$, $v \in U$ and $w \in V(G) - U$, then $vw \in B_u$.*

Proof. Since $v \in U$, v has a base node in $T_0 \cup \{u\}$. Since $w \notin U$, w has a base node in $T_1 \cup \{u\}$. Hence the path from the base node of v to the base node of w must travel through u and so $vw \in B_u$. □

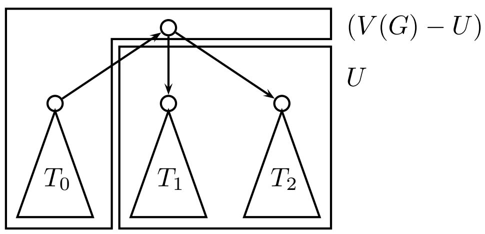

Figure 5.2: An edge with exactly one endpoint in U must be in u. Here we have the case when U contains exactly the base vertices of $V(T_1) \cup V(T_2)$.

Showing that $|B_u|$ is large is sufficient to prove our lower bounds on $\mathsf{tw}(L(G))$ and $\mathsf{pw}(L(G))$. Consider the tree decomposition case first. If u has Type 1, then each v in U has at least $\delta(G)$ neighbours, but at most $\lceil \frac{1}{3}\delta(G) \rceil - 1$ neighbours are also in U, since $|U| = \lceil \frac{1}{3}\delta(G) \rceil$. Hence for each $v \in U$ there are at least $\lfloor \frac{2}{3}\delta(G) \rfloor + 1$ vertices w such that $w \notin U$ and $vw \in E(G)$. Hence by Lemma 5.7, $|B_u| \geq \lceil \frac{1}{3}\delta(G) \rceil (\lfloor \frac{2}{3}\delta(G) \rfloor + 1) > \frac{1}{3}\delta(G)\frac{2}{3}\delta(G) \geq \frac{2}{9}\delta(G)^2$.

Alternatively, u has Type 2. We can say that $|U| = (\frac{1}{2} + \epsilon)\delta(G)$ where $|\epsilon| \leq \frac{1}{6}$. (Note ϵ may be negative.) Now each $v \in U$ has at least $(\frac{1}{2} - \epsilon)\delta(G) + 1$ neighbours in $V(G) - U$. Hence by Lemma 5.7, $|B_u| \geq (\frac{1}{2} + \epsilon)\delta(G)((\frac{1}{2} - \epsilon)\delta(G) + 1) > (\frac{1}{2} + \epsilon)(\frac{1}{2} - \epsilon)\delta(G)^2 \geq \frac{2}{9}\delta(G)^2$, since $\epsilon^2 \leq \frac{1}{36}$. This proves Theorem 1.5.

Finally we consider the path decomposition case. Each vertex v in U has at least $\lfloor \frac{1}{2}\delta(G) \rfloor + 1$ neighbours in $V(G) - U$. Hence by Lemma 5.8, $|B_u| \geq \lceil \frac{1}{2}\delta(G) \rceil (\lfloor \frac{1}{2}\delta(G) \rfloor + 1) > \frac{1}{4}\delta(G)^2$. This proves Theorem 1.6.

5.3 The General Upper Bound and Extensions

Călinescu *et al.* [11] and Atserias [3] independently proved the following upper bound on $\mathsf{tw}(L(G))$:

Theorem 5.9 (Atserias [3], Călinescu *et al.* [11])**.** *Let $\Delta(G)$ be the maximum degree of a graph G. Then for every graph G,*

$$\mathsf{tw}(L(G)) \leq (\mathsf{tw}(G)+1)(\Delta(G)) - 1.$$

Proof. Take a minimum width tree decomposition of G. Then replace every vertex in a given bag with every edge incident to that vertex. This is a tree decomposition of $L(G)$. The width of this tree decomposition is at most $(\mathsf{tw}(G)+1)(\Delta(G)) - 1$. □

A similar result holds for the pathwidth as well.

Corollary 5.10. *Let $\Delta(G)$ be the maximum degree of a graph G. Then*

$$\mathsf{pw}(L(G)) \leq (\mathsf{pw}(G)+1)(\Delta(G)) - 1.$$

Given the format of Theorem 5.9, we might hope that some analogous lower bound exists in terms of minimum degree and treewidth. Consider the following: there exist some constants $c, c' > 0$ such that for every graph G with minimum degree $\delta(G)$,

$$\mathsf{tw}(L(G)) \geq c\,\mathsf{tw}(G)\delta(G) \tag{5.3a}$$

$$\mathsf{pw}(L(G)) \geq c'\,\mathsf{pw}(G)\delta(G). \tag{5.3b}$$

(5.3a) and (5.3b) would be a strengthening of Theorems 1.5 and 1.6 respectively, since $\mathsf{pw}(G) \geq \mathsf{tw}(G) \geq \delta(G)$. However, (5.3a) and (5.3b) do not hold. In some sense, this implies Theorems 1.5 are best possible in the sense that we probably cannot replace $\delta(G)$ with anything stronger. We now provide a proof of this fact—thanks to Bruce Reed for this example.

For positive integers n, k construct the following graph $H_{n,k}$. Begin with the $n \times n$ grid, and for each vertex v of the grid, construct $k - \deg(v)$ cliques of size $k+1$. For each clique, add a single edge from a single vertex of the clique to the corresponding vertex v of the grid. Every vertex of this graph has degree k, except those vertices of the cliques which are adjacent to vertices on the grid, which have degree $k+1$. Hence the minimum degree $\delta(H_{n,k}) = k$. Since $H_{n,k}$ contains an $n \times n$ grid, it follows that $\mathsf{tw}(H_{n,k}) \geq n$. (In fact, it can be shown that $\mathsf{tw}(H_{n,k}) = n$ when $n \geq k+1$, but we omit this proof. Also, see Lemma 2.23 for a proof that the grid has treewidth n.)

Lemma 5.11. $\mathsf{pw}(L(H_{n,k})) \leq 2n + k + \binom{k+1}{2} - 1.$

Proof. First construct a path decomposition for the line graph of the $n \times n$ grid. Label the rows of the grid $1, \dots, n$ from top to bottom. Now any edge either has both endpoints in the same row, or an endpoint in two sequential rows. So label the edges $1, \dots, 2n(n-1)$ in the following fashion. First label all of the edges in the row 1 in the from left to right. Then label all edges with an endpoint in both row 1 and row 2 from left to right. Continue with the edges in row 2, then the edges between rows 2 and 3, and so on and so forth. Note that if two edges i, j are incident and $i < j$, then $j \leq 2n + i - 1$. Hence define our path decomposition $(P, \mathcal{X})$ as follows. Let P be a path with $2n(n-1)$ nodes. For the i^{th} node in the path, let the bag indexed by this node (which we denote by X_i) contain edges $\{i, \dots, 2n + i - 1\}$. Note that the edges are really acting as vertices of the line graph, but we refer to them as edges for simplicity. Also, for large values of i, not all of these edges exist—in that case simply place the defined edges into the bag. Finally, note that currently the largest bag has size $2n$. All that remains is to extend this path decomposition for the grid into a path decomposition for $L(H_{n,k})$.

For every vertex v of the $n \times n$ grid, there is a bag X_i that contains all edges of the grid incident to v. (In some cases there may be several legitimate choices of X_i, if so, choose one of them arbitrarily.) Refer to such an X_i as the bag corresponding to v. To this bag, add all edges incident to v that are not in the grid. (Recall there are $k - \deg(v)$ such edges.) Currently the largest bag has size $\leq 2n + k$. For each bag corresponding to some v, duplicate it $k - \deg(v) - 1$ times. (By duplicate, we mean to subdivide an edge incident to its node, and add a copy of X_i as the bag for that node.) Clearly, this maintains all path decomposition properties. Then for each clique corresponding to v, place all of the edges of that clique into exactly one of the copies of X_i; a different copy of X_i for each clique. This gives a path decomposition for $H_{n,k}$. The largest bag has size at most $2n + k + \binom{k+1}{2}$, which is sufficient to prove our result. $\square$

Now if either (5.3a) or (5.3b) hold, then $\mathsf{pw}(L(G)) \geq c\,\mathsf{tw}(G)\delta(G)$ for some fixed constant $c > 0$. However, consider $H_{n,k}$ and set $n \geq k + \binom{k+1}{2} - 1$ and $k > \frac{3}{c}$. Then $\mathsf{pw}(L(H_{n,k})) \leq 3n$ by Lemma 5.11 and $c\,\mathsf{tw}(H_{n,k})\delta(H_{n,k}) \geq cnk > 3n$, which is a contradiction. Hence neither (5.3a) nor (5.3b) hold.

Chapter 6

Treewidth of the Kneser Graph and the Erdős-Ko-Rado Theorem

6.1 Introduction

Recall the following definitions. The set $[n] = \{1, \ldots, n\}$. For any set $S \subseteq [n]$, a subset of S of size k is called a *k-set*, or occasionally a k-set in S. Let $\binom{S}{k}$ denote the set of all k-sets in S. We say two sets *intersect* when they have non-empty intersection. The *Kneser graph* $\mathsf{Kneser}(n,k)$ is the graph with vertex set $\binom{[n]}{k}$, such that two vertices are adjacent if they are disjoint.

In this chapter, we prove the following result.

Theorem 1.7. *Let $G = \mathsf{Kneser}(n,k)$ with $n \geq 4k^2 - 4k + 3$ and $k \geq 3$. Then*

$$\mathsf{tw}(G) = \binom{n-1}{k} - 1.$$

This gives an exact answer for the treewidth of the Kneser graph when n is sufficiently large. In order to prove this, we show that $\binom{n-1}{k} - 1$ is both an upper bound and lower bound on the treewidth. We construct a tree decomposition directly in Section 6.3 to prove an upper bound. In Section 6.4 we prove the lower bound by using the relationship between treewidth and separators, which was previously discussed in Chapter 2.

We also prove the following more precise result when $k = 2$.

Theorem 1.8. *Let $G = \mathsf{Kneser}(n,2)$. Then*

$$\mathsf{tw}(G) = \begin{cases} 0 & \text{if } n \leq 3 \\ 1 & \text{if } n = 4 \\ 4 & \text{if } n = 5 \\ \binom{n-1}{2} - 1 & \text{if } n \geq 6. \end{cases}$$

The upper bounds for Theorem 1.8 are proved in Section 6.3, and the lower bounds in Section 6.5.

In the process of proving Theorem 1.7, we prove the following generalisation of the Erdős-Ko-Rado Theorem (Theorem 6.2 in Section 6.2), which says that if $n \geq 2k$ and H is a complete subgraph in the complement of $\mathsf{Kneser}(n,k)$ then $|H| \leq \binom{n-1}{k-1}$. We prove the same bound for balanced complete multipartite graphs.

Theorem 1.10. *Say $c \in [\frac{2}{3}, 1)$ and $n \geq \max\{4k^2 - 4k + 3, \frac{1}{1-c}(k^2 - 1) + 2\}$. If H is a complete multipartite subgraph of the complement of $\mathsf{Kneser}(n,k)$ such that no colour class contains more than $c|H|$ vertices, then $|H| \leq \binom{n-1}{k-1}$.*

Note that similar, but incomparable, generalisations of the Erdős-Ko-Rado Theorem have recently been explored in [36, 37, 99]. Theorem 1.10 is proven in Section 6.4, since it follows almost directly from our proof of the lower bound on the treewidth of a Kneser graph.

Finally, in Section 6.6, we are able to obtain a weaker result on the lower bound of $\mathsf{Kneser}(n,k)$ for much smaller values of n. We do this by generalising the techniques used to prove Theorem 1.7.

Theorem 1.9. *Let $G = \mathsf{Kneser}(n,k)$ with $n \geq \frac{1}{2}(\sqrt{5k^2 - 12k + 8} + 3k + 2)$ and $k \geq 3$. Then*

$$\binom{n-1}{k} - \binom{n-1}{k-1} - 1 \leq \mathsf{tw}(G) \leq \binom{n-1}{k} - 1.$$

Recall that since $k \geq 3$, Theorem 1.9 holds when $n \geq 3k - 1$.

6.2 Basic Definitions and Preliminaries

From now on, we refer to the graph $\mathsf{Kneser}(n,k)$ as G, with n and k implicit.

Let $\Delta(H)$ be the maximum degree of a graph H and $\delta(H)$ be the minimum degree of a graph H. Also let $\alpha(H)$ be the size of the largest independent set of H, where an *independent set* is a set of pairwise non-adjacent vertices. If $k = 1$, then G is a complete graph. If $n < 2k$ then G contains no edges. If $n = 2k$ then G is an induced matching. From now on, we shall assume that $n \geq 2k + 1$ and $k \geq 2$, since the treewidth is trivial in the other cases.

In order to prove a lower bound on the treewidth of the Kneser graph, we the relationship between treewidth and separators. Recall the definition of a (k, S, c)-separator in Section 2.5. For this chapter, always set $S = V(G)$ and choose $c \in [\frac{2}{3}, 1)$, rather than

$[\frac{1}{2}, 1)$. We do this to ensure that if X is a $(|X|, S, c)$-separator, then $G - X$ can be partitioned into two parts with less than $c|G - X|$ vertices and no edges between them. This (essentially) follows from Corollary 2.6. This gives the following lemma.

Lemma 6.1. *Let X be a $(|X|, V(G), c)$-separator where $c \in [\frac{2}{3}, 1)$. Then $V(G - X)$ can be partitioned into two parts A and B, with no edge between A and B, such that*

- $(1 - c)|G - X| \leq |A| \leq \frac{1}{2}|G - X|$,
- $\frac{1}{2}|G - X| \leq |B| \leq c|G - X|$.

We use a few important well-known combinatorial results. Recall the following from Chapter 1.

Theorem 6.2 (Erdős-Ko-Rado [28, 50])**.** *Let $G = \mathsf{Kneser}(n, k)$ for some $n \geq 2k$. Then*

$$\alpha(G) = \binom{n-1}{k-1}.$$

If $n \geq 2k + 1$ and $\mathcal{A}$ is an independent set such that $|\mathcal{A}| = \binom{n-1}{k-1}$, then $\mathcal{A} = \{v | i \in v\}$ for a fixed element $i \in [n]$.

The original Erdős-Ko-Rado Theorem defines $\mathcal{A}$ as a set of k-sets in $[n]$, such that the k-sets of $\mathcal{A}$ pairwise intersect. Our formulation in terms of vertices in the Kneser graph is clearly equivalent. We will use Theorem 6.2 when determining an upper bound for $\mathsf{tw}(G)$.

The second major result is by Pyber [80]. Let $\mathcal{A}$ and $\mathcal{B}$ be sets of vertices of the Kneser graph G, such that for all $v \in \mathcal{A}$ and $w \in \mathcal{B}$ the pair vw is not an edge. Then we say the pair $(\mathcal{A}, \mathcal{B})$ are *cross-intersecting families*.

Theorem 6.3 (Erdős-Ko-Rado for Cross-Intersecting Families [78, 80])**.** *Let $n \geq 2k$ and let $(\mathcal{A}, \mathcal{B})$ be cross-intersecting families in $G = \mathsf{Kneser}(n, k)$. Then*

$$|\mathcal{A}||\mathcal{B}| \leq \binom{n-1}{k-1}^2.$$

If $n \geq 2k + 1$ and $(\mathcal{A}, \mathcal{B})$ are cross-intersecting families such that $|\mathcal{A}||\mathcal{B}| = \binom{n-1}{k-1}^2$, then $\mathcal{A} = \mathcal{B} = \{v | i \in v\}$ for a fixed element $i \in [n]$.

As with Theorem 6.2, the original formulation by Pyber of Theorem 6.3 is more general. We have given the result in an equivalent form that is sufficient for our requirements. The first statement in the theorem was originally proven by Pyber [80]. Matsumoto and Tokushige [78] proved the statement regarding the maximum choice of $\mathcal{A}$ and $\mathcal{B}$.

Let X be a $(|X|, V(G), \frac{2}{3})$-separator and A,B the parts of the vertex partition of $G - X$ as in Lemma 6.1. Now for all $v \in A$ and $w \in B$, v and w are in different components and as such are non-adjacent. So (A, B) are cross-intersecting families. We know $|A| = c|G - X|$ where $\frac{1}{3} \leq c \leq \frac{1}{2}$. By Theorem 6.3, it follows that $c(1-c)|G-X|^2 \leq \binom{n-1}{k-1}^2$. Thus $|G - X| \leq \sqrt{\frac{1}{c(1-c)}}\binom{n-1}{k-1}$. Since $\sqrt{\frac{1}{c(1-c)}}$ is maximised when $c = \frac{1}{3}$, it follows that $|G - X| \leq \frac{3}{\sqrt{2}}\binom{n-1}{k-1}$. This gives a lower bound on $|X|$, and as such a lower bound on the treewidth (by Lemma 2.7). Hence $\mathsf{tw}(G) \geq \binom{n}{k} - \frac{3}{\sqrt{2}}\binom{n-1}{k-1} - 1$.

However, note that the parts A and B of $V(G - X)$ are vertex disjoint, but that the definition of a pair of cross-intersecting families does not require this. In fact, Theorem 6.3 shows that in the case where $|\mathcal{A}||\mathcal{B}|$ is maximised, $\mathcal{A} = \mathcal{B}$. We show we can do better than the above naïve lower bound on $\mathsf{tw}(G)$ when $\mathcal{A}$ and $\mathcal{B}$ are disjoint.

Before considering our final preliminary, we provide the following definitions. Consider all of the a-sets in $[b]$. Define the *colexicographic* or *colex* ordering on the a-sets as follows: if x and y are distinct a-sets, then $x < y$ when $\max(x - y) < \max(y - x)$. This is a strict total order. A set X of a-sets in $[b]$ is *first* if X consists of the first $|X|$ a-sets in the colex ordering of all the a-sets in $[b]$.

Now consider the colex ordering of a-sets in $[b]$. All of the a-sets in $[i]$ (where $i < b$) come before any a-set containing an element greater than or equal to $i + 1$. To see this, note if x is an a-set in $[i]$ and y is an a-set with $j \in y$ such that $j \geq i + 1$, then $\max(x - y) \leq \max(x) \leq i$, and $\max(y - x) \geq j \geq i + 1$ since $j \in y - x$. We will use this when determining the make-up of first sets in Section 6.4.

Let X be a set of a-sets in $[b]$. For $p \leq a$, the *p-shadow of X* is the set $\{x : |x| = p,$ and $\exists y \in X$ such that $x \subseteq y\}$. That is, the p-shadow contains all p-sets that are contained within a-sets of X. If x is an a-set in $[b]$, let the *complement* of x be the $(b - a)$-set $y = [b] - x$. If X is a set of a-sets on $[b]$, then the *complement* of X is $\overline{X} := \{y : y$ is the complement of some $x \in X\}$. Note $|X| = |\overline{X}|$.

Lemma 6.4 (A first set minimises the shadow [49, 64] (see Frankl [31] for a short proof))**.** *Let X be a set of a-sets on $[b]$, $p \leq a$ and S be the p-shadow of X. Suppose $|X|$ is fixed but X is not. Then $|S|$ is minimised when X is first.*

This idea is also used by Pyber [80] and Matsumoto and Tokushige [78]. Intuitively, the shadow S should be minimised whenever the a-sets of X "overlap" as much as possible, so that each p-set in S is a subset of as many a-sets as possible.

6.3 Upper Bound for Treewidth

This section proves the upper bounds on $\mathsf{tw}(G)$ in Theorems 1.7 and 1.8.

In both Theorems 1.7 and 1.8, the upper bound is almost always $\binom{n-1}{k} - 1$. The only exceptions are the trivial cases (when $n \leq 2k$), and the case when $k = 2$ and $n = 5$, which is the Petersen graph. The Petersen graph is well known to have treewidth 4 ([75], for example, or below).

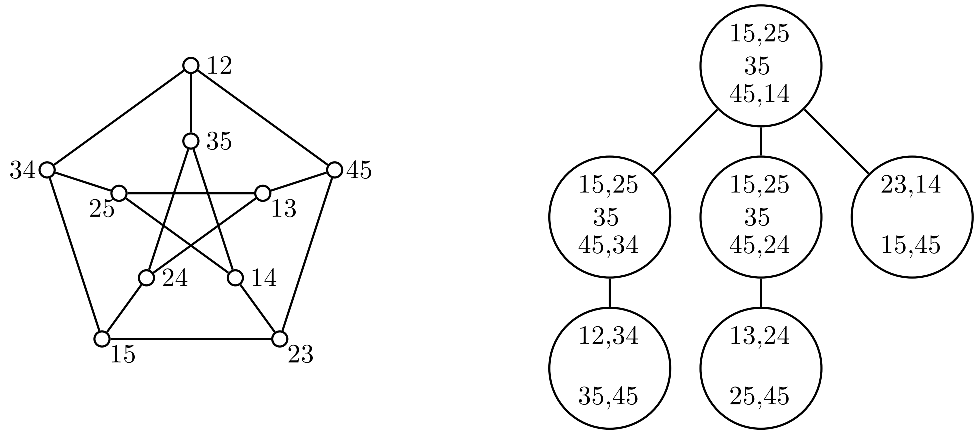

Figure 6.1: The Petersen graph $\mathsf{Kneser}(5, 2)$, together with a minimum width tree decomposition of $\mathsf{Kneser}(5, 2)$.

What follows is a general upper bound on the treewidth of any graph, which is sufficient to prove the remaining cases.

Lemma 6.5. *If H is any graph, then* $\mathsf{tw}(H) \leq \max\{\Delta(H), |V(H)| - \alpha(H) - 1\}$.

Proof. Let $\alpha := \alpha(H)$. We shall construct a tree decomposition with underlying tree T, where T is a star with $\alpha(H)$ leaves. Let R be the bag indexed by the central node of T, and label the other bags $B_1, \ldots, B_\alpha$. Let $X := \{x_1, \ldots x_\alpha\}$ be a maximum independent set in H. Let $R := V(H) - X$ and $B_i := N(x_i) \cup \{x_i\}$ for all $i \in \{1, \ldots, \alpha\}$. We now show this is a tree decomposition:

Any vertex not in X is contained in R. Given the structure of the star, any induced subgraph containing the central node is connected. Alternatively, if a vertex is in X, then it appears only in bags indexed by leaves. However, since X is an independent set, $x_i \in X$ appears only in B_i, not in any other bag B_j. A single node is obviously connected. If vw is an edge of H, then at most one of v and w is in X. Say $v = x_i \in X$. Then v, w both appear in the bag B_i. Otherwise neither vertex is in X, and both vertices appear in R.

So this is a tree decomposition. The size of R is $|V(H)| - \alpha(H)$. The size of B_i is the degree of x_i, plus one, which is at most $\Delta(H) + 1$. From here our lemma is proven. $\square$

Note that we can do slightly better than the above result if the vertices in the leaf bags are all known to have smaller than maximum degree. This is not an improvement with regards to Kneser graphs, since they are regular, but is helpful in Lemma 2.1.

We now consider Lemma 6.5 for the Kneser graph itself.

Lemma 6.6. *If G is a Kneser graph with $k \geq 2$ and $n \geq 2k+1$, then* $\mathsf{tw}(G) \leq \binom{n}{k-1} - 1$.

Proof. By Lemma 6.5 and Theorem 6.2, and since $n \geq 2k+1$,

$$\mathsf{tw}(G) \leq \max\{\Delta(G), |V(G)| - \alpha(G) - 1\} = \max\left\{\binom{n-k}{k}, \binom{n}{k} - \binom{n-1}{k-1} - 1\right\}.$$

Since $k \geq 2$, $\mathsf{tw}(G) \leq \binom{n-1}{k} - 1$, as required. □

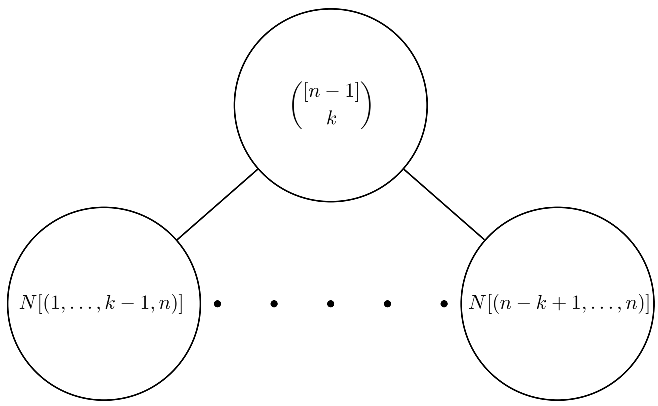

Figure 6.2: The tree decomposition described by Lemma 6.5 for $\mathsf{Kneser}(n,k)$. Recall that the closed neighbourhood $N[v] = N(v) \cup \{v\}$.

6.4 Separators in the Kneser Graph

To complete the proof of Theorem 1.7, it is sufficient to prove a lower bound on the treewidth. The following lemma, together with Lemma 2.7, provides this. It is the heart of the proof of Theorem 1.10.

Lemma 6.7. *Let X be a $(|X|, V(G), c)$-separator of the Kneser graph G where $c \in [\frac{2}{3}, 1)$. If $n \geq \max\{4k^2 - 4k + 3, \frac{1}{1-c}(k^2-1) + 2\}$, then $|X| \geq \binom{n-1}{k}$.*

Proof. Assume, for the sake of a contradiction, that $|X| < \binom{n-1}{k}$. Then $|G - X| > \binom{n-1}{k-1}$. By Lemma 6.1, $V(G-X)$ can be partitioned into two parts A and B such that $(1-c)|G-X| \leq |A| \leq \frac{1}{2}|G-X|$ and $\frac{1}{2}|G-X| \leq |B| \leq c|G-X|$ and no edge has an endpoint in both A and B.

For a given element $i \in [n]$, let $A_i := \{v \in A : i \in v\}$. Also define $A_{-i} := \{v \in A : i \notin v\}$. So A_i and A_{-i} partition the set A, for any choice of i. Define analogous sets for B.

Claim 1. There exists some i such that $|B_i| \geq \frac{1}{k}|B|$.

Proof. Since $|A| \geq (1-c)|G-X| > 0$, there is a vertex $v \in A$. Without loss of generality, $v = \{n-k+1, \dots, n\}$. Each $w \in B$ is not adjacent to v, and so w and v intersect. Thus each w must contain at least one of $n-k+1, \dots, n$. Hence at least one of these elements appears in at least $\frac{1}{k}|B|$ of the vertices of B, as required. $\square$

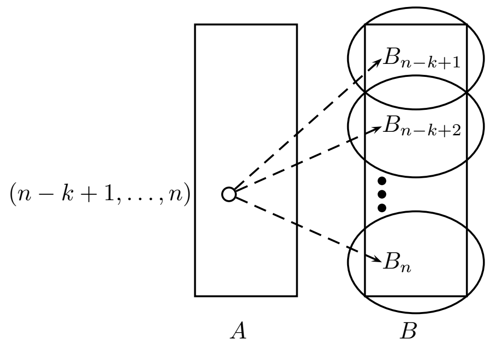

Figure 6.3: Diagram for Claim 1.

Without loss of generality, $|B_n| \geq \frac{1}{k}|B|$.

Claim 2. $|B_n| > \binom{n-3}{k-2} + \binom{n-2}{k-2}$.

Proof. Recall $|B| \geq \frac{1}{2}|G-X| \geq \frac{1}{2}\binom{n-1}{k-1}$. Then by Claim 1 and our subsequent assumption, $|B_n| \geq \frac{1}{k}|B| \geq \frac{1}{2k}|G-X| \geq \frac{1}{2k}\binom{n-1}{k-1}$. Assume for the sake of a contradiction that $|B_n| \leq \binom{n-3}{k-2} + \binom{n-2}{k-2}$. So

$$\frac{1}{2k}\binom{n-1}{k-1} \leq \binom{n-3}{k-2} + \binom{n-2}{k-2}.$$

Thus

$$(n-1)! \leq 2k(k-1)((n-k)(n-3)! + (n-2)!).$$

Hence

$$n^2 - 3n + 2 = (n-1)(n-2) \leq 2k(k-1)(2n-k-2) = 4k^2n - 4kn - 2k^3 - 2k^2 + 4k.$$

So $n^2 + (4k - 4k^2 - 3)n + 2k^3 + 2k^2 - 4k + 2 \leq 0$. Since $n \geq 4k^2 - 4k + 3$, it follows $2k^3 + 2k^2 - 4k + 2 \leq 0$. Given that $k \geq 1$, this provides our desired contradiction. $\square$

Consider the set $\overline{A_{-n}}$, that is, the complements of the vertices in A that do not contain n. So every set in $\overline{A_{-n}}$ contains n. Let $\overline{A_{-n}}^* := \{\overline{v} - n : \overline{v} \in \overline{A_{-n}}\}$. That is, remove n from each set in $\overline{A_{-n}}$. There is clearly a one-to-one correspondence between $(n-k)$-sets in $\overline{A_{-n}}$ and $(n-k-1)$-sets in $\overline{A_{-n}}^*$.

Similarly, define $B_n^* := \{v - n : v \in B_n\}$. That is, remove from each vertex of B_n the element n, which they all contain. The resultant sets are $(k-1)$-sets in $[n-1]$.

Claim 3. If $v^* \in B_n^*$ and $\overline{w}^* \in \overline{A_{-n}}^*$, then $v^* \not\subseteq \overline{w}^*$.

Proof. Assume, for the sake of a contradiction, that $v^* \subseteq \overline{w}^*$. Then it follows that $v \subset \overline{w}$, by re-adding n to both sets. Thus v and w are adjacent. However, $v \in B_n \subset B$ and $w \in A_n \subset A$, which is a contradiction. □

Let S be the $(k-1)$-shadow of $\overline{A_{-n}}^*$. Hence if $v \in B_n^*$, then $v \notin S$, by Claim 3. So, it follows that

$$B_n^* \subseteq \binom{[n-1]}{k-1} - S.$$

Hence we have an upper bound for $|B_n^*|$ when we take $|S|$ to be minimised. By Lemma 6.4, $|S|$ is minimised when $\overline{A_{-n}}^*$ is first.

Claim 4. $|A_{-n}| \leq \binom{n-3}{k-2}$.

Proof. $|A_{-n}| = |\overline{A_{-n}}| = |\overline{A_{-n}}^*|$, so it is sufficient to show that $|\overline{A_{-n}}^*| \leq \binom{n-3}{k-2}$. Assume for the sake of contradiction that $|\overline{A_{-n}}^*| \geq \binom{n-3}{k-2} = \binom{n-3}{n-k-1}$.

Firstly, we show that $|S| \geq \binom{n-3}{k-1}$. It is sufficient to prove this lower bound when $|S|$ is minimised. Hence we can assume that $\overline{A_{-n}}^*$ is first, and contains the first $\binom{n-3}{n-k-1}$ $(n-k-1)$-sets in the colexicographic ordering. That is, it contains all $(n-k-1)$-sets on $[n-3]$. This is because there are $\binom{n-3}{n-k-1}$ such sets, and they come before all other sets in the ordering. In that case, S contains all $(k-1)$-sets in $[n-3]$. Since all of the $(k-1)$-sets in $[n-3]$ are in S, it follows that $|S| \geq \binom{n-3}{k-1}$, as required.

Then it follows that $|B_n^*| \leq \binom{n-1}{k-1} - \binom{n-3}{k-1} = \binom{n-3}{k-2} + \binom{n-2}{k-2}$. However, $|B_n^*| = |B_n| > \binom{n-3}{k-2} + \binom{n-2}{k-2}$ by Claim 2. This provides our desired contradiction. □

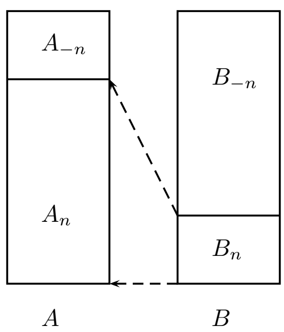

Figure 6.4: Diagram for Claim 4.

The basic idea is as follows. A large proportion of the vertices of B use the element n. If $v \in A$, then v must intersect all vertices of B, including all those that use element n. To do this, v can either use element n itself (which is, in some sense, the "easy" way), or v can intersect each vertex of B_n in another element (the "hard" way). The important fact, shown in Claim 4, it is not possible to have too many vertices intersect the vertices of B_n the "hard" way. This forces the proportion of the vertices of A using element n to be large.

Claim 5. $|A_n| \geq \frac{k}{k+1}|A|$.

Proof. First we show that $|A_n| \geq k|A_{-n}|$. Suppose otherwise, for the sake of a contradiction. By Claim 4, $|A| = |A_n| + |A_{-n}| < (k+1)|A_{-n}| \leq (k+1)\binom{n-3}{k-2}$. But $|A| \geq (1-c)|G - X|$. Hence $(1-c)\binom{n-1}{k-1} < (k+1)\binom{n-3}{k-2}$. Thus $(n-1)(n-2) < \frac{1}{1-c}(k+1)(k-1)(n-k) \leq \frac{1}{1-c}(k+1)(k-1)(n-2)$. Thus $n < \frac{1}{1-c}(k^2-1)+1$, which contradicts our lower bound on n. Then $|A_n| \geq k|A_{-n}| = k(|A| - |A_n|)$. So $(k+1)|A_n| \geq k|A|$ as required. $\square$

Given that a large proportion of the vertices of A use element n, the same principle holds for the vertices of B here as held for A above. By repeatedly using similar arguments, we force the proportion of the vertices using n in A and B to increase until all vertices of $A \cup B$ use element n.

Claim 6. $B_n = B$.

Proof. Suppose, for the sake of a contradiction, that there exists some vertex $v \in B$ such that $n \notin v$. So each $w \in A_n$ contains n (by definition) and some element of v (which is not n), since vw is not an edge. Any vertex of A_n can be constructed as follows—take element n, choose one of the k elements of v, and choose the remaining $k-2$ elements from the remaining $n-2$ elements of $[n]$. Thus

$$|A_n| \leq 1 \cdot k\binom{n-2}{k-2}.$$

Note this is actually a weak upper bound, since we have counted some of the vertices of A_n more than once. Recall $|A| \geq (1-c)|G-X| \geq (1-c)\binom{n-1}{k-1}$. So by Claim 5,

$$\frac{(1-c)k}{(k+1)}\binom{n-1}{k-1} \leq \frac{k}{k+1}|A| \leq k\binom{n-2}{k-2}.$$

Thus $\frac{n-1}{k-1} \leq \frac{1}{1-c}(k+1)$ and $n \leq \frac{1}{1-c}(k^2-1)+1$, which contradicts our lower bound on n. $\square$

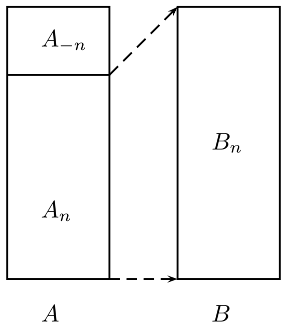

Figure 6.5: Diagram for Claim 6.

Claim 7. $A_n = A$.

Proof. This follows by essentially the same argument as Claim 6. Assume our claim does not hold and there exists $v \in A$ such that $n \notin v$. By Claim 6, $|B_n| = |B| \geq \frac{1}{2}\binom{n-1}{k-1}$. There is an upper bound on $|B_n|$ equal to the upper bound on $|A_n|$ in the proof of Claim 6. Then

$$\frac{1}{2}\binom{n-1}{k-1} \leq |B| = |B_n| \leq k\binom{n-2}{k-2},$$

and so $n \leq 2k(k-1)+1$. This contradicts our lower bound on n. □

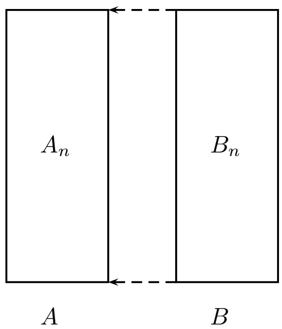

Figure 6.6: Diagram for Claim 7.

Claims 6 and 7 show that every vertex in $G - X = A \cup B$ contains n. Thus $|G - X| \leq \binom{n-1}{k-1}$ and $|X| \geq \binom{n-1}{k}$, our desired contradiction. □

By Lemma 6.7, if X is a $(|X|, V(G), \frac{2}{3})$-separator of the Kneser graph G and $n \geq 4k^2 - 4k + 3$, then $|X| \geq \binom{n-1}{k}$. Hence by Lemma 2.7, $\mathsf{tw}(G) \geq \binom{n-1}{k} - 1$. This proves Theorem 1.7.

Also, Lemma 6.7 allows us to prove Theorem 1.10.

Proof of Theorem 1.10. Let $C_1, \dots, C_r$ be the colour classes of H and recall $G = \mathsf{Kneser}(n, k)$. Let $X := V(\overline{G}) - V(H)$, so that $X, C_1, \dots, C_r$ is a partition of the vertex set of $\overline{G}$ (and also G). In G there are no edges between any pair C_i, C_j, and $|C_i| \leq c|H| = c|G - X|$ for each i. So X is a $(|X|, V(G), c)$-separator of G, and $|X| \geq \binom{n-1}{k}$ by Lemma 6.7. Hence $|H| \leq \binom{n-1}{k-1}$. □

6.5 Lower Bound for Treewidth when $k = 2$

To complete our proof of Theorem 1.8, we need to obtain a lower bound on the treewidth when $k = 2$. If $n \leq 4$, then Theorem 1.8 is trivial. When $n = 5$, then G is the Petersen graph, which contains a K_5-minor forcing $\mathsf{tw}(G) \geq 4$. Hence we may assume that $n \geq 6$.

Assume, for the sake of a contradiction that $\mathsf{tw}(G) < \binom{n-1}{2} - 1$. Let $(T, (B_x : x \in V(T)))$ be a minimum width tree decomposition for G, and normalise the tree decomposition as allowed by Lemma 2.2. By Lemma 2.7, there exists some $(\binom{n-1}{2} - 1, V(G), \frac{2}{3})$-separator X. In fact, by the proof of Lemma 2.7, we can go further and assert that X is a subset of a bag of $(B_x : x \in V(T))$.

Now $|G - X| = \binom{n}{2} - |X| > \binom{n-1}{1} = n - 1$. By Lemma 6.1, $V(G - X)$ can be partitioned into two parts A and B such that $\frac{1}{3}|G - X| \leq |A|, |B| \leq \frac{2}{3}|G - X|$ and there is no edge with an endpoint in A and B. (Note that this bound on $|A|$ and $|B|$ is slightly weaker than in Lemma 6.1, but has the benefit of being the same on both parts.) Since $n \geq 6$, it follows that $|A|, |B| \geq 2$. By Theorem 6.2, $V(G - X)$ is too large to be an independent set, and so it contains an edge, with both endpoints in A or both endpoints in B.

Without loss of generality this edge is $\{1, 2\}\{3, 4\} \in A$. Since every vertex in B must intersect both endpoints, $B \subseteq \{\{1, 3\}, \{1, 4\}, \{2, 3\}, \{2, 4\}\}$. If B contains an edge, then any other vertex in A or B must contain two elements of $\{1, 2, 3, 4\}$. So $V(G - X) \subseteq \{\{1, 2\}, \{1, 3\}, \{1, 4\}, \{2, 3\}, \{2, 4\}, \{3, 4\}\}$ and has maximum order 6. Otherwise, without loss of generality, $B = \{\{1, 3\}, \{1, 4\}\}$ and $A = \{\{3, 4\}, \{1, i\} | i \notin \{1, 3, 4\}\}$, so $|G - X| = n$. (Note A must be exactly that set, or $|G - X|$ is too small.)

If $n \geq 7$, then $|G - X| \geq 7$ and the first case cannot occur. However in the second case, $|B| = 2 < \frac{1}{3} \cdot 7 \leq \frac{1}{3}n$. So neither case can occur, and we have forced a contradiction on either $|G - X|$ or $|B|$. This completes the proof when $n \geq 7$. Hence, let $n = 6$, and note $|G - X| = 6$ in either case. In the first case, $G - X$ contains three disjoint matching of three edges, so that, without loss of generality, the endpoints of one edge is in A and then endpoints of the other two are in B. In the second case, the subgraph induced by A forms a star. In either case, A is connected.

Now we use the fact that X is a subset of some bag B_x. Now for all $x \in V(T)$, $|B_x| \leq \binom{5}{2} - 1 = 9$. Since $|G - X| = 6$, it follows $|X| = 9$. Hence X is exactly a bag of maximum order. Since A is a connected component for either choice of $G - X$, there is some subtree of $T - x$ that contains all vertices of A. Let y be the node of this subtree adjacent to x. Also note, for either choice of $G - X$, that each vertex of X has a neighbour in A. (In the first case, each vertex in X contains at most one element from $\{1, 2, 3, 4\}$ and so is adjacent to one of $\{1, 2\}, \{3, 4\}$. In the second case, no vertex of X uses element 1, and as such is adjacent to one of $\{1, 2\}, \{1, 5\}$ and $\{1, 6\}$.) So every vertex of B_x is also in bag B_y, which contradicts our normalisation.

Thus, if $n \geq 6$, then $\mathsf{tw}(G) \geq \binom{n-1}{2} - 1$. This completes the proof of Theorem 1.8.

6.6 A Weaker Lower Bound for Treewidth

We now extend our proof technique from Theorem 1.7 to prove Theorem 1.9.

The upper bound for Theorem 1.9 follows directly from Lemma 6.6. To prove the lower bound, we follow the same process as in Section 6.4 and show that a large separator is required. The following lemma is sufficient.

Lemma 6.8. *Let X be a $(|X|, V(G), \frac{2}{3})$-separator of the Kneser graph G. If $n \geq \frac{1}{2}(\sqrt{5k^2 - 12k + 8} + 3k + 2)$ and $k \geq 3$, then $|X| \geq \binom{n-1}{k} - \binom{n-1}{k-1}$.*

Proof. We assume, for the sake of a contradiction, that $|X| < \binom{n-1}{k} - \binom{n-1}{k-1}$. Recall by Lemma 6.1 that $V(G - X)$ can be partitioned into two parts A and B with no edge between A and B such that $\frac{1}{3}|G - X| \leq |A| \leq \frac{1}{2}|G - X|$ and $\frac{1}{2}|G - X| \leq |B| \leq \frac{2}{3}|G - X|$. Since $|X| < \binom{n-1}{k} - \binom{n-1}{k-1}$, it follows that $|G - X| > \binom{n}{k} - \binom{n-1}{k} + \binom{n-1}{k-1} = 2\binom{n-1}{k-1}$.

Now $|B| \geq \binom{n-1}{k-1}$. If $|A| > \binom{n-1}{k-1}$, then $|A||B| > \binom{n-1}{k-1}^2$, which contradicts Theorem 6.3. It follows that $|A| \leq \binom{n-1}{k-1}$.

As discussed in Lemma 6.7, if $w \in A$ and $v \in B$, then $w \nsubseteq \overline{v}$, since otherwise w and v do not intersect and there is an edge between A and B. Let S be the k-shadow of $\overline{A}$. Then

$$B \subseteq V(G) - S.$$

First, we consider some results about the order of A. Define the sequence of integers $t_2, \ldots, t_{n-k+2}$ as follows:

- $t_2 := \binom{n-2}{k-2}$,
- $t_i := t_{i-1} + \binom{n-i}{k-2}$.

Claim 1. $|A| > t_3$.

Proof. Recall that $|A| \geq \frac{1}{3}|G - X| > \frac{2}{3}\binom{n-1}{k-1}$. Assume for the sake of a contradiction that $|A| \leq t_3 = \binom{n-2}{k-2} + \binom{n-3}{k-2}$. Thus

$$\frac{2}{3}\binom{n-1}{k-1} < \binom{n-2}{k-2} + \binom{n-3}{k-2}$$

$$\frac{2}{3}(n-1)(n-2) < (k-1)(n-2) + (k-1)(n-k)$$

$$\frac{2}{3}(n^2 - 3n + 2) < 2kn - 2n - k - k^2 + 2.$$

$$n < \frac{1}{2}(\sqrt{3k^2 - 6k + 4} + 3k).$$

Hence we have our required contradiction, since this forces n to be smaller than its lower bound. □

Claim 2. $|A| \leq t_{n-k+2}$.

Proof. Recall that

$$|A| \leq \binom{n-1}{k-1} = \binom{n-2}{k-1} + \binom{n-2}{k-2} = \binom{n-2}{k-1} + t_2.$$

$$|A| \leq \binom{n-3}{k-1} + \binom{n-3}{k-2} + t_2 = \binom{n-3}{k-1} + t_3.$$

$$\vdots$$

$$|A| \leq \binom{n-(n-k+2)}{k-1} + t_{n-k+2} = \binom{k-2}{k-1} + t_{n-k+2} = t_{n-k+2}.$$

□

Since $t_3 < |A| \leq t_{n-k+2}$, and since $t_3 < \cdots < t_{n-k+2}$, we can fix i such that $t_{i-1} < |A| \leq t_i$. Then $i \geq 4$.

Claim 3. $|S| \geq \binom{n-2}{k} + \binom{n-3}{k-1} + \cdots + \binom{n-(i-1)}{k-(i-1)+2}$.

Proof. Since $|A| > t_{i-1}$ and $|A| = |\overline{A}|$, it follows that $|\overline{A}| > t_{i-1}$. Now S is minimised when $\overline{A}$ is *first*, by Lemma 6.4. Now we consider what $\overline{A}$ must contain when it is first. Note $t_{i-1} = \binom{n-2}{k-2} + \ldots \binom{n-(i-1)}{k-2} = \binom{n-2}{n-k} + \cdots + \binom{n-(i-1)}{n-k-i-1}$. The set $\overline{A}$ is a set of $(n-k)$-subsets, so it must contain all $(n-k)$-subsets on $\{1, \ldots, n-2\}$, of which there are $\binom{n-2}{n-k}$. Next in the ordering are the $(n-k)$ sets using $(n-1)$ but not n or $(n-2)$. There are $\binom{n-3}{n-k-1}$ of these, since $(n-1)$ is fixed in all of these sets and the remaining elements are chosen from $\{1, \ldots, n-3\}$. Subsequently, $\overline{A}$ contains the sets using $(n-1)$ and $(n-2)$ but not n or $(n-3)$, of which there are $\binom{n-4}{n-k-2}$. This follows in a logical fashion. Each of these sets corresponds to a element in t_{i-1} (where t_{i-1} is represented as a sum). Given this, we determine what S contains when $\overline{A}$ is first. From the first set of $\binom{n-2}{n-k}$ $(n-k)$-sets, the shadow contains all k-sets on $[n-2]$. There are $\binom{n-2}{k}$ of these sets. Next, given all $(n-k)$-sets using $(n-1)$ but not n or $(n-2)$, the shadow contains all k-sets using $(n-1)$ with the rest of the elements chosen from $\{1,\dots,n-3\}$, of which there are $\binom{n-3}{k-1}$. Note here we have only counted what is new to the shadow. Hence, when $\overline{A}$ is first, it follows that $|S| \geq \binom{n-2}{k} + \binom{n-3}{k-1} + \cdots + \binom{n-(i-1)}{k-(i-1)+2}$. This lower bound for S holds in general. $\square$

Thus it follows that

$$|B| \leq \binom{n}{k} - \binom{n-2}{k} - \binom{n-3}{k-1} - \cdots - \binom{n-(i-1)}{k-(i-1)+2}.$$

Note that

$$\begin{aligned}\binom{n-1}{k} &= \binom{n-2}{k} + \binom{n-2}{k-1}\\ &= \binom{n-2}{k} + \binom{n-3}{k-1} + \binom{n-3}{k-2}\\ &= \binom{n-2}{k} + \cdots + \binom{n-(i-1)}{k-(i-1)+2} + \binom{n-(i-1)}{k-(i-1)+1}.\end{aligned}$$

Hence it follows

$$\begin{aligned}|B| \leq &\binom{n-1}{k-1} + \binom{n-2}{k} + \cdots + \binom{n-(i-1)}{k-(i-1)+2} + \binom{n-(i-1)}{k-(i-1)+1}\\ &- \binom{n-2}{k} - \binom{n-3}{k-1} - \cdots - \binom{n-(i-1)}{k-(i-1)+2}\\ =&\binom{n-1}{k-1} + \binom{n-(i-1)}{k-(i-1)+1}.\end{aligned}$$

Also recall that

$$\binom{n-1}{k-1} - \binom{n-i}{k-1} = \binom{n-2}{k-2} + \cdots + \binom{n-i}{k-2} = t_i.$$

So $|A| \leq \binom{n-1}{k-1} - \binom{n-i}{k-1}$.

Thus $|A| + |B| \leq 2\binom{n-1}{k-1} + \binom{n-i+1}{k-i+2} - \binom{n-i}{k-1}$. Thus the following claim is sufficient to give our required contradiction.

Claim 4. $\binom{n-i+1}{k-i+2} - \binom{n-i}{k-1} \leq 0$.

Proof. Assume otherwise. Thus

$$\binom{n-i+1}{k-i+2} > \binom{n-i}{k-1}.$$

If $i \geq k+3$, then $\binom{n-i+1}{k-i+2} = 0$, which is a contradiction. Hence $i \leq k+2$. Thus we obtain from the above equation

$$(k-1)\dots(k-i+3)(n-i+1) > (n-k-1)\dots(n-i-k+2).$$

And so

$$(k-i+3)(n-i+1) > (n-k-1)(n-k-2)\frac{n-k-3}{k-1}\dots\frac{n-k-(i-2)}{k-(i-4)}.$$

Since $i \geq 4$, we get $(k-1)(n-3) > (n-k-1)(n-k-2)$. So $n < \frac{1}{2}(\sqrt{5k^2-12k+8}+3k+2)$, giving a contradiction with our lower bound on n. □

This completes the proof of Lemma 6.8. □

Thus we have achieved our desired contradiction and shown that $\mathsf{tw}(G) \geq \binom{n-1}{k} - \binom{n-1}{k-1} - 1$ when $n \geq \frac{1}{2}(\sqrt{5k^2-12k+8}+3k+2)$. This completes the proof of Theorem 1.9. However, recall that the minimum degree is a naïve lower bound on the treewidth of any graph. In our case, this gives $\mathsf{tw}(G) \geq \binom{n-k}{k}$. We now show that the lower bound we have constructed is actually an improvement.

Lemma 6.9. *If $n \geq 2k+1$ and $k \geq 3$, then $\binom{n-1}{k} - \binom{n-1}{k-1} - 1 \geq \binom{n-k}{k}$.*

Proof. Assume for the sake of a contradiction that $\binom{n-1}{k} - \binom{n-1}{k-1} \leq \binom{n-k}{k}$. This gives

$$(n-1)\dots(n-k) - k(n-1)\dots(n-k+1) \leq (n-k)\dots(n-2k+1),$$

$$(n-1)\dots(n-k+1)(n-2k) \leq (n-k)\dots(n-2k+1),$$

By cancelling $(n-2)$ with $(n-k)$, $(n-3)$ with $(n-(k+1))$ and so forth, we get the following.

$$(n-1)(n-2k) \leq (n-(2k-2))(n-(2k-1)),$$

$$(k-2)n \leq 2k^2-4k+1.$$

By substituting the lower bound for n, we get

$$2k^2-3k-2 \leq 2k^2-4k+1.$$

This gives a contradiction for $k > 3$. By a slightly longer calculation which we omit, there is also a contradiction when $k = 3$, as required. □

6.7 Open Questions

We conjecture that Theorem 1.7 should also hold for smaller values of n.

Conjecture 6.10. *Let G be a Kneser graph with $n \geq 3k$ and $k \geq 2$. Then $\mathsf{tw}(G) = \binom{n-1}{k} - 1$.*

This conjecture follows directly from Theorem 1.8 when $k = 2$. The Petersen graph also shows that $n \geq 3k$ is a tight bound when $k = 2$.

In general, we can determine a slightly better tree decomposition when $n \leq 3k - 2$. Let $X = \{v \in V(G) : 1 \in v\}$, and let W be an independent set in $V(G) - X$ such that no two vertices of W have a common neighbour in X. We define a tree decomposition for G with underlying tree T as follows. Let r denote the root node of T, and let r have one child node for each vertex in W and each vertex in X adjacent to no vertex in W. Label each of these child nodes by their associated vertex of G. Let each node labelled by a vertex $w \in W$ have one child node for each vertex of $N(w) \cap X$. Label each of those child nodes by their associated vertex of G, and note that since every vertex of X has at most one neighbour in W, no vertex of G labels more than one node of T.

Define the bag indexed by r to be $V(G) - W - X$. Note this bag contains less than $\binom{n-1}{k}$ vertices when $W \neq \emptyset$. If a node is labelled by a vertex $v \in X$, let the corresponding bag be $N(v) \cup \{v\}$. These bags contain $\binom{n-k}{k} + 1$ vertices. If a node is labelled by a vertex $w \in W$, let the corresponding bag be $\{w\} \cup \{u : uw \in E(G), 1 \notin u\} \cup \{u : ux \in E(G)$ where $xw \in E(G)$ and $1 \in x\}$. These bags contain less than $\binom{n-1}{k}$ vertices whenever $|W| \geq 2$, since they contain no vertex in X, and each contains only one vertex from W. This is a valid tree decomposition, but we omit the proof. When $|W| \geq 2$, the width of this tree decomposition is less than the width given by Lemma 6.5.

However, when $|W| \leq 1$, this tree decomposition has the same width as given by Lemma 6.5. We can construct W such that $|W| \geq 2$ iff $n < 3k - 1$. For example, let $W = \{\{2, \ldots, (k+1)\}, \{(k+1), \ldots, 2k\}\}$. If $n \leq 3k - 2$, then any vertex of X must be non-adjacent to at least one vertex of W. Alternatively, if $n \geq 3k - 1$ and $|W| \geq 2$, then there exists two vertices $x, y \in W$ such that $|x \cup y| \leq 2k - 1$. Then X contains a vertex adjacent to both x and y. Hence, for general n, we cannot improve the lower bound on n in Theorem 1.7 to $3k - 2$ or below. This does leave a question about what may occur for $n = 3k - 1$. It is possible that Theorem 1.7 holds for $n \geq 3k - 1$, with the Petersen graph as a single exception.

We now discuss possible strategies for proving results in the direction of Conjecture 6.10. Let X be a $(|X|, V(G), c)$-separator (where $c \in [\frac{2}{3}, 1)$) and say A and B are the parts of $V(G - X)$ by Lemma 6.1. In Lemma 6.7, we showed that $B_n^* \subseteq \binom{[n-1]}{k-1} - S$, where S is the $(k-1)$-shadow of $\overline{A_{-n}}^*$. Similarly, we can argue that $B \subseteq \binom{[n]}{k} - S$, where S is the k-shadow of $\overline{A}$. This is a key point in the papers by Pyber [80] and Matsumoto and Tokushige [78]. However, since A and B are disjoint, we could say that $B \subseteq \binom{[n]}{k} - (S \cup A)$, since B cannot contain a vertex in A, or a neighbour of such a vertex. However, we do not know what choice of A minimises $S \cup A$. (It can be seen that it is not the first $|A|$

sets under the colex order.)

If we improved Lemma 6.7 by determining the optimal A, then we would improve both Theorems 1.7 and 1.10. We believe that this should be possible, and that the optimal choice of A that minimises $S \cup A$ should be similar to (but not exactly) the first $|A|$ sets under colex order. This would hopefully allow an argument similar to that of Pyber, and give a lower bound on n that is linear in terms of k (ideally with a small constant).

Part II

Graph Minors

Chapter 7

Finding a Minor Quickly in Graphs with High Average Degree

7.1 Introduction

This chapter presents a linear-time algorithm for finding an H-minor in a graph with high average degree. Recall from Chapter 1 that

$$g(H) = \inf\{D : \text{every graph } G \text{ with average degree } d(G) \geq D \text{ contains an } H\text{-minor}\}.$$

We prove the following theorem.

Theorem 1.11. *For every fixed t-vertex graph H, there exists a $O(n)$ time algorithm that, given an n-vertex graph G with $d(G) \geq 2(g(H) + t)$, finds an H-minor in G.*

Given a t-vertex graph H, $g(H) \geq t-2$, since $d(K_{t-1}) = t-2$ but K_{t-1} cannot contain an H-minor. As a result, $2(g(H) + t)$ is bound above by a small constant factor of $g(H)$.

7.2 Algorithm

Given a vertex v of a graph G, we denote by $\deg_G(v)$ and $N_G(v)$ the degree and neighbourhood of v in G, respectively. We drop the subscript when G is clear from the context. Define a *matching* $M \subseteq E(G)$ to be a set of edges such that no two edges in M share an endpoint. Let $V(M)$ be the set of endpoints of the edges in M. An *induced matching* in G is a matching such that any two vertices x, y of $V(M)$ are only adjacent in G when $xy \in M$. Given a matching M in G, let G/M be the graph formed by contracting each edge of M in G.

We may assume that $t \geq 3$, since finding an H-minor efficiently is trivial when $t \leq 2$. Consider the following algorithm that takes as input a graph given as a list of vertices and a list of edges. The implicit output of the algorithm is the sequence of contractions and deletions that produce an H-minor.

Algorithm 1 FindMinor (input: n-vertex graph G with $d(G) \geq 2(g(H)+t)$)

1: Delete edges of G so that $2(g(H)+t) \leq d(G) \leq 2(g(H)+t)+1$.

2: Delete vertices of low degree so that the minimum degree $\delta(G) > \frac{1}{2}d(G)$.

3: Let $S := \{v \in V(G) : \deg(v) \leq d(G)^2\}$, and let $B := \{v \in V(G) : \deg(v) > d(G)^2\}$.
[Note that B is possibly empty, and that S and B partition $V(G)$.]

4: Say an edge $vw \in E(G)$ is *good* if $v, w \in S$ and $|N(v) \cap N(w)| \leq \frac{1}{2}(d(G)-2)$. Greedily construct a maximal matching M of good edges.
[Note that it is possible that no edges are good, in which case $M = \emptyset$.]

5: If $|M| > \frac{1}{8d(G)}n$, then greedily construct a maximal induced submatching M' of M. That is, initialise $M' := \emptyset$ and $Q := M$, and repeat the following algorithm until $Q = \emptyset$: pick an edge $vw \in Q$, add vw to M', and delete from Q the edge vw and every edge with an endpoint adjacent to v or w.
Let $G' := G/M'$. Run FindMinor(G') and stop.

6: Now assume $|M| \leq \frac{1}{8d(G)}n$. Let $B' := B \cup V(M)$ and $S' := S - V(M)$.
[Note that, similarly to Step 3, S' and B' partition $V(G)$.]

7: Greedily compute a maximal subset A of S' such that each vertex $u \in A$ is *assigned* to a pair of vertices in $N(u) \cap B'$, and each pair of vertices in B' has at most one vertex in A assigned to it.

8: If $2|A| \geq d(G)|B'|$ and $B' \neq \emptyset$, then let G' be the graph obtained from G as follows: For each pair of distinct vertices $x, y \in B'$ with an assigned vertex $z \in A$, contract the edge xz.
Run FindMinor($G'[B']$) and stop.

9: Now assume $2|A| < d(G)|B'|$ or $B' = \emptyset$. Choose $v \in S' - A$.
[We prove below that $S' - A \neq \emptyset$. Since v is not assigned, for every pair x, y of vertices in $N(v) \cap B$ some vertex $z \in A$ is assigned to x, y.]

10: If $|N(v) \cap B'| \geq t$, then let G' be the graph obtained from G as follows: For each pair of distinct vertices $x, y \in N(v) \cap B'$, if z is the vertex in A assigned to x and y, then contract xz into x (so that the new vertex is in B'). Then $G'[N(v) \cap B'] \supseteq K_t \supseteq H$. Stop.

11: Otherwise let $G' := G[\{v\} \cup (N_G(v) \cap S')]$ and run an exhaustive search to find an H-minor in G'.
[Below we prove that $d(G') > g(H)$ and $|V(G')| \leq d(G)^2 + 1$.]

7.3 Correctness of Algorithm

First, we prove that FindMinor(G) does output an H-minor. Define $m := |E(G)|$. We must ensure the following: that FindMinor finds an H-minor in Steps 5 and 8; that $S' - A \neq \emptyset$ in Step 9; that the graph constructed in Step 10 contains a K_t subgraph; and that our exhaustive search in Step 11 finds an H-minor of G.

Consider Step 5. Assume that FindMinor finds an H-minor in any graph G' with $|V(G')| < n$ where $d(G') \geq 2(g(H) + t)$. Consider the induced matching M'. Contracting any single edge vw of M' does not lower the average degree, since we only lose $|N(v) \cap N(w)| + 1 \leq \frac{1}{2}d(G)$ edges and one vertex. Since the matching is induced, contracting every edge in M' does not lower the average degree. Since $|M| > \frac{1}{8d(G)}n$, M is not empty and M' is not empty. Thus $d(G') \geq d(G) \geq 2(g(H) + t)$ and $|V(G')| < |V(G)| = n$. Thus, by induction, running the algorithm on G' finds an H-minor, and as such we find one for G.

If we recurse at Step 8, then $2|A| \geq d(G)|B'|$ and $B' \neq \emptyset$. Now $|V(G'[B'])| = |B'|$ and $|E(G'[B'])| \geq |A|$, since every assigned vertex corresponds to an edge of $G'[B']$. Thus

$$d(G'[B']) = \frac{2|E(G'[B'])|}{|V(G'[B'])|} \geq \frac{2|A|}{|B'|} \geq d(G).$$

Also, $|V(G'[B'])| = |B'| < n$, since otherwise $A = S' = \emptyset$, contradicting $2|A| \geq d(G)|B'| > 0$. Hence, by assumption, the algorithm will find an H-minor in $G'[B']$. Thus the algorithm finds an H-minor for G.

Now we show that $|S'| > |A|$ in Step 9. We have $2|A| < d(G)|B'|$ or $B' = \emptyset$. First consider the case when $2|A| < d(G)|B'|$. Note that $2m = d(G)n$, and that $d(G)^2|B| < \sum_{v \in B} \deg(v) \leq 2m$, and so $|B| < \frac{2m}{d(G)^2} = \frac{1}{d(G)}n$. Now $|S'| = |S| - 2|M| \geq |S| - \frac{1}{4d(G)}n$ by Step 6. Thus,

$$|S'| \geq |S| - \frac{1}{4d(G)}n = (n - |B|) - \frac{1}{4d(G)}n > n - \frac{1}{d(G)}n - \frac{1}{4d(G)}n = \frac{4d(G) - 5}{4d(G)}n.$$

By Step 9 and Step 6,

$$|A| < \frac{d(G)}{2}|B'| = \frac{d(G)}{2}(|B| + 2|M|) < \frac{d(G)}{2}\left(\frac{1}{d(G)}n + \frac{1}{4d(G)}n\right) = \frac{5}{8}n.$$

Thus, if $|S'| \leq |A|$ then $\frac{4d(G)-5}{4d(G)}n < \frac{5}{8}n$, so $3d(G) < 10$. This is a contradiction since $d(G) \geq 2(g(H) + t) \geq 2t \geq 4$. Hence, $|S'| > |A|$. Now consider the case that $B' = \emptyset$. Then $|S'| = n$ and $A = \emptyset$, since the vertices of A are assigned to pairs of vertices in B'. Hence $|S'| > |A|$.

Now consider Step 10. The subgraph $G'[N(v) \cap B']$ contains at least t vertices by assumption. Each pair of distinct vertices x, y in $N(v) \cap B'$ has an assigned vertex in A, since otherwise v would have been assigned to x and y. Hence the vertex z exists, and

x and y are adjacent after contracting xz. Therefore all pairs of vertices in $N(v) \cap B'$ become adjacent, and $G'[N(v) \cap B']$ is a complete graph. Hence we have found a K_t-minor in G, and our desired H-minor is simply a subgraph of this K_t-minor.

Finally consider Step 11. Since G' is an induced subgraph of G, if we can find H as a minor in G', we have an H-minor in G. We use an exhaustive search, so all we need to ensure is that G' does contain an H-minor. Thus, we simply need to ensure that $d(G') > g(H)$. By Step 1 and Step 2, $\deg_G(v) > \frac{1}{2}d(G) \geq g(H) + t \geq t$. Since Step 10 was not applicable, v has at most $t-1$ neighbours in B'. Thus v has some neighbour in S'. Let w be a vertex of $G' - v$. Thus vw is an edge and $v, w \in S'$. Since neither v nor w was matched by M, and since M is maximal, vw is not good. Since $v, w \in S' \subseteq S$, this means that $|N(v) \cap N(w)| > \frac{1}{2}(d(G) - 2)$. Since v has at most $t-1$ neighbours in B', we have $|N(v) \cap N(w) \cap S'| > \frac{1}{2}(d(G) - 2) - (t-1)$. Every common neighbour of v and w in S' is a neighbour of w in G', by definition, so $\deg_{G'}(w) > \frac{1}{2}(d(G) - 2) - (t-1)$. Since v is dominant in G', $d(G') \geq \delta(G') > \frac{1}{2}(d(G) - 2) - (t-1) \geq \frac{1}{2}(2(g(H) + t) - 2) - (t-1) \geq (g(H) + t - 1) - (t-1) = g(H)$, as required.

7.4 Time Complexity

Now that we have shown that FindMinor will output an H-minor, we must ensure it does so in $O(n)$ time (for fixed H).

First, suppose FindMinor runs without recursing. Recall that our input graph G is given as a list of vertices and a list of edges, from which we will construct adjacency lists as it is read in. Since our goal in Step 1 is to ensure that $m \leq \frac{1}{2}(2(g(H)+t)+1)n$, we can do this by taking, at most, the first $\frac{1}{2}(2(g(H)+t)+1)n$ edges, and ignoring the rest. This can be done in $O(n)$ time, and from now on we may assume that $m \in O(n)$. In Step 2, since we are only deleting vertices of bounded degree, this can be done in $O(n)$ time. Clearly, Steps 3, 6 and 9 can be implemented in $O(n)$ time. By definition, the degree of any vertex in S or S' is at most $(2(g(H)+t)+1)^2$. Hence Steps 4, 5, 7, 8 and 10 take $O(n)$ time. Finally, for Step 11 note that $|V(G')| \leq d(G)^2 + 1$, so exhaustive search runs in $O(1)$ time for fixed H. Hence the algorithm without recursion runs in $O(n)$ time.

Should FindMinor recurse, we need to ensure that the order of the graph we recurse on is a constant factor less than n. Then the overall time complexity is $O(n)$ (by considering the sum of a geometric series). In Step 5, the endpoints of edges in M have degree less than or equal to $d(G)^2$, and so $|M'| \geq \frac{1}{2d(G)^2}|M| \geq \frac{1}{16d(G)^3}n$. This ensures that $|V(G')| \leq (1 - \frac{1}{16d(G)^3})n$, as desired. In Step 8, the order of $G'[B']$ is at most $\frac{2|A|}{d(G)} \leq \frac{2n}{d(G)}$. Hence it follows that the overall time complexity is $O(n)$.

Chapter 8

Hadwiger's Conjecture for Circular Arc Graphs

8.1 Introduction

Recall a circular arc graph G is an intersection graph where the vertex set is a collection of arcs on a circle. Also recall the cover number $\beta(G)$ is the size of the smallest set of arcs in $V(G)$ which cover the entire circle. If no set of arcs cover the entire circle, then $\beta(G) = \infty$, and G is an interval graph. A normal Helly circular arc graph G is a circular arc graph for which $\beta(G) > 3$.

In this chapter, we prove the following weakening of Hadwiger's Conjecture:

Theorem 1.12. *For a normal Helly circular arc graph G, $\mathsf{had}(G) \geq \chi(G) - 1$.*

To prove this, we let G be a vertex-minimum counterexample, that is, $\mathsf{had}(G) < \chi(G) - 1$ and $\mathsf{had}(G') \geq \chi(G') - 1$ for every circular arc graph G' with $\beta(G') > 3$ and less vertices than G. Given that (as shown in Chapter 1) Hadwiger's Conjecture holds for interval graphs, we can assume that $\beta(G)$ is finite. (If $\beta(G) = \infty$, then there is a point on the circle with no arcs—we can "cut" the circle at this point and convert G into an interval graph, for which Hadwiger's Conjecture holds.)

We shall show that either $\mathsf{had}(G) \geq \chi(G) - 1$ or that we can colour G with less than $\chi(G)$ colours, either of which form a contradiction.

8.2 Preliminaries

For a circular arc graph G, recall the maximum load $\mathcal{L}(G)$ is the maximum number of arcs at any point of the circle. For simplicity, let $\mathcal{L} := \mathcal{L}(G)$. Since all of these arcs intersect, this forces the existence of a clique of order $\mathcal{L}(G)$, and so $\chi(G) \geq \mathcal{L}(G)$. Let q be a point of maximum load, and let Q be the vertex set of the clique at q. Define the interval graph H such that $H := G - Q$.

From now on, we choose to think of G as like an interval graph in the following sense: "cut" the circle at point q and straighten it to obtain a line. Then G is the interval graph H, plus the vertices of Q, each of which is represented by two intervals—one starting at $-\infty$ and one ending at $+\infty$. For such a vertex $u \in Q$, call the interval at starting $-\infty$ the *left interval of* u and the interval ending at $+\infty$ the *right interval of* u. Denote the vertices of Q as *Q-vertices* and the remaining vertices as *H-vertices*. For an H-vertex v, define $l(v)$ to be the left endpoint of the interval, and $r(v)$ to be the right endpoint. For a Q-vertex u, let $r(u)$ be the right endpoint of the left interval of u, and $l(u)$ be the left endpoint of the right interval. The left endpoint of the left interval is always $-\infty$ and the right endpoint of the right interval is $+\infty$, so we do not need to denote these specifically. It is well known that we can assume all intervals have distinct endpoints (except for the endpoints $-\infty$ and $+\infty$), since we can perturb the endpoints of an interval to ensure this. For two points p and r on the line, denote p is left of r by writing $p < r$. We say an H-vertex v *covers* an H-vertex w if $l(v) < l(w)$ and $r(w) < r(v)$. A Q-vertex v *covers* an H-vertex w if $r(w) < r(v)$ or $l(v) < l(w)$. A Q-vertex v *covers* a Q-vertex w if $r(w) < r(v)$ and $l(v) < l(w)$.

Define a *small vertex* v of G to be a vertex such that there is no vertex w covered by v. Then call any other vertex *large*. For each large vertex v there is a small vertex w such that w is covered by v.

Define $k := \chi(G) - \mathcal{L}(G)$.

A graph is *colour critical* if any vertex deletion causes the chromatic number to decrease. This concept was first introduced by Dirac [24].

Lemma 8.1. *If G is the vertex-minimum counterexample to Theorem 1.12, then G is colour critical.*

Proof. Say that G is not colour critical. Then there exists $v \in V(G)$ such that $\chi(G - v) = \chi(G)$. Then note that $\mathsf{had}(G - v) \leq \mathsf{had}(G)$, since vertex deletion is a valid operation when constructing a minor. Also, $\beta(G - v) \geq \beta(G) > 3$, since any set of vertices covering the circle in $G - v$ also covers the circle in G. So $\mathsf{had}(G - v) \leq \mathsf{had}(G) < \chi(G) - 1 = \chi(G - v) - 1$. Thus $G - v$ is a smaller counterexample, contracting our assumption that G is a vertex-minimum counterexample. $\square$

Every colour critical graph has minimum degree $\delta(G) \geq \chi(G) - 1$. So Lemma 8.1 implies the following:

Corollary 8.2. *If G is the vertex-minimum counterexample to Theorem 1.12, then the minimum degree of G is $\delta(G) \geq \chi(G) - 1 = \mathcal{L} + k - 1$.*

If u is a Q-vertex and v is an H-vertex, then we say that u is a *left-Q-neighbour* of v if $uv \in E(G)$ and the interval of v intersects the left interval of u. In this case we also say v has a left-Q-neighbour. Similarly we define *right-Q-neighbour.* Lemma 8.3 through to Corollary 8.5 prove some basic but important results about the structure of the neighbourhood of a vertex of H.

Lemma 8.3. *Let G be the vertex-minimum counterexample to Theorem 1.12, let Q be the set of vertices at a point of maximum load q and let $H := G - Q$. Then no H-vertex v has both a left-Q- and right-Q-neighbour.*

Proof. Say $v \in V(H)$ has a left-Q-neighbour u and a right-Q-neighbour w. Then $\{v, u, w\}$ cover the entire circle, contradicting $\beta(G) > 3$. □

Lemma 8.4. *Let G, Q, H be as in Lemma 8.3. If v is an H-vertex with no right-Q-neighbour, then v has at least k H-neighbours that are at $r(v)$ and not at $l(v)$, which we call right-only-H-neighbours.*

Proof. First suppose v is small. Then every neighbour of v is either at $l(v)$ and/or $r(v)$. Since $\deg(v) \geq (\mathcal{L}+k-1)$ by Corollary 8.2, and since v can have at most $(\mathcal{L}-1)$ neighbours at $l(v)$, v has at least k neighbours at $r(v)$ that are not at $l(v)$. Every Q-neighbour of v, if there are any, is at $l(v)$ since left-Q-neighbours come from $-\infty$. Thus v has k H-neighbours at $r(v)$ and not at $l(v)$.

Alternatively, v is large. Let u be the vertex covered by v with rightmost left endpoint. Now u is small—if u covers some w, then v covers w and $l(w) > l(u)$, contradicting the choice of u. Similarly, each vertex w at $r(u)$ and not at $l(u)$ has $l(w) > l(u)$, so w must be at $r(v)$. Since u has k right-only-H-neighbours, so does v. □

By symmetry we have:

Corollary 8.5. *Let G, Q, H be as in Lemma 8.3. If v is an H-vertex with no left-Q-neighbour, then v has at least k H-neighbours that are at $l(v)$ and not at $r(v)$, which we call left-only-H-neighbours.*

Definition Let $s_1, \ldots, s_{k-1}$ be the first $(k-1)$ vertices of H by left endpoint, where $l(s_1) < l(s_2) < \cdots < l(s_{k-1})$. Call this set S. Let $t_1, \ldots, t_{k-1}$ be the last $(k-1)$ vertices by right endpoint, where $r(t_1) < r(t_2) < \cdots < r(t_{k-1})$. Call this set T.

In Section 8.3, we shall attempt to construct a series of paths starting in S and ending in T. Before doing so, we prove a series of useful results about S and T.

Lemma 8.6. *Let G, Q, H be as in Lemma 8.3. If u is a vertex in Q, then u is adjacent to at least one of s_i and t_i for all $i \in \{1, \ldots, k-1\}$.*

Proof. Say $u \in Q$ is not adjacent to s_i and t_i. Then the left interval of u is adjacent to at most $s_1, \ldots, s_{i-1}$ from H, and the right interval of u is adjacent to at most $t_{i+1}, \ldots, t_{k-1}$ from H. The vertex u also has $\mathcal{L} - 1$ neighbours in Q. Thus, $\deg(u) \leq (i-1) + ((k-1)-i) + (\mathcal{L}-1) = \mathcal{L} + k - 3 < \mathcal{L} + k - 1 \leq \delta(G)$ by Corollary 8.2, which is the desired contradiction. □

Lemma 8.7. *Let G, Q, H be as in Lemma 8.3. Each vertex of S has a left-Q-neighbour, and each vertex of T has a right-Q-neighbour.*

Proof. Say s_i has no left-Q-neighbour. Then s_i has k left-only-H-neighbours by Corollary 8.5. However, the only plausible left-only-H-neighbours are $s_1, ..., s_{i-1}$, of which there are less than k. A similar argument holds for t_i. □

By Lemma 8.7 and Lemma 8.3, we get the following:

Corollary 8.8. *Let G, Q, H be as in Lemma 8.3. The sets S and T are vertex disjoint.*

Lemma 8.9. *Let G, Q, H be as in Lemma 8.3. Both S and T are cliques.*

Proof. Say $s_i s_j \notin E(G)$, for some $i < j$. By Lemma 8.7 and Lemma 8.3, s_i has no right-Q-neighbour, so s_i has k right-only-H-neighbours by Lemma 8.4. However, since s_j is not adjacent to s_i, the only possible right-only-H-neighbours of s_i are $s_{i+1}, ..., s_{j-1}$, of which there are less than k. A similar argument holds for T. □

Label the vertices of Q twice, as follows. First, label using $q_1, \ldots, q_{\mathcal{L}}$ such that $r(q_1) < r(q_2) < \cdots < r(q_{\mathcal{L}})$ with respect to the left interval of each of these vertices. Second, label using $q'_1, \ldots, q'_{\mathcal{L}}$ such that $l(q'_{\mathcal{L}}) > l(q'_{\mathcal{L}-1}) > \cdots > l(q'_1)$ with respect to the right interval of each of these vertices.

Define the sets $S_i = \{q_1, \ldots, q_i\}$ and $T_i = \{q'_{\mathcal{L}}, q'_{\mathcal{L}-1}, \ldots, q'_{\mathcal{L}-(k-1)+i}\}$. No vertex of S_i intersects s_i and no vertex of T_i intersects t_i, as shown below.

Lemma 8.10. *Let G, Q, H be as in Lemma 8.3. No vertex in S_i intersects s_i, and no vertex in T_i intersects t_i.*

Proof. Since S is a clique by Lemma 8.9, we know that at $l(s_i)$, the vertices $s_1, \ldots, s_i$ are all present. Hence, at most $(\mathcal{L} - i)$ Q-vertices are at this point, else the load here is greater than the maximum load. Thus, if some $q_j \in S_i$ is at this point, then so are $q_{j+1}, ..., q_{\mathcal{L}}$ by how we have chosen our labels, and so the load at $l(s_i)$ is greater than $\mathcal{L}$. Therefore no vertex in S_i is at $l(s_i)$, and since s_i has a left-Q-neighbour by Lemma 8.7 and no H-vertex has both a left- and right-Q-neighbour by Lemma 8.3, $S_i \cap N(s_i) = \emptyset$. Again, a similar argument holds for t_i. $\square$

By Lemma 8.6 and Lemma 8.10, we get the following:

Corollary 8.11. *Let G, Q, H be as in Lemma 8.3. No vertex of Q is in S_i and T_i*

8.3 Special Path Sets

Recall, our goal is to show that vertex-minimum counterexample G either contains a complete minor of order $\chi(G) - 1 = \mathcal{L} + k - 1$, or has a re-colouring with one less colour (that is, with at most $\mathcal{L} + k - 1$ colours). Firstly, we shall try to construct a minor. We describe this minor using the terminology of models, which can be found in Section 2.2. In our intended model, our branch sets will be the $\mathcal{L}$ vertices of our clique Q, and a set of $(k-1)$ paths which link s_i and t_i by travelling around the circle. If successful, this will give us a minor of order $\mathcal{L} + k - 1$; see Lemma 8.12. We construct these paths as follows:

Definition A *special path set* $\mathcal{P}$ is a set of paths $P_1, \ldots, P_{k-1}$ in H (where $k = \chi(G) - \mathcal{L}(G)$) that satisfy the following properties:

(P1) $P_1, ..., P_{k-1}$ are pairwise vertex-disjoint.
(P2) $s_i \in P_i$ for all i, and $t_j \notin P_i$ for all $i \neq j$.

If $t_i \in P_i$, then call P_i *finished*, otherwise it is *unfinished*.

Lemma 8.12. *Let G be the vertex-minimum counterexample to Theorem 1.12, let Q be the set of vertices at a point of maximum load q and let $H := G - Q$. Then there does not exist a special path set $\mathcal{P}$ such that each $P_i \in \mathcal{P}$ is finished.*

Proof. Assume, for the sake of a contradiction, such a $\mathcal{P}$ does exist. To find our contradiction, we show $\mathsf{had}(G) \geq \mathcal{L} + k - 1 = \chi(G) - 1$. Our $K_{\mathcal{L}+k-1}$-model will contain the following branch sets:

- $Q_i := \{q_i\}$ for all $i \in \{1, ..., \mathcal{L}(G)\}$
- P_i from $\mathcal{P}$ for $i \in \{1, ..., (k-1)\}$

Each of these sets is clearly connected, so it remains to show that they are pairwise adjacent. Since $q_i \in Q_i$ and $q_j \in Q_j$, and $q_i, q_j \in Q$, which is by definition a clique, these sets are clearly adjacent. Also, since $s_i \in P_i$ and $s_j \in P_j$ and S is a clique by Lemma 8.9, P_i and P_j are adjacent. Finally, consider some P_i and some Q_j. Now, q_j is adjacent to either s_i or t_i by Lemma 8.6, so since P_i is finished and $t_i \in P_i$, it follows that Q_j and P_i are adjacent. $\square$

So, if all of our paths are finished, G is not a counterexample. Thus, we can assume at least one path in $\mathcal{P}$ is unfinished. We use the fact that $\mathcal{P}$ contains an unfinished path to determine some key facts about the structure of G, which will help with the recolouring.

Define the point p on the line to be the last point (from left to right) that has all $(k-1)$ paths of $\mathcal{P}$ present. That is, p is the endpoint of the first path in $\mathcal{P}$ to "stop". Choose our $\mathcal{P}$ so that:

(C1) The point p, as defined above, is as far to the right as possible.
(C2) Subject to (C1), $|V(P_1) \cup \cdots \cup V(P_{k-1})|$ is minimised. When there is no ambiguity, we shall denote $|V(P_1) \cup \cdots \cup V(P_{k-1})| = V(\mathcal{P})$.
(C3) Subject to (C1) and (C2), for each $P_i \in \mathcal{P}$ the j^{th} vertex v in P_i has $l(v)$ as far to the left as possible.

Call a vertex on some $P_i \in \mathcal{P}$ a *path vertex*. Any remaining vertex of H that is not some t_i we shall call a *free vertex*.

Lemma 8.13. *Let G, Q, H be as in Lemma 8.12, and let $\mathcal{P}$ be a special path set chosen with respect to (C1),(C2) and (C3). Every $P_i \in \mathcal{P}$ is an induced path, and no vertex of P_i is covered by any other vertex of P_i.*

Proof. Say P_i is not induced. Then we can clearly take an induced path in P_i from s_i to the last vertex of P_i by right endpoint. Call this path P_i^*, and construct a special path set $\mathcal{P}^*$ from $\mathcal{P}$ by replacing P_i with P_i^*. Since $s_i \in P_i^*$, and we only removed vertices from P_i^*, $\mathcal{P}^*$ is a special path set. The path P_i^* travels as far along the interval as P_i, so the position of p has not changed. However $P_i^* \subsetneq P_i$, so $|V(\mathcal{P}^*)| < |V(\mathcal{P})|$, contradicting our choice of $\mathcal{P}$.

Similarly, if we have some $v, w \in P_i$ such that v is covered by w, then since P_i is induced, either v follows directly after w in P_i or vice versa. However, since any neighbour of v is a neighbour of w, v has no other neighbours in P_i, so v is either the first vertex or the last vertex of P_i. It is not the case that $v = s_i$ or $v = t_i$, since such vertices can not be covered by another vertex of the same path, by definition. Hence v is the last vertex of P_i, but $v \neq t_i$. If we remove v from P_i then, as above, p has not moved but the number of vertices on P_i has fallen. As before, this contradicts our choice of $\mathcal{P}$. □

Lemma 8.14. *Let G, Q, H and $\mathcal{P}$ be as in Lemma 8.13. Then p is at the end of an unfinished path.*

Proof. The point p is certainly at the end of some path P_i by the second half of Lemma 8.13, so say P_i is finished. Thus $p = r(t_i)$. Since not all of these paths are finished, there is an unfinished path P_j that ends at a vertex a_j. Thus $r(a_j) > p$, but this means that $a_j \in T$. Hence either $a_j = t_j$, and P_j is finished, or $a_j \neq t_j$ and this is not a special path set. □

From this point forward, we say that p is the endpoint of an unfinished path P_i, and a_i is the last vertex of P_i; that is, $p = r(a_i)$. In Lemma 8.15 to Lemma 8.20 we prove some useful facts about the structure of the paths of $\mathcal{P}$ and the free vertices. The basic idea behind these results is that, if they did not hold, then it would be possible to construct a "better" set of paths $\mathcal{P}$. These results will help with recolouring in Section 8.4.

Lemma 8.15. *Let G, Q, H and $\mathcal{P}$ be as in Lemma 8.13. There is neither a free vertex nor the vertex t_i at p.*

Proof. Let u be a vertex at p such that u is either a free vertex or $u = t_i$. Hence u is not on any path, by the definition of a free vertex, and by (P2). Then add u on to path P_i and rename it P_i^*. Keep all the other paths the same, and call this new set $\mathcal{P}^*$. Now $\mathcal{P}^*$ satisfies (P1) since we only added vertices to P_i that were not on any other path. Also $\mathcal{P}^*$ satisfies (P2) since we did not place any $t_j (j \neq i)$ onto P_i. Thus $\mathcal{P}^*$ is a special path set. However, P_i now ends at $r(u)$, which is to the right of p. All other paths still end right of p. Thus $\mathcal{P}^*$ contradicts our choice of $\mathcal{P}$. □

Lemma 8.16. *Let G, Q, H and $\mathcal{P}$ be as in Lemma 8.13. The number of H-vertices at point p is at least $(k+1)$.*

Proof. Define $p+\epsilon$ to be the point on the line immediately to the right of p, that is, before any other endpoint of any other interval. The load at $p+\epsilon$ is one less than the load at p, and every vertex at $p+\epsilon$ is at p. Let v be the last vertex by left endpoint such that $r(v) < p+\epsilon$. Let w be the first vertex by right endpoint such that $l(w) > p+\epsilon$. Such a w exists since $l(t_i) > p+\epsilon$. Then either v has no right-Q-neighbour or w has no left-Q-neighbour—otherwise both w and v have both left- and right-Q-neighbours, contradicting Lemma 8.3. Say that v has no right-Q-neighbour. Then v has k right-only-H-neighbours, by Lemma 8.4. Each of these vertices must be at $p+\epsilon$ by the choice of v, similar to the second half of Lemma 8.4. There are k H-vertices at $p+\epsilon$, and thus $k+1$ at point p including a_i. If w has no left-Q-neighbours, the result also follows by symmetry, using Corollary 8.5. □

We say a path P_j of our special path set *appears twice at point r* if there are two vertices of P_j at point r. Note that since all paths in $\mathcal{P}$ are induced, there can be at most two vertices of P_j at any given point. (If there are three vertices of P_j at a point, then path P_j contains a triangle.)

Corollary 8.17. *Let G, Q, H and $\mathcal{P}$ be as in Lemma 8.13. Some path $P_j \neq P_i$ appears twice at point p.*

Proof. By Lemma 8.16, there are at least $(k+1)$ H-vertices at p, but from Lemma 8.15, none of these vertices can be free, or t_i. Since there are only $(k-1)$ paths, this means some path must appear more than once at p, and it cannot be P_i. (In fact, there are at least two paths that appear twice at p.) □

Lemma 8.18. *Let G, Q, H and $\mathcal{P}$ be as in Lemma 8.13. If P_j appears twice at point p, then one of the two vertices of P_j at p is t_j.*

Proof. Say P_j appears twice at p and let a_j and b_j be the two vertices of P_j at p, such that a_j is before b_j on P_j, and $b_j \neq t_j$. Now remove b_j and anything later from P_j, and add b_j to P_i. Rename this path P_i^*. So P_i^* remains connected since b_j is at p. Also relabel P_j as P_j^*, now that it only goes as far as a_j. Label by $\mathcal{P}^*$ the path set $\mathcal{P}$ with P_i, P_j replaced by P_i^*, P_j^* respectively. Now, $\mathcal{P}^*$ satisfies (P1) trivially—we only deleted vertices after and including b_j from P_j before adding b_j to P_i^*. Also, since $b_j \neq t_j$, and s_j is before a_j on P_j, (P2) still holds. Now P_i^* ends at $r(b_j)$ and P_j^* ends at $r(a_j)$, both of which are after p. All other paths remain unchanged and end after p. Thus $\mathcal{P}^*$ contradicts the choice of $\mathcal{P}$. □

Corollary 8.19. *Let G, Q, H and $\mathcal{P}$ be as in Lemma 8.13. There is a finished path P_j that appears twice at p, one of the vertices of P_j is t_j.*

From now on, we shall refer to a path satisfying Corollary 8.19 as a *blocking path*, since it blocks point p from being any further to the right.

Definition Consider the graph formed from G by deleting Q, all of $\mathcal{P}$, and all remaining t_i, leaving only free vertices. Call a component in this graph a *free component*. A free component is a connected subgraph of H.

Lemma 8.20. *Let G, Q, H and $\mathcal{P}$ be as in Lemma 8.13. There is no free component that is adjacent to two vertices from P_i and two vertices from $P_j - t_j$, where P_j is a blocking path.*

Proof. Assume some free component U is adjacent to two vertices from P_i and two vertices from $P_j - t_j$. Now, since U covers a connected part of the line, it follows that U is adjacent to connected subsets of P_i and $P_j - t_j$, of order greater than one. Let x_i be the first vertex of P_i adjacent to U, and y_i the vertex after it in P_i. Define x_j, y_j similarly for $P_j - t_j$. As before, let a_i be the last vertex of P_i and a_j the last vertex of $P_j - t_j$; both of these vertices and t_j are at p by Lemma 8.14 and the definition of a blocking path. To show a contradiction, we again construct a better path set than $\mathcal{P}$. Now since U defines a connected part of the line the same way that a vertex does, we can consider the point $l(U)$, the left end point of U. Since no vertex of S is free, $l(U) > l(s_{k-1})$. Hence x_i, x_j are at $l(U)$. Firstly, partition $P_j - t_j$ into the following subpaths: $P_j^1 := (s_j, ..., x_j)$ and $P_j^2 := (y_j, ..., a_j)$. Similarly let $P_i^1 := (s_i, ..., x_i)$ and $P_i^2 := (y_i, ..., a_i)$. There are two cases to consider, depending on the relationship between the right endpoints of x_i and x_j:

Case 1: $r(x_i) < r(x_j)$. Note that in this case x_j and y_i are adjacent. Define P_j^* and P_i^* as follows:

- $P_j^* = P_j^1 \cup P_i^2 \cup t_j$.
- $P_i^* = P_i^1 \cup U^* \cup P_j^2$, where U^* is a path through U from a vertex adjacent to x_i to one adjacent to y_j.

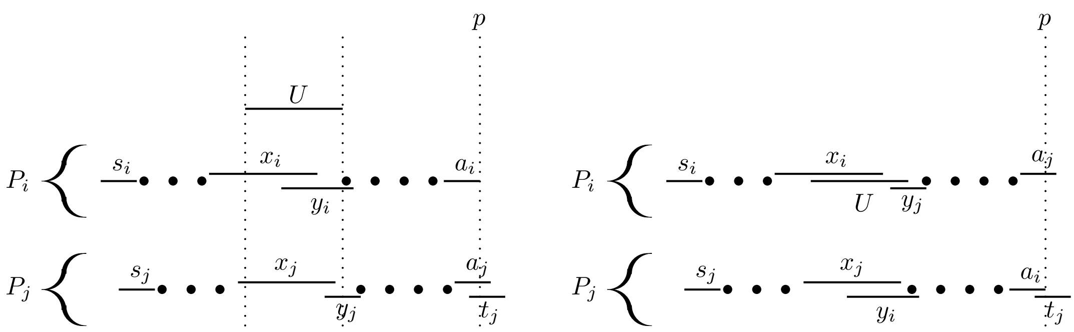

Figure 8.1: An illustration of Case 1, showing the paths before and after our rearrangement. The "second halves" of P_i and P_j swap with each other, and vertices of U are used to maintain connectivity. (In this figure U is just a single vertex, but the principle holds when U is a connected subgraph.) After this change, all paths extend past the point p.

Now, since x_j is adjacent to y_i and a_i is adjacent to t_j at p, P_j^* is a connected path. Similarly since U is adjacent to x_i and y_j, P_i^* is a connected path. Now let $\mathcal{P}^*$ be the path set formed by replacing $P_i, P_j \in \mathcal{P}$ with P_i^*, P_j^*. (P1) holds since we partitioned $P_j - t_j$ and P_i, meaning no vertex appears in both P_j^* and P_i^*. (P2) holds since $s_i \in P_i^1 \subset P_i^*$ and $s_j \in P_j^1 \subset P_j^*$, and since P_i is unfinished and we only considered $P_j - t_j$, we ensure no vertex of T is placed on the wrong path. Thus $\mathcal{P}^*$ is a special path set. Since P_i^* ends at either $r(a_j)$ and P_j^* ends at $r(t_j)$, all of which are further right than p, it follows that $\mathcal{P}^*$ contradicts the choice of $\mathcal{P}$.

Case 2: $r(x_j) < r(x_i)$. Note that here x_i is adjacent to y_j. Define P_j^* and P_i^* as follows:

- $P_j^* = P_j^1 \cup U^* \cup P_i^2 \cup t_j$, where U^* is a path through U from a vertex adjacent to x_j to one adjacent to y_i.
- $P_i^* = P_i^1 \cup P_j^2$.

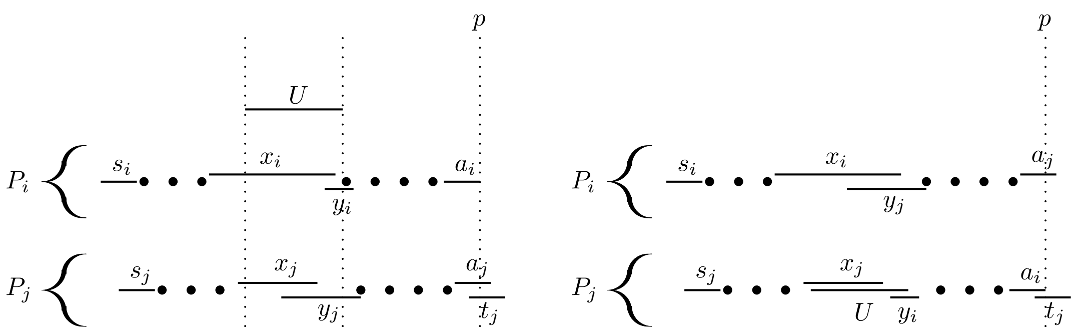

Figure 8.2: Case 2 is similar to Case 1, except the free vertices are used in the other path.

Define $\mathcal{P}^*$ as before. By an almost identical argument to the above, we find $\mathcal{P}^*$ contradicts the choice of $\mathcal{P}$. □

Lemma 8.21. *Let G, Q, H and $\mathcal{P}$ be as in Lemma 8.13. Every vertex entirely after p is in T.*

Proof. By Lemma 8.15, there are no free vertices at point p. Now let u be the first H-vertex by right endpoint entirely after p.

We claim u has no left-Q-neighbour. Otherwise, there is a left-Q-neighbour of u at p. By Corollary 8.19 there is some blocking path P_j, and thus t_j is at p. Now the left-Q-neighbour of u is also a left-Q-neighbour of t_j. However, t_j also has a right-Q-neighbour by Lemma 8.7, contradicting Lemma 8.3. Hence u has no left-Q-neighbour.

So u has k left-only-H-neighbours by Corollary 8.5. All of these are at p, by our choice of u, so none of them are free vertices, by Lemma 8.15. Recall P_i is the path that ends at p and a_i is the last vertex of P_i. Then a_i is not one of these k vertices, since $p = r(a_i)$. Hence there is some path P_j such that two vertices of P_j are left-only-H-neighbours of u. (In fact, there are at least two.) These two vertices are also at p by choice of u. Hence P_j is a blocking path by Lemma 8.18. Thus t_j is a left-only-H-neighbour of u. But then u must also be a vertex of T, since these vertices are the last $k-1$ vertices by right endpoint, and

$r(u) > r(t_j)$. Then, if v is any vertex entirely after p, then $r(v) > r(u)$ by the definition of u, so $v \in T$ by the same argument. □

Finally, Lemma 8.22 through to Lemma 8.26 prove a series of results about the structure of the free components of G. Again, these results will help with recolouring in Section 8.4.

Lemma 8.22. *Let G, Q, H and $\mathcal{P}$ be as in Lemma 8.13. If U is a free component and P_j is a blocking path, then there exists a vertex $w \in P_i \cup P_j$ such that the connected part of the line defined by U is entirely covered by the interval of w, that is, $l(w) < l(U) < r(U) < r(w)$.*

Proof. First suppose that U is not adjacent to t_j and suppose this lemma does not hold. Then, let x_i, x_j be vertices on P_i, P_j respectively at point $l(U)$. Then since there is no vertex w on either of these paths covering U, there is some other vertex y_i, y_j on $P_i, P_j - t_j$ respectively at $r(U)$. This contradicts Lemma 8.20.

Alternatively, suppose that U is adjacent to t_j. First we claim U is entirely left of p. Since no vertex of T is free, U cannot be entirely after p by Lemma 8.21, and U cannot be at p by Lemma 8.15. Thus U is entirely to the left of p.

So t_j is at $r(U)$ since t_j is at p. Then if a_j is the penultimate vertex on P_j, then a_j is also at p by Corollary 8.19. Since $a_j < t_j$, a_j is at $r(U)$. Hence if a_j does not cover U, there is some other vertex of P_j at $l(U)$. But then U is adjacent to two vertices of $P_j - t_j$, and by Lemma 8.20, U is adjacent to only one vertex of P_i, which must cover U. □

Now, we generalise our definition of left-Q-neighbours and right-Q-neighbours to free components—a free component U has a left-Q-neighbour if there is a vertex in Q whose left interval intersects U. Similarly we define right-Q-neighbour of U.

Corollary 8.23. *Let G, Q, H and $\mathcal{P}$ be as in Lemma 8.13. A free component U does not have both a left-Q-neighbour and a right-Q-neighbour.*

Proof. From Lemma 8.22 and Corollary 8.19, U is covered by some vertex w. Since w does not have both left- and right-Q-neighbours by Lemma 8.3, neither does U. □

Lemma 8.24. *Let G, Q, H and $\mathcal{P}$ be as in Lemma 8.13. A free vertex u is adjacent to at most three vertices from any path P_f.*

Proof. Say u is adjacent to $r \geq 4$ vertices of P_f. Label them $x_1, ..., x_r$ in their order in P_f. Then consider the path P_f^* with $x_2, ..., x_{r-1}$ replaced by u. This path still ends at the same place, but since we have removed at least two vertices and replaced them with only one, the number of vertices on the path has decreased. Thus if $\mathcal{P}^*$ is $\mathcal{P}$ with P_f replaced by P_f^*, then $\mathcal{P}^*$ contradicts condition (C2) in the choice of $\mathcal{P}$. □

Lemma 8.25. *Let G, Q, H and $\mathcal{P}$ be as in Lemma 8.13. If a free vertex u is adjacent to distinct vertices $x, y, z \in P_f$ for some P_f, where x, y, z is the order of the vertices in the path, then x and y are both at $l(u)$.*

Proof. Some vertex in P_f is at $l(u)$ and some vertex in P_f is at $r(u)$. Since u is adjacent to only $x, y, z \in P_f$, then x is at $l(u)$ and z is at $r(u)$. If y is not at $l(u)$, then construct P_f^* from P_f by replacing y with u. Now P_f^* is still a path and it ends at the same place as before, but $l(u) < l(y)$, contradicting (C3) in the choice of $\mathcal{P}$. □

Lemma 8.26. *Let G, Q, H and $\mathcal{P}$ be as in Lemma 8.13. A free component U has a right-Q-neighbour.*

Proof. Let u be the vertex of U such that $r(u) = r(U)$. The free component U has a right-Q-neighbour if and only if u has a right-Q-neighbour, so suppose u does not. Thus by Lemma 8.4 u has k right-only-H-neighbours. By our choice of u, there are no free vertices other than u at $r(u)$. Hence all of these k vertices are path vertices.

We claim that for each P_f, there is at most one vertex of P_f which is a right-only-H-neighbour of u. By Lemma 8.24, $|N(u) \cap P_f| \leq 3$. If $|N(u) \cap P_f| = 1$, then this vertex must be at $l(u)$ and $r(u)$, so it is not a right-only-H-neighbour. If $|N(u) \cap P_f| = 2$, then one of these two vertices is at $l(u)$, so there is at most one right-only-H-neighbour of u in P_f. Finally $|N(u) \cap P_f| = 3$, then by Lemma 8.25, two of these vertices are at $l(u)$, so again P_f contributes at most one right-only-H-neighbour of u.

However, there are only $(k-1)$ paths, and each path contributes at most one right-only-H-vertex, meaning together there are at most $(k-1)$ vertices, which is a contradiction. □

All of these lemmas together give enough of an idea of the structure of G, forcing enough non-adjacency amongst the vertices to allow us to determine a colouring of G with less than $\chi(G)$ colours.

8.4 Colouring G

Given all the facts we have proven about our graph, we will now show how to colour G with less that $\chi(G) = \mathcal{L}(G) + k$ colours, proving that there are no counterexamples. We will colour the graph in three parts.

First, we colour the clique Q with $\mathcal{L}$ colours, such that the vertex q_i is coloured i. Then, we colour the vertices of our special path set $\mathcal{P}$ and the remaining vertices of T. Finally we colour the free vertices. Call this colouring c. For $X \subset V(G)$, $c(X)$ will denote the set of colours assigned to the vertices in X.

For $f \in \{1, \dots, k-1\}$, define $\tau_f = P_f \cup t_f$. If P_f is finished, then $\tau_f = P_f$. If not, then τ_f can be thought of as a path from s_f to t_f with the final edge missing. We shall now colour all of the τ_f.

We want to define $(k-1)$ disjoint sets of colours $E_{k-1}, \dots, E_1$, where E_f is the *palette* for τ_f, that is, a set of colours we shall allow use of for τ_f. By doing this, we ensure that there can be no monochromatic edge from a vertex in one path to a vertex in another—the sets of colours we use for each path is different.

Recall that $S_f = \{q_1, \dots q_f\}$ and $T_f = \{q'_{\mathcal{L}}, \dots q'_{\mathcal{L}-(k-1)+f}\}$. Let $\mathcal{A} := c(S_{k-1})$ and $\mathcal{B} := c(T_1)$. Note that $|\mathcal{A}| = |\mathcal{B}| = (k-1)$.

Construct $E_{k-1}, \dots, E_1$ in that order by the following algorithm. A colour in some E_f is said to be *used*.

Say we are constructing E_f.

- Now $f \in \mathcal{A}$. If $f \in \mathcal{B}$, then set $E_f = \{f, \alpha_f\}$, where α_f is a new colour. Say E_f has Type 1.

Otherwise $f \notin \mathcal{B}$. Then select a colour from $c(T_f)$ that has not been used before by any palette, and call it b_f. (In Lemma 8.27 below, we prove that this is always possible.) Since $T_f \subseteq T_1$, we have $b_f \in \mathcal{B}$.

- If $b_f \in \mathcal{A}$, then set $E_f = \{b_f, \alpha_f\}$, where α_f is a new colour. Say E_f has Type 2.
- Finally, if $f \notin \mathcal{B}$ and $b_f \notin \mathcal{A}$, then let $E_f = \{f, \alpha_f, b_f\}$, where α_f is a new colour. Say E_f has Type 3.

We have used $(k-1)$ new colours $\alpha_1, \dots, \alpha_{k-1}$. However, we need to ensure that the above algorithm constructs well-defined sets of distinct colours.

Lemma 8.27. *Let G be the vertex-minimum counterexample to Theorem 1.12, let Q be the set of vertices at a point of maximum load q and let $H := G - Q$. Let $\mathcal{P}$ be a special path set chosen with respect to (C1),(C2) and (C3). It is possible to construct a set of palettes as described above, and these palettes are pairwise disjoint.*

Proof. It is sufficient to prove the following stronger statement by induction: For all $f \in \{(k-1), \dots 1\}$, the sets $E_{k-1}, \dots, E_f$ are well-defined and pairwise disjoint, and $\bigcup_{j=f}^{k-1} E_j$ contains $(k-f)$ colours in $\mathcal{B}$, and $\{1, \dots, f-1\} \cap (\bigcup_{j=f}^{k-1} E_j) = \emptyset$.

First we show it is possible to construct E_{k-1}, the base case:

- Note that $(k-1) \in \mathcal{A}$. If $(k-1) \in \mathcal{B}$, then $E_{k-1} = \{(k-1), \alpha_{k-1}\}$. So far we have used one colour from $\mathcal{B}$, and we have not used $1, \dots, (k-2)$.

If $(k-1) \notin \mathcal{B}$, then the algorithm selects a colour b_{k-1} from $c(T_{k-1}) = \{c(q'_{\mathcal{L}})\}$ that has not been used before—since we have not used any colours yet, this is fine.

- If $b_{k-1} \in \mathcal{A}$, then $E_{k-1} = \{b_{k-1}, \alpha_{k-1}\}$. So far we have used one colour from $\mathcal{B}$. The vertex $q'_{\mathcal{L}}$ cannot be in S_{k-1} and T_{k-1} by Corollary 8.11. Since $b_{k-1} \in c(T_{k-1})$, $b_{k-1} \notin c(S_{k-1})$ and we have not used $1, \ldots, (k-2)$.
- If $b_{k-1} \notin \mathcal{A}$, then $E_{k-1} = \{k-1, \alpha_{k-1}, b_{k-1}\}$. So, again, we have used one colour from $\mathcal{B}$, and we have not used $1, \ldots, (k-2)$.

Now, since we can find b_{k-1} if it is required, E_{k-1} is well-defined, and since we have only defined one palette, all defined palettes are disjoint trivially. Also note that we have used one colour from $\mathcal{B}$ and we have not used $1, \ldots, (k-2)$, satisfying our other requirements.

Now, by induction, say we have constructed $E_{k-1}, \ldots E_{f+1}$, and that we have used $(k-f-1)$ colours from $\mathcal{B}$ and we have not used $1, \ldots, f$. We show it is possible to construct E_f as required:

- If $f \in \mathcal{B}$, then $E_f = \{f, \alpha_f\}$. Hence we use one more colour from $\mathcal{B}$, and $1, \ldots, (f-1)$ remain unused.

If $f \notin \mathcal{B}$, then the algorithm selects an unused colour b_f from the set $c(T_f) = \{c(q'_{\mathcal{L}}), \ldots, c(q'_{\mathcal{L}-(k-1)+f})\}$. There are $(k-f)$ colours in this set, and they are all in $\mathcal{B}$. However, we have only used $(k-f-1)$ colours in $\mathcal{B}$. Hence since $c(T_f) \subseteq \mathcal{B}$, there is at least one unused colour $b_f \in c(T_f)$.

- If $b_f \in \mathcal{A}$, then $E_f = \{b_f, \alpha_f\}$. Hence we use one more colour from $\mathcal{B}$, and $1, \ldots, (f-1)$ remain unused, since $b_f \notin c(S_f)$ by Corollary 8.11, similar to the base case.
- If $b_f \notin \mathcal{A}$, then $E_f = \{f, \alpha_f, b_f\}$. So, again, we have used one more colour from $\mathcal{B}$ and $1, \ldots, (f-1)$ all remain unused.

Thus E_f is well-defined. Since $1, \ldots, f$ are not used in $E_{k-1}, \ldots, E_{f+1}$, and b_f was chosen such that it had not been used, all defined palettes are pairwise disjoint. Also, $1, \ldots, (f-1)$ remain unused and only $k-f$ colours from $\mathcal{B}$ have been used, so our induction holds. □

Now we have enough to colour G with only less than $\chi(G)$ colours and obtain our contradiction.

Theorem 8.28. *Let G be a vertex-minimum counterexample to Theorem 1.12. Then G can be coloured with $\mathcal{L} + k - 1 = \chi(G) - 1$ colours.*

Proof. Recall that we can colour Q with $\mathcal{L}$ colours such that $c(q_f) = f$. All that remains to do is to colour τ_f for $f \in \{1, \ldots, k-1\}$ and the free components.

Firstly we colour $\tau_{k-1}, ..., \tau_1$ using the sets $E_{k-1}, ..., E_1$ we constructed in Lemma 8.27. We claim we can colour each τ_f with the colours of E_f without creating a monochromatic edge.

First note that since the colours of the palettes are pairwise disjoint, there will be no monochromatic edge between a vertex of τ_f and τ_g for $f \neq g$, and since the paths are induced, there will be no monochromatic edge inside τ_f as long as we ensure there are no monochromatic edges between consecutive vertices on τ_f. The only thing to be careful of is monochromatic edges between τ_f and Q.

Say E_f has Type 1. Note we can use α_f for any vertex of τ_f since it is a new colour, that is, a colour not used by Q. We now show that f can be used for any vertex other than t_f. Since $f \in c(S_f)$, the left interval of the vertex q_f does not reach $l(s_f)$ by Lemma 8.10, and since $f \in \mathcal{B}$, the right interval of q_f does not reach $r(t_1)$. Since there is only one vertex of T in τ_f, q_f intersects only one vertex of this set. Thus, τ_f can be coloured by assigning t_f the colour α_f, and alternating between f and α_f from right to left along the path.

Say E_f has Type 2. This case is very similar to the previous one; α_f is available for any vertex, and b_f is available for any vertex other than s_f, by the mirror of the previous argument. Colour s_f by α_f and alternate b_f and α_f along the rest of the path from left to right.

Say E_f has Type 3. Partition τ_f into two subpaths τ_f^1, the vertices of τ_f without q_f as a right-Q-neighbour, and τ_f^2, the remaining vertices of τ_f that do have q_f as a right-Q-neighbour. Now the vertices of τ_f^1 can all use either α_f or f, since they do not intersect the right interval of q_f by definition, and they do not intersect the left interval of q_f by Lemma 8.10. Since the vertices of τ_f^2 have a right-Q-neighbour they do not have a left-Q-neighbour by Lemma 8.3. Hence none of these vertices are adjacent to the left interval of the vertex of Q coloured b_f, and since $b_f \in c(T_f)$, none are adjacent to the right interval of this Q-vertex by Lemma 8.10. Hence we can use b_f or α_f for any vertex of τ_f^2. So colour τ_f^1 such that the last vertex has colour f, that is, colour it f and alternate α_f, f back towards s_f, and colour the first vertex of τ_f^2 by α_f and alternate b_f, α_f from left to right along τ_f^2.

Hence we can colour each τ_f as required. So far we have used $\mathcal{L} + k - 1$ colours. We need to colour the rest of the graph without using any new colours. Only the free vertices remain. Let U be a free component. Then we claim we can colour U with the existing colours.

By Lemma 8.26, U has a right-Q-neighbour, and by Corollary 8.23 it has no left-Q-neighbour. We shall colour the vertices of U by traversing the component from right to left, and colouring a vertex u when we reach its right endpoint, with any available colour. It suffices to show that when we come to colour u that one of the $\mathcal{L} + k - 1$ colours is not assigned to a vertex adjacent to u. Since U has no left-Q-neighbour, neither does u. Thus, if there is a Q-vertex adjacent to u, it is at $r(u)$. Similarly, by the order in which we colour the free vertices, any coloured free vertex in the neighbourhood of u is at $r(u)$. So if a colour d appears in the neighbourhood of u but not on a vertex at $r(u)$, then it must appear on a path vertex.

We claim that for a path P_f, there is at most one colour on a vertex in $P_f \cap N(u)$ that is not at $r(u)$. For each path P_f, there is at least one vertex of that path at $r(u)$, so there is at least one colour of $c(P_f \cap N(u))$ at $r(u)$. Hence, if $|c(P_f \cap N(u))| \leq 2$, there is at most one colour of $c(P_f \cap N(u))$ not at $r(u)$, as required. Since $P_f \subseteq \tau_f$, and τ_f is coloured with E_f, we have $|c(P_f \cap N(u))| \leq |E_f|$. Hence if $|E_f| = 2$ (that is, if E_f has Type 1 or Type 2), our claim holds.

Otherwise, $|E_f| = 3$, and E_f has Type 3. When we coloured τ_f with three colours, we actually 2-coloured each of the two subpaths of τ_f. Hence if $N(u)$ only intersects one subpath of τ_f, then $|c(P_f \cap N(u))| \leq 2$, and we get the same result as above. If $|c(P_f \cap N(u))| = 3$, then $N(u)$ contains vertices from both subpaths. By Lemma 8.24, $|P_f \cap N(u)| \leq 3$, so $|c(P_f) \cap N(u)| = 3$ and $N(u)$ contains either one vertex from τ_f^1 and two from τ_f^2, or vice versa. In the first case, $c(P_f \cap N(u)) = \{f, \alpha_f, b_f\}$. Let the three vertices of $P_f \cap N(u)$ be x, y, z by their order in P_f. By Lemma 8.25, x and y are at $l(u)$, and z is at $r(u)$. If y is also at $r(u)$, then we have only one colour of this set not at $r(u)$. Otherwise, $r(y) < r(u)$ and thus any right-Q-neighbour of y is a right-Q-neighbour of u. Since $y \in \tau_f^2$, y has $q_f \in Q$ as a right-Q-neighbour where $c(q_f) = f$. But then, f is at $r(u)$, so again there is only one colour of $c(P_f \cap N(u))$ not at $r(u)$. In the second case, $c(P_f \cap N(u)) = \{\alpha_f, f, \alpha_f\}$, so $|c(P_f \cap N(u))| = 2$, and our claim holds.

Now, there are $(k-1)$ paths, and for each path there is at most one colour on that path not at $r(u)$. Since these are the only colours not at $r(u)$, there are only $(k-1)$ colours adjacent to u that are not at $r(u)$. There are at most $(\mathcal{L} - 1)$ colours at $r(u)$ since $\mathcal{L}$ is the maximum load. Thus there are at most $(\mathcal{L} - 1) + (k - 1) = \mathcal{L} + k - 2$ colours adjacent to u. We are colouring with $\mathcal{L} + k - 1$ colours, so there is always a colour free for u. Thus our claim holds. $\square$

Thus, we can colour G with only $\mathcal{L} + k - 1$ colours, which contradicts $\chi(G) = \mathcal{L} + k$. This completes the proof of Theorem 1.12.

8.5 Extensions

The question remains as to whether Theorem 1.12 can be improved, either by removing the -1, and thus proving Hadwiger's Conjecture for circular arc graphs when $\beta(G) > 3$, or by removing the requirement on the cover number.

We believe that it should be possible to prove that $\mathsf{had}(G) \geq \chi(G)$ whenever G is a circular arc graph with $\beta(G) > 3$, and that this should be able to be proven using techniques similar to those found in this chapter. Note the following facts. A special path set can be expanded to contain k paths rather than $k-1$ paths without too much trouble. All of the important results about the nature of special path sets still hold if we make this adjustment. However, there are some issues with the colouring arguments that need to be fixed. For example, if we place a new colour in each palette $E_1, \ldots, E_k$, then there are k new colours. Together with the $\mathcal{L}$ colours of Q, this give $\chi(G)$ colours in total, which does not give the improvement we require. Our initial attempts to fix this issue involved colouring one of the paths with only existing colours, possibly using the fact that at least one of the paths is unfinished. If P_f is unfinished, then τ_f is a path with an edge missing, which should be easier to colour than a normal path.

On the other hand, proving $\mathsf{had}(G) \geq \chi(G)$ when $\beta(G) \leq 3$ is likely to be much harder. First note that the assumption $\beta(G) > 3$ is used repeatedly in the proofs of this chapter, and there is no obvious way to get similar results without this assumption. It is likely that when $\beta(G) \leq 3$, the constructed minor will need to contain two sets of paths, rather than a clique and a single special path set. Note, however, that if $\beta(G) = 1$, then G contains a *dominating vertex*, a vertex adjacent to all other vertices. A vertex minimum counterexample does not contain such a vertex, since adding a dominating vertex increases both the chromatic number and the Hadwiger number by 1. Thus removing such a vertex lowers both parameters by 1, and gives a smaller counterexample. Thus we would only need to prove the conjecture when $\beta(G) \in \{2, 3\}$.

Chapter 9

Linkages in Interval Graphs

9.1 Introduction

Recall that there is some slight overlap in terminology between two different concepts of "linked". In this chapter, we shall not use the concept of linkedness presented in Chapter 2, so there should be no confusion. Recall $2k$ distinct vertices $s_1, \ldots, s_k, t_1, \ldots, t_k$ can be *linked* if there exists a set of k pairwise vertex disjoint paths $P_1, \ldots, P_k$ such that P_i starts at s_i and ends at t_i. The paths $P_1, \ldots, P_k$ are called a *linkage*. For a graph G, if $|V(G)| \geq 2k$ and if any $2j$ distinct vertices (where $j \leq k$) can be linked, then we say the graph G is *k-linked*. (Note that if any $2k$ distinct vertices can be linked, then so can any $2j$ distinct vertices for $j < k$.) We call the vertices $s_1, \ldots, s_k$ *sources* and $t_1, \ldots, t_k$ *targets*.

The power of a path, P_n^k, is the graph formed by taking the n-vertex path and adding edges between any two vertices u, v where the distance $d(u, v) \leq k$. That is, if we label the vertices of the path $1, \ldots, n$, then there is an edge between i and j whenever $|i - j| \leq k$. We shall always label the vertices in this fashion, and we say vertices are *left* or *right* of each other with respect to this labelling. It is clear that P_n^k is k-connected whenever $n > k$. The graph P_n^k is the interval graph of $\{[i, i + k] : 1 \leq i \leq n\}$. (Given that no interval is completely covered by another, P_n^k is a *proper* interval graph.)

The interval graphs are a subset of the chordal graphs. To see this, take a cycle of length at least 4 and consider the intervals of this cycle in the real line. Let v be the vertex with the leftmost right endpoint; v has two neighbours in the cycle, but they are both at the right endpoint of v and this forces a triangle as an induced subgraph of the cycle. Thus every cycle of length at least 4 has a chord. Recall the following linkage result for chordal graphs.

Lemma 9.1 (Böhme *et al.* [9])**.** *If G is a $(2k-1)$-connected chordal graph, then G is*

k-linked.

We now provide an alternate proof of Lemma 9.1.

Proof. As discussed in Section 2.4, if G is a chordal graph, then it is possible to find a tree decomposition T of G such that every bag of the tree decomposition is a clique. Since G is $(2k-1)$-connected, for any two adjacent bags X, Y of T, it follows that $|X \cap Y| \geq 2k-1$, since the set $X \cap Y$ is a cut-set. Similarly, any cut set must contain $X \cap Y$ for some pair of adjacent bags X, Y; otherwise it is possible to find a path in G between any pair of vertices x, y by taking vertices along the path in T between a bag containing x and a bag containing y.

Perform induction on $k \geq 1$. The base case is clear since a 1-connected graph is 1-linked. Now consider the case for general k. Given $2k$ distinct vertices $s_1, \ldots, s_k, t_1, \ldots, t_k$, let $G' := G - \{s_2, \ldots, s_k, t_2, \ldots, t_k\}$. Since G is $(2k-1)$-connected and only $2(k-1) = 2k-2$ vertices have been deleted, the graph G' is connected. Let P_1 be an induced path from s_1 to t_1 in G'. The path P_1 is also an induced path from s_1 to t_1 in G which avoids $s_2, \ldots, s_k, t_2, \ldots, t_k$. Consider $G - P_1$. Since P_1 is induced, it contains at most two vertices in every bag of T, or else P_1 contains a triangle. Hence $G - P_1$ is $(2k-3)$-connected, since deleting P_1 deleted at most two vertices from $X \cap Y$ for any pair of adjacent bags X, Y. Since $2k - 3 = 2(k-1) - 1$, the graph $G - P_1$ is $(k-1)$-linked by induction. Thus, let $P_2, \ldots, P_k$ be the linkage for $s_2, \ldots, s_k, t_2, \ldots, t_k$ in $G - P_1$. Thus the set $P_1, P_2, \ldots, P_k$ is the required linkage for $s_1, \ldots, s_k, t_1, \ldots, t_k$ in G. $\square$

We now show that Lemma 9.1 is tight, even for powers of a path. (Böhme *et al.* proved a similar result, but not in this fashion.) Consider P_{3k-1}^{2k-2}; it is sufficient to show that this graph is not k-linked. Let $s_i = i$ and $t_i = 2k - 1 + i$ for $1 \leq i \leq k$. Note that s_i is not adjacent to t_i for any choice of i. Thus each path from s_i to t_i contains at least three vertices. However, k pairwise vertex disjoint paths, each containing at least three vertices requires a total of at least $3k$ vertices. Hence, P_{3k-1}^{2k-2} is not k-linked.

This above example can be extended to powers of paths with $n > 3k - 1$ as long as we ensure our choice of sources and targets are close to one another (that is, they induce a subgraph that is a power of a path of length $3k - 1$).

Before proving Theorem 1.13, we prove a more powerful result for the class of powers of a path.

9.2 "Selection Sort" Paths in the Power of a Path

Consider the power of a path P_n^{k+r}. We desire to show this graph is something like k-linked for a small integer $r > 0$. In order to do this, we need to restrict the sources and the targets, otherwise this is impossible. First, given any source-target pair s_i, t_i, we declare that the vertex on the left will always be the source, and the vertex on the right will always be the target. More drastically, we need to ensure that the minimum distance between *any* source s_i and target t_j is at least $\frac{k}{r}$. By ensuring that the sources and targets are "far apart", there is enough "room" to be able to organise the paths such that each reaches the correct target. Also note that this restriction ensures that the example proving the tightness of Lemma 9.1 is not an obstruction to this alternate result.

Lemma 9.2. *Let $s_1, \ldots, s_k, t_1, \ldots, t_k$ be distinct vertices in the graph P_n^{k+r} (with $r > 0$) such that s_i is left of t_i for all i, and such that the distance between* any *source and* any *target in P_n^{k+r} is at least $\frac{k}{r}$. Then $s_1, \ldots, s_k, t_1, \ldots, t_k$ can be linked.*

Given the restrictions in Lemma 9.2, note the following. Say s_i is the rightmost source and t_j is the leftmost target. Now the distance between s_i and t_j is at least $\frac{k}{r}$. (In fact, that the distance between s_i and t_j is at least $\frac{k}{r}$ is sufficient to prove all other sources and targets are far enough apart.) Given that P_n^{k+r} is k-connected, it is possible to take paths starting at $\{s_1, \ldots, s_k\} - s_i$ and ending at the $k-1$ vertices immediately left of s_i, that do not use any vertex right of s_i. (This last part follows from the fact that we can ensure these paths are actually paths in P_n^k, and the k consecutive vertices ending at s_i are a cut set of P_n^k.)

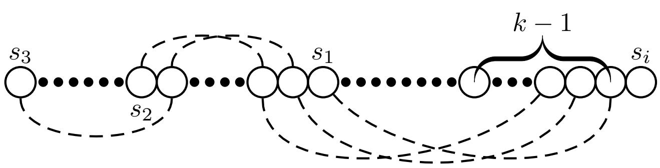

Figure 9.1: It is sufficient to take paths from the sources to the $k-1$ vertices immediately left of s_i, and then take paths from those k vertices (that is, including s_i) to the correct corresponding targets.

A similar fact holds concerning the vertices immediately right of t_j. Thus, if we can link the k vertices ending at s_i with the k vertices starting at t_j (for any possible choice of the k pairs), then this is sufficient to prove Lemma 9.2. So the following lemma is sufficient.

Lemma 9.3. *Let $n \geq 2k$, and let $s_1, \ldots, s_k$ be the first k vertices of the graph P_n^{k+r} (in any order) and $t_1, \ldots, t_k$ be the last k vertices (again in any order). If $r > 0$ and the distance between the rightmost source and the leftmost target is at least $\frac{k}{r}$, then $s_1, \ldots, s_k, t_1, \ldots, t_k$ can be linked.*

We can actually go further than this. Partition the vertices of P_n^{k+r} into blocks of size k, such that the first k vertices form the first block, the second k vertices the second, and so on. Note that the first block contains exactly the sources; call this the source block, and label it 0. Label the subsequent blocks consecutively. All vertices in the j^{th} block are distance j from the final source. Whenever $n \not\equiv 0 \mod k$, the targets will not all be in the same block, and will be split over two consecutive blocks. Given that the distance between sources and targets is at least $\lceil \frac{k}{r} \rceil$, it follows that the first block that contains targets will be labelled $\lceil \frac{k}{r} \rceil$ or higher. From the targets in the second target block take the set of edges into the corresponding non-target vertices of the first target block. This is possible given that the targets and corresponding non-targets are distance k in P_n. If we construct the correct paths from the sources into the vertices of the first target block, then those paths with the above edges added give the desired linkage.

Also note all vertices of the first target block are still at a distance of at least $\lceil \frac{k}{r} \rceil$ from the sources. If the distance is greater than $\lceil \frac{k}{r} \rceil$, then take a set of edges from each vertex in the first target block to the corresponding vertex in the previous block. If we construct a set of paths from the sources to the block before the first target block, then we can add these edges to get the desired linkage. By doing this repeatedly, we can ensure we only need paths from the source block to the block $\lceil \frac{k}{r} \rceil$.

Finally note that since both sources and targets appear in any order, we can permute the labels of the targets so that they appear in the obvious order (that is, t_1 before t_2 and so on), as long as we perform the corresponding permutation on the sources. Taken together, this means it is sufficient to prove Lemma 9.4 in order to prove Lemma 9.2.

Lemma 9.4. *Consider the graph P_n^{k+r} where $r > 0$ and $n = \lceil \frac{k}{r} \rceil k + k$. Label the last k vertices $t_1, \ldots, t_k$, such that t_k is the last vertex, t_{k-1} the second last and so on. Label the first k vertices $s_1, \ldots, s_k$. Regardless of the order of $s_1, \ldots, s_k$, the vertices $s_1, \ldots, s_k, t_1, \ldots, t_k$ can be linked.*

Proof. Recall that we have partitioned the vertices of P_n^{k+r} into $\lceil \frac{k}{r} \rceil + 1$ blocks of k vertices each, labelled from 0 to $\lceil \frac{k}{r} \rceil$. Each path from a source to a target will intersect each block exactly once. Label each vertex of a block by one of $p_1, \ldots, p_k$ to denote which path it is on. Label each source s_i with p_i. As long as

- the vertex labelled with p_i in one block is adjacent to the vertex labelled p_i in the previous block, and
- each t_i is labelled p_i,

then the required linkage has been constructed.

Let $1 \leq j \leq \lceil \frac{k}{r} \rceil$. In the j^{th} block, label the first $\min\{jr, k\}$ vertices by $p_1, \ldots, p_{\min\{jr,k\}}$ in that order, and call this set of labels the first subset. When $\min\{jr, k\} = jr$, we may still have the labels $p_{jr+1}, \ldots, p_k$ to place. Place these remaining labels on the remaining $k - jr$ vertices of the j^{th} block in the same order as the labels appear in the 0^{th} block. Call this the second subset. The final block is labelled $\lceil \frac{k}{r} \rceil$, and so for this block all vertices are labelled by the first subset, and so each t_i is labelled by p_i. Hence this labelling satisfies the second bulleted requirement above.

It remains to check that the first requirement holds. Consider the j^{th} block. It is easy to see that the vertices labelled $p_1, \ldots, p_{(j-1)r}$ are adjacent to the equivalently labelled vertices in the previous block. (In fact, this fact would hold even if the graph was P_n^k.) The remaining labels may have appeared on the corresponding vertex in the previous block, or on a vertex at most r vertices before the corresponding vertex, given the r new labels in the first subset. However, in P_n^{k+r} each vertex is adjacent to $k + r$ vertices before it in the ordering, so the first bulleted requirement holds. This proves the lemma. $\square$

We call these "selection sort" paths since if $r = 1$, then the labels in each block appear to be undergoing selection sort—each block in sequence is another pass over the linked list and another label is placed in the correct position. Finally at the last step, all labels are in the correct position, and the sources are linked up to the correct targets. When $r > 1$, this is equivalent to selection sort with more labels moved at every given step. This means less steps are required. We need to ensure that the sources and targets are far enough apart so that "selection sort" has the "time" (by which we mean, the space in the vertex set) to run completely. Given that each vertex is adjacent to $k + r$ vertices after it in the ordering, a label can only be moved up by at most r places, so this process cannot obviously be improved. Because of this, Lemma 9.4 is best possible in the sense that if $r = 0$ and the ordering of the sources is not simply $s_1, \ldots, s_k$, then it is impossible to link the sources and targets; as we try to build the paths from the sources to the targets, sequential vertices in a given path must be corresponding vertices in sequential blocks, or some other path cannot be extended. However, this means no kind of "rearrangement" is possible, and as such the sources and the targets will not link up correctly. So we cannot ensure k-connectivity gives even our weakened version of k-linked.

We discuss a use of this "selection sort" technique in Section 9.4.

9.3 Improved Linkages in Interval Graphs

The result of Section 9.2 is of some interest, but it would be far more preferable to extend Lemma 9.2 to a more general class of graphs. Here we extend a limited version of this result to the class of interval graphs.

Theorem 1.13. *Let G be a $\lceil \frac{3k}{2} \rceil$-connected interval graph, and let $s_1, \ldots, s_k, t_1, \ldots, t_k$ be $2k$ pairwise distinct vertices, such that no source s_i and no target t_j are adjacent, and such that s_i is left of t_i for all i. Then $s_1, \ldots, s_k, t_1, \ldots, t_k$ can be linked.*

Note that this is essentially Lemma 9.2 generalised for the interval graph when $r = \frac{k}{2}$. Many of the techniques used in the proof of Theorem 1.13 are similar to those used in Chapter 8. Also note that since s_i and t_i are not adjacent, the requirement that s_i is left of t_i is unambiguous.

First, we recall the following basic facts about an interval graph G (which are quite similar to the basic facts about circular arc graphs). Given that every vertex of the interval graph corresponds to an interval on the real line, we often treat the vertex and its corresponding interval interchangeably. (Only in rare cases do we need to be more explicit.) Thus every vertex v has a *left endpoint* denoted $l(v)$ and a *right endpoint* denoted $r(v)$. By perturbing the endpoints of all the intervals, we can ensure that no point is an endpoint (left or right) of more than one interval. Connected induced subgraphs also have corresponding intervals on the real line (that is, the union of all the intervals corresponding to vertices of the subgraph); hence if U is a connected induced subgraph we define $l(U)$ and $r(U)$ to be its endpoints. If we consider all of the right endpoints of the vertices in G, define the leftmost of these points to be the *start point*. Also define the rightmost left endpoint of a vertex to be the *final point*. The *minimum load* $\ell(G)$ is the minimum number of vertices at any point between the start and final points. (We make this restriction, otherwise the minimum load might be 0.) It is well known and easily seen that the minimum load is equal to the connectivity of G.

Define a *special path set* $\mathcal{P}$ to be a set of paths $P_1, \ldots, P_k$ in G that satisfy the following properties.

(P1) $P_1, \ldots, P_k$ are pairwise vertex-disjoint.

(P2) $s_i \in P_i$ for all i, and $t_j \notin P_i$ when $j \neq i$.

This is essentially identical to the definition of a special path set in Chapter 8, however note that there are k (and not $(k-1)$) paths in $\mathcal{P}$. If $t_i \in P_i$ we say that P_i is *finished*, otherwise it is *unfinished*. If all paths are finished, then we say $\mathcal{P}$ is itself finished, otherwise it is unfinished. Let p_i be the rightmost point of P_i, for all P_i.

Obviously, if $\mathcal{P}$ is finished then it is a linkage, and Theorem 1.13 holds. Otherwise, we can assume that every special path set is unfinished. Let $X \subseteq \{p_1, \dots, p_k\}$ be the set of right endpoints from unfinished paths, and let p be the leftmost point in X (note $X \neq \emptyset$). Choose $\mathcal{P}$ so that

(C1) The point p, as defined above, is as far right as possible.
(C2) Subject to (C1), $|V(P_1) \cup \cdots \cup V(P_k)|$ is minimised.

This choice of $\mathcal{P}$ (specifically the definition of p) *is* different than the choice of $\mathcal{P}$ in Chapter 8. The definition of p is done this way since now we cannot assume that the sources and the targets form cliques at the start and end of the interval respectively, as was the case in Chapter 8. Also, the requirement (C3) is no longer required since it was previously included to assist with colouring, which is not relevant here.

We will prove a contradiction by showing that $\mathcal{P}$ can be modified so that p is moved further to the right, which gives a better choice of $\mathcal{P}$.

Say P_i is the unfinished path such that $p_i = p$, and say P_j is a path (finished or unfinished) such that $p_j > p$. If replace P_i with a path that is either finished, or unfinished with a right endpoint further right than p, then p itself has been moved further right. (It is possible p may no longer refer to the same p_i, but the point it refers to will be right of the old p given that p is leftmost in X.) If, while replacing P_i with this new path, we also replace P_j with another shorter path, then as long as the right endpoint of the new P_j is right of the original p, the new p is still further right than the old p. This follows even if P_j has gone from finished to unfinished. This means we can replace P_i (and perhaps P_j) in a way that is essentially identical to Chapter 8 in order to construct a better choice of $\mathcal{P}$. As a result, for many of our basic results we shall simply cite the appropriate previous result.

Lemma 9.5. *Let G be a $\lceil\frac{3k}{2}\rceil$-connected interval graph, and let $s_1, \dots, s_k, t_1, \dots, t_k$ be $2k$ distinct vertices, such that no source s_i and no target t_j are adjacent, and such that s_i is left of t_i for all i. Let $\mathcal{P}$ be the special path set chosen with respect to (C1) and (C2). Then every $P_i \in \mathcal{P}$ is an induced path, and no vertex of P_i other than s_i or t_i is covered by any other vertex of P_i.*

Proof. This follows from Lemma 8.13. However note a slight weakening—it is possible that the source or the target on P_i is covered by another vertex. This is because s_i and t_i are not necessarily at the start and end of the interval. □

As before, a vertex of G that is in no path of $\mathcal{P}$ and is not a target is called a *free vertex*.

Lemma 9.6. *Let $G, s_1, \ldots, s_k, t_1, \ldots t_k$, and $\mathcal{P}$ be as in Lemma 9.5. Let P_i be the unfinished path such that $p_i = p$. Then there is neither a free vertex nor the vertex t_i at p, and as such t_i is right of p.*

Proof. The first part of this lemma follows from Lemma 8.15. The second part follows from the self-evident fact that the unfinished path P_i is not adjacent to t_i at any point, and from the relative positions of s_i and t_i on the real line. □

By Lemma 9.5, every path P_j of $\mathcal{P}$ is induced, so it follows that at any point on the real line at most two intervals in P_j are present. If a path P_j contains two vertices at a point r, we say P_j *appears twice at point* r. If P_i is the path such that $p_i = p$, then P_i does not appear twice at point p, since p is the right endpoint of an end vertex of the P_i.

Lemma 9.7. *Let $G, s_1, \ldots, s_k, t_1, \ldots t_k$, and $\mathcal{P}$ be as in Lemma 9.5. If P_j appears twice at point p, then one of the two vertices of P_j at p is the vertex t_j.*

Proof. This follows from Lemma 8.18. □

Now, consider the point $p + \epsilon$, the point just after p but before the left endpoint of any other vertex. At $p + \epsilon$ there are at least $\lceil \frac{3k}{2} \rceil$ vertices, since $p + \epsilon$ is left of $l(t_i)$ by Lemma 9.6 and thus left of the final point. Hence at p itself there are at least $\lceil \frac{3k}{2} \rceil + 1$ vertices. By Lemmas 9.6 and 9.7, all of the vertices at p are on paths of $\mathcal{P}$, and any path that appears twice at point p includes a target vertex. There must be $\lceil \frac{k}{2} \rceil + 1$ paths appearing twice at point p, otherwise there are insufficient vertices at p. Hence there is at least one target at p, and so no source at p.

If P_i is the unfinished path such that $p = p_i$, then P_i contains at least two vertices since the vertex at p is not s_i. Let y_i denote the rightmost vertex of P_i such that no target at p is also at $l(y_i)$. We show that such a vertex is well-defined. If v is the vertex of P_i adjacent to s_i, then either $l(v)$ is inside s_i, or it is left of $l(s_i)$. In either case, any target at both p and $l(v)$ must be adjacent to s_i, contradicting our assumption about G. Hence there exists some vertex of P_i such that there is no target at both its left endpoint and at p, and so y_i is well-defined. It also follows that $y_i \neq s_i$, so let x_i denote the neighbour of y_i in P_i that is before y_i in the path.

Recall that a maximal connected induced subgraph of G containing only free vertices, is called a *free component*. Denote the free component at $l(y_i)$ by U, and say $U = \emptyset$ if there is no such free component. We define the *reverse point* q to be $l(U)$ if $U \neq \emptyset$, or $l(y_i)$ if $U = \emptyset$. It is possible that the reverse point is left of the start point; if it is not, we say it has *type 1* and if it is we say it has *type 2*. When q has type 1, we have a situation similar to Lemma 8.20, but when q has type 2 our proof is rather different.

Lemma 9.8. *Let $G, s_1, \dots, s_k, t_1, \dots t_k$, and $\mathcal{P}$ be as in Lemma 9.5. Let q be the reverse point. Then q does not have type 1.*

Proof. Since q is right of the start point and q is the endpoint of some vertex, it follows that there are at least $\lceil \frac{3k}{2} \rceil + 1$ vertices at q (for the same reason there are at least that many at p). By the maximality of U, at most one of these vertices is a free vertex, so it follows that at least $\lceil \frac{k}{2} \rceil$ paths of $\mathcal{P}$ appear twice at q. Given that $\lceil \frac{k}{2} \rceil + 1$ paths of $\mathcal{P}$ appear twice at p, it follows that some path P_j appears twice at q and at p. Let P_i be the path such that $p = p_i$; it follows $i \neq j$, since P_i does not appear twice at point p. Denote the vertices of P_j at q by a_j, b_j such that $r(a_j) < r(b_j)$. Denote the vertices of P_j at p by c_j, t_j such that $r(c_j) < r(t_j)$. It is possible that $b_j = c_j$, but otherwise a_j, b_j, c_j, t_j are pairwise disjoint and $r(b_j) < r(c_j)$. This follows since there are no targets at both q and p, by choice of y_i and since $q \leq l(y_i) < p$.

Given this, the path P_j has the form $(s_j, \dots, a_j, b_j, \dots, c_j, t_j)$. Denote the vertex of P_i at p by a_i. Then P_i has the form $(s_i, \dots, x_i, y_i, \dots, a_i)$. It is possible that $s_i = x_i$ and/or $y_i = a_i$. Partition these paths into the following subpaths: $P_j^1 := (s_j, \dots, a_j)$, $P_j^2 := (b_j, \dots, c_j)$, $P_j^3 := (t_j)$ and $P_i^1 := (s_i, \dots, x_i)$, $P_i^2 := (y_i, \dots, a_i)$.

Consider $P_j^1 \cup U \cup P_i^2 \cup P_j^3$. We can replace P_j with this path which travels from s_j to t_j, however, we will also need to replace P_i, otherwise P_i^2 is contained into two paths of $\mathcal{P}$. Fortunately, the set P_j^2 is not longer being used in P_j, and this subpath covers the real line from q to p, and as such covers $l(y_i)$, a point which contains x_i. Hence replace P_i with $P_i^1 \cup P_j^2$.

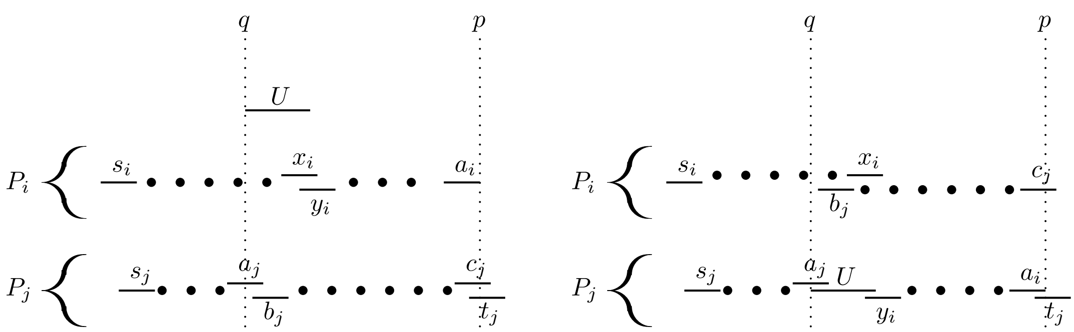

Figure 9.2: The subpath P_j^2 is placed into P_i, and P_i^2 is placed into P_j. It may be necessary to place vertices of U into P_j to maintain connectivity. Note the similarity to Figure 8.2.

The right end of each path remains unchanged except for P_i, and $p_i \geq r(c_i) > p$. Hence p has been moved further right, contradicting our choice of $\mathcal{P}$. □

In the proof of Lemma 9.8, we essentially "swap" the middle sections of P_i and P_j in such a way that P_j is still finished, but P_i is now able to travel further right. This is the key idea in the proof of Theorem 1.13. However, we also must deal with a few other cases.

Lemma 9.9. *Let $G, s_1, \ldots, s_k, t_1, \ldots t_k$, and $\mathcal{P}$ be as in Lemma 9.5. Let q be the reverse point. Then q does not have type 2.*

Proof. In order to prove this, we also define a *second reverse point*. Let P_i be the unfinished path such that $p = p_i$. Let the free component at $r(x_i)$ be denoted U' (and let $U' = \emptyset$ if there is no such free component). Then let the second reverse point $q' := r(U')$ (or $r(x_i)$ if $U' = \emptyset$). Since P_i does not appear twice at p and since there are no free vertices at p, it follows that $q' < p$. As with q in Lemma 9.8, there are $\lceil \frac{k}{2} \rceil$ paths appearing twice at q', and as such there exists a P_j that appears twice at q' and p.

Let a_i be the vertex of P_i at p and partition the path P_i as in Lemma 9.8: $P_i^1 := (s_i, \ldots, x_i)$ and $P_i^2 := (y_i, \ldots, a_i)$. Given that q is left of the start point, it follows that $P_i^2 \cup U$ covers the entire real line from the start point to p. Hence, since P_j contains a target at p, it follows that $r(s_j) < p$ and as such s_j is adjacent to $P_i^2 \cup U$.

If P_i^1 is adjacent to $P_j - s_j$, then replace P_j with $\{s_j\} \cup P_i^2 \cup U \cup \{t_j\}$ and P_i with $P_i^1 \cup (P_j - \{s_j, t_j\})$. In this case p is further right, since the new P_i ends further right than p due to P_j appearing twice at p.

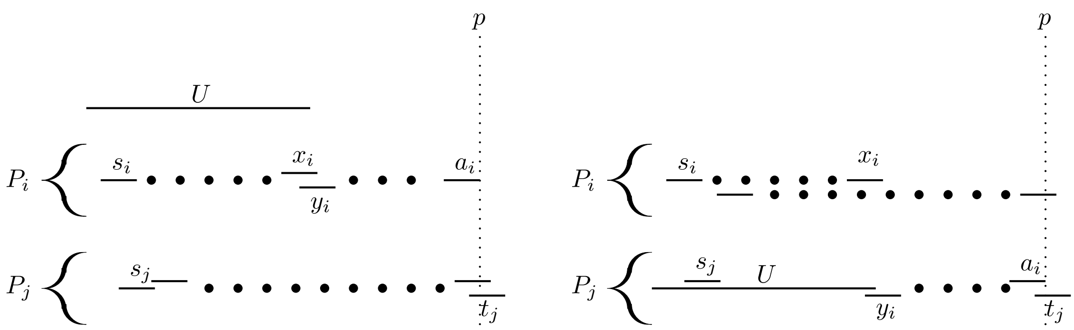

Figure 9.3: Given that q is left of the start point, it is easy to place $P_i^2 \cup U$ into P_j. This leaves $P_j - \{s_j, t_j\}$ available for P_i. When $P_j - s_j$ is adjacent to P_i^1, this is sufficient. Also note the second reverse point is not used in this case.

However, it is possible that P_i^1 is not adjacent to $P_j - s_j$. This means that $l(P_j - s_j) > r(x_i)$, and since s_j is at $l(P_j - s_j)$, it follows s_j is adjacent to some vertex of P_i^2. Since P_j appears twice at q', the subpath $P_j - s_j$ is adjacent to $P_i^1 \cup U'$. Hence replace P_j with

$\{s_j\} \cup P_i^2 \cup \{t_j\}$ and replace P_i with $P_i^1 \cup U' \cup (P_j - \{s_j, t_j\})$. The point p has once again been moved further right, contradicting our choice of $\mathcal{P}$. $\square$

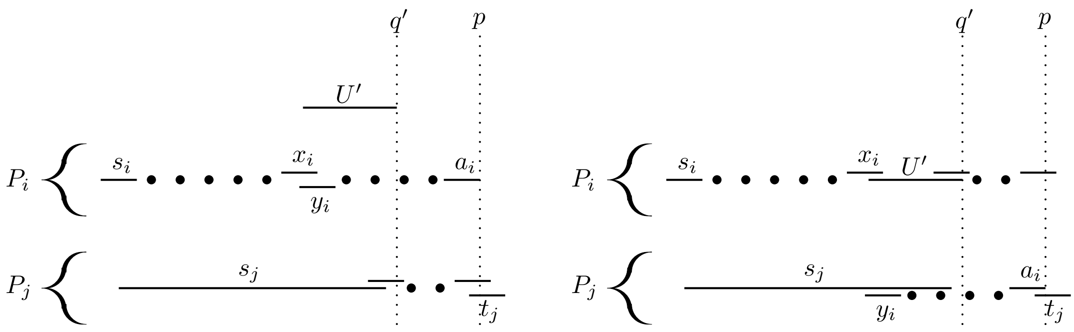

Figure 9.4: Using the second reverse point, it is possible to place P_i^2 into P_j and then $P_j - \{s_j, t_j\}$ into P_i, maintaining connectivity with U'.

Given that reverse point q must exist and have either type 1 or type 2, Lemma 9.8 and Lemma 9.9 are sufficient to prove that a finished $\mathcal{P}$ must exist. This proves Theorem 1.13.

The connectivity requirement in Theorem 1.13 is tight. Let k be even, and consider $G := P_n^{\frac{3k}{2}-1}$, where $n = 2k + \frac{3k}{2} - 1$. Label the vertices (in order) $s_1, \ldots, s_k, a_1, \ldots, a_{\frac{3k}{2}-1}, t_{\frac{k}{2}+1}, \ldots, t_k, t_1, \ldots, t_{\frac{k}{2}}$. The graph G is not $\frac{3k}{2}$-connected, and no s_i and t_j are adjacent, so if we show $s_1, \ldots, s_k, t_1, \ldots, t_k$ cannot be linked, then Theorem 1.13 is tight.

Suppose for the sake of a contradiction there is a linkage $P_1, \ldots, P_k$. Each P_i in the linkage contains s_i and t_i, but all other vertices in P_i are labelled by some a_j. Since no source is adjacent to no target, each P_i contains at least three vertices.

Say there exists some a_j adjacent to both s_i and t_i where $i \leq \frac{k}{2}$. Now since a_j is adjacent to s_i, it follows that $j \leq \frac{k}{2} + i - 1$. Since a_j is adjacent to t_i, it follows that $j \geq \frac{k}{2} + i$. So there is no such a_j. As a result, each P_i contains four vertices when $i \leq \frac{k}{2}$. Thus the linkage contains at least $3\frac{k}{2} + 4\frac{k}{2}$ vertices in total, but this is greater than $n = 2k + \frac{3k}{2} - 1$, the number of vertices in G. Hence $s_1, \ldots, s_k, t_1, \ldots, t_k$ cannot be linked.

Theorem 1.13 is a restricted extension of Lemma 9.2 for the broader class of interval graphs. We believe that it should be possible to extend the entire lemma, and as such make the following conjecture.

Conjecture 9.10. *Let $s_1, \ldots, s_k, t_1, \ldots, t_k$ be $2k$ distinct vertices in $(k + r)$-connected interval graph G (with $r > 0$), such that s_i is left of t_i for all i, and such that the distance* *between* any *source and* any *target in G is at least $\frac{k}{r}$. Then $s_1, \ldots, s_k, t_1, \ldots, t_k$ can be linked.*

Our motivation for Theorem 1.13 comes from Theorem 1.12. Recall that given a circular arc graph, we can delete the vertices at a point and obtain an interval graph. Constructing a linkage along this interval was a key step in trying to construct a complete minor in our proof of Theorem 1.12. However, the requirement that source and target vertices be non-adjacent (or even further apart) essentially forces the cover number of the circular arc graph, β, to be large. (This is to say, if there is a small vertex cover of the circular arc graph, then there is probably a short path between any source and any target, depending on where we "cut" the circle.) As a result of this Conjecture 9.10, while interesting, would not really assist in Chapter 8. Any attempt to prove Hadwiger's Conjecture for circular arc graphs when β is small will require some alternate approach.

Finally, it is also worth asking whether Theorem 1.13 can be extended to the more general class of chordal graphs. Unfortunately, the answer to this question is no. We provide a counterexample that is $(2k-2)$-connected but not k-linked, even when sources and targets are not adjacent. (This is essentially a more general example that Lemma 9.1 is tight.)

Given the connection between chordal graphs and treewidth as seen in Section 2.4 and the proof of Lemma 9.1, we shall exhibit this graph as a tree decomposition. The vertices of G are exactly the vertices that appear in at least one bag of the tree decomposition, and $E(G)$ contains all acceptable edges for the tree decomposition—if two vertices share a bag, then they are adjacent. Recall that as long as every two adjacent bags in the tree decomposition contain at least $2k-2$ vertices in common, then the chordal graph G that arises from this tree decomposition is $(2k-2)$-connected, since deleting any set of less than $2k-2$ vertices leaves a vertex present in each pair of adjacent bags which is enough to ensure connectivity.

Let the underlying tree T be the star with k leaves, such that each edge has been subdivided $k-1$ times. The vertices of G are as follows:

- sources $s_1, \ldots, s_k$,
- targets $t_1, \ldots, t_k$,
- $a_1, \ldots, a_{k-1}$, called the a-vertices,
- $b^i_1, \ldots, b^i_k$ for all $1 \leq i \leq k$.

In the bag indexed by the centre node, which we label C, place the vertices $\{s_1, \ldots, s_k, a_1, \ldots, a_{k-1}\}$. Denote the path in T from the centre node to the i^{th} leaf as the i^{th} path. Let B_i denote the bag adjacent to C in the i^{th} path. The bag B_i contains all vertices of C except s_i, and also includes vertex b_1^i. For each subsequent bag on the i^{th} path, remove a source vertex and add a vertex from $b_1^i, \ldots, b_k^i$, until the final bag on the i^{th} path contains $\{a_1, \ldots, a_{k-1}, b_1^i, \ldots, b_k^i\}$. To this final bag, also add t_i. Given the way we have constructed these bags, it is clear that adjacent bags contain at least $2k-2$ vertices in common. Hence the chordal graph that arises from this tree decomposition is $(2k-2)$-connected.

However, we cannot link each s_i to t_i in this graph G. Note that s_i does not appear in any bag along the i^{th} path, and t_i only appears at the end of the i^{th} path. Thus any path from s_i to t_i contains a vertex in $C \cap B_i = \{s_1, \ldots, s_k, a_1, \ldots, a_{k-1}\} - \{s_i\}$, since this set of vertices separates the graph such that s_i and t_i are in different components. Given that this path cannot use another source, it must contain an a-vertex. But there are k paths, and only $(k-1)$ a-vertices, and as such we cannot link these sources and targets, even though they are not adjacent.

9.4 Hadwiger Number of the Power of a Cycle

Recall the k^{th}-*power of a cycle* C_n^k is the graph formed by taking a cycle and adding edges between any two vertices at distance at most k.

In Chapter 1, we proved a lower bound on $\mathsf{had}(C_n^k)$ when $n \equiv 1 \mod k$ and $n \geq 2k+1$. As promised, we prove a lower bound that is independent of the modulus of n, using the results of Section 9.2.

Lemma 9.11. *If $n \geq k^2 + 2k$, then $\mathsf{had}(C_n^k) \geq 2k$.*

Proof. Label the vertices of the power of the cycle $1, \ldots, n$ clockwise. The graph $C_n^k - \{1, \ldots, k+1\}$ is isomorphic to the power of a path P_{n-k-1}^k. Given Lemma 9.4, we construct a set of $k-1$ paths in P_{n-k-1}^k (and thus in C_n^k) from $k+2, \ldots, 2k$ to $n-k+2, \ldots, n$ respectively. These paths form $k-1$ branch sets of the K_{2k} model. (It is clear that these branch sets are adjacent to each other.) The remaining $k+1$ branch sets are singleton branch sets $\{\{1,\}, \ldots, \{k+1\}\}$. (Note these vertices form a clique.) Each vertex $i \in 1, \ldots, k$ is adjacent to $\{n-k+i, \ldots, n\}$ and $\{k+2, \ldots, k+i\}$, and as such is adjacent to every one of the paths. Finally, the vertex $k+1$ is adjacent to $k+2, \ldots, 2k+1$ and so is also adjacent to each one of the paths. This gives a complete model with $k+1+k-1 = 2k$ branch sets, as required. $\square$

Note that the above minor is very similar to the minor we attempted to construct in Chapter 8, but much stronger given what we know about the power of a cycle.

Recall we proved in Section 1.4 that $\mathsf{pw}(C_n^k) \leq 2k$. For large n, this means

$$2k - 1 \leq \mathsf{had}(C_n^k) - 1 \leq \mathsf{tw}(C_n^k) \leq \mathsf{pw}(C_n^k) \leq 2k.$$

Thus Lemma 9.11 is almost best possible.

Bibliography

[1] Albertson, M. O., Chappell, G. G., Kierstead, H. A., Kündgen, A., and Ramamurthi, R. (2004). Coloring with no 2-colored P_4's. *Electron. J. Combin.*, **11 #R26**.

[2] Appel, K. and Haken, W. (1976). A proof of the four color theorem. *Discrete Math.*, **16**(2), 179–180.

[3] Atserias, A. (2008). On digraph coloring problems and treewidth duality. *European J. Combin.*, **29**(4), 796–820.

[4] Belkale, N. and Chandran, L. S. (2009). Hadwiger's conjecture for proper circular arc graphs. *European J. Combin.*, **30**(4), 946–956.

[5] Bellenbaum, P. and Diestel, R. (2002). Two short proofs concerning tree-decompositions. *Combin. Probab. Comput.*, **11**(6), 541–547.

[6] Bodlaender, H. L. (1993). A tourist guide through treewidth. *Acta Cybernet.*, **11**(1-2), 1–21.

[7] Bodlaender, H. L. (1998). A partial k-arboretum of graphs with bounded treewidth. *Theoret. Comput. Sci.*, **209**(1-2), 1–45.

[8] Bodlaender, H. L., Grigoriev, A., and Koster, A. M. C. A. (2008). Treewidth lower bounds with brambles. *Algorithmica*, **51**(1), 81–98.

[9] Böhme, T., Gerlach, T., and Stiebitz, M. (2008). Ordered and linked chordal graphs. *Discuss. Math. Graph Theory*, **28**(2), 367–373.

[10] Bollobás, B. and Thomason, A. (1996). Highly linked graphs. *Combinatorica*, **16**(3), 313–320.

[11] Călinescu, G., Fernandes, C., and Reed, B. A. (2003). Multicuts in unweighted graphs and digraphs with bounded degree and bounded tree-width. *J. Algorithms*, **48**(2), 333–359.

[12] Catlin, P. A. (1979). Hajós' graph-coloring conjecture: variations and counterexamples. *J. Combin. Theory Ser. B*, **26**(2), 268–274.

[13] Chambers, J. (2002). *Hunting for Torus Obstructions*. M.Sc. thesis, Department of Computer Science, University of Victoria, Canada.

[14] Chekuri, C. and Chuzhoy, J. (2013). Polynomial bounds for the grid-minor theorem. arXiv: 1305.6577.

[15] Chudnovsky, M. and Ovetsky Fradkin, A. (2008). Hadwiger's conjecture for quasi-line graphs. *J. Graph Theory*, **59**(1), 17–33.

[16] Chudnovsky, M., Reed, B. A., and Seymour, P. D. (2011). The edge-density for $K_{2,t}$ minors. *J. Combin. Theory Ser. B*, **101**(1), 18–46.

[17] Colin de Verdière, Y. (1998). Multiplicities of eigenvalues and tree-width of graphs. *J. Combin. Theory Ser. B*, **74**, 121–146.

[18] Curtis, A., Lin, M. C., McConnell, R., Nussbaum, Y., Soulignac, F. J., Spinrad, J., and Szwarcfiter, J. L. (2013). Isomorphism of graph classes related to the circular-ones property. *Discrete Math. Theor. Comput. Sci.*, **15**(1).

[19] DeVos, M., Dvořák, Z., Fox, J., McDonald, J., Mohar, B., and Scheide, D. (2014). A minimum degree condition forcing complete graph immersion. *Combinatorica, in print*.

[20] Di Giacomo, E., Didimo, W., Liotta, G., and Montecchiani, F. (2012). h-Quasi planar drawings of bounded treewidth graphs in linear area. In *Proceedings of the 38th International Conference on Graph-Theoretic Concepts in Computer Science*, WG'12, pages 91–102, Berlin. Springer-Verlag.

[21] Díaz, J., Serna, M., and Wormald, N. (2007). Bounds on the bisection width for random d-regular graphs. *Theoret. Comput. Sci.*, **382**(2), 120–130.

[22] Diestel, R. (2000). *Graph Theory*. Springer.

[23] Diestel, R., Jensen, T. R., Gorbunov, K. Y., and Thomassen, C. (1999). Highly connected sets and the excluded grid theorem. *J. Combin. Theory Ser. B*, **75**(1), 61–73.

[24] Dirac, G. A. (1952). A property of 4-chromatic graphs and some remarks on critical graphs. *J. London Math. Soc.*, **27**, 85–92.

[25] Dujmović, V., Morin, P., and Wood, D. R. (2005). Layout of graphs with bounded tree-width. *SIAM J. Comput.*, **34**(3), 553–579.

[26] Dujmović, V., Harvey, D. J., Joret, G., Reed, B. A., and Wood, D. R. (2013). A linear-time algorithm for finding a complete graph minor in a dense graph. *SIAM J. Discrete Math.*, **27**(4), 1770–1774.

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

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

[29] Flum, J. and Grohe, M. (2006). *Parameterized Complexity Theory*. Texts in Theoretical Computer Science. An EATCS Series. Springer.

[30] Fox, J. (2011). Constructing dense graphs with sublinear Hadwiger number. *J. Combin. Theory Ser. B* (to appear). arXiv: 1108.4953.

[31] Frankl, P. (1984). A new short proof for the Kruskal-Katona theorem. *Discrete Math.*, **48**(2-3), 327–329.

[32] Fulkerson, D. R. and Gross, O. A. (1965). Incidence matrices and interval graphs. *Pacific J. Math.*, **15**, 835–855.

[33] Ganian, R., Hliněnỳ, P., Langer, A., Obdržálek, J., Rossmanith, P., and Sikdar, S. (2014). Lower bounds on the complexity of MSO1 model-checking. *J. Comput. System Sciences*, **80**(1), 180–194.

[34] Gavril, F. (1974a). Algorithms on circular-arc graphs. *Networks*, **4**(4), 357–369.

[35] Gavril, F. (1974b). The intersection graphs of subtrees in trees are exactly the chordal graphs. *J. Combin. Theory Ser. B*, **16**, 47–56.

[36] Gerbner, D., Lemons, N., Palmer, C., Patkós, B., and Szécsi, V. (2012). Almost intersecting families of sets. *SIAM. J. Discrete Math.*, **26**(4), 1657–1669.

[37] Gerbner, D., Lemons, N., Palmer, C., Pálvölgyi, D., Patkós, B., and Szécsi, V. (2013). Almost cross-intersecting and almost cross-Sperner pairs of families of sets. *Graph. Combin.*, **29**, 489–498.

[38] Grippo, L. N. and Safe, M. D. (2012). On circular-arc graphs having a model with no three arcs covering the circle. In *Anais do XLIV Simpósio Brasileiro de Pesquisa Operacional*, pages 4093–4104. arXiv: 1402.2641.

[39] Grohe, M. and Marx, D. (2009). On tree width, bramble size, and expansion. *J. Combin. Theory Ser. B*, **99**(1), 218–228.

[40] Hadwiger, H. (1943). Über eine Klassifikation der Streckenkomplexe. *Vierteljschr. Naturforsch. Ges. Zürich*, **88**, 133–142.

[41] Halin, R. (1976). S-functions for graphs. *J. Geometry*, **8**(1-2), 171–186.

[42] Harvey, D. J. and Wood, D. R. (2012). Treewidth of line graphs. arXiv: 1210.8205.

[43] Harvey, D. J. and Wood, D. R. (2013). Parameters tied to treewidth. arXiv: 1312.3401.

[44] Harvey, D. J. and Wood, D. R. (2014a). Treewidth of the Kneser graph and the Erdős-Ko-Rado theorem. *Electron. J. Combin.*, **21**(1), P1.48.

[45] Harvey, D. J. and Wood, D. R. (2014b). Treewidth of the line graph of a complete graph. *J. Graph Theory* (to appear).

[46] Jensen, T. R. and Toft, B. (1995). *Graph Coloring Problems*. John Wiley.

[47] Jung, H. A. (1970). Eine Verallgemeinerung des n-fachen Zusammenhangs für Graphen. *Math. Ann.*, **187**, 95–103.

[48] Karapetjan, I. A. (1980). Coloring of arc graphs. *Akad. Nauk Armyan. SSR Dokl.*, **70**(5), 306–311.

[49] Katona, G. O. H. (1968). A theorem of finite sets. In *Theory of Graphs (Proc. Colloq., Tihany, 1966)*, pages 187–207. Academic Press, New York.

[50] Katona, G. O. H. (1972). A simple proof of the Erdős-Chao Ko-Rado theorem. *J. Combin. Theory Ser. B*, **13**, 183–184.

[51] Kawarabayashi, K. (2007). On the connectivity of minimum and minimal counterexamples to Hadwiger's Conjecture. *J. Combin. Theory Ser. B*, **97**(1), 144 – 150.

[52] Kawarabayashi, K. and Kobayashi, Y. (2012). Linear min-max relation between the treewidth of H-minor-free graphs and its largest grid minor. In *29th International Symposium on Theoretical Aspects of Computer Science*, volume 14 of *LIPIcs. Leibniz Int. Proc. Inform.*, pages 278–289. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern.

[53] Kawarabayashi, K. and Mohar, B. (2007). Some recent progress and applications in graph minor theory. *Graphs Combin.*, **23**(1), 1–46.

[54] Kawarabayashi, K., Kostochka, A. V., and Yu, G. (2006). On sufficient degree conditions for a graph to be k-linked. *Combin. Probab. Comput.*, **15**(5), 685–694.

[55] Kawarabayashi, K., Kobayashi, Y., and Reed, B. A. (2012). The disjoint paths problem in quadratic time. *J. Combin. Theory Ser. B*, **102**(2), 424–435.

[56] Kneser, M. (1955). Aufgabe 360. *Jahresber. Deutsch. Math.-Verein.*, **58:27**.

[57] Kostochka, A. V. (1982). The minimum Hadwiger number for graphs with a given mean degree of vertices. *Metody Diskret. Analiz.*, (38), 37–58.

[58] Kostochka, A. V. (1984). Lower bound of the Hadwiger number of graphs by their average degree. *Combinatorica*, **4**(4), 307–316.

[59] Kostochka, A. V. and Prince, N. (2008). On $K_{s,t}$-minors in graphs with given average degree. *Discrete Math.*, **308**(19), 4435–4445.

[60] Kostochka, A. V. and Prince, N. (2012). On $K_{s,t}$-minors in graphs with given average degree, II. *Discrete Math.*, **312**(24), 3517–3522.

[61] Kreutzer, S. (2012). On the parameterized intractability of monadic second-order logic. *Log. Methods Comput. Sci.*, **8**(1), 1:27, 35.

[62] Kreutzer, S. and Tazari, S. (2010a). Lower bounds for the complexity of monadic second-order logic. In *Proc. 25th Annual IEEE Symposium on Logic in Computer Science (LICS '10)*, pages 189–198. IEEE.

[63] Kreutzer, S. and Tazari, S. (2010b). On brambles, grid-like minors, and parameterized intractability of monadic second-order logic. In *Proc. 21st Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 354–364. SIAM.

[64] Kruskal, J. B. (1963). The number of simplices in a complex. In *Mathematical Optimization Techniques*, pages 251–278. Univ. of California Press, Berkeley, Calif.

[65] Kündgen, A. and Pelsmajer, M. J. (2008). Nonrepetitive colorings of graphs of bounded tree-width. *Discrete Math.*, **308**(19), 4473–4478.

[66] Kuratowski, K. (1930). Sur le probléme des courbes gauches en topologie. *Fund. Math.*, **16**, 271–283.

[67] Larman, D. G. and Mani, P. (1970). On the existence of certain configurations within graphs and the 1-skeletons of polytopes. *Proc. London Math. Soc. (3)*, **20**, 144–160.

[68] Li, D. and Liu, M. (2007). Hadwiger's conjecture for powers of cycles and their complements. *European J. Combin.*, **28**(4), 1152–1155.

[69] Lin, M. C., Soulignac, F. J., and Szwarcfiter, J. L. (2013). Normal helly circular-arc graphs and its subclasses. *Discrete Appl. Math.*, **161**(78), 1037 – 1059.

[70] Lipton, R. J. and Tarjan, R. E. (1979). A separator theorem for planar graphs. *SIAM J. Appl. Math.*, **36**(2), 177–189.

[71] Lovász, L. (1978). Kneser's conjecture, chromatic number, and homotopy. *J. Combin. Theory Ser. A*, **25**(3), 319–324.

[72] Lucena, B. (2007). Achievable sets, brambles, and sparse treewidth obstructions. *Discrete Appl. Math.*, **155**(8), 1055–1065.

[73] Mader, W. (1967). Homomorphieeigenschaften und mittlere Kantendichte von Graphen. *Math. Ann.*, **174**, 265–268.

[74] Mader, W. (1968). Homomorphiesätze für Graphen. *Math. Ann.*, **178**, 154–168.

[75] Marchal, L. (2012). *Treewidth*. Ph.D. thesis, Maastricht University, Netherlands.

[76] Markov, I. and Shi, Y. (2011). Constant-degree graph expansions that preserve treewidth. *Algorithmica*, **59**(4), 461–470.

[77] Marx, D. (2010). Can you beat treewidth? *Theory Comput.*, **6**, 85–112.

[78] Matsumoto, M. and Tokushige, N. (1989). The exact bound in the Erdős-Ko-Rado theorem for cross-intersecting families. *J. Combin. Theory Ser. A*, **52**(1), 90–97.

[79] Myers, J. S. and Thomason, A. (2005). The extremal function for noncomplete minors. *Combinatorica*, **25**(6), 725–753.

[80] Pyber, L. (1986). A new generalization of the Erdős-Ko-Rado theorem. *J. Combin. Theory Ser. A*, **43**(1), 85–90.

[81] Reed, B. A. (1992). Finding approximate separators and computing tree width quickly. In *Proceedings of the twenty-fourth annual ACM Symposium on Theory of Computing*, STOC '92, pages 221–228, New York, USA. ACM.

[82] Reed, B. A. (1997). Tree width and tangles: a new connectivity measure and some applications. In *Surveys in Combinatorics*, volume 241 of *London Math. Soc. Lecture Note Ser.*, pages 87–162. Cambridge Univ. Press.

[83] Reed, B. A. and Seymour, P. D. (1998). Fractional colouring and Hadwiger's conjecture. *J. Combin. Theory Ser. B*, **74**(2), 147–152.

[84] Reed, B. A. and Seymour, P. D. (2004). Hadwiger's conjecture for line graphs. *European J. Combin.*, **25**(6), 873–876.

[85] Reed, B. A. and Wood, D. R. (2009). A linear-time algorithm to find a separator in a graph excluding a minor. *ACM Trans. Algorithms*, **5**(4), Art. 39.

[86] Reed, B. A. and Wood, D. R. (2012). Polynomial treewidth forces a large grid-like-minor. *European J. Combin.*, **33**(3), 374–379.

[87] Reed, B. A. and Wood, D. R. (2014). Forcing a sparse minor. arXiv: 1402.0272.

[88] Robertson, N. and Seymour, P. D. (1983–2012). Graph minors I–XXIII. *J. Combin. Theory Ser. B*.

[89] Robertson, N. and Seymour, P. D. (1983). Graph minors. I. Excluding a forest. *J. Combin. Theory Ser. B*, **35**(1), 39–61.

[90] Robertson, N. and Seymour, P. D. (1986a). Graph minors. II. Algorithmic aspects of tree-width. *J. Algorithms*, **7**(3), 309–322.

[91] Robertson, N. and Seymour, P. D. (1986b). Graph minors. V. Excluding a planar graph. *J. Combin. Theory Ser. B*, **41**(1), 92–114.

[92] Robertson, N. and Seymour, P. D. (1991). Graph minors. X. Obstructions to tree-decomposition. *J. Combin. Theory Ser. B*, **52**(2), 153–190.

[93] Robertson, N. and Seymour, P. D. (1995). Graph minors. XIII. The disjoint paths problem. *J. Combin. Theory Ser. B*, **63**(1), 65–110.

[94] Robertson, N. and Seymour, P. D. (2003). Graph minors. XVI. Excluding a non-planar graph. *J. Combin. Theory Ser. B*, **89**(1), 43–76.

[95] Robertson, N., Seymour, P. D., and Thomas, R. (1993). Hadwiger's conjecture for K_6-free graphs. *Combinatorica*, **13**(3), 279–361.

[96] Robertson, N., Seymour, P. D., and Thomas, R. (1994). Quickly excluding a planar graph. *J. Combin. Theory Ser. B*, **62**(2), 323–348.

[97] Scheinerman, E. R. and Ullman, D. H. (1997). *Fractional Graph Theory*. Wiley.

[98] Schrijver, A. (2003). *Combinatorial Optimization: Polyhedra and Efficiency*. Algorithms and combinatorics. Springer.

[99] Scott, A. and Wilmer, E. (2013). Hypergraphs of bounded disjointness. arXiv: 1306.4236.

[100] Seymour, P. D. and Leaf, A. (2012). Treewidth and planar minors. `https://web.math.princeton.edu/~pds/papers/treewidth/paper.pdf`.

[101] Seymour, P. D. and Thomas, R. (1993). Graph searching and a min-max theorem for tree-width. *J. Combin. Theory Ser. B*, **58**(1), 22–33.

[102] Tazari, S. and Müller-Hannemann, M. (2009). Shortest paths in linear time on minor-closed graph classes, with an application to Steiner tree approximation. *Discrete Appl. Math.*, **157**(4), 673–684.

[103] Thomas, R. and Wollan, P. (2005). An improved linear edge bound for graph linkages. *European J. Combin.*, **26**(3-4), 309–324.

[104] Thomason, A. (1984). An extremal function for contractions of graphs. *Math. Proc. Cambridge Philos. Soc.*, **95**(2), 261–265.

[105] Thomason, A. (2001). The extremal function for complete minors. *J. Combin. Theory Ser. B*, **81**(2), 318–338.

[106] Thomassen, C. (2005). Some remarks on Hajós' conjecture. *J. Combin. Theory Ser. B*, **93**(1), 95–105.

[107] Toft, B. (1996). A survey of Hadwiger's conjecture. *Congr. Numer.*, **115**, 249–283.

[108] Tucker, A. (1971). Matrix characterizations of circular-arc graphs. *Pacific J. Math.*, **39**, 535–545.

[109] Tucker, A. (1975). Coloring a family of circular arcs. *SIAM J. Appl. Math.*, **29**(3), 493–502.

[110] Valencia-Pabon, M. (2003). Revisiting Tucker's algorithm to color circular arc graphs. *SIAM J. Comput.*, **32**(4), 1067–1072.

[111] van der Holst, H. (1996). *Topological and Spectral Graph Characterizations*. Ph.D. thesis, Amsterdam University, Netherlands.

[112] Wagner, K. (1937a). Über eine Eigenschaft der ebene Komplexe. *Math. Ann.*, **114**, 570–590.

[113] Wagner, K. (1937b). Über eine Erweiterung des Satzes von Kuratowski. *Deutsche Math.*, **2**, 280–285.

[114] Wood, D. R. (2011). Clique minors in Cartesian products of graphs. *New York J. Math.*, **17**, 627–682.

[115] Wulff-Nilsen, C. (2010). Faster shortest path algorithm for H-minor free graphs with negative edge weights. *CoRR*. arXiv: 1008.1048.

[116] Wulff-Nilsen, C. (2011). Separator theorems for minor-free and shallow minor-free graphs with applications. In *Proc. 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 37–46. IEEE.

[117] Yuster, R. and Zwick, U. (2007). Maximum matching in graphs with an excluded minor. In *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 108–117, New York. ACM.

[118] Ziegler, G. M. (2001). Generalized Kneser coloring theorems with combinatorial proofs. *Inventiones Math*, **147**, 671–691.

Index