Local properties of graphs with large chromatic number

by

James Davies

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Combinatorics and Optimization

Waterloo, Ontario, Canada, 2022

Examining Committee Membership

The following served on the Examining Committee for this thesis. The decision of the Examining Committee is by majority vote.

External Examiner: Alex Scott
Professor, Mathematical Institute,
University of Oxford

Supervisor(s): Jim Geelen
Professor, Dept. of Combinatorics and Optimization,
University of Waterloo

Internal Member: Sophie Spirkl
Assistant Professor, Dept. of Combinatorics and Optimization,
University of Waterloo

Internal Member: Penny Haxell
Professor, Dept. of Combinatorics and Optimization,
University of Waterloo

Internal-External Member: Anna Lubiw
Professor, Cheriton School of Computer Science,
University of Waterloo

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

This thesis deals with problems concerning the local properties of graphs with large chromatic number in hereditary classes of graphs.

We construct intersection graphs of axis-aligned boxes and of lines in $\mathbb{R}^3$ that have arbitrarily large girth and chromatic number. We also prove that the maximum chromatic number of a circle graph with clique number at most ω is equal to $\Theta(\omega \log \omega)$. Lastly, extending the χ-boundedness of circle graphs, we prove a conjecture of Geelen that every proper vertex-minor-closed class of graphs is χ-bounded.

Acknowledgements

First and foremost, I would like to thank my supervisor, Jim Geelen for his mentorship. His enthusiasm and constant encouragement made this thesis possible. I am also grateful for the numerous interesting problems Jim would share, despite knowing how distracted a good new problem could make me.

I would also like to thank my committee members, Alex Scott, Sophie Spirkl, Penny Haxell, and Anna Lubiw for their careful reading and comments, which improved the thesis.

I thank Bartosz Walczak, Tomasz Krawczyk, Michał Pilipczuk, Marcin Briański, and Rose McCarty for various fun visits working in Poland. I also thank all my other collaborators that I have been fortunate enough to have the joy and privilege of working with during my time at Waterloo.

The friends I have made here the last four years have made this truly a good time in my life. Sabrina Lato in particular has been a constant source of companionship and support, and, despite everything, made the lockdowns enjoyable. I also thank Tim Miller, Lukas Nabergall, Joshua Gunter, Rose McCarty, Nick Olsen-Harris, Martin Pei, Soffía Árnadóttir, Tina Chen, Ben Moore, Matthew Kroeker, Amena Assem, Fernanda Rivera, Mathieu Rundström, Lise Turner, Kelvin Chan, Alan Wong, Cynthia Rodríguez, Kazuhiro Nomoto, Rutger Campbell, and many others.

Initially, I was hesitant to start a PhD so far from home. But I am so glad I did, and I thank for Dan Král' for his encouragement to do so.

Last but not least, I thank my family for their love and support and for encouraging my love of maths, although they may have gone further than intended there.

Dedication

This thesis is dedicated to my parents and grandparents.

Table of Contents

Chapter 1

Introduction

What causes a graph to have large chromatic number? It is well understood what makes a graph 2-colourable; these are exactly the graphs without an odd cycle. However already what distinguishes 3-colourable graphs and graphs with huge chromatic number is not well understood. Certainly we can not expect any answer as simple as for 2-colourable graphs. Even distinguishing whether a graph is either 1000-colourable or not 3-colourable is an NP-complete problem [46].

There are however some simple obstructions a graph may have to being 3-colourable. For example it is straightforward to check whether or not a graph contains a 4-vertex clique, and since the vertices of a clique must all receive distinct colours, this verifies that the graph is not 3-colourable. More generally the chromatic number of a graph is lower bounded by its clique number. However, the clique number of a graph is not enough to determine, or even to obtain an upper bound for the chromatic number of a graph; Tutte [31, 32] constructed triangle-free graphs with arbitrarily large chromatic number.

Despite this, it's well understood when the clique number does exactly determine chromatic number for a graph and its induced subgraphs. A ***perfect graph*** is a graph for which every induced subgraph has its chromatic number exactly equal to its clique number. In addition to odd cycles of length at least 5, their complements also provide minimal examples of non-perfect graphs. The strong perfect graph theorem of Chudnovsky, Robertson, Seymour, and Thomas [19] states that these are the only minimal non-perfect graphs. Beyond this, it is already much less well understood when the chromatic number of a graph is at most one more than the clique number. However, it is at least known that this class of graphs does contain all line graphs by Vizing's theorem [106].

The property of when the clique number provides an upper bound for the chromatic

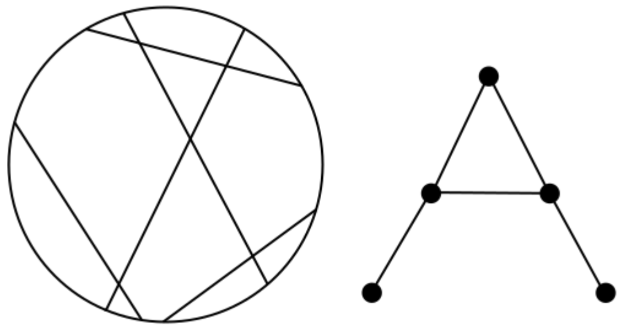

Figure 1.1: A chord diagram and its corresponding circle graph.

number of a graph in a class is formalized by the notion of χ-boundedness as introduced by Gyárfás [57]. For a graph G, the chromatic number and clique number are denoted by $\chi(G)$ and $\omega(G)$ respectively. A class of graphs $\mathcal{G}$ is ***χ-bounded*** if there exists a function $f\colon \mathbb{N} \to \mathbb{N}$ such that $\chi(G) \leq f(\omega(G))$ for all $G \in \mathcal{G}$. We call such an f the ***χ-bounding function***. We remark that χ-bounded classes of graphs have sometimes instead been called "near perfect graphs" due to the fact that such classes generalize perfect graphs. Since every graph is a subgraph of a sufficiently large clique, it is most natural to consider χ-boundedness of classes that are closed under taking induced subgraphs rather than subgraphs. We call such classes ***hereditary***.

A ***circle graph*** is a graph whose vertices correspond to chords of a circle and where two vertices are adjacent whenever their corresponding chords intersect (see Figure 1.1 for an example). In addition to perfect graphs and line graphs, circle graphs provide another example of a χ-bounded class of graphs. This was proved by Gyárfás [56] and is one the most classical results in the study of both χ-boundedness and geometric intersection graphs.

For $k \geq 3$ it is NP-complete to determine if a graph is k-colourable [46], and this holds even when restricted to many hereditary classes of graphs such as line graphs [61] and circle graphs [45]. On the other hand, for many classes such as line graphs and circle graphs [47], the clique number can be computed in polynomial time. For χ-bounded classes, this can allow for a polynomial time approximation algorithm for the chromatic number based on the clique number and the known χ-bounding function. For line graphs this gives a very good polynomial time approximation algorithm since the clique number and chromatic number differ by at most one. While this level of accuracy is not always possible, it is still desirable to have a good, say polynomial approximation for the chromatic number. This can be achieved if the χ-bounding function is at most some polynomial. We say that such classes are ***polynomially χ-bounded***. Not all χ-bounded classes are polynomially

χ-bounded [12], and polynomial χ-bounding functions are often much harder to obtain.

Circle graphs are generalized by intersection graphs of line segments in the plane. Answering a question of Erdős, Pawlik et al [85] proved that segment intersection graphs are not χ-bounded. Despite this, in some sense they are not far off from being χ-bounded. The notion of χ-boundedness effectively captures the idea of when huge chromatic number is essentially an "extremely local" property, but what if we relax our notion of "local" here? Then we may hope to find that despite our class of graphs not being χ-bounded, huge chromatic number is still a "local" property.

Scott [94] introduced the notion of ρ-control which effectively captures the idea of when huge chromatic number is a local property for an entire class of graphs. For a graph G and positive integer ρ, we let $\chi^{(\rho)}(G)$ denote the maximum chromatic number of an induced subgraph of G with radius at most ρ. A class of graphs $\mathcal{G}$ is ***ρ-controlled*** if there exists a function f such that $\chi(G) \leq f(\chi^{(\rho)}(G))$ for all $G \in \mathcal{G}$. It is easy to prove inductively that a hereditary class of graphs is 1-controlled if and only if it is χ-bounded.

Chudnovsky, Scott, and Seymour [21] proved that segment intersection graphs are 2-controlled, so despite not being χ-bounded [85], huge chromatic number is still a local property for segment intersection graphs. This is a non-trivial result in general because Erdős [36] proved that there are graphs with arbitrarily large girth and chromatic number, where the ***girth*** of a graph is equal to the length of its shortest cycle. If G has girth at least $2r+2$, then every induced subgraph of radius at most r is a tree. So in particular, for every positive integer ρ, there are graphs G with arbitrarily large chromatic number such that $\chi^{(\rho)}(G) = 2$. Already ρ-control is a notion of its own independent interest, but the idea of first proving ρ-control as an intermediate step towards obtaining χ-boundedness has also been used to resolve a number of conjectures on χ-boundedness in recent years [97].

In the next few sections we introduce and discuss the contributions of this thesis. The highlights are briefly summarized as follows.

- In Chapter 2 we construct intersection graphs of axis-aligned boxes and of lines in $\mathbb{R}^3$ that have arbitrarily large girth and chromatic number. These are the first non-trivial examples of geometric intersection graphs with large girth and chromatic number.
- In Chapter 3 we improve the known χ-bounding function for circle graphs to $O(\omega \log \omega)$, and within a small constant factor of the $\Omega(\omega \log \omega)$ lower bound construction of Kostochka [67]. We also improve this lower bound.
- In Chapter 4, generalizing the χ-boundedness of circle graphs, we prove a conjecture of Geelen (see [34]) that every proper vertex-minor-closed class of graphs is χ-bounded. As an intermediate result, we also prove that they are linearly 2-controlled.

1.1 Constructions

Of course not all classes of graphs are χ-bounded, and it is important to amass an arsenal of graph constructions that can be used to show that particular classes of graphs are not χ-bounded. Constructions of triangle-free graphs with arbitrarily large chromatic number include Tutte's [31, 32], Zykov's [107], Mycielski's [79] constructions, Burling graphs [13], shift graphs [37], and Kneser graphs [74]. Tutte [31,32] gave the first such construction, and his construction is particularly malleable, allowing for a number of versatile modifications. So let us present Tutte's construction.

Tutte's construction builds a $(k+1)$-chromatic triangle-free graph G_{k+1} from a k-chromatic triangle-free graph G_k, as follows. The graph G_{k+1} has a stable set S such that each connected component of $G_{k+1} - S$ is isomorphic to G_k; we refer to these components as the copies of G_k. To ensure that G_{k+1} is triangle-free we require that each vertex in S is adjacent to at most one vertex in each copy of G_k. We also require that each vertex in each copy of G_k is adjacent to exactly one vertex in S. In fact, not only does G_{k+1} remain triangle-free, but if G_k has girth at least 6, then so does G_{k+1}. Now to get $\chi(G_{k+1}) > \chi(G_k)$ we want that for each k-colouring of S there is a copy of G_k whose neighbourhood in S is monochromatic. This will ensure that G_{k+1} has no k-colouring since every vertex in a copy of G_k has a neighbour in S and G_k has no $(k-1)$-colouring. Tutte achieves this by choosing $|S| > k(|V(G_k)| - 1)$, and adding a copy of G_k whose neighbourhood is $X \subseteq S$ for every $|V(G_k)|$-vertex subset X of S. This is clearly sufficient, since by the pigeonhole principle, any k-colouring of S will contain a monochromatic set of $|V(G_k)|$ vertices.

One malleable part of Tutte's construction is in the choice of $|S|$ and how copies of G_k are added by choosing their $|V(G_k)|$-vertex neighbourhoods in S. As observed by Toft [100], one way to view this step in Tutte's construction is that you consider an auxiliary complete $|V(G_k)|$-uniform hypergraph H on vertex set S, then add a copy of G_k for each hyperedge of H. Again, when $|S| > k(|V(G_k)| - 1)$, any k-colouring of S will have a monochromatic set of $|V(G_k)|$ vertices which corresponds to a monochromatic hyperedge of H for which we added a copy of G_k. So the important feature of the auxiliary $|V(G_k)|$-uniform hypergraph H used here is that every k-colouring of H has a monochromatic hyperedge, or in other words, H is not k-colourable. Tutte's construction guarantees this in the simplest way by taking H to be a large complete hypergraph. By additionally choosing H to be a hypergraph with girth at least $g/3$, Nešetřil and Rödl [81] observed that if G_k has girth at least g, then G_{k+1} will also have girth at least g (for comparison here, note that a complete uniform hypergraph has girth equal to 2). Taking as an input a graph G and an auxiliary $|V(G)|$-uniform hypergraph H, we let $\mathrm{T}(G, H)$ be the operation of performing the hypergraph variation of Tutte's construction on the graph G, using H as the auxiliary hypergraph. By the above discussion, we have the following.

Lemma 1.1.1. *Let G be a graph with chromatic number at least k and girth at least g, and let H be a $|V(G)|$-uniform hypergraph with chromatic number greater than k and girth at least $g/3$. Then any graph G^* obtainable from the operation $\mathrm{T}(G, H)$ has chromatic number greater than k and girth at least g.*

By further modifying Tutte's construction to also work for n-uniform hypergraphs, rather than just 2-uniform hypergraphs (graphs), Nešetřil and Rödl [81] used this to give a short constructive proof that there are graphs with arbitrarily large girth and chromatic number.

Surprisingly, it appears that this freedom in choosing the auxiliary hypergraph had not been exploited beyond this. We explore this further by choosing our hypergraphs to be ones based on Gallai's theorem [88], a generalization of Van der Waerden's theorem [105] on monochromatic arithmetic progressions. By further using a spare variant of Gallai's Theorem due to Prömel and Voigt [86] we can even ensure that the hypergraphs arising from Gallai's theorem have large girth. This modification appears to be particularly well suited to geometric constructions; in this thesis we give two applications to box and line intersection graphs in $\mathbb{R}^3$. These provide the first non-trivial examples of geometric intersection graphs with arbitrarily large girth and chromatic number. Let us call a class of graphs containing graphs with arbitrarily large girth and chromatic number ***χ-amorphous***; these are the classes where large chromatic is in no sense a local property.

We discuss each of these applications further, but before that, let us mention that there are also other malleable parts of Tutte's construction that we do not explore in this thesis. One example of this that has recently received a lot of attention is that Tutte's construction is well suited for adding in edges to ensure that induced subgraphs with large chromatic number contain certain induced subgraphs [12,15,52,63]. For instance this was used in the recent proofs that there are K_4-free graphs with arbitrarily large chromatic number and with no 5-chromatic triangle-free induced subgraph [15], and also that there are hereditary χ-bounded classes of graphs that are not polynomially χ-bounded [12].

The first application concerns intersection graphs of axis-aligned boxes. The ***intersection graph*** of a collection of sets $\mathcal{C}$ is the graph with vertex set $\mathcal{C}$ where two elements C, C' of $\mathcal{C}$ are adjacent if C and C' intersect. In 1948, Bielecki [4] asked if triangle-free intersection graphs of axis-aligned rectangles in the plane have bounded chromatic number. This was answered positively in 1960 by Asplund and Grünbaum [2], who proved more generally that intersection graphs of axis-aligned rectangles with clique number ω are $O(\omega^2)$-colourable, making this one of the first results on χ-boundedness. Recently

Chalermsook and Walczak [16] improved the bound to $O(\omega \log \omega)$. The best known lower bound is that there are rectangle intersection graphs with clique number ω and chromatic number 3ω [66].

Surprisingly, the situation is different in $\mathbb{R}^3$. In 1965, Burling [13] proved that there are axis-aligned boxes in $\mathbb{R}^3$ whose intersection graphs are triangle-free and have arbitrarily large chromatic number. In light of Burling's construction, the problem of whether intersection graphs of axis-aligned boxes in $\mathbb{R}^3$ with large girth have bounded chromatic number was raised by Kostochka and Perepelitsa [71]. Later Kostochka [66] further speculated, for all positive integers d, that intersection graphs of axis-aligned boxes in $\mathbb{R}^d$ with girth at least 5 should have bounded chromatic number.

We prove that there are intersection graphs of axis-aligned boxes in $\mathbb{R}^3$ with arbitrarily large girth and chromatic number. In other words, we prove that these graphs are χ-amorphous, and so there is no hope for any weakening of χ-boundedness such as ρ-control to hold.

Theorem 1.1.2. *There are intersection graphs of axis-aligned boxes in* $\mathbb{R}^3$ *with arbitrarily large girth and chromatic number.*

One pleasing consequence of the connection to arithmetic Ramsey theory used in the proof of Theorem 1.1.2 is that we may obtain Burling's [13] classical result as an application of Van der Waerden's theorem [105] on arithmetic progressions.

Despite not being χ-bounded [85], segment intersection graphs having huge chromatic number is still in some sense a local property. Kostochka and Nešetřil [69] proved that if the intersection graph of n segments in the plane has girth at least 5, then the graph has only $O(n)$ edges, consequently such graphs have bounded chromatic number. Fox and Pach [44] asked if this could be extended to segments in $\mathbb{R}^3$. As a second application, we answer this question with an defnatic no; even the subclass of line intersection graphs is χ-amorphous.

Theorem 1.1.3. *There are intersection graphs of straight lines in* $\mathbb{R}^3$ *with arbitrarily large girth and chromatic number.*

The problem of whether intersection graphs of lines in $\mathbb{R}^3$ are χ-bounded had been raised explicitly by Pach, Tardos, and Tóth [84], but had been circulating in the community a few years prior. Norin (see [22]) answered this in the negative by showing that double shift graphs (which are girth-4 graphs with large chromatic number [37]) are line intersection graphs. Thus Theorem 1.1.3 extends Norin's result.

1.2 Circle graphs

A ***circle graph*** is an intersection graph of chords on a circle (see Figure 1.1). Gyárfás [56] proved that circle graphs are χ-bounded with χ-bounding function $2^\omega(2^\omega - 2)\omega^2$ and asked [57, 58] for improved χ-bounding functions. In particular Gyárfás [57] originally asked if a linear χ-bounding function was possible. This was answered in the negative by Kostochka [67] who gave the superlinear lower bound of $\frac{1}{2}\omega(\ln\omega - 2)$. Kostochka [67] also improved the χ-bounding function to $2^\omega\omega(\omega+2)$ and later Kostochka and Kratochvíl [68] improved this further to $50 \cdot 2^\omega - 32\omega - 64$ for the more general class of ***polygon-circle graphs***, which are the intersection graphs of polygons inscribed in a circle.

Until very recently, a major open problem of Esperet [40] was to determine whether or not every hereditary χ-bounded class of graphs is polynomially χ-bounded. With Briański and Walczak [12], we very recently answered this in the negative. Before this, it was believed that circle graphs were a good candidate for a counter-example. Indeed, for over 30 years they had certainly been resilient to attempts to improve the χ-bounding function beyond the exponential 2^ω barrier. Introducing new techniques, with McCarty [29] we dispelled this widely held belief by proving that circle graphs are polynomially χ-bounded with a χ-bounding function of $\omega^2 + 2\omega\lceil 2\log_2\omega\rceil + 8\omega \leq 7\omega^2$.

In this thesis we further extend and refine the techniques introduced in [29] to improve the χ-bounding function for circle graphs to within a constant factor of Kostochka's [67] lower bound construction.

Theorem 1.2.1. *Every circle graph with clique number at most ω has chromatic number at most $2\omega\log_2(\omega) + 2\omega\log_2(\log_2(\omega)) + 10\omega$.*

Circle graphs and their representations are fundamental objects that appear in a diverse range of study. Some examples include knot theory [3,5], bioinformatics [60], quantum field theory [76], quantum computing [11, 104], and data structures [43]. On the more combinatorial side, in addition to discrete and computational geometry, circle graphs and their representations also appear in the study of continued fractions [103], vertex-minors [51], matroid representation [10], and in various sorting problems [54]. Circle graphs are also deeply related to planar graphs; the fundamental graphs of planar graphs are exactly the class of bipartite circle graphs [30]. Direct applications of colouring circle graphs include finding the minimum number of stacks needed to obtain a given permutation [42], solving routing problems such as in VLSI physical design [98], and finding stack layouts of graphs, which also has a number of additional applications of its own (see [33]).

With these applications in mind, it is desirable to have an efficient algorithm for colouring circle graphs. While their clique numbers can be found in polynomial time [47], unfortunately the problem of determining if a circle graph is k-colourable is NP-complete [45]. So the best that can be hoped for is an efficient approximation algorithm for the chromatic number. The proof of Theorem 1.2.1 is constructive and yields a practical polynomial time algorithm for colouring circle graphs with a colouring that is optimal up to at most a logarithmic factor of the chromatic number.

For completeness, we also provide a new simple lower bound construction for the χ-bounding function of circle graphs. As a bonus it improves Kostochka's [67] lower bound by a factor of 2.

Theorem 1.2.2. *For every positive integer ω there is a circle graph with clique number at most ω and chromatic number at least $\omega(\ln \omega - 2)$.*

For large clique number this leaves a constant factor of about $2\log_2(e) \approx 2.8854$ between the upper and lower bounds. These new upper and lower bounds are remarkably tight, but the difference between the logarithmic bases used in the upper and lower bounds is certainly curious.

For small clique number just one non-trivial tight bound is known; Kostochka [67] proved that triangle-free circle graphs are 5-colourable, and Ageev [1] constructed a triangle-free circle graph with chromatic number 5. In the next case, the best known upper bound is due to Nenashev [80] who proved that K_4-free circle graphs have chromatic number at most 30. The best lower bound for K_4-free circle graphs that we are aware of is a 6-chromatic graph arising from a modification of Ageev's [1] construction. By optimizing the proof of Theorem 1.2.1 to the $\omega = 3$ case, we improve the upper bound for K_4-free circle graphs to 19. As with Theorem 1.2.1, its proof can be made algorithmic.

Theorem 1.2.3. *Every K_4-free circle graph is 19-colourable.*

Other classes generalising circle graphs are now known to have polynomially χ-bounding functions. Krawczyk and Walczak [72] proved that if $f(\omega)$ is a χ-bounding function for circle graphs, then $f(\omega)\cdot\binom{\omega+1}{2}$ is a χ-bounding function for polygon-circle graphs and (more generally) interval filament graphs. So combining Theorem 1.2.1 with this result improves the best known χ-bounding function for interval filaments graphs to $O(\omega^3 \log \omega)$. Krawczyk and Walczak [72] proved that there are interval filament graphs with clique number ω and chromatic number $\binom{\omega+1}{2}$. An ***L-shape*** consists of a vertical segment and a horizontal segment that meet at their lowermost and leftmost endpoints respectively, so they form a

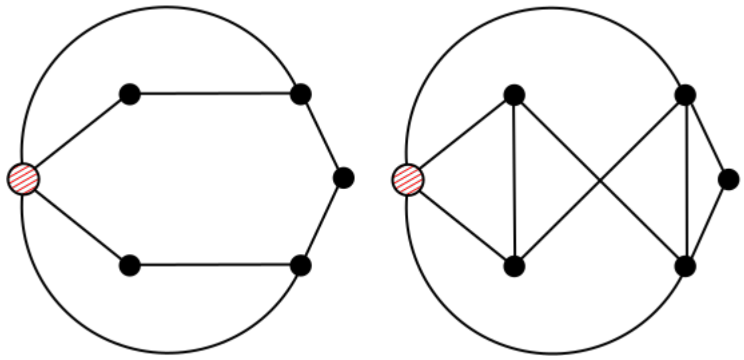

Figure 1.2: The effect of local complementation on the dashed red vertex.

"L". An L-shape is ***grounded*** if its uppermost point is on the x-axis. A ***grounded L-graph*** is an intersection graph of grounded L-shapes. Grounded L-graphs generalize circle graphs; it can be shown that circle graphs are exactly the intersection graphs of grounded L-shapes whose rightmost endpoints are on the $y = -x$ line [78]. Grounded L-graphs were introduced by McGuinness [78] who proved that they are χ-bounded with a χ-bounding function of $2^{O(4^{\omega})}$. With Krawczyk, McCarty and Walczak [28], again, by extending ideas of [29], we improved this doubly exponential bound to the polynomial bound of $17\omega^4$. Theorem 1.2.2 gives the best known lower bound construction for grounded L-graphs.

There are also many classes that contain circle graphs and are known to be χ-bounded (although they usually have extremely large χ-bounding functions). In terms of intersection graphs of grounded geometric objects, the most general are ***outer-string*** graphs, which are the intersection graphs of strings in the upper half-plane that have an endpoint on the x-axis. Rok and Walczak [93] proved that outer-string graphs are χ-bounded with χ-bounding function $2^{O(2^{\omega(\omega-1)/2})}$. Optimistically, we may hope that the ideas introduced in [29] and further developed here an in [28] could be pushed further to eventually prove that outer-string graphs are polynomially χ-bounded.

For further results on χ-boundedness of classes of graphs generalizing circle graphs, see [21, 24, 25, 27, 92, 95]. Two of these results; proper classes of graphs that are closed under taking either vertex-minors [25], or more generally pivot-minors [24] are the focus of the next section.

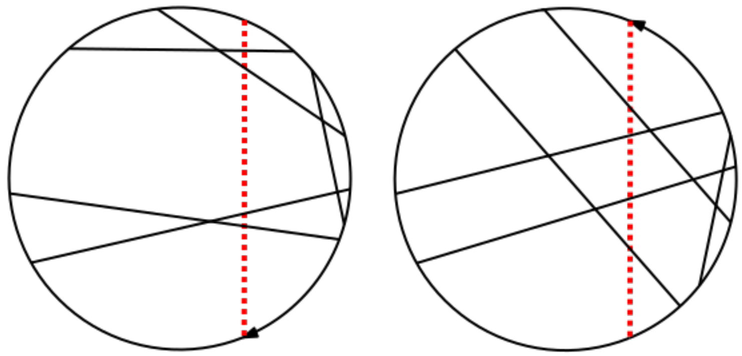

Figure 1.3: An illustration of how local complementation is performed on the corresponding chord diagram of a circle graph. The circle graphs of these two chord diagrams are exactly the graphs in Figure 1.2.

1.3 Vertex-minors

The action of performing ***local complementation*** at a vertex v in a graph G replaces the induced subgraph on $N(v)$ by its complement (see Figure 1.2 for an example). We denote the resulting graph by $G * v$. We say that a graph H is a ***vertex-minor*** of a graph G if H can be obtained from G by a sequence of vertex deletions and local complementations. There are a number of parallels between vertex-minors and graph minors, so vertex-minors can be thought of as a analogue of the graph minor relation for dense classes of graphs. One striking parallel is highlighted by the grid minor theorem [90] and a recent grid theorem for vertex-minors [51]. Vertex-minors were given their descriptive name by Oum [83], but they had appeared earlier in Bouchet's [8] work on isotropic systems. Since then, local complementation and vertex-minors have appeared in other areas such as bioinformatics [35] and quantum computing [104].

Circle graphs provide a natural class of graphs that are closed under taking vertex-minors. Vertex deletion is clear as deleting a vertex of a circle graph simply corresponds to deleting its corresponding chord in its chord diagram. Local complementation is less obvious, but to locally complement at a vertex v of our circle graph, we first look at one of the two arcs between the end points of its corresponding chord, then we reverse the order that the end points of other chords appear along this arc. The intersection graph of the resulting chord diagram is then the graph obtained by locally complementing at v (see Figure 1.3 for an illustration of this operation).

We prove a conjecture of Geelen (see [34]) that every proper vertex-minor-closed class of graphs is χ-bounded.

Theorem 1.3.1. *Every proper vertex-minor-closed class of graphs is χ-bounded.*

Several special cases of Theorem 1.3.1 have been proved in the past. As previously discussed, Gyárfás [56] proved that circle graphs are χ-bounded. Another important vertex-minor-closed class of graphs is those with bounded rank-width; Dvořák and Král' [34] proved that such graphs are χ-bounded. Geelen, Kwon, McCarty and Wollan [51] proved the grid theorem for vertex-minors, which can be applied with the result of Dvořák and Král' [34]. More precisely, they proved that if H is a circle graph, then the class of graphs with no H vertex-minor has bounded rank-width and so is also χ-bounded. Let W_n denote the wheel graph consisting of an n-cycle and a single additional dominating vertex. Two of the three minimal forbidden vertex-minors for circle graphs are the wheels W_5 and W_7 [7], and so generalizing Gyárfás's result that circle graphs are χ-bounded, Choi, Kwon, Oum and Wollan [17] proved that for each n, the class of graphs with no W_n vertex-minor is also χ-bounded. Kostochka [67] proved that the complements of circle graphs are χ-bounded. This class of graphs is not vertex-minor-closed, however its closure under vertex-minors can be shown to be χ-bounded as an extension of Kostochka's result [48].

Building on Geelen's conjecture, Kim, Kwon, Oum and Sivaraman [65] further asked if all proper vertex-minor-closed classes of graphs are polynomially χ-bounded. Recently there have also been significant developments on this problem. As discussed in the previous section, with McCarty [29] we proved that circle graphs are polynomially χ-bounded. Bonamy and Pilipczuk [6] proved that graphs of bounded rank-width are polynomially χ-bounded. As a result, it also follows that if H is a circle graph, then the class of graphs with no H vertex-minor is polynomially χ-bounded [6, 51]. To prove this conjecture in general, it is likely that the full strength of a general structure theorem for proper vertex-minor closed classes of graphs would be required. The starting point towards such a theorem is the aforementioned grid theorem of Geelen, Kwon, McCarty and Wollan [51] and there has since been significant further progress towards a general structure theorem [77]. The polynomial χ-boundedness conjecture would follow from the conjectured structure theorem [77], a decomposition theorem in [6], and the fact that circle graphs are polynomially χ-bounded [29].

One advantage of our proof of Theorem 1.3.1 over a potential proof using the conjectured structure theorem is that it is of course much shorter and simpler than such a proof (although our proof is neither short nor simple). Our other motivation for proving Theorem 1.3.1 without a structure theorem was to develop techniques that could be used for proving the more general conjecture of Choi, Kwon and Oum [18] that proper pivot-minor-closed classes of graphs are χ-bounded (we delay the definition of pivot-minors until Section 4.1). Very recently [24] we have been able to extend the techniques used in the proof of Theorem 1.3.1 to prove the pivot-minor χ-boundedness conjecture in full.

Theorem 1.3.2. *Every proper pivot-minor-closed class of graphs is χ-bounded.*

The proof of Theorem 1.3.2 is significantly more complicated than the proof of Theorem 1.3.1, so for simplicity, we omit the proof of Theorem 1.3.2 in this thesis and just prove Theorem 1.3.1. We will however in Section 4.7 discuss differences in their proofs as well as ideas used to prove Theorem 1.3.2. The full proof of Theorem 1.3.2 will appear in an upcoming paper [24]. The proof of Theorem 1.3.1 in this thesis differs a bit from in [25]; it incorporates some of the ideas used to prove Theorem 1.3.2 to simply the proof.

To prove Theorem 1.3.1 we use the idea of ρ-controlled classes of graphs as introduced by Scott [94]. Let G be a graph. Recall that $\chi^{(\rho)}(G)$ denotes the maximum chromatic number of an induced subgraph of G with radius at most ρ, and we say that a class of graphs $\mathcal{G}$ is ρ-controlled if there exists a function f such that $\chi(G) \leq f(\chi^{(\rho)}(G))$ for all $G \in \mathcal{G}$. For this natural weakening of χ-boundedness, we obtain improved bounds. We say that a class of graphs $\mathcal{G}$ is ***linearly ρ-controlled*** if there exists a constant $c \geq 1$ such that $\chi(G) \leq c\chi^{(\rho)}(G)$ for all $G \in \mathcal{G}$.

Theorem 1.3.3. *Every proper vertex-minor-closed class of graphs is linearly 2-controlled.*

This was conjectured in [25] where we only proved that proper vertex-minor closed classes are linearly 9-controlled. Note that in general not all vertex-minor-closed classes of graphs have linear χ-bounding functions, or are even linearly 1-controlled. For instance, induced subgraphs of circle graphs with radius 1 are permutation graphs, which are perfect (see Lemma 3.1.1), while by Theorem 1.2.2 or the earlier construction of Kostochka [66,67], circle graphs have no linear χ-bounding function. We remark that in general, the bounding function for ρ-controlled classes of graphs may need to grow arbitrarily fast (even for any $\rho' > \rho$). The proof of this is a simple combination of ideas in [12] and [52].

Chapter 2

Boxes and lines

In this chapter we prove Theorems 1.1.2 and 1.1.3 that there are intersection graphs of axis-aligned boxes and of lines in $\mathbb{R}^3$ with arbitrarily large girth and chromatic number. These results also have corollaries concerning the average degree of bipartite box and line intersections with arbitrarily large girth.

Tomon and Zakharov [102] recently showed that there exist bipartite box intersection graphs with girth at least 6 that have a super-linear number of edges. This resolved a problem of Kostochka [66]. We extend this showing that bipartite box intersection graphs with arbitrarily large girth can have a super-linear number of edges.

Kwan, Letzter, Sudakov, and Tran [73] proved that every triangle-free graph with minimum degree at least c contains an induced bipartite subgraph of minimum degree at least $\Omega(\ln c/\ln\ln c)$. So as a corollary of Theorem 1.1.2, we obtain the following.

Corollary 2.0.1. *There are bipartite intersection graphs of axis-aligned boxes in* $\mathbb{R}^3$ *with arbitrarily large girth and minimum degree.*

It is also possible to modify the construction of Theorem 1.1.2 to prove Corollary 2.0.1 directly. Of course, we can also obtain the analogous corollary for intersection graphs of lines in $\mathbb{R}^3$. Since our construction for Theorem 1.1.2 repeatedly applies the sparse Gallai's theorem, the number of boxes used in Theorem 1.1.2 grows extremely rapidly with the chromatic number. Consequently, the minimum degree of our bipartite intersection graphs of axis-aligned boxes in $\mathbb{R}^3$ in Corollary 2.0.1 might grow extremely slowly with the number of boxes. Very recently Tomon [101] improved on Corollary 2.0.1 by proving that for fixed $g \geq 4$, there are bipartite intersection graphs of n axis-aligned boxes in $\mathbb{R}^3$ with girth g

and minimum degree at least $\Omega_g(\log \log n)$. In fact, their graphs are incidence graphs of points and axis-aligned rectangles in $\mathbb{R}^2$.

This new method of combining Tutte's [31,32] construction with Gallai's [88] theorem and its variations appears to be particularity well suited to constructing graphs with large chromatic number in geometric settings. Indeed, with Keller, Kleist, Smorodinsky, and Walczak [26], we have recently successfully applied the method in another geometric setting. A ***constellation*** is a finite collection of circles in the plane in which no three circles are tangent at the same point. The ***tangency graph*** $G(\mathcal{C})$ of a constellation $\mathcal{C}$ is the graph with vertex set $\mathcal{C}$ and edges comprising of the pairs of tangent circles in $\mathcal{C}$. Ringel's circle problem [62] from 1959 asked for the maximum chromatic number of tangency graphs of constellations. With Keller, Kleist, Smorodinsky, and Walczak [26], we resolved Ringel's circle problem in a strong sense by proving that these graphs are χ-amorphous; they contain graphs with arbitrarily large girth and chromatic number.

We remark that our choice of hypergraph based on Gallai's theorem appears to be crucial in these constructions; with the exception of box intersection graphs in $\mathbb{R}^3$, we suspect that the graph classes in these applications do not contain the graphs from Tutte's original construction. Even in the case of boxes, using Gallai's theorem allows us to prove the theorem for the subclass of so called grounded square box graphs (see Theorem 2.2.1).

A natural subclass of the intersection graphs of axis-aligned boxes in $\mathbb{R}^3$, are those whose boxes have disjoint interiors. Using a necessarily more relaxed variation of Tutte's construction, Reed and Allwright [89] showed that these graphs are not χ-bounded either. Magnant and Martin [75] further showed that there are such intersection graphs where the boxes only intersect on their top and bottom faces (which in particular means they are triangle-free), and have arbitrarily large chromatic number.

We remark that while Tutte's construction can be sparsified, these variations used for boxes with disjoint interiors cannot be; an intersection graph of n axis-aligned boxes in $\mathbb{R}^3$ with disjoint interiors, and with girth at least 5, has at most $24n$ edges and thus has bounded chromatic number.

Theorem 2.0.2. *Every intersection graph of n axis-aligned boxes in $\mathbb{R}^3$ with disjoint interiors, and with girth at least 5, has at most $24n$ edges*

Proof. If two boxes with disjoint interiors intersect, then there is a plane that contains a pair of intersecting faces, one from each of the two boxes. Glebov [53] proved that every intersection graph of r axis-aligned rectangles in the plane with girth at least 5 has at most $4r$ edges. Each of the 6 faces of a box is contained in a unique plane, so we obtain the bound of $24n$ by applying Glebov's result to each plane contained in $\mathbb{R}^3$ and the rectangular faces that they contain. □

This can be strengthened to show that such graphs are 2-controlled by using the fact that intersection graphs of axis-aligned rectangles are 2-controlled [21] in place of Glebov's [53] result.

2.1 Gallai's theorem

Before proving Theorems 1.1.2 and 1.1.3, we must first introduce our main tool; Gallai's theorem [88] and its sparse variant [86].

A cornerstone of Ramsey theory is Van der Waerden's theorem on monochromatic arithmetic progressions.

Theorem 2.1.1 (Van der Waerden [105])**.** *Every finite colouring of $\mathbb{N}$ contains arbitrarily long arithmetic progressions.*

Gallai [88] proved a generalization concerning monochromatic homothetic copies of finite sets $T \subset \mathbb{R}^d$. A ***homothetic map*** $f : \mathbb{R}^d \to \mathbb{R}^d$ is one of the form $f(x) = sx + \bar{c}$ for some $s \in \mathbb{R}_{>0}$ and $\bar{c} \in \mathbb{R}^d$. In other words, a homothetic map is a composition of uniform scaling and a translation. A set $T' \subset \mathbb{R}^d$ is a ***homothetic copy*** of a set $T \subset \mathbb{R}^d$ if there is a homothetic map f such that $f(T) = T'$.

Theorem 2.1.2 (Gallai [88])**.** *Let $T \subset \mathbb{R}^d$ be a finite set. Then any finite colouring of $\mathbb{R}^d$ contains a monochromatic homothetic copy of T.*

Note that Gallai's theorem is equivalent to the multidimensional van der Waerden's theorem; simply consider the lattice $\mathbb{Z}[T]$.

Gallai's theorem would be enough to construct graphs as in Theorems 1.1.2 and 1.1.3 with girth at least 6, however for larger girth we require a sparse version proven by Prömel and Voigt [86].

We say that a collection of distinct sets $T_1, \ldots, T_k$ form a cycle C of length k if there exist distinct elements $x_1, \ldots, x_k$ such that for all $i \in \{1, \ldots, k-1\}$ we have $x_i \in T_i \cap T_{i+1}$ and $x_k \in T_k \cap T_1$. The ***girth*** of a hypergraph is equal to the length of its shortest cycle. A triangle-free intersection graph of a collection of sets $\mathcal{T}$ has a cycle of length k if and only if there is a collection of k sets contained in $\mathcal{T}$ that form a cycle of length k. We often find it convenient to consider colourings and cycles of objects in $\mathbb{R}^3$ directly rather than in their intersection graphs.

We may now state the sparse version of Gallai's theorem [88] due to Prömel and Voigt [86] that we require.

Theorem 2.1.3 (Prömel and Voigt [86])**.** *Let $T \subset \mathbb{R}^d$ be a finite subset containing at least three elements and let g and k be positive integers. Then there exists a finite set $X \subset \mathbb{R}^d$ such that every k-colouring of X contains a monochromatic homothetic copy T' of T and no set of at most $g-1$ homothetic copies of T form a cycle.*

Prömel and Voigt [87] also proved sparse versions of other theorems including the Hales–Jewett theorem [59] and the Graham–Rothschild theorem [55]. A special case of Theorem 2.1.3 is also equivalent to a sparse van der Waerden's theorem.

As discussed in Section 1.1, the graphs we shall construct are based on a hypergraph variation of Tutte's [31, 32] construction of triangle-free graphs with arbitrarily large chromatic number that has been considered several times before [70, 81, 100].

Note that Theorem 2.1.3 is implicitly a theorem concerning the existence of certain hypergraphs with large girth and chromatic number; the (implicit) hypergraph has vertex set X and hyperedge set being the homothetic copies of T in X. The proofs of Theorems 1.1.2 and 1.1.3 both go by showing that we can build up graphs with larger and larger chromatic number with each new graph created as in the hypergraph variation of Tutte's construction, with the input auxiliary hypergraph being the one implicit in an application of Theorem 2.1.3. So by Lemma 1.1.1, it will just be enough to ensure that the intersections between our geometric objects witness this construction.

2.2 Boxes

Before proving Theorem 1.1.2, we need to introduce a more restricted setting to facilitate the inductive argument. We say that an axis-aligned box $B \subseteq \mathbb{R}^3$ is a ***square box*** if its top and bottom faces are squares. Let $P_{x=y} \subset \mathbb{R}^3$ be the $x = y$ plane, and let $L_{x=y}$ be the line given by projecting $P_{x=y}$ into the (x, y)-plane. We say that a square box $B \subset \mathbb{R}^3$ is ***grounded*** if B intersects $P_{x=y}$, and the rest of B is contained in the $x > y$ side of $\mathbb{R}^3$ (as divided by $P_{x=y}$). In other words, B is grounded if the top left corner of the square given by projecting B into the (x, y)-plane is contained in the $x = y$ line $L_{x=y}$. Notice that if B is grounded, then $B \cap P_{x=y}$ is a vertical line segment, an edge of the box, and its projection

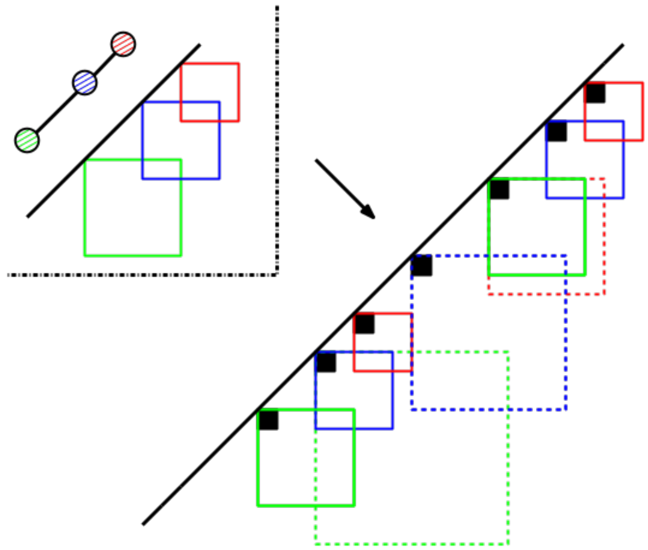

Figure 2.1: An illustration from above of part of the construction applied to three boxes forming a path. The solid black boxes are tall thin boxes of $\mathcal{B}_X$. The figure includes three homothetic copies of $\mathcal{B}_{g,k}$, each corresponding to a different homothetic copy of T. Two correspond to translations of T in X, and the one with dashed sides corresponds to a translation and uniform scaling by a factor of 2. We use the unseen z-axis to avoid boxes corresponding to different homothetic copies of T intersecting, in particular, the boxes corresponding to the dotted rectangles are higher.

onto the (x, y)-plane is a single point of the line $L_{x=y}$. A graph is a ***grounded square box*** graph if it is the intersection graph of a collection of grounded square boxes.

We prove Theorem 1.1.2 for the subclass of grounded square box graphs. While the grounded condition is used to ease the inductive argument, achieving square boxes rather than general grounded axis-aligned boxes is simply a bonus we get for free from the proof.

Theorem 2.2.1. *For every $g \geq 3$ and $k \geq 1$, there exists a grounded square box graph with girth at least g and chromatic number at least k.*

Proof. This is trivially true for $k \leq 3$, as all odd cycles are grounded square box graphs. So we fix g and proceed inductively on k. Let $\mathcal{B}_{g,k}$ be a collection of grounded square boxes whose intersection graph has girth at least g and chromatic number at least k. We wish to apply the 1-dimensional (sparse) Gallai's theorem to a finite set T contained in the line $L_{x=y}$. Each point of T will correspond to a box $\mathcal{B}_{g,k}$. We choose the point t_B of T

corresponding a box $B \in \mathcal{B}_{g,k}$ to be the point given by projecting $P_{x=y} \cap B$ onto the (x, y)-plane. The points of T should all be distinct to maintain a one to one correspondence, this can be achieved by slightly perturbing the grounded square boxes of $\mathcal{B}_{g,k}$ along $P_{x=y}$ (while preserving their intersections).

Now we apply Theorem 2.1.3 to $T \subset L_{x=y}$ to obtain a finite set $X \subset L_{x=y}$ such that every k-colouring of X contains a monochromatic homothetic copy T' of T, and no set of at most $\lfloor \frac{g}{3} \rfloor$ homothetic copies of T contained in X form a cycle. Let $\mathcal{T}$ be the set of homothetic copies of T in X. With G being the intersection graph of $\mathcal{B}_{g,k}$, and H being the hypergraph with vertex set X and hyperedge set $\mathcal{T}$, we now aim to witness a graph G^* obtainable from the operation $\mathrm{T}(G, H)$ (as in Lemma 1.1.1) as an intersection graph of grounded square boxes. Throughout the rest of the proof the reader may wish to refer to Figure 2.1 for an illustration of the construction.

We begin with constructing the boxes corresponding to the vertices $V(H) = X$ of G^*. For each $x \in X$, we choose a very thin and very tall grounded square box B_x whose intersection with $P_{x=y}$ projects to the point $x \in L_{x=y}$. Each of these boxes are chosen to contain the same (large) interval Z when projected onto the z-axis. Let $\mathcal{B}_X = \{B_x : x \in X\}$.

Now we need to construct the boxes corresponding to the copies of G, one for each for each $T' \in E(H) = \mathcal{T}$. For each $T' \in \mathcal{T}$, we choose a set of grounded square boxes $\mathcal{B}_{T'}$ homothetic to $\mathcal{B}_{g,k}$, so that the points of $\mathcal{B}_{T'}$ intersecting $P_{x=y}$ project onto the subset $T' \subseteq X$. By appropriately translating these homothetic copies along the z-axis, we can ensure that the boxes of each $\mathcal{B}_{T'}$ intersect exactly their corresponding box of $\{B_t : t \in T'\} \subseteq \mathcal{B}_X$, and also ensure that there are no intersection between boxes of distinct homothetic copies $\mathcal{B}_{T_1}, \mathcal{B}_{T_2}$ of $\mathcal{B}_{g,k}$.

Let $\mathcal{B}_{g,k+1}$ be the union of the grounded square boxes $\mathcal{B}_X$ and $(\mathcal{B}_{T'} : T' \in \mathcal{T})$. Now the intersection graph of $\mathcal{B}_{g,k+1}$ is a graph obtainable from the operation $\mathrm{T}(G, H)$, so the theorem follows by Lemma 1.1.1. $\square$

It only takes a minor modification of this proof to obtain, as an application of Van der Waerden's theorem, a proof of Burling's [13] classical result that there are triangle-free intersection graphs of axis-aligned boxes with arbitrarily large chromatic number. In fact, here we can already achieve a girth of 6. One just has to arrange at the start that the boxes of $\mathcal{B}_{6,k}$ are perturbed in such a way that their intersections with $P_{x=y}$ project down to a set of rational points T in $L_{x=y}$. Then there is some arithmetic progression in $L_{x=y}$ that contains T, and so we may use Van der Waerden's theorem in place of Theorem 2.1.3. For larger girth we could also choose to use the sparse Van der Waerden's theorem [87] in place of Theorem 2.1.3.

2.3 Lines

Next we prove Theorem 1.1.3, that there are intersection graphs of lines in $\mathbb{R}^3$ with arbitrarily large girth and chromatic number. The proof is similar to that of Theorem 2.2.1, we just have to modify the geometric arguments for this setting. We restate Theorem 1.1.3 slightly differently for convenience.

Theorem 2.3.1. *For every $g \geq 3$ and $k \geq 1$, there exist lines in $\mathbb{R}^3$ whose intersection graph has girth at least g and chromatic number at least k.*

Proof. The theorem is trivially true for $k \leq 3$, as all odd cycles are intersection graphs of lines in $\mathbb{R}^3$. So we fix g and proceed inductively on k. Let $\mathcal{S}_{g,k}$ be a collection of lines in $\mathbb{R}^3$ whose intersection graph has girth at least g and chromatic number at least k.

Choose a plane $P \subset \mathbb{R}^3$ such that every line of $\mathcal{S}_{g,k}$ intersects P in a single point, and no two lines of $\mathcal{S}_{g,k}$ intersect P at the same point. Let T be the set of points of P contained in a line of $\mathcal{S}_{g,k}$, in particular there is a one to one correspondence between points of T and lines in $\mathcal{S}_{g,k}$.

By Theorem 2.1.3, there exists a finite set $X \subset P$ such that every k-colouring of X contains a monochromatic homothetic copy T' of T, and no set of at most $\lfloor \frac{g}{3} \rfloor$ homothetic copies of T contained in X form a cycle. Let $\mathcal{T}$ be the set of homothetic copies of T in X. With G being the intersection graph of $\mathcal{S}_{g,k}$, and H being the hypergraph with vertex set X and hyperedge set $\mathcal{T}$, we now aim to witness a graph G^* obtainable from the operation $\mathrm{T}(G, H)$ (as in Lemma 1.1.1) as an intersection graph of lines.

For each $T' \in \mathcal{T}$, let $\mathcal{S}_{T'}$ be a homothetic copy of $\mathcal{S}_{g,k}$ such that $\mathcal{S}_{T'} \cap P = T'$. Now choose a direction $\vec{d}$ parallel to P, but not parallel to a line between any two points of X, so that the homothetic copies $(\mathcal{S}_{T'} : T' \in \mathcal{S})$ of $\mathcal{S}_{g,k}$ can each be translated in the $\vec{d}$ direction so as to avoid intersections between lines of distinct homothetic copies $\mathcal{S}_{T_1}, \mathcal{S}_{T_2}$ of $\mathcal{S}_{g,k}$. Perform such translations to each $(\mathcal{S}_{T'} : T' \in \mathcal{S})$.

Now, for each $x \in X$, let S_x be the line containing x that is parallel to $\vec{d}$. Let $\mathcal{S}_X = \{S_x : x \in X\}$. The lines of $\mathcal{S}_X$ are all parallel and non-intersecting, and for each $T' \in \mathcal{T}$, each line of $\mathcal{S}_{T'}$ intersects exactly its corresponding line of $\{S_{t'} : t' \in T'\} \subseteq \mathcal{S}_X$.

Let $\mathcal{S}_{g,k+1}$ be the union of the lines of $\mathcal{S}_X$ and $(\mathcal{S}_{T'} : T' \in \mathcal{T})$. We have the desired intersections between lines of $\mathcal{S}_{g,k+1}$, so the theorem now follows from Lemma 1.1.1. $\square$

Chapter 3

Circle graphs

In this chapter we prove Theorems 1.2.1 and 1.2.2 narrowing the optimal χ-bounding function for circle graphs to within a small constant factor. Theorem 1.2.2 has a simple construction and is proved in Section 3.6, while Sections 3.1-3.4 are dedicated to proving Theorem 1.2.1. In Section 3.5 we optimize the proof of Theorem 1.2.1 to the K_4-free case to prove Theorem 1.2.3, that K_4-free circle graphs are 19-colourable.

The proof of Theorem 1.2.1 essentially goes by proving a stronger statement on being able to extend certain well-structured partial pre-colourings. This better facilitates an inductive argument and is an idea most famously used in Thomassen's [99] proof that planar graphs are 5-choosable. As in our proof that circle graphs are $O(\omega^2)$-colourable with McCarty [29], we use what we call a pillar assignment to colour our circle graphs. The reason for this is two-fold: pillar assignments provide a convenient way to describe the possible pre-colourings, and they also act as a useful tool for extending the pre-colourings. However we require a definition of pillar assignments that is different to that of [29].

In [29] we used pillar assignments to obtain an improper colouring such that every monochromatic component was a permutation graph. By exploiting the structure of our improper colouring and using a natural Turán-type lemma on permutation graphs, we were able to bound the number of colours needed in this improper colouring. Then finally a proper colouring was obtained by refining the improper colouring.

Although significantly easier, obtaining a proper colouring by first going via this improper colouring appears to present a degree of inefficiency in minimizing the number of colours used. So the most significant difference with the notion of pillar assignment that we use is that it provides a proper colouring of the circle graph directly. This involves colouring certain induced permutation subgraphs in a particular well-structured way. The purpose of this additional structure in the colouring is to allow for a new Turán-type lemma (see Lemma 3.3.3). Although this lemma is less natural, it is much more specialized to our notion of pillar assignments. With this new notion of pillar assignment and its tailor-made Turán-type lemma, we are then able to obtain the improved bounds with an inductive argument on extending pillar assignments.

As a step towards proving our required tailor-made Turán-type lemma, we actually prove a tight (and somewhat more abstract) version of the Turán-type lemma on permutation graphs used in [29] (see Theorem 3.3.2).

An old problem of Kostochka and Kratochvíl [68] asks whether the optimal χ-bounding function for polygon-circle graphs is within a constant factor of the optimal χ-bounding function for circle graphs. By Theorems 1.2.1 and 1.2.2, circle graphs are now known to have an optimal χ-bounding function of $\Theta(\omega \log \omega)$. So two reasonable approaches to this problem would be to either improve the lower bound construction, or to extend the proof of Theorem 1.2.1 to polygon-circle graphs. We believe that polygon-circle graphs are also $O(\omega \log \omega)$-colourable, but somewhat surprisingly the proof of Theorem 1.2.1 does not appear to easily extend to polygon-circle graphs. On the other hand, our original proof with McCarty [29] that circle graphs are polynomially χ-bounded can be extended with very little changes to prove the same $O(\omega^2)$ χ-bounding function for polygon-circle graphs (see [23]). The key part of the proof of Theorem 1.2.1 that does not appear to translate well to polygon-circle graphs is Lemma 3.1.1 which gives a colouring of permutation graphs with desirable properties. For polygon-circle graphs the analogue is trapezoid graphs, but they appear to not have such colourings.

3.1 Preliminaries

For convenience of proving Theorem 1.2.1 we use an interval overlap representation of our circle graphs rather than a chord diagram representation. An ***interval system*** is a collection of open intervals in $(0, 1)$ such that no two share an endpoint. Two distinct intervals I_1, I_2 ***overlap*** if they have non-empty intersection, and neither is contained in the other. The ***overlap graph*** of an interval system $\mathcal{I}$ is the graph with vertex set $\mathcal{I}$ where two vertices are adjacent whenever their corresponding intervals overlap. By cutting the circle and deforming it onto the real line, it can easily be checked that circle graphs are exactly overlap graphs of interval systems. Similarly, ***permutation graphs*** are exactly the overlap graphs of interval systems $\mathcal{I}$ such that there exists a $p \in (0, 1)$ with $p \in I$ for all intervals $I \in \mathcal{I}$ [54].

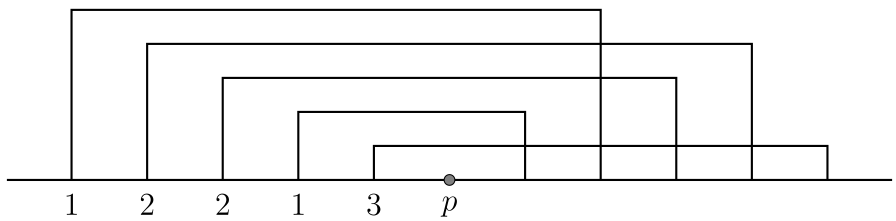

Figure 3.1: The colouring ϕ_p of an interval system whose intervals all contain p in the case that $C = \{1, 2, 3\}$. The colour that the intervals receive is the number appearing below their leftmost endpoint.

It is often more convenient to examine properties of a circle graph as equivalent properties of their interval systems. Note that sets of pairwise non-overlapping intervals in an interval system correspond to stable sets in the overlap graph, and sets of pairwise overlapping intervals correspond to cliques. Given an interval system $\mathcal{I}$, we let $\omega(\mathcal{I})$ be equal to the size of the largest set of pairwise overlapping intervals contained in $\mathcal{I}$. Equivalently, $\omega(\mathcal{I})$ is equal to the clique number of the overlap graph of $\mathcal{I}$. Similarly we consider colourings of an interval system with a notion equivalent to that of colourings of their overlap graphs. A ***proper partial colouring*** of an interval system $\mathcal{I}$ is an assignment of colours to a subset of the intervals of $\mathcal{I}$ so that no pair of overlapping intervals receive the same colour. We say that a proper partial colouring of $\mathcal{I}$ is ***complete*** if every interval of $\mathcal{I}$ is assigned a colour.

For an interval $I \subseteq (0, 1)$, let $\ell(I)$ be its leftmost endpoint and let $r(I)$ be its rightmost endpoint. For two intervals $I_1, I_2 \subseteq (0, 1)$, we use $I_1 < I_2$ to denote that $r(I_1) < \ell(I_2)$, and similarly $I_1 > I_2$ to denote that $\ell(I_1) > r(I_2)$. Given a finite partially ordered set $(X, \preceq)$, and some $x \in X$, the ***height*** $h(x)$ of x in the partial order is equal to the maximum length of a chain ending in x. For a positive integer k, we let $[k] = \{1, \ldots, k\}$.

We finish this section with a lemma on colouring permutation graphs that is used in our definition of pillar assignments. In addition to the bound on the number of colours required, we also make use of an additional property of this colouring.

Given an interval system $\mathcal{I}$ with $\omega(\mathcal{I}) = \omega$ such that all intervals of $\mathcal{I}$ contain some given point $p \in \mathbb{R}$, we can naturally define a partial order on $\mathcal{I}$ as follows. Let $\preceq$ be the partial order of $\mathcal{I}$ such that $I \preceq I'$ whenever $\ell(I) < \ell(I')$ and $r(I) < r(I')$ (or $I = I'$). Notice that two intervals overlap exactly when they are comparable in the partially ordered set $(\mathcal{I}, \preceq)$. So we can obtain a colouring ϕ_p of $\mathcal{I}$ by colouring the intervals according to their height in the partially ordered set $(\mathcal{I}, \preceq)$. Since comparable elements overlap, the longest chain in $(\mathcal{I}, \preceq)$ has length ω, so ϕ_p is a ω-colouring of $\mathcal{I}$. For an example of such a colouring ϕ_p, see Figure 3.1.

With this colouring, whenever we find a sequence of intervals whose left (or right) endpoints are increasing (with respect to the natural ordering of $\mathbb{R}$), and whose colours as given by ϕ_p are strictly increasing (as a subset of $\mathbb{N}$), although these intervals need not overlap, we can find the same number of intervals between them that do overlap. This is a useful property of the colouring since it gives us a way to indirectly find sets of overlapping intervals. Let us justify this property formally.

Suppose that $I_1, \ldots, I_k \in \mathcal{I}$ are such that $\phi_p(I_1) < \cdots < \phi_p(I_k)$ and $\ell(I_1) < \cdots < \ell(I_k)$. Let $I_k^* = I_k$ and for each $j < k$ in decreasing order, let I_j^* be the interval with $\ell(I_j^*)$ maximum, subject to $I_j^* \prec I_{j+1}^*$ and $\phi_p(I_j^*) = \phi_p(I_j)$. Such intervals I_j^* must exist with $\ell(I_j) \leq \ell(I_j^*) < \ell(I_{j+1}^*)$ by the choice of colouring ϕ_p as if $I' \in \mathcal{I}$ were an interval with $I' \prec I_{j+1}^*$, $\phi_p(I') = \phi_p(I_j)$, and $\ell(I') < \ell(I_j)$, then I' and I_j would be non-overlapping (since they receive the same colour), and so I_j would overlap with I_{j+1}^* and hence precede I_{j+1}^* in the partial order. Then $I_1^* \prec \cdots \prec I_k^*$, and so the intervals $I_1^*, \ldots I_k^*$ are pairwise overlapping with $\ell(I_1^*), \ldots, \ell(I_k^*) \in [\ell(I_1), \ell(I_k)]$.

Similarly, if we instead had that $r(I_1) < \cdots < r(I_k)$, then the analogous property for right endpoints also holds. So by the above discussion, we obtain the following lemma. Since the given colours are important for the lemma's applications, we allow our colours to be an arbitrary set $C \subset \mathbb{N}$ with $|C| = \omega$.

Lemma 3.1.1. *Let $\mathcal{I}$ be an interval system with $\omega(\mathcal{I}) = \omega$ such that all intervals of $\mathcal{I}$ contain some given point $p \in \mathbb{R}$, and let $C \subset \mathbb{N}$ have $|C| = \omega$. Then there is a proper colouring $\phi_p : \mathcal{I} \to C$ such that if $I_1, \ldots, I_k \in \mathcal{I}$ are intervals with $\phi_p(I_1) < \cdots < \phi_p(I_k)$ and $\ell(I_1) < \cdots < \ell(I_k)$ (or $r(I_1) < \cdots < r(I_k)$), then there exist k pairwise overlapping intervals $I_1^*, \ldots I_k^* \in \mathcal{I}$ with $\ell(I_1^*), \ldots, \ell(I_k^*) \in [\ell(I_1), \ell(I_k)]$ (or $r(I_1^*), \ldots, r(I_k^*) \in [r(I_1), r(I_k)]$ respectively).*

3.2 Pillar assignments

We start this section by defining our notion of pillar assignments, the tool we use to colour circle graphs. The colouring in Lemma 3.1.1 is crucial to the definition of pillar assignments and thus crucial for colouring our circle graphs. Afterwards we examine some properties of pillar assignments.

A ***pillar*** of an interval system $\mathcal{I}$ is a point within $(0, 1)$ that is distinct from the endpoints of the intervals of $\mathcal{I}$. For totally ordered pillars $(P, \preceq)$ (where $\preceq$ is not necessarily

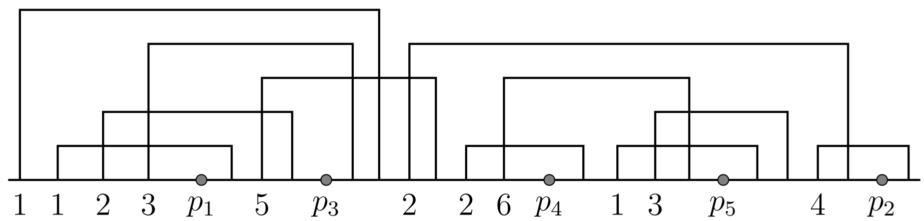

Figure 3.2: The colouring $\psi_{(P,\preceq)}$ of an interval system $\mathcal{I}$ for a collection of totally ordered pillars $(P, \preceq)$ with $P = \{p_1, p_2, p_3, p_4, p_5\}$ and $p_1 \prec p_2 \prec p_3 \prec p_4 \prec p_5$. The colour that the intervals receive is the number appearing below their leftmost endpoint. In this example we get that $C_{p_1} = \{1, 2, 3\}$, $C_{p_2} = \{4\}$, $C_{p_3} = \{5\}$, $C_{p_4} = \{2, 6\}$, and $C_{p_5} = \{1, 3\}$.

the natural ordering of $\mathbb{R}$), we say that an interval $I \in \mathcal{I}$ is ***assigned*** to a pillar $p \in P$ if $p \in I$ and there is no pillar $p' \in P$ such that $p' \in I$ and $p' \prec p$. So every interval is assigned to at most one pillar. For each pillar $p \in P$, we let $\mathcal{I}_p$ be the intervals of $\mathcal{I}$ that are assigned to p. The ***foundation*** F_p of a pillar $p \in P$ is the open interval containing p that has its endpoints in $\{p' \in P : p' \prec p\} \cup \{0, 1\}$ and contains no pillar $p' \in P$ with $p' \prec p$.

Next we show how to obtain a proper partial colouring of an interval system $\mathcal{I}$ from a collection of totally ordered pillars $(P, \preceq)$. We refer the reader to Figure 3.2 for an illustration of a pillar assignment and the colouring obtained from ordered pillars. For each pillar $p \in P$ in order, we assign a set of colours $C_p \subset \mathbb{N}$ to p and a C_p-colouring $\phi_p : \mathcal{I}_p \to C_p$ of the intervals assigned to p as follows.

If p is the first pillar in the total order $\preceq$, then let $C_p = \{1, \ldots, \omega(\mathcal{I}_p)\}$, and let $\psi_p = \phi_p$ be a C_p-colouring of $\mathcal{I}_p$ as in Lemma 3.1.1.

Otherwise let p^* be the pillar immediately preceding p in the total order $\preceq$. Then let $\mathcal{F}_p$ be the intervals of $\mathcal{I}$ that have exactly one endpoint in F_p. Let C_p be the set of the smallest $\omega(\mathcal{I}_p)$ positive integers that are not contained in $\psi_{p^*}(\mathcal{F}_p)$. Then let ϕ_p be a C_p-colouring of $\mathcal{I}_p$ as in Lemma 3.1.1. Let $\psi_p = \psi_{p^*} \cup \phi_p$. Note that ψ_p remains a proper partial colouring of $\mathcal{I}$ as the intervals of $\mathcal{I}_p$ are all contained in the foundation F_p, and so do not overlap with any of the intervals $\psi_{p^*}^{-1}(C_p)$ by the choice of C_p.

Then for the last pillar q in the total order $\preceq$, we let $\psi_{(P,\preceq)} = \psi_q$. Another convenient equivalent definition which we often use is $\psi_{(P,\preceq)} = \bigcup_{p \in P} \phi_p$.

A ***pillar assignment*** of an interval system $\mathcal{I}$ is a triple $(P, \preceq, \psi)$ such that P is a set of pillars, $\preceq$ is a total ordering of P, and ψ is the proper partial colouring $\psi_{(P,\preceq)}$ of $\mathcal{I}$ as described above. A pillar assignment $(P, \preceq, \psi)$ is ***complete*** if every interval of $\mathcal{I}$ contains some pillar of P (or equivalently if ψ colours every interval of $\mathcal{I}$). For a pillar assignment $(P, \preceq, \psi)$, let $\chi(P, \preceq, \psi)$ be equal to $\left|\bigcup_{p \in P} C_p\right| = |\psi(\mathcal{I})|$, in other words $\chi(P, \preceq, \psi)$ is the number of colours that the pillar assignment $(P, \preceq, \psi)$ uses to colour its interval system $\mathcal{I}$. So if $(P, \preceq, \psi)$ is a complete pillar assignment, then $\chi(\mathcal{I}) \leq \chi(P, \preceq, \psi)$. By the above definition and discussion, we have the following.

Lemma 3.2.1. *Let $(P, \preceq, \psi)$ be a pillar assignment of an interval system $\mathcal{I}$. Then ψ is a proper partial colouring of $\mathcal{I}$ that colours every interval containing a pillar of P, and furthermore if the pillar assignment is complete then ψ is a complete proper colouring of the interval system $\mathcal{I}$, and so $\chi(\mathcal{I}) \leq \chi(P, \preceq, \psi)$.*

Next we analyze the endpoints of intervals assigned to a given pillar, and also the endpoints of intervals with a given colour in a pillar assignment. These two lemmas are used in the proof of our tailor-made Turán-type lemma in Section 3.3. An ***arch*** of a pillar assignment $(P, \preceq, \psi)$ is an open interval with endpoints in $\{0, 1\} \cup P$ that contains no pillar of P.

Lemma 3.2.2. *Let $(P, \preceq, \psi)$ be a pillar assignment of an interval system $\mathcal{I}$, let K be an arch of $(P, \preceq, \psi)$, and let $\mathcal{I}_K$ be the intervals of $\mathcal{I}$ with exactly one endpoint in K. Let $\bigcup_{p \in P} \mathcal{I}_{(K,p)}$ be the partition of $\mathcal{I}_K$ where for each $p \in P$, the intervals $\mathcal{I}_{(K,p)}$ are exactly the intervals of $\mathcal{I}_K$ that are assigned to pillar p. Then there is a collection of disjoint intervals $\{K_p : p \in P\}$ contained in $(0, 1)\backslash K$ such that for every $p \in P$, the intervals of $\mathcal{I}_{(K,p)}$ have an endpoint within K_p.*

Proof. Let $K = (k^-, k^+)$. First observe that for each pillar $p \in P$, the intervals of $\mathcal{I}_{(K,p)}$ must all be contained in either $(k^-, 1)$ or $(0, k^+)$ depending on whether the pillar p is contained in $[k^+, 1)$ or $(0, k^-]$.

Now suppose for the sake of contradiction that no such collection of disjoint intervals $\{K_p : p \in P\}$ exist. Then there must exist two distinct pillars p, p' and distinct intervals $I_1, I_2 \in \mathcal{I}_{(K,p)}$, $I' \in \mathcal{I}_{(K,p')}$ such that the endpoints $e_1, e_2.e'$ of I_1, I_2, I' respectively that are contained in $(0, 1)\backslash K$ are such that $e_1 < e' < e_2$, and either $e_1 < e' < e_2 < k^-$, or $k^+ < e_1 < e' < e_2$.

Suppose in the first case that $e_1 < e' < e_2 < k^-$. Then $I'\backslash K$ must contain p', and furthermore both $I_1\backslash K$ and $I_2\backslash K$ must contain p. Hence $I_1\backslash K$ contains both p and p'. As I_1 is assigned to p, we see that $p \prec p'$. Then I' does not contain p as I' is assigned to p' and $p \prec p'$. But this contradicts the fact that $I_2\backslash K \subset I'\backslash K$ contains p. The second case that $k^+ < e_1 < e' < e_2$ is argued similarly and we conclude that such a collection of disjoint intervals $\{K_p : p \in P\}$ exists. □

Lemma 3.2.3. *Let $(P, \preceq, \psi)$ be a pillar assignment of an interval system $\mathcal{I}$, let K be an arch of $(P, \preceq, \psi)$, and let $\mathcal{I}_K$ be the intervals of $\mathcal{I}$ with exactly one endpoint in K. Let $\bigcup_{p \in P} \mathcal{I}_{(K,p)}$ be the partition of $\mathcal{I}_K$ where for each $p \in P$, the intervals $\mathcal{I}_{(K,p)}$ are exactly the intervals of $\mathcal{I}_K$ that are assigned to pillar p. For each $c \in \psi(\mathcal{I})$, let $\mathcal{I}_{(K,c)}$ be the intervals of $\mathcal{I}_K$ that are coloured c by ψ. Then for each $c \in \psi(\mathcal{I})$, there is a pillar $p \in P$, such that $\mathcal{I}_{(K,c)} \subseteq \mathcal{I}_{(K,p)}$.*

Proof. Suppose not, then there must exist intervals $I_1, I_2 \in \mathcal{I}_K$ such that $\psi(I_1) = \psi(I_2)$ and I_1, I_2 are assigned to distinct pillars p_1, p_2. In particular this means that $\phi_{p_1}(I_1) = \phi_{p_2}(I_2)$ as in the definition of the colouring $\psi = \psi_{(P,\preceq)}$.

Without loss of generality, we may assume that $p_1 \prec p_2$. Then the foundation F_{p_2} of p_2 must contain I_2, and so F_{p_2} contains K as well. Hence the foundation F_{p_2} contains an endpoint of I_1. But this now contradicts the choice of C_{p_2} by the definition of the colouring $\psi = \psi_{(P,\preceq)}$. □

Next we define a notion for the degree of an interval J contained within an arch of a pillar assignment $(P, \preceq, \psi)$ of some interval system $\mathcal{I}$. It is this notion of degree that our tailor-made Turán-type lemma is based on.

For an interval J within an arch of a pillar assignment $(P, \preceq, \psi)$ of an interval system $\mathcal{I}$, the ***degree*** $d_{(P,\preceq,\psi)}(J)$ of J is equal to the number of colours that intervals of $\mathcal{I}$ with an endpoint in J receive from ψ. As an example, for the pillar assignment $(P, \preceq, \psi)$ depicted in Figure 3.2, $d_{(P,\preceq,\psi)}(p_3, p_4) = 5$, and $d_{(P,\preceq,\psi)}((p_3, p_4), \{p_1, p_2, p_3\}) = 2 + 0 + 1 = 3$. When the pillar assignment is clear from context we often omit the subscript on the degrees.

A pillar assignment $(P^*, \preceq^*, \psi^*)$ ***extends*** a pillar assignment $(P, \preceq, \psi)$ if $P \subset P^*$, every pillar of P precedes every pillar of $P^* \backslash P$ in $\preceq^*$ and $(P^*, \preceq^*)|_P = (P, \preceq)$, and ψ^* is a proper partial colouring that extends ψ. We remark that by definition, the last condition that ψ^* extends ψ is implied by the conditions on $(P^*, \preceq^*)$, because for the interval system $\mathcal{I}$, the colourings $\psi = \psi_{(P,\preceq)}$ and $\psi^* = \psi_{(P^*,\preceq^*)}$ are determined solely by the totally ordered pillars $(P, \preceq)$ and $(P^*, \preceq^*)$ respectively.

We require a divide and conquer lemma that under favourable conditions allows for a certain extension of a pillar assignment that maintains a low total number of colours used and low degree arches.

Lemma 3.2.4. *Let $(P, \preceq, \psi)$ be a pillar assignment of an interval system $\mathcal{I}$ with $\omega(\mathcal{I}) = \omega$, let K be an arch, let t be a positive integer, and let $Q \subset K$ be a finite collection of pillars such that $d_{(P,\preceq,\psi)}(J) \leq t$ for every interval J contained in $K \backslash Q$. Then there is a pillar assignment $(P^*, \preceq^*, \psi^*)$ extending $(P, \preceq, \psi)$ such that:*

- $P^* = P \cup Q$,
- $d_{(P^*,\preceq^*,\psi^*)}(J) \leq t + \omega\lceil \log_2(|Q|+1)\rceil$ *for every interval* $J \subset K \backslash Q$, *and*
- $\chi(P^*,\preceq^*,\psi^*) \leq \max\{\chi(P,\preceq,\psi),\ d_{(P,\preceq,\psi)}(K) + \omega\lceil \log_2(|Q|+1)\rceil\}$.

Proof. Firstly the result is trivially true if $|Q| = 0$. So from here we argue inductively on $|Q|$.

Let the endpoints of K be q_0 and q_n, with $q_0 < q_n$ and $n = |Q| + 1$. Then let the elements of Q be $\{q_1, \dots, q_{n-1}\}$ where $q_1 < \cdots < q_{n-1}$. Consider the pillar assignment $(P', \preceq', \psi')$ extending $(P, \preceq, \psi)$ that is obtained by adding the pillar $q_{\lceil \frac{n-1}{2} \rceil}$ immediately after all the pillars of P in the total ordering $(P, \preceq)$.

Then with respect to the pillar assignment $(P', \preceq', \psi')$, the interval K contains exactly two arches; $K_1 = (q_0, q_{\lceil \frac{n-1}{2} \rceil})$, and $K_2 = (q_{\lceil \frac{n-1}{2} \rceil}, q_n)$. By considering the colouring $\phi_{q_{\lceil \frac{n-1}{2} \rceil}} = \psi_{(P',\preceq')} \backslash \psi_{(P,\preceq)}$, we can observe that:

- $d_{(P',\preceq',\psi')}(K_1),\ d_{(P',\preceq',\psi')}(K_2) \leq d_{(P,\preceq,\psi)}(K) + \omega$,
- $\chi(P', \preceq', \psi') \leq \max\{\chi(P,\preceq,\psi),\ d_{(P,\preceq,\psi)}(K) + \omega\}$, and
- $d_{(P',\preceq',\psi')}(q_i, q_{i-1}) \leq t + \omega$ for every $i \in [n]$.

Next note that with respect to the colouring, extending the pillar assignment $(P', \preceq', \psi')$ within each of the arches K_1 and K_2 is independent of the other. So we may apply the result of the inductive hypothesis twice, once to the pillars $\{q_1, \dots, q_{\lceil \frac{n-1}{2} \rceil - 1}\} \subset K_1$, and then to the pillars $\{q_{\lceil \frac{n-1}{2} \rceil + 1}, \dots q_{n-1}\} \subset K_2$, to obtain a new pillar assignment $(P^*, \preceq^*, \psi^*)$ extending $(P', \preceq', \psi')$ (and so also extending $(P, \preceq, \psi)$). Furthermore the resulting pillar assignment $(P^*, \preceq^*, \psi^*)$ is such that: $P^* = P' \cup \{q_1, \dots, q_{\lceil \frac{n-1}{2} \rceil - 1}\} \cup \{q_{\lceil \frac{n-1}{2} \rceil + 1}, \dots q_{n-1}\} = P \cup Q$, and for each $i \in [n]$,

$$
\begin{aligned}
d_{(P^*,\preceq^*,\psi^*)}(q_{i-1}, q_i) &\leq t + \omega + \omega \left\lceil \log_2 \left(\max \left\{ \left\lceil \frac{n-1}{2} \right\rceil,\ n - \left\lceil \frac{n-1}{2} \right\rceil \right\} \right) \right\rceil \\
&\leq t + \omega \lceil \log_2(n) \rceil \\
&= t + \omega \lceil \log_2(|Q|+1) \rceil,
\end{aligned}
$$

and lastly,

$$\chi(P^*, \preceq^*, \psi^*) \leq \max \left\{ \begin{array}{c} \chi(P', \preceq', \psi'), \\ d_{(P', \preceq', \psi')}(K_1) + \omega \left\lceil \log_2 \left(\left\lceil \frac{n-1}{2} \right\rceil \right) \right\rceil, \\ d_{(P', \preceq', \psi')}(K_2) + \omega \left\lceil \log_2 \left(n - \left\lceil \frac{n-1}{2} \right\rceil \right) \right\rceil \end{array} \right\}$$
$$\leq \max \left\{ \chi(P, \preceq, \psi),\ d_{(P,\preceq,\psi)}(K) + \omega \lceil \log_2(n) \rceil \right\}$$
$$= \max \left\{ \chi(P, \preceq, \psi),\ d_{(P,\preceq,\psi)}(K) + \omega \lceil \log_2(|Q|+1) \rceil \right\}.$$

Hence $(P^*, \preceq^*, \psi^*)$ provides the desired pillar assignment. □

Lemma 3.2.4 is enough for proving Theorem 1.2.1, but to obtain the improved bound of 19 for K_4-free circle graphs as in Theorem 1.2.3, we need slightly better bounds than those given in Lemma 3.2.4. This can be obtained by adding the technical additional condition that for every $q \in Q$, there is an interval $I_q \in \bigcup_{p \in P} \mathcal{I}_p$ that overlaps with every interval in $\mathcal{I} \setminus \left(\bigcup_{p \in P} \mathcal{I}_p \right)$ that contains q. Then in any pillar assignment $(P^*, \preceq^*, \psi^*)$ extending $(P, \preceq, \psi)$ such that $P^* = P \cup Q$ and for any $q \in Q$, we would have that $\omega(\mathcal{I}_q) \leq \omega(\mathcal{I}) - 1 = \omega - 1$. With this additional condition and observation, we may replace every occurrence of "ω" in the resulting bound (and in the proof) with "$\omega - 1$". Thus we obtain the following.

Lemma 3.2.5. *Let $(P, \preceq, \psi)$ be a pillar assignment of an interval system $\mathcal{I}$ with $\omega(\mathcal{I}) = \omega$, let K be an arch, let t be a positive integer, and let $Q \subset K$ be a finite collection of pillars such that $d_{(P,\preceq,\psi)}(J) \leq t$ for every interval J contained in $K \backslash Q$, and for every $q \in Q$, there is an interval $I_q \in \bigcup_{p \in P} \mathcal{I}_p$ that overlaps with every interval in $\mathcal{I} \setminus \left(\bigcup_{p \in P} \mathcal{I}_p \right)$ that contains q. Then there is a pillar assignment $(P^*, \preceq^*, \psi^*)$ extending $(P, \preceq, \psi)$ such that:*

- $P^* = P \cup Q$,
- $d_{(P^*,\preceq^*,\psi^*)}(J) \leq t + (\omega - 1)\lceil \log_2(|Q| + 1) \rceil$ *for every interval* $J \subset K \backslash Q$, *and*
- $\chi(P^*, \preceq^*, \psi^*) \leq \max\{\chi(P, \preceq, \psi),\ d_{(P,\preceq,\psi)}(K) + (\omega - 1)\lceil \log_2(|Q| + 1) \rceil\}$.

Of course while going from "ω" to "$\omega - 1$" is significant when $\omega = 3$, it is only a very minor improvement when ω is large.

3.3 Extremal results

In this section we prove the Turán-type lemma that is tailor-made for our notion of pillar assignment and degree (Lemma 3.3.3). The purpose of this is to enable prudent usage of Lemma 3.2.4 in the proof of Theorem 3.4.1, a strengthening of Theorem 1.2.1 that concerns extending pillar assignments. The idea of Lemma 3.3.3 is based on a Turán-type theorem of Capoyleas and Pach [14] for circle graphs. For an interval system $\mathcal{I}$ and a collection of disjoint intervals $\mathcal{J}$, the Turán-type theorem of Capoyleas and Pach [14] bounds (in terms of $\omega(\mathcal{I})$ and $|\mathcal{J}|$) the number of pairs of distinct intervals $J_1, J_2 \in \mathcal{J}$ such that there is an interval of $\mathcal{I}$ with an endpoint in both J_1 and J_2.

First we need to prove Theorem 3.3.2, a theorem in a similar style to that of the Erdős-Szekeres theorem [39]. This result may be of independent interest. Indeed it can be shown that Theorem 3.3.2 is equivalent to a tight version of the Turán-type lemma used in [29], so it can also be considered an exact permutation graph analogue of the aforementioned Turán-type theorem of Capoyleas and Pach [14]. With a bit more care one can even characterise the extremal examples.

Before stating and proving Theorem 3.3.2, we first require two definitions and a simple lemma. Given some $S \subseteq \mathbb{R}^d$, we define the ***strong dominance partial ordering*** $\preceq_{sd}$ of S to be the partial order such that $u \preceq_{sd} v$ exactly when each coordinate of v is greater than the corresponding coordinate of u (or $u = v$). Given two sets A and B, we let $A \times B$ denote the ***Cartesian product*** $\{(a, b) : a \in A \text{ and } b \in B\}$, of A and B.

Lemma 3.3.1. *Let a, b be positive integers, and let $\preceq_{sd}$ be the strong dominance partial ordering of $[a] \times [b]$. Then the maximum length of an antichain in $([a] \times [b], \preceq_{sd})$ is equal to $a + b - 1$.*

Proof. Let A be the antichain $\{(a, j) : j \in [b]\} \cup \{(i, b) : i \in [a - 1]\}$, then $|A| = a + b - 1$. For each integer k with $1 - b \leq k \leq a - 1$, let $C_k = \{(x, y) \in [a] \times [b] : x - y = k\}$. Then $C_{1-b}, \ldots, C_{a-1}$ is a chain cover of $[a] \times [b]$ of size $a + b - 1$. An antichain contains at most one element of every chain in a chain cover. Hence the maximum length of an antichain in $([a] \times [b], \preceq_{sd})$ is equal to $|A| = a + b - 1$ as required. □

Theorem 3.3.2. *Let a, b, n be positive integers with $n \leq a, b$, and let $\prec_{sd}$ be the strong dominance partial ordering of $[a] \times [b]$. Let $S \subseteq [a] \times [b]$ be a set containing no chain of length greater than n. Then $|S| \leq n(a + b - n)$.*

Proof. Let m be the maximum length of a chain contained in $(S, \preceq_{sd})$, and let $A_1, \ldots, A_m$ be the antichain cover of S where $A_k = \{(x, y) \in S : h_{(S, \preceq_{sd})}(x, y) = k\}$ for each $k \in [m]$.

Then for each $k \in [m]$, and each $(x, y) \in A_k$, there exists a chain $C_{(x,y)}$ of length k ending in (x, y). This implies that $x, y \geq k$. So for each $k \in [m]$, the antichain A_k is contained in the grid $([a]\backslash[k-1]) \times ([b]\backslash[k-1])$. Then by Lemma 3.3.1, $|A_k| \leq (a-k+1)+(b-k+1)-1 = a + b - 2k + 1$ for every $k \in [m]$.

Lastly

$$|S| = \sum_{k=1}^{m} |A_k| \leq \sum_{k=1}^{m} (a + b - 2k + 1) = m(a + b - m) \leq n(a + b - n)$$

as desired. □

The bound in this theorem is tight: one extremal example is $\{(x, y) \in [a] \times [b] : x \leq n \text{ or } y \leq n\}$, which contains no chain of length greater than n. The theorem can also be generalised to higher dimensional grids with essentially the same proof. We anticipate that Theorem 3.3.2 will likely also find further applications in improving χ-bounding functions for other classes of geometric intersection graphs.

We now proceed with applying Theorem 3.3.2 to prove our tailor-made Turán-type lemma.

Lemma 3.3.3. *Let $(P, \preceq, \psi)$ be a pillar assignment of an interval system $\mathcal{I}$ with $\omega(\mathcal{I}) = \omega$, let K be an arch such that $d(K) \geq \omega$, and let $\mathcal{J}$ be a collection of disjoint open intervals contained within K such that $|\mathcal{J}| \geq \omega$. Then*

$$\sum_{J \in \mathcal{J}} d(J) \leq \omega(d(K) + |\mathcal{J}| - \omega).$$

Proof. Let $\mathcal{I}_K$ be the intervals of $\mathcal{I}$ with exactly one endpoint in K. Let $\bigcup_{p \in P} \mathcal{I}_{(K,p)}$ be the partition of $\mathcal{I}_K$ where for each $p \in P$, the intervals $\mathcal{I}_{(K,p)}$ are exactly the intervals of $\mathcal{I}_K$ that are assigned to pillar p. Let P' be the set of pillars $p \in P$ such that $\mathcal{I}_{(K,p)}$ is non-empty. Then by Lemma 3.2.2 there is a collection of disjoint intervals $\{K_p : p \in P'\}$ contained in $(0, 1)\backslash K$ such that for every $p \in P'$, the intervals of $\mathcal{I}_{(K,p)}$ have an endpoint within K_p. Let $\preceq_K$ be the ordering of $\{K_p : p \in P'\}$ so that $K_{p_1} \prec_K K_{p_2}$ exactly when either $K_{p_1} < K_{p_2} < K$, or $K < K_{p_1} < K_{p_2}$, or $K_{p_2} < K < K_{p_1}$. Since the intervals $\{K_p : p \in P'\}$ are pairwise disjoint, $\preceq_K$ is a total ordering of $\{K_p : p \in P'\}$. The key property of this total order is that if p, p' are distinct pillars of P' with $p \prec p'$, and $I, I' \in \mathcal{I}_K$ are intervals assigned to p and p' respectively, then I overlaps with I' if the endpoint of I in that is contained in K precedes the endpoint of I_2 that is contained in K.

Let $C' = \psi(\mathcal{I}_K)$. For each $c \in C'$, let $\mathcal{I}_{(K,c)}$ be the intervals of $\mathcal{I}_K$ that are coloured c by ψ. By Lemma 3.2.3, for each $c \in C'$, there exists a pillar $p \in P'$ such that $\mathcal{I}_{(K,c)} \subseteq \mathcal{I}_{(K,p)}$, and in particular every interval of $\mathcal{I}_{(K,c)}$ has an endpoint in K_p. Now let $\preceq_C$ be the total ordering of C' such that for every $c_1, c_2 \in C'$, we have that $c_1 \preceq c_2$ exactly when either there exists a pillar $p \in P'$ such that $\mathcal{I}_{(K,c_1)}, \mathcal{I}_{(K,c_2)} \subseteq \mathcal{I}_{(K,p)}$, and $c_1 \leq c_2$, or there exists distinct pillars $p_1, p_2 \in P'$ such that $\mathcal{I}_{(K,c_1)} \subseteq \mathcal{I}_{(K,p_1)}$, $\mathcal{I}_{(K,c_2)} \subseteq \mathcal{I}_{(K,p_2)}$ and $p_1 \prec_K p_2$. Let $f : C' \to [d(K)]$ be the bijection such that $f(c_1) \leq f(c_2)$ exactly when $c_1 \preceq_C c_2$.

Now let $\mathcal{J} = \{J_1, \ldots, J_{|\mathcal{J}|}\}$ where $J_1 < \cdots < J_{|\mathcal{J}|}$. Next let S be the set of all elements $(x, y) \in [d(K)] \times [|\mathcal{J}|]$ such that there is an interval I of $\mathcal{I}$ that is coloured $f^{-1}(x)$ by ψ, and has an endpoint in J_y. Note that $|S| = \sum_{J \in \mathcal{J}} d(J)$.

Suppose now for the sake of contradiction that

$$\sum_{J \in \mathcal{J}} d(J) > \omega(d(K) + |\mathcal{J}| - \omega).$$

Then by Theorem 3.3.2, there is a chain W contained in S of length at least $\omega + 1$. Since this chain is contained in S, there must exist colours $c_1 \prec_C \cdots \prec_C c_{\omega+1}$ contained in C', and integers $1 \leq x_1 < \cdots < x_{\omega+1} \leq |\mathcal{J}|$ so that for each $j \in [\omega + 1]$, there is an interval $I_j \in \mathcal{I}$ with an endpoint e_j contained in J_{x_j} and $\psi(I_j) = c_j$. Since $J_{x_1} < \cdots < J_{x_{\omega+1}}$, we have that $e_1 < \cdots < e_{\omega+1}$.

Let $p_1, \ldots, p_n \in P'$ be the collection of distinct pillars such that for some integers $0 = a_0 < \cdots < a_{n-1} < a_n = \omega + 1$, we have that for each p_i, the intervals $\mathcal{I}_{p_i} = \{I_j : a_{i-1} < j \leq a_i\}$ are all assigned to pillar p_i. Note that $p_1 \prec_K \cdots \prec_K p_n$ by the definition of the total ordering $(C', \preceq_C)$.

Then by Lemma 3.1.1 and the definition of the pillar assignment $(P, \preceq, \psi)$, for each $i \in [n]$, there exist pairwise overlapping intervals $I^*_{a_{i-1}+1}, \ldots, I^*_{a_i}$ that are all assigned to the pillar p_i, and all have an endpoint contained in $[e_{a_{i-1}+1}, e_{a_i}]$. Since $p_1 \prec_K \cdots \prec_K p_n$, and $[e_1, e_{a_1}] < \cdots < [e_{a_{n-1}+1}, e_{\omega+1}]$, we see that any two intervals of $\{I^*_1, \ldots, I^*_{\omega+1}\}$ that are assigned to distinct pillars also overlap. Hence the intervals $I^*_1, \ldots, I^*_{\omega+1}$ pairwise overlap, a contradiction to the fact that $\omega(\mathcal{I}) = \omega$. □

3.4 Colouring circle graphs

By Lemma 3.2.1, the following theorem strengthens and so implies Theorem 1.2.1, that every circle graph with clique number at most ω has chromatic number at most $2\omega \log_2(\omega) + 2\omega \log_2(\log_2(\omega)) + 10\omega$.

Theorem 3.4.1. *Let $\omega \geq 2$ be an integer, let G be a circle graph with clique number at most ω, and let $\mathcal{I}$ be an interval system with overlap graph G. Let $(P, \preceq, \psi)$ be a pillar assignment of $\mathcal{I}$ such that $\chi(P, \preceq, \psi) \leq 2\omega\log_2(\omega) + 2\omega\log_2(\log_2(\omega)) + 10\omega$, and $d(K)_{(P,\preceq,\psi)} \leq \omega\log_2(\omega) + \omega\log_2(\log_2(\omega)) + 6\omega$ for every arch K of $(P, \preceq, \psi)$. Then there is a complete pillar assignment $(P^*, \preceq^*, \psi^*)$ of $\mathcal{I}$ extending $(P, \preceq, \psi)$ with $\chi(P^*, \preceq^*, \psi^*) \leq 2\omega\log_2(\omega) + 2\omega\log_2(\log_2(\omega)) + 10\omega$.*

Proof. The theorem is trivially true if $(P, \preceq, \psi)$ is a complete pillar assignment, so we proceed by induction on the number of intervals that are not coloured by $\psi = \psi_{(P,\preceq)}$.

Let K be an arch of $(P, \preceq, \psi)$ that contains some interval I of $\mathcal{I}$. Then I is not coloured by ψ. Let q^* be a pillar contained in I. Let $q_0 = \ell(K)$. Now for each integer $i \geq 1$ in increasing order, if the pillar q_{i-1} was chosen and $d_{(P,\preceq,\psi)}(q_{i-1}, r(K)) > 2\omega$, then we choose the next pillar $q_i \in K$ so that $q_i > q_{i-1}$ and $d_{(P,\preceq,\psi)}(q_{i-1}, q_i) = 2\omega$. Note that such a q_i can always be chosen if $d_{(P,\preceq,\psi)}(q_{i-1}, r(K)) > 2\omega$ as incrementally increasing some $q > q_{i-1}$ increases the degree of (q_{i-1}, q) by at most 1. Let n be equal to the largest i such that the pillar q_i is chosen, and let $Q = \{q_1, \ldots, q_n, q^*\}$. Then $d_{(P,\preceq,\psi)}(q_{n-1}, r(K)) > 2\omega$.

Let $\mathcal{J} = \{(q_0, q_1), \ldots, (q_{n-1}, r(K))\}$. Then $\sum_{J \in \mathcal{J}} d(J) > 2\omega n$. So by Lemma 3.3.3,

$$2\omega n < \sum_{J \in \mathcal{J}} d(J) \leq \omega(d(K) + n - \omega) \leq \omega(\omega\log_2(\omega) + \omega\log_2(\log_2(\omega)) + 5\omega + n).$$

Hence $n < \omega\log_2(k) + \omega\log_2(\log_2(\omega)) + 5\omega$, and so $|Q| < \log_2(\omega) + \omega\log_2(\log_2(\omega)) + 5\omega + 1$.

Then by Lemma 3.2.4 there is a pillar assignment $(P', \preceq', \psi')$ extending $(P, \preceq, \psi)$ such that $P' = P \cup Q$, and for every arch K' of $(P', \preceq', \psi')$ contained in K,

$$\begin{aligned} d_{(P',\preceq',\psi')}(K') &\leq 2\omega + \omega\lceil\log_2(|Q| + 1)\rceil \\ &< 3\omega + \omega\log_2(\omega\log_2(\omega) + \omega\log_2(\log_2(\omega)) + 5\omega + 2) \\ &\leq 3\omega + \omega\log_2(8\omega\log_2(\omega)) \\ &= \omega\log_2(\omega) + \omega\log_2(\log_2(\omega)) + 6\omega, \end{aligned}$$

and furthermore

$$\begin{aligned} \chi(P', \preceq', \psi') &\leq \max\left\{\chi(P, \preceq, \psi),\ d_{(P,\preceq,\psi)}(K) + \omega\lceil\log_2(|Q| + 1)\rceil\right\} \\ &\leq \max\left\{\chi(P, \preceq, \psi),\ 2\omega\log_2(\omega) + 2\omega\log_2(\log_2(\omega)) + 10\omega\right\} \\ &= 2\omega\log_2(\omega) + 2\omega\log_2(\log_2(\omega)) + 10\omega. \end{aligned}$$

Hence $(P', \preceq', \psi')$ satisfies the inductive hypothesis. Since $q^* \in Q \subseteq P'$, the interval I is coloured by ψ'. As I is not coloured by ψ, the number of intervals of $\mathcal{I}$ that are not coloured by ψ' is strictly less than the number of intervals of $\mathcal{I}$ that are not coloured by ψ. Hence by induction there exists a complete pillar assignment $(P^*, \preceq^*, \psi^*)$ extending $(P', \preceq', \psi')$ (and thus $(P, \preceq, \psi)$) with $\chi(P^*, \preceq^*, \psi^*) \leq 2\omega \log_2(\omega) + 2\omega \log_2(\log_2(\omega)) + 10\omega$ as required. □

We remark that with more careful arguments it is possible to improve the lower order terms slightly, but we are not aware of a way to improve the leading constant.

Let us briefly sketch how to obtain a $O(n^3)$ time algorithm for obtaining the colouring given by the proof of Theorem 1.2.1 (once the clique number has already been determined, which can be done in $O(n^2)$ time [47]).

Let $q_1, \ldots, q_r$ be pillars added to the pillar assignment $(P, \preceq, \psi)$ as in the proof of Theorem 3.4.1. With a bit more care, we can ensure that each new pillar will have at least one interval assigned to it in the extended pillar assignment $(P', \preceq', \psi')$. Each new pillar q_i can be found in $O(n^2)$ time since by incrementally increasing q, one at a time, a new endpoint of the intervals of $\mathcal{I}$ is included in the interval (q_{i-1}, q_i), and for each new endpoint we just need to check the colour (if it is coloured) it has already received from $(P, \preceq, \psi)$. The new pillars $q_1, \ldots, q_r$ are added to the pillar assignment accordingly to a divided and concur argument ordering (given by Lemma 3.2.4), which can be found in $O(r \log r)$ time. For each new pillar p in the order given by $(P', \preceq')$, the intervals $\mathcal{I}_p$ assigned to p and the colour palette C_p can both be found in $O(n^2)$ time. The colouring of the intervals $\mathcal{I}_p$ is given by Lemma 3.1.1, which is just a first fit colouring according to the ordering of $\mathcal{I}_p$ given by their leftmost endpoints. This can be done in $O(|\mathcal{I}_p|^2)$ time. Putting this altogether, the extended pillar assignment $(P', \preceq', \psi')$ can be found from $(P, \preceq, \psi)$ in $O(n^2 m)$ time, where m is the number of intervals of $\mathcal{I}$ coloured by $(P', \preceq', \psi')$, but not $(P, \preceq, \psi)$. This gives a $O(n^3)$ time algorithm for obtaining the colouring given by Theorem 1.2.1.

The above sketch has been optimized for simplicity rather than efficiency. We suspect that the arguments could be optimized to significantly improve the run time of the algorithm.

3.5 Colouring K_4-free circle graphs

Now by specializing the proof of Theorem 3.4.1 to the case $\omega = 3$, and using Lemma 3.2.5 in place of Lemma 3.2.4, we prove the following which strengthens and so implies Theorem 1.2.3, that every K_4-free circle graph is 19-colourable.

Theorem 3.5.1. *Let G be a circle graph with clique number at most 3, and let $\mathcal{I}$ be an interval system with overlap graph G. Let $(P, \preceq, \psi)$ be a pillar assignment of $\mathcal{I}$ such that $\chi(P, \preceq, \psi) \leq 19$, and $d(K)_{(P,\preceq,\psi)} \leq 13$ for every arch K of $(P, \preceq, \psi)$. Then there is a complete pillar assignment $(P^*, \preceq^*, \psi^*)$ of $\mathcal{I}$ extending $(P, \preceq, \psi)$ with $\chi(P^*, \preceq^*, \psi^*) \leq 19$.*

Proof. As before, the theorem is trivially true if $(P, \preceq, \psi)$ is a complete pillar assignment, so we proceed by induction on the number of intervals that are not coloured by $\psi = \psi_{(P,\preceq)}$.

Let K be an arch of $(P, \preceq, \psi)$ that contains some interval I of $\mathcal{I}$. Then I is not coloured by ψ. Let $q_0 = \ell(K)$. Now for each integer $i \geq 1$ in increasing order, if the pillar q_{i-1} was chosen and $d_{(P,\preceq,\psi)}(q_{i-1}, r(K)) > 7$, then we choose the next pillar $q_i \in K$ so that $q_i > q_{i-1}$ and $d_{(P,\preceq,\psi)}(q_{i-1}, q_i) = 7$. Let n be equal to the largest i such that the pillar q_i is chosen, and let $Q = \{q_1, \ldots, q_n\}$. Then $d_{(P,\preceq,\psi)}(q_{n-1}, r(K)) > 7$.

Let $\mathcal{J} = \{(q_0, q_1), \ldots, (q_{n-1}, r(K))\}$. Then $\sum_{J \in \mathcal{J}} d(J) > 7n$. So by Lemma 3.3.3,

$$7n < \sum_{J \in \mathcal{J}} d(J) \leq 3(d(K) + n - 3) \leq 30 + 3n.$$

Hence $|Q| = n \leq 7$.

Note that by the choice of $q_1, \ldots, q_n$, for every $q \in Q$, there is an interval $I_q \in \bigcup_{p \in P} \mathcal{I}_p$ that overlaps with every interval in $\mathcal{I} \setminus \left(\bigcup_{p \in P} \mathcal{I}_p\right)$ that contains q. So by Lemma 3.2.4 there is a pillar assignment $(P', \preceq', \psi')$ extending $(P, \preceq, \psi)$ such that $P' = P \cup Q$, and for every arch K' of $(P', \preceq', \psi')$ contained in K,

$$\begin{aligned} d_{(P',\preceq',\psi')}(K') &\leq 7 + (3-1)\lceil \log_2(|Q| + 1) \rceil \\ &\leq 7 + 2\lceil \log_2(8) \rceil \\ &= 13, \end{aligned}$$

and furthermore

$$\begin{aligned} \chi(P', \preceq', \psi') &\leq \max\left\{\chi(P, \preceq, \psi),\ d_{(P,\preceq,\psi)}(K) + (3-1)\lceil \log_2(|Q| + 1) \rceil\right\} \\ &\leq \max\{\chi(P, \preceq, \psi),\ 13 + 6\} \\ &= 19. \end{aligned}$$

Hence $(P', \preceq', \psi')$ satisfies the inductive hypothesis. If ψ' does not colour an interval that is uncoloured by ψ, then note that, by the choice of Q, we would have $d_{(P',\preceq',\psi')}(q_0, q_1), \ldots, d_{(P',\preceq',\psi')}(q_{n-1}, q_n), d_{(P',\preceq',\psi')}(q_n, r(K)) \leq 7$. In this case we could simply extend the pillar assignment $(P', \preceq', \psi')$ by picking one extra pillar contained in the interval I, and the resulting pillar assignment would colour I and still satisfy the inductive hypothesis. So we may assume that ψ' colours an interval that is uncoloured by ψ. Hence by induction there exists a complete pillar assignment $(P^*, \preceq^*, \psi^*)$ extending $(P', \preceq', \psi')$ (and thus $(P, \preceq, \psi)$) with $\chi(P^*, \preceq^*, \psi^*) \leq 19$ as required. □

Despite this improvement to 19, we suspect that Theorem 1.2.3 is still far from best possible. It is likely that more specialized arguments could significantly improve this bound. Using the ideas of Ageev's [1] construction, the best lower bound we have been able to find is that there are K_4-free circle graphs with chromatic number at least 6. Although 6 seems closer to the correct answer, we make no conjecture on what the maximum chromatic number of a K_4-free circle graph could be.

The proof of Theorem 1.2.3 can of course be made algorithmic in the same way as Theorem 1.2.1.

3.6 Lower bound

In this section we give a simple construction to prove Theorem 1.2.2. We find it more convenient to use a chord diagram representation of our circle graphs, rather than the interval overlap representations that were used to prove Theorems 1.2.1 and 1.2.3 in the previous sections. We allow chords to coincide and consider the chords to be open, so two chords that share an endpoint only intersect if they share both their endpoints. By slightly perturbing open chords so that no two share an endpoint and so that exactly the same pairs intersect, it can be shown that circle graphs are exactly intersection graphs of open chords on a circle where chords can coincide.

The construction is inspired by those given by Kostochka [67] for both circle graphs and their complements, as well as his proof that the complements of circle graphs are χ-bounded. With essentially the same arguments, our construction also yields a new proof that there are complements of circle graphs with clique number at most ω and chromatic number at least $\omega(\ln \omega - O(1))$.

For positive integers ω and n with $n > 3\omega - 3$ we define a chord diagram $\mathcal{D}_{n,\omega}$ as follows. Let $p_1, q_1, p_2, q_2, \ldots, p_n, q_n$ be points on a circle in cyclic clockwise order. Now for each $i \in [n]$, and $j \in [\omega - 1]$, let $\mathcal{C}_{i,j}$ consist of exactly $\left\lfloor \frac{\omega}{j+1} \right\rfloor$ coinciding open chords with endpoints p_i, q_{i+j} (taking $i + j$ modulo n). Then let $\mathcal{D}_{n,\omega} = \bigcup_{i \in [n]} \bigcup_{j \in [\omega-1]} \mathcal{C}_{i,j}$. For an example, see Figure 3.3, which illustrates the chord diagram $\mathcal{D}_{17,6}$.

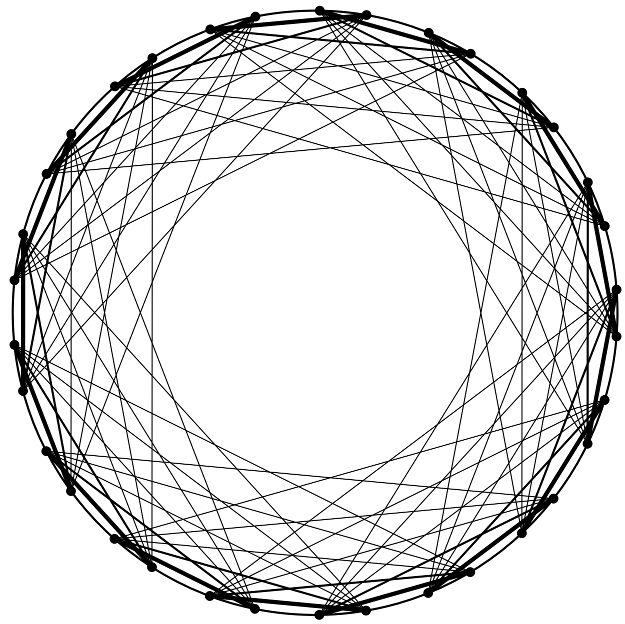

Figure 3.3: The chord diagram of $\mathcal{D}_{17,6}$. The thickness of the chords corresponds to the number of chords that coincide.

Now we prove some bounds on the number of chords contained in $\mathcal{D}_{n,\omega}$, as well as the size of the largest set of pairwise intersecting and pairwise disjoint chords in $\mathcal{D}_{n,\omega}$.

Lemma 3.6.1. *The set of chords $\mathcal{D}_{n,\omega}$ has size greater than $n\omega(\ln\omega - 2)$.*

Proof. By definition,

$$|\mathcal{D}_{n,\omega}| = \left|\bigcup_{i\in[n]}\bigcup_{j\in[\omega-1]}\mathcal{C}_{i,j}\right| = \sum_{i=1}^{n}\sum_{j=1}^{\omega-1}|\mathcal{C}_{i,j}| = n\sum_{j=1}^{\omega-1}\left\lfloor\frac{\omega}{j+1}\right\rfloor = n\sum_{j=1}^{\omega}\left\lfloor\frac{\omega}{j}\right\rfloor - n\omega.$$

Then observe that:

$$n\sum_{j=1}^{\omega}\left\lfloor\frac{\omega}{j}\right\rfloor - n\omega \geq n\sum_{j=1}^{\omega}\frac{\omega}{j} - 2n\omega = n\omega\sum_{j=1}^{\omega}\frac{1}{j} - 2n\omega > n\omega\ln\omega - 2n\omega = n\omega(\ln\omega - 2).$$

Hence the lemma follows. □

Lemma 3.6.2. *There are no $\omega + 1$ pairwise intersecting chords contained in $\mathcal{D}_{n,\omega}$.*

Proof. Let $C \subseteq \mathcal{D}_{n,\omega}$ be a set of pairwise intersecting chords. Let P be the set of endpoints of chords in C that are contained in $\{p_1, \dots p_n\}$, and similarly for Q. Then $|P| = |Q|$ as a pair of open chords that share an endpoint only intersect if they share both their endpoints. Furthermore after possibly rotating the chords of C around the circle, we can assume without loss of generality that $P = \{p_{a_1}, \dots, p_{a_\ell}\}$ and $Q = \{q_{b_1}, \dots, q_{b_\ell}\}$ with $a_1 < \cdots < a_\ell \leq b_1 < \cdots < b_\ell$, and that every chord of C has one of $\{p_{a_1}, q_{b_1}\}, \dots, \{p_{a_\ell}, q_{b_\ell}\}$ as its endpoints. For each $i \in [\ell]$, there are exactly $\left\lfloor \frac{\omega}{(b_i - a_i - 1) + 1} \right\rfloor = \left\lfloor \frac{\omega}{b_i - a_i} \right\rfloor$ chords with endpoints contained in $\{p_{a_1}, q_{b_1}\}$. Therefore

$$|C| \leq \sum_{i=1}^{\ell} \left\lfloor \frac{\omega}{b_i - a_i} \right\rfloor \leq \sum_{i=1}^{\ell} \frac{\omega}{b_i - a_i} \leq \sum_{i=1}^{\ell} \frac{\omega}{\ell} = \omega.$$

□

Lemma 3.6.3. *There is no set of n pairwise disjoint chords contained in $\mathcal{D}_{n,\omega}$.*

Proof. Let S be a set of pairwise disjoint chords of $\mathcal{D}_{n,\omega}$. Now consider an auxiliary directed graph G on vertex set $\{v_1, \dots, v_n\}$ where there is an edge directed from v_i to v_j whenever S contains a chord with endpoints p_i and q_j. First note that $|S| = |E(G)|$ as all the chords with endpoints p_i and q_j intersect.

Now observe that G is outerplanar, with the natural embedding of $v_1, \dots, v_n$ being on the circle in clockwise order and all edges of G directed in the clockwise direction. In a directed outerplanar graph with such an embedding, all cycles contain a directed path of length 2 in the clockwise direction. However G has no directed path of length 2 as such a path with internal vertex v_i would imply that S contains a chord with an endpoint p_i, and another with the endpoint q_i, a contradiction since all such chords of $\mathcal{D}_{n,\omega}$ intersect. Hence G is a forest, and so $|S| = |E(G)| < |V(G)| = n$ as required. □

We now prove Theorem 1.2.2, that for every positive integer k there is a circle graph with clique number at most k and chromatic number at least $k(\ln k - 2)$.

Proof of Theorem 1.2.2. For a positive integer ω, choose some $n > 3\omega - 3$. Let $G_{n,\omega}$ be the intersection graph of the chord diagram $\mathcal{D}_{n,\omega}$, so $G_{n,\omega}$ is a circle graph. By Lemma 3.6.2, the graph $G_{n,\omega}$ has clique number at most ω. By Lemma 3.6.1, $|V(G_{n,\omega})| > n\omega(\ln \omega - 2)$, and by Lemma 3.6.3, the stable sets of $G_{n,\omega}$ all have size less than n. Hence $\chi(G_{n,\omega}) > \frac{n\omega(\ln \omega - 2)}{n} = \omega(\ln \omega - 2)$ as desired. □

Chapter 4

Vertex-minors

In this chapter we prove Theorem 1.3.1 that proper vertex-minor-closed classes of graphs are χ-bounded. Along the way we also prove Theorem 1.3.3 that they are also linearly 2-controlled. At the end of the chapter in Section 4.7, we will also discuss some of the details of the proof of Theorem 1.3.2 that proper pivot-minor-closed classes of graphs are χ-bounded.

Let us recall the definitions of local complementation and vertex-minors. The action of performing local complementation at a vertex v in a graph G replaces the induced subgraph on $N(v)$ by its complement Figure 1.2). The resulting graph is denoted by $G * v$. We say that a graph H is a vertex-minor of a graph G if H can be obtained from G by a sequence of vertex deletions and local complementations.

Something that is conceptually simpler than finding a given graph H as a vertex-minor is finding an induced subdivision of H. In a graph H, the act of replacing an edge uw with a vertex v adjacent to u and w only is known as ***subdividing*** the edge uw. A graph G is a ***subdivision*** of a graph H if G can be obtained from H by a sequence of subdivisions. If v is a vertex of degree two in a graph G and v is adjacent to two non-adjacent vertices u and w then we say that the graph obtained by removing the vertex v and adding an edge between u and w is the graph obtained from G by ***smoothing*** the vertex v. Observe that the graph obtained from G by smoothing a vertex v is $(G * v) - v$, and so in particular is a vertex-minor of G. So more generally, by repeated smoothing of vertices, a graph H is a vertex-minor of any subdivision of H. We let G^k denote the graph obtained from G by subdividing each edge k times.

Scott [94] conjectured that for every graph H, the graphs containing no induced subdivision of H are χ-bounded. This is known to be true when H is a tree [94], a cycle [20], and generalizing both trees and cycles, a banana tree [95]. However Scott's conjecture is false in general; Pawlik et al. [85] proved that segment intersection graphs (which contain no induced subdivision of K_5^1) are not χ-bounded. Theorem 1.3.1 recovers one possible weakening of Scott's conjecture.

To prove Theorem 1.3.1 we use the idea of ρ-controlled classes of graphs. Recall that $\chi^{(\rho)}(G)$ denotes the maximum chromatic number of an induced subgraph of G with radius at most ρ, and we say that a class of graphs $\mathcal{G}$ is ρ-controlled if there exists a function f such that $\chi(G) \leq f(\chi^{(\rho)}(G))$ for all $G \in \mathcal{G}$.

With the idea of ρ-control in mind, one may naturally split the problem of proving that a class of graphs $\mathcal{G}$ is χ-bounded into subproblems. The first is to show that for some $\rho \geq 2$, $\mathcal{G}$ is ρ-controlled. The next is to reduce control and show that $\mathcal{G}$ is 2-controlled. The final subproblem is to make use of the fact that $\mathcal{G}$ is 2-controlled to prove χ-boundedness. So intuitively the strategy is essentially to first show that huge chromatic number is (at least in this looser sense) a local property. From there we then argue that it is actually a more and more local property until eventually getting down to the extremely local property of containing a large clique. The framework of this strategy was first introduced by Scott [94].

For vertex-minors, the first step turns out to be the most challenging. In our proof of Theorem 1.3.1 in [25], the first step was to prove that proper vertex-minor-closed classes are (linearly) 9-controlled. For the second step we then applied a theorem of Chudnovsky, Scott, and Seymour [21] to quickly reduce "9" down to "2", although at the cost of not retaining linearity. The proof of Theorem 1.3.1 in this thesis very roughly follows the same framework as that in [25], but we incorporate some ideas that will be used to prove Theorem 1.3.2 in [24]. This simplifies the proof somewhat with the most significant change being that we prove 2-control directly. An advantage of this is that we are able to retain linearity and thus prove Theorem 1.3.3 that vertex-minor-closed classes are linearly 2-controlled. Let us remark that despite new ideas from [24] allowing us to simplify the proof of Theorem 1.3.1, the proof of Theorem 1.3.2 that pivot-minor-closed classes of graphs are χ-bounded still remains significantly more complicated than that of Theorem 1.3.1 (see Section 4.7 for further discussion on this).

In Section 4.1 we cover some necessary preliminaries and discuss simple universal graphs for vertex-minors. Aiming to find particular universal graphs as vertex-minors is much more convenient than aiming to find an arbitrary graph H as a vertex-minor. However, there is a balance to be struck between how simple the universal graphs that we aim to find are, and how easy they are to find as vertex-minors. In Section 4.2 we then find more complicated universal graphs that are easier to find as vertex-minors. The main result of this section (Theorem 4.2.7) is later used to finish off each of the steps in the ρ-control strategy. Sections 4.3–4.5 are then dedicated proving Theorem 1.3.3 that vertex-minor-closed classes are (linearly) 2-controlled. Each these sections are very roughly dedicated to finding different parts of the universal graphs that we seek. We will give a more detailed overview of these sections at the end of Section 4.2, after we have introduced these universal graphs. Then Section 4.6 handles the last step of the ρ-control strategy by using Theorem 1.3.3 to finally prove Theorem 1.3.1 that vertex-minor-closed class of graphs are χ-bounded. After all this, in Section 4.7, we further discuss pivot-minors and the proof of Theorem 1.3.2.

4.1 Preliminaries

Given two sets A and B, we let $A - B$ denote the subset of A obtained by removing the elements of $A \cap B$.

Given a vertex v of a graph G, we let $N(v)$ denote its ***neighbourhood***, which is the set of vertices adjacent to v. More generally given a set A of vertices in a graph G, we let $N(A)$ be the set of vertices in $V(G) - A$ that are adjacent to a vertex of A. If the graph is not clear from context, we use $N_G(v)$ or $N_G(A)$. Given an integer $t \geq 0$, we let $N_t(A)$ be the set of vertices at distance exactly t from A, and we let $N_t[A]$ be the set of vertices at distance at most t from A. We may denote the ***closed neighbourhood*** $N_1[A]$ by $N[A]$.

We say that two sets of vertices A and B in a graph G are ***complete*** to each other if for all $a \in A$ and $b \in B$, we have $ab \in E(G)$. Similarly A and B are ***anti-complete*** if for all $a \in A$ and $b \in B$, we have $ab \notin E(G)$. If for all $b \in B$, there exists a vertex $a \in A$ that is adjacent to b, then we say that A ***dominates*** B. For a simple example observe that if v is a vertex of a graph G, then $N_{t-1}(v)$ dominates $N_t(v)$. Recall that for a positive integer n, we let $[n]$ denote the set $\{1, 2, \ldots, n\}$.

Given a set C of vertices on a graph G, we denote the ***induced subgraph*** of G on vertex set C by $G[C]$. For convenience we often use $\chi(C)$ for $\chi(G[C])$. Given a set A of vertices of a graph G, we let $G - A$ be the graph obtained from G by deleting the vertices A. Similarly for a set F of edges of G, we let $G - F$ be the graph obtained from G by deleting the edges F. For a set F of edges in G, the graph obtained from G by contracting each edge of F (and then removing any resulting loops or multiple edges) is denoted by G/F. Given two disjoint sets A and B of vertices in a graph G, we let $E(A, B)$ denote the set of edges between A and B. For a set A of vertices in a graph G, we let $E(A)$ denote the set of edges between vertices of A.

For an edge uv of a graph G, let $V_1 = N(u) - N[v]$, $V_2 = N(v) - N[u]$, and $V_3 =$

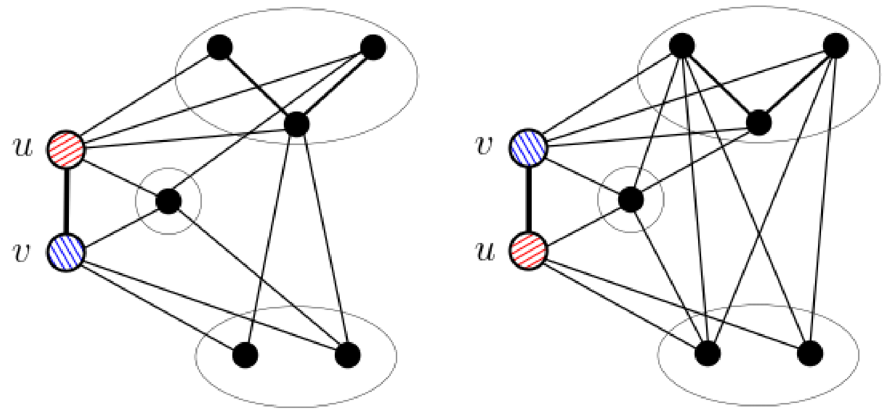

Figure 4.1: The effect of pivoting an edge uv.

$N(u) \cap N(v)$. ***Pivoting*** the edge uv of G is the act of first complementing the edges between each of the three pairs of vertex sets $(V_1, V_2), (V_2, V_3)$, and (V_1, V_3), and then swapping the vertex labels of u and v (see Figure 4.1). We denote this graph by $G \wedge uv$. A graph H is a ***pivot-minor*** of a graph G if H can be obtained from G by a sequence of vertex deletions and pivots. Pivot-minors are closely related to vertex-minors; it can be shown that $G \wedge uv = G * u * v * u = G * v * u * v$ (see [83] for a proof). In particular, notice that this implies that if H is a pivot-minor of G, then H is also a vertex-minor of G. We will use this fact repeatedly without reference.

For positive integers n, m, we let $K_{n,m}$ denote the complete bipartite graph whose vertices can be partitioned into two stable sets of size n and m that are complete to each other. We prove a motivating lemma.

Lemma 4.1.1. *The graph* $K^1_{n,\binom{n}{2}}$ *contains every* n*-vertex graph as a vertex-minor.*

Proof. First observe that by smoothing and deleting vertices we may obtain K^1_n as a vertex-minor. Now given an n-vertex graph G, we may associate its vertices with the non-subdivision vertices of K^1_n. Now for each pair of distinct vertices u and v of G we may do one of two things. If uv is an edge of G then we may simply smooth the corresponding degree-2 vertex of K^1_n. If uv is not an edge then we may just delete the corresponding degree-2 vertex of K^1_n. Doing this for each such pair u, v results in the desired vertex-minor G. □

So $K^1_{n,m}$ provides suitable universal graphs for vertex-minors.

4.2 Beyond $K^1_{n,m}$

In this section we extend Lemma 4.1.1 to find universal graphs that are easier to find as vertex-minors. Lemma 4.2.2 will show that $K^1_{n,m}$ remains universal even when certain additional edges are allowed to be added. Since $K^2_{n,m}$ is a subdivision of $K^1_{n,m}$, these also provide universal graphs for vertex-minors. Lemma 4.2.3 again shows that these graphs are universal even when certain additional edges can be added to $K^2_{n,m}$. Lastly the main result of this section (Theorem 4.2.7) unifies and slightly extends both Lemma 4.2.2 and Lemma 4.2.3. This is the result that we will eventually use to conclude the proof of Theorem 1.3.3 and Theorem 1.3.1.

All the proofs in section have a common theme of appropriately applying the multicolour Ramsey theorem to reduce to key cases, and then proceeding with an appropriate sequence of local completions and vertex deletions to obtain either $K^1_{n,m}$ or $K^2_{n,m}$ as a vertex-minor. For positive integers t, n, we let $R^*(n;t)$ denote the ***t*-colour Ramsey number of *n***. In other words, the multicolour Ramsey theorem tells us that for every t-colouring of the edges of $K_{R^*(n;t)}$, there is a monochromatic clique with n vertices.

We begin with taking Lemma 4.1.1 a step further by showing that $K^1_{n,m}$ remains universal even when certain additional edges are added. A graph G with vertex-set $\{w_1, \ldots, w_n\} \cup \{y_{i,j} : i \in [n], j \in [m]\} \cup \{z_1, \ldots, z_m\}$ is an ***interfered*** $\boldsymbol{K^1_{n,m}}$ if

- for each $i \in [n], j \in [m]$, $w_i y_{i,j} \in E(G)$ and $y_{i,j} z_j \in E(G)$,
- all other edges of G are contained in $\{w_k y_{i,j} : i, k \in [n], j \in [m] \text{ with } i < k\}$.

See Figure 4.2 for an example of an interfered $K^1_{3,3}$.

Lemma 4.2.1. *Let G be an interfered $K^1_{R^*(n;2^{2m}),2m}$. Then G contains $K^1_{n,m}$ as a pivot-minor.*

Proof. Consider an auxiliary complete graph A on vertex set $[R^*(n;2^{2m})]$. For each $1 \leq i < k \leq R^*(n;2^{2m})$, colour the edge ik of A according to which of the $2m$ edges of $\{w_k y_{i,j} : j \in [2m]\}$ are in G. This is a 2^{2m}-edge colouring of A. So by the multicolour Ramsey theorem, there exists a $N \subseteq [R^*(n;2^{2m})]$ with $|N| = n$ such that for all $i, i', k, k' \in N$ and $j \in [2m]$ with $i < k$ and $i' < k'$, $y_{i,j}$ is adjacent to w_k if and only if $y_{i',j}$ is adjacent to $w_{k'}$.

If there is a $M \subseteq [2m]$ with $|M| = m$ such that for all $i, k \in N$ and $j \in M$ with $i < k$, $y_{i,j}$ is not adjacent to w_k, then we obtain $K^1_{n,m}$ as an induced subgraph on the vertex set $\{w_i : i \in N\} \cup \{y_{i,j} : i \in N, j \in M\} \cup \{z_j : j \in M\}$.

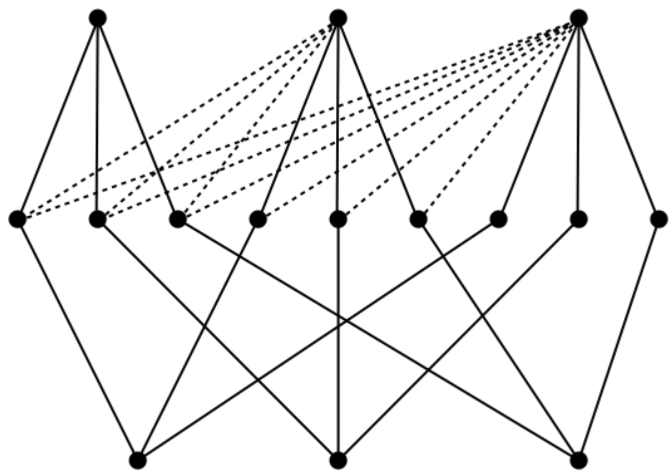

Figure 4.2: An interfered $K^1_{3,3}$. Dashed lines indicate possible additional edges.

So, we may assume that there exists a $M \subseteq [2m]$ with $|M| = m+1$ such that for all $i, k \in N$ and $j \in M$ with $i < k$, $y_{i,j}$ is adjacent to w_k. Let $N = \{a_1, \ldots, a_n\}$ where $a_1 < \ldots < a_n$ and choose some $b \in M$. Now let

$$H = G \wedge w_{a_1}y_{a_1,b} \wedge w_{a_2}y_{a_2,b} \wedge \cdots \wedge w_{a_n}y_{a_n,b}.$$

For each $1 \leq \ell \leq n$ in order, on the vertex set $\{w_{a_k} : \ell \leq k \leq n\} \cup \{y_{i,j} : i \in N, j \in M\} \cup \{z_j : j \in M \backslash \{b\}\}$, pivoting the edge $w_{a_\ell}y_{a_\ell,b}$ removes the edges between vertex sets $\{y_{a_\ell,j} : j \in M \backslash \{b\}\}$ and $\{w_{a_k} : \ell < k \leq n\}$, and swaps the labels of w_{a_ℓ} and $y_{\ell,b}$. Thus, the induced subgraph of H on vertex set $\{y_{i,b} : i \in N\} \cup \{y_{i,j} : i \in N, j \in M \backslash \{b\}\} \cup \{z_j : j \in M \backslash \{b\}\}$ provides the desired pivot-minor of G. □

As a consequence of Lemma 4.1.1 and Lemma 4.2.1, we get the following.

Lemma 4.2.2. *For every positive integer n, there exists a pair of positive integers q, h such that every interfered $K^1_{q,h}$ contains every n-vertex graph as a vertex-minor.*

Now we will examine supergraphs of $K^2_{n,m}$ that retain the property of being universal graphs for vertex-minors. For the rest of this section we will assume that $V(K^2_{n,m}) = \{w_1, \ldots, w_n\} \cup \{x_{i,j} : i \in [n], j \in [m]\} \cup \{y_{i,j} : i \in [n], j \in [m]\} \cup \{z_1, \ldots, z_m\}$ where for each $i \in [n], j \in [m]$, $w_i x_{i,j}, x_{i,j}y_{i,j}, y_{i,j}z_j \in E(K^2_{n,m})$.

A graph G is an ***interfered*** $\boldsymbol{K^2_{n,m}}$ if it contains $K^2_{n,m}$ as a subgraph (with vertex set as above), and all other edges of G are contained in the union of the two sets

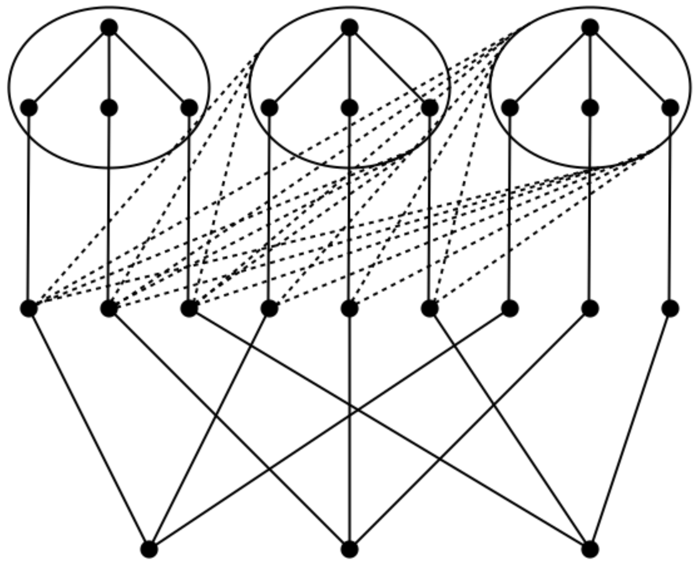

Figure 4.3: An interfered $K^2_{3,3}$. Dashed lines indicate possible additional edges.

1. $E_T = \{y_{i,j}w_k : i, k \in [n], j \in [m] \text{ with } i < k\}$, and
2. $E_M = \{y_{i,j}x_{k,j'} : i, k \in [n], j, j' \in [m] \text{ with } i < k\}$.

We call the edges of $E(G) \cap E_T$ the ***tip edges*** of G, and the edges of $E(G) \cap E_M$ the ***middle edges*** of G. See Figure 4.3 for an example of an interfered $K^2_{3,3}$. These are also universal graphs for vertex-minors.

Lemma 4.2.3. *For every positive integer n, there exists a pair of positive integers q, h such that every interfered $K^2_{q,h}$ contains every n-vertex graph as a vertex-minor.*

To prove Lemma 4.2.3 we will aim to find $K^2_{n,\binom{n}{2}}$ or $K^1_{n,\binom{n}{2}}$ as a pivot-minor by clearing out the additional edges in each of the two sets added to some $K^2_{q,h}$ to create G. Let us first take care of the tip edges.

Lemma 4.2.4. *Let G be an interfered $K^2_{R^*(3n-2;2^{m+1}),m+1}$. Then G contains an interfered $K^2_{n,m}$ as a pivot-minor that has no tip edges.*

Proof. By the multicolour Ramsey theorem, there exists a $N \subseteq [R^*(3n-2; 2^{m+1})]$ with $|N| = 3n - 2$ such that for all $i, i', k, k' \in N$ and $j \in [m+1]$ with $i < k$ and $i' < k'$, $y_{i,j}$ is adjacent to w_k if and only if $y_{i',j}$ is adjacent to $w_{k'}$.

If there is no $b \in [m+1]$ such that for all $i, k' \in N$ with $i < k$, $y_{i,b}$ is adjacent to w_k, then G contains an interfered $K^2_{3n-2,m+1}$ without any tip edges as an induced subgraph on the vertex set $\{w_i : i \in N\} \cup \{x_{i,j} : i \in N, j \in [m+1]\} \cup \{y_{i,j} : i \in N, j \in [m+1]\} \cup \{z_j : j \in [m+1]\}$.

So we may assume that there exists a $b \in [m+1]$ such that for all $i, k' \in N$ with $i < k$, $y_{i,b}$ is adjacent to w_k. Without loss of generality, we may assume that $b = m+1$. Let $N = \{a_1, \ldots, a_{3n-2}\}$ where $a_1 < \ldots < a_{3n-2}$. Now let

$$H = G \wedge y_{a_2,m+1}w_{a_3} \wedge y_{a_5,m+1}w_{a_6} \wedge \cdots \wedge y_{a_{3n-4},m+1}w_{a_{3n-3}}.$$

Note that for each $1 \leq \ell \leq n-1$ in order, two consequences of pivoting the edge $y_{a_{3\ell-1},m+1}w_{a_{3\ell}}$ is the removal of the edges between vertex sets $\{y_{a_{3\ell-2},j} : j \in [m]\}$ and $\{w_{a_i} : 3\ell < i \leq 3n-2\}$, and a possible alteration of the edges between $\{y_{a_{3\ell-2},j} : j \in [m]\}$ and $\{x_{a_i,j} : 3\ell < i \leq 3n-2, j \in [m]\}$. These pivots alter the induced subgraph H' of H on vertex set $\{w_{a_{3i-2}} : i \in [n]\} \cup \{x_{a_{3i-2},j} : i \in [n], j \in [m]\} \cup \{y_{a_{3i-2},j} : i \in [n], j \in [m]\} \cup \{z_j : j \in [m]\}$ in no other way. Thus, H' provides the desired pivot-minor of G. □

It now just remains to eliminate the middle edges. These are trickier to remove. Let G be an interfered $K^2_{n,m}$, and let H be a bipartite graph on vertex set $\{y_1, \ldots, y_m\} \cup \{x_1, \ldots, x_m\}$. We say that a graph G is an H-interfered $K^2_{n,m}$ if G is an interfered $K^2_{n,m}$ without any tip edges, and for all $i, k \in [n]$ and $j, j' \in [m]$ with $i < k$, $y_{i,j}$ is adjacent to $x_{k,j'}$ in G if and only if y_j is adjacent to $x_{j'}$ in H. We apply the multicolour Ramsey theorem to find a H-interfered $K^2_{n,m}$ for some H.

Lemma 4.2.5. *Let G be an interfered $K^2_{R^*(n;2^{m^2}),m}$ without any tip edges. Then G contains a H-interfered $K^2_{n,m}$ as an induced subgraph for some bipartite graph H on vertex set $\{y_1, \ldots, y_m\} \cup \{x_1, \ldots, x_m\}$.*

Proof. Consider an auxiliary complete graph on vertex set $[R^*(n;2^{m^2})]$, for each pair $1 \leq i < k \leq R^*(n;2^{m^2})$, we colour the edge ik according to the bipartite subgraph of G with vertex set $\{y_{i,j} : j \in [m]\} \cup \{x_{k,j} : j' \in [m]\}$. The lemma now follows by applying the multicolour Ramsey theorem. □

It is not always possible to eliminate the middle edges as we may hope, however in this case we can still find a $K^1_{n,m}$ as a pivot-minor. So for the next step we aim to find either $K^1_{n,m}$ or $K^2_{n,m}$ as a pivot-minor. Our application of the multicolour Ramsey theorem in the proof of the following lemma is essentially the same as using a Ramsey theorem for ordered bipartite graphs.

Lemma 4.2.6. *Let G be a H-interfered $K^2_{3n+1,2R^*(3m-2;4)}$ with $m \geq 2$. Then G contains either $K^1_{n,m}$ or $K^2_{n,m}$ as a pivot-minor.*

Proof. Consider an auxiliary complete graph A on vertex set $[2R^*(3m-2;4)]$. For each pair $1 \leq j < \ell \leq [2R^*(3m-2;4)]$ we colour the edge $j\ell$ of A one of four colours corresponding to whether or not y_jx_ℓ is an edge of H and whether or not $y_\ell x_j$ is an edge of H. Furthermore we colour the vertices v of A one of two colours corresponding to whether or not y_vx_v is an edge of H. By pigeon-hole principle, there exists a $B \subseteq [2R^*(3m-2;4)]$ with $|B| = R^*(3m-2;4)$ so that either y_vx_v is an edge of H for every $v \in B$, or y_vx_v is not an edge of H for every $v \in B$. By the multicolour Ramsey theorem, there exists a $D \subseteq B$ with $|D| = 3m-2$ such that for every $j, j'\ell, \ell' \in D$ with $j < \ell$ and $j' < \ell'$, y_jx_ℓ is an edge of H if and only if $y_{j'}x_{\ell'}$ is an edge of H, and $y_\ell x_j$ is an edge of H if and only if $y_{\ell'}x_{j'}$ is an edge of H. So after possibly relabelling vertices (so that D is mapped to $[3m-2]$ with either the same or reversed ordering), we obtain an induced subgraph G' of G that is a H'-interfered $K^2_{3n+1,3m-2}$ where either

(1) for every $i \in [3m-2]$, $N_{H'}(y_i) = \{x_i\}$, in which case we say that H' is a ***coupled matching***, or

(2) for every $i \in [3m-2]$, $N_{H'}(y_i) = \{x_1, \ldots x_{3m-2}\}$, in which case we say that H' is ***complete***, or

(3) for every $i \in [3m-2]$, $N_{H'}(y_i) = \{x_1, \ldots x_i\}$, in which case we say that H' is a ***coupled half-graph***, or

(4) for every $i \in [3m-2]$, $N_{H'}(y_i) = \emptyset$, in which case we say that H' is ***anti-complete***, or

(5) for every $i \in [3m-2]$, $N_{H'}(y_i) = \{x_1, \ldots x_{3m-2}\} \backslash \{x_i\}$, in which case we say that H' is a ***anti-coupled matching***, or

(6) for every $i \in [3m-2]$, $N_{H'}(y_i) = \{x_1, \ldots x_{i-1}\}$, in which case we say that H' is a ***uncoupled half-graph***.

We must now handle these six cases. Case (4) that H' is anti-complete is trivial, since in this case G' is isomorphic to $K^2_{3n+1,3m-2}$. Case (1) is also straightforward since then G' contains an induced subgraph isomorphic to $K^1_{3n,3m-2}$ (with vertices $w_2, \ldots, w_{3n+1}, y_{1,1}, \ldots y_{1,3m-2}$ forming the vertices degree at least three). By using pivot-minors, we can essentially reduce cases (2) and (6) to case (4), and cases (3) and (5) to case (1).

Claim 4.2.6.1. *If H' is either complete or a anti-coupled matching, then G' contains a H^*-interfered $K^2_{n+1,3m-4}$ as a pivot-minor where H^* is either anti-complete or a coupled matching respectively.*

Proof. Let

$$G'' = G' \wedge y_{2,3m-2}x_{3,3m-3} \wedge y_{5,3m-2}x_{6,3m-3} \wedge \cdots \wedge y_{3n-1,3m-2}x_{3n,3m-3}.$$

Note that for each $1 \leq \ell \leq n$ in order, a consequence of pivoting the edge $y_{3\ell-1,3m-2}x_{3\ell,3m-3}$ is that the edges between $\{y_{3\ell-2,j} : j \in [3m-4]\}$ and $\{x_{3i-2,j} : 3\ell < i \leq 3n+1, j \in [3m-4]\}$ are complemented. Let G^* be the induced subgraph of G'' on vertices $\{w_{3i-2} : i \in [n+1]\} \cup \{x_{3i-2,j} : i \in [n+1], j \in [3m-4]\} \cup \{y_{3i-2,j} : i \in [n+1], j \in [3m-4]\} \cup \{z_j : j \in [3m-4]\}$. The pivots alter the induced subgraph G^* in no other way. Thus, G^* provides the desired pivot-minor of G'. □

So in particular, if H' is complete, then G' contains $K^2_{n+1,3m-4}$ as a pivot-minor, and if H' is a anti-coupled matching, then as before, G' contains $K^1_{n,3m-4}$ as a pivot-minor. This handles cases (2) and (5) respectively.

Claim 4.2.6.2. *If H' is either a coupled or uncoupled half-graph, then G' contains a H^*-interfered $K^2_{n+1,m}$ as a pivot-minor where H^* is either a coupled matching, or anticomplete respectively.*

Proof. Let

$$\begin{aligned} G^*_0 = G' &- \{w_2, w_3, w_5, w_6, \ldots, w_{3n-1}, w_{3n}\} \\ &- \{z_2, z_3, z_5, z_6, \ldots, z_{3n-1}, z_{3n}\} \\ &- \{x_{3i-2,3j-1} : i \in [n+1], j \in [m-1]\} \\ &- \{x_{3i-2,3j} : i \in [n+1], j \in [m-1]\} \\ &- \{y_{3i-2,3j-1} : i \in [n+1], j \in [m-1]\} \\ &- \{y_{3i-2,3j} : i \in [n+1], j \in [m-1]\}. \end{aligned}$$

For each $1 \leq \ell \leq n-1$ in order, let

$$\begin{aligned} G^*_\ell = (&G^*_{\ell-1} \wedge y_{3\ell-1,3}x_{3\ell,2} \wedge y_{3\ell-1,6}x_{3\ell,5} \wedge \cdots \wedge y_{3\ell-1,3m-3}x_{3\ell,3m-4}) \\ &- \{x_{i,j} : i \in \{3\ell-1, 3\ell\}, j \in [3m-2]\} \\ &- \{y_{i,j} : i \in \{3\ell-1, 3\ell\}, j \in [3m-2]\}. \end{aligned}$$

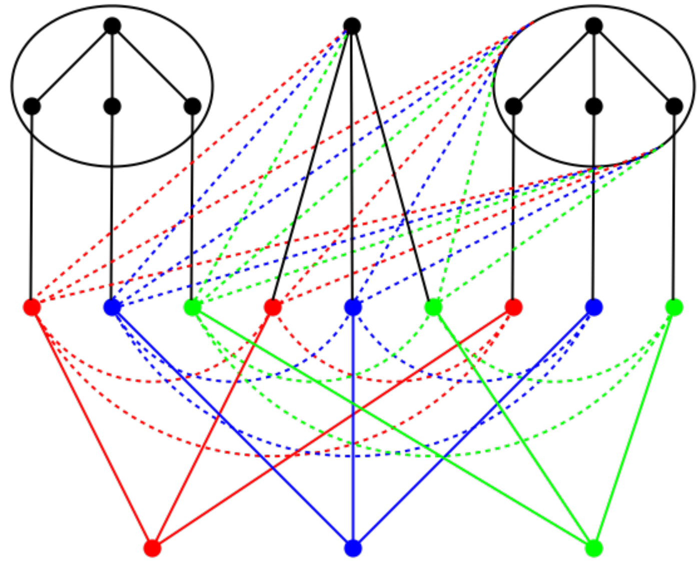

Figure 4.4: A $(3,3)$-frame. Dashed lines indicate possible additional edges. The vertices and certain edges incident to them are coloured differently for each of $N[z_1], N[z_2], N[z_3]$.

Observe that for fixed ℓ, for each $1 \leq r \leq m-1$ in order, pivoting the edge $y_{3\ell-1,3r}x_{3\ell,3r-1}$ of $G^*_{\ell-1}$ removes all the edges between the two vertex sets $\{y_{3\ell-2,3j-2} : r < j \leq m\} \cup \{y_{3\ell-1,3j} : r < j \leq m-1\}$ and $\{x_{i,j} : 3\ell \leq i \leq 3n+1,\ 1 \leq j \leq 3r\} \cap V(G^*_{\ell-1})$. So for each $j \in [m]$, if H' is coupled half-graph, then $N_{G^*_\ell}(y_{3\ell-2,3j-2}) = \{x_{3\ell-2,3j-2}, z_{3j-2}\} \cup \{x_{i,3j-2} : 3\ell \leq i \leq 3n+1\}$, otherwise if H' is a uncoupled half-graph, then $N_{G^*_\ell}(y_{3\ell-2,3j-2}) = \{x_{3\ell-2,3j-2}, z_{3j-2}\}$. So we see that G^*_{n-1} provides the desired pivot-minor of G'. □

As before, if H' is a coupled half-graph, then G' contains $K^1_{n,m}$ as a pivot-minor, and if H' is a uncoupled half-graph, then G' contains $K^2_{n+1,m}$ as a pivot-minor. This handles cases (3) and (6) respectively, thus completing the proof. □

Since $K^2_{n,m}$ contains $K^1_{n,m}$ as a vertex-minor, Lemma 4.2.3 now follows from applying Lemma 4.2.4, Lemma 4.2.5, Lemma 4.2.6, and Lemma 4.1.1 in order.

We say that a graph G' is a ***(q, h)-frame*** if there is a graph G that is an interfered $K^2_{q,h}$ and a $X \subseteq [q]$ such that G' can be obtained from G by first contracting the edges $w_i x_{i,j}$ such that $i \in X$ and $j \in [h]$, and then possibly adding edges to each of the induced subgraphs on vertex sets $N(z_1), \ldots, N(z_h)$. See Figure 4.4 for an example of a $(3,3)$-frame. In particular, compare Figure 4.4 with Figures 4.2 and 4.3, and observe that (q,h)-frames generalise both the notion of an interfered $K^1_{q,h}$, and the notion of an interfered $K^2_{q,h}$.

The following is the main result of this section, it unifies and slightly extends both Lemma 4.2.2 and Lemma 4.2.3.

Theorem 4.2.7. *For every positive integer n, there exists a pair of positive integers q, h such that every (q, h)-frame contains every n-vertex graph as a vertex-minor.*

Proof. Let q_1, h_1 be as in Lemma 4.2.2, and let q_2, h_2 be as in Lemma 4.2.3. Let $q' = q_1 + q_2 - 1$. Then let $h = \max\{h_1, h_2\}$ and $q = R^*(q'; 2h)$.

By the multicolour Ramsey theorem, there exists some $Q \subseteq [q]$ with $|Q| = q'$ such that for each $j \in [h]$, $\{y_{i,j} : i \in Q\}$ is either a clique or a stable set. By possibly locally complementing on vertices of $\{z_j : j \in [h]\}$, we may assume that each is a stable set. Then by pigeon-hole principle, we may find either an interfered $K^1_{q_1,h_1}$ or an interfered $K^2_{q_2,h_2}$ as an induced subgraph (and so as a vertex-minor after the possible local complementations). The result then follows by Lemma 4.2.2 and Lemma 4.2.3 respectively. □

To prove Theorem 1.3.3, by Theorem 4.2.7 it now suffices to show that if $\chi^{(2)}(G)$ is bounded and $\chi(G)$ is sufficiently huge, then G contains a large (q, h)-frame as a vertex-minor. This is what Sections 4.3–4.5 are dedicate to. Once we have Theorem 1.3.3, in Section 4.6 we then prove Theorem 1.3.1 that proper vertex-minor-closed class of graphs are χ-bounded. We finish this section with a sketch of the strategy for proving Theorem 1.3.3 and give a overview of these next three sections.

Roughly speaking, a (q, h)-frame has two halves. The "bottom" half consists connected components on vertex sets $N[z_1], \dots, N[z_h]$, where for each $1 \leq j \leq h$, the induced subgraph on vertex set $N[z_j]$ has a spanning star with centre z_j. The "top" half consists is the subgraph obtained by removing vertices $z_1, \dots, z_h$. This is a collection of stars and 1-subdivided stars with centres $w_1, \dots, w_h$ and possible additional edges between these subgraphs. It is not too difficult to find either half individually, even as induced subgraphs rather than as vertex-minors. However, in order to create our (q, h)-frame vertex-minor, we need to the find top half in such a way that it is "compatible" with the bottom half that we then find later.

Instead of aiming to find the stars or 1-subdivided stars of the top half of our (q, h)-frame directly as induced subgraphs, we will use vertex-minors. The advantage of this is that we can instead aim to find as an induced subgraph a more general structure where each star or 1-subdivided star can instead be a more general tree-like structure with many leaves. With such a induced graph, we can then use vertex-minors to "simulate" a contraction-like operation on these trees to reduce them down to 1-subdivided stars as we desired for the top half of our (q, h)-frame. So, Section 4.3 is dedicated to proving a lemma that will allow us to find these tree-like structures (see Lemma 4.3.6). Then, Section 4.4 is dedicated to proving a technical lemma that will allow us to suitably simulate this contraction-like operation on at least one of these tree-like structures (see Lemma 4.4.4).

Lastly, in Section 4.5, we put everything together. The section starts by using the fact that $\chi^{(2)}(G)$ is bounded and $\chi(G)$ is huge to find an induced subgraph that both contains what will become the bottom half of our (q, h)-frame, but also contains a suitable structure from which will be able build the compatible top half (see Lemma 4.5.3). After this we then use the results of Sections 4.3 and 4.4 to extract our tree-like structures that we need for the top half of the (q, h)-frame, and then to carefully "contract" them down one at a time. Putting it all together, this will finally complete our (q, h)-frame and prove Theorem 1.3.3 that proper vertex-minor-closed classes of graphs are (linearly) 2-controlled.

4.3 Large induced bloated trees

This section is devoted to the analysis of large induced tree-like structures. From these structures we shall later obtain the "top" half of the (q, h)-frame vertex-minor that we seek.

If T is a tree then we say that a vertex of degree at most 1 is a ***leaf*** and that a vertex of degree at least 3 is a ***branching vertex***. The degree sum of an n-vertex tree is $2n - 2$, therefore if a tree has $\ell > 1$ leaves, then it has at most $\ell - 2$ branching vertices. Likewise, a tree with $b \geq 1$ branching vertices has at least $b + 2$ leaves. We use these two facts repeatedly without reference.

We call maximal cliques (with respect to vertex inclusion) of size at least three ***big cliques***, or ***big k-cliques*** when we wish to refer to their size. We say that a graph G is a ***bloated tree*** if

- every edge is contained in at most one big clique,
- the vertices of every big clique of size $k \geq 3$ have degree at most k, and
- the graph obtained by contracting each big clique is a tree.

Note that a bloated tree is not necessarily a tree, but all trees are bloated trees. An alternative definition for bloated trees is that they are block graphs such that for each $k \geq 3$, the vertices that are contained in the big k-cliques have degree at most k. We say that a vertex of a bloated tree is a ***leaf*** if it has degree at most 1. A vertex of a bloated tree is ***branching*** if it has degree at least 3 and is not contained in a triangle.

Erdős, Saks and Sós [38] proved that for each $r \geq 3$, there exists an increasing function $t_r : \mathbb{N} \to \mathbb{N}$ such that $\lim_{n\to\infty} t_r(n) = \infty$, and every connected K_r-free graph with at least n vertices contains an induced tree on at least $t_r(n)$ vertices.

We require a version for bloated trees, and for convenience we may make this independent of the clique number. We do not attempt to optimize the bounds. Letting $f(1) = 1, f(2) = 2$, and $f(n) = \max\{r : t_r^{-1}(r) \leq n\}$ for $n \geq 3$, we see that f is increasing and $\lim_{n\to\infty} f(n) = \infty$. Since a clique is itself a bloated tree, we obtain the following version of the theorem of Erdős, Saks and Sós [38].

Theorem 4.3.1 (Erdős, Saks and Sós [38])**.** *There exists an increasing function $f : \mathbb{N} \to \mathbb{N}$ such that $\lim_{n\to\infty} f(n) = \infty$, and if G is a connected graph on at least n vertices, then it contains an induced bloated tree T on at least $f(n)$ vertices.*

We further require a suitable version of Theorem 4.3.1 in which we seek an induced bloated tree containing many vertices from a (much larger) set of distinguished vertices. First we need a lemma on cut vertices and bridges in maximal induced bloated trees of a graph. A vertex v of a connected graph G is a ***cut vertex*** if $G - v$ is disconnected. Similarly, an edge e of a connected graph G is a ***bridge*** if $G - e$ is disconnected.

Lemma 4.3.2. *Let G be a connected graph and T a maximal induced bloated tree of G. If u and v are adjacent vertices that have degree two in T, and both u and v are cut vertices in G, then uv is a bridge of G.*

Proof. As u is a cut vertex of G and has degree two in the maximal induced bloated tree T, we observe that, $G - u$ must have exactly two connected components. In particular the two vertices that u is adjacent to in T are in separate connected components of $G - u$. So u is not contained in a big clique of T. Similarly for v.

Suppose that uv is not a bridge. Then there exists an induced cycle C of G containing the edge uv. No vertex of $C - \{u, v\}$ is adjacent to any vertex of $T - \{u, v\}$ as this would contradict the fact that both $G - u$ and $G - v$ have exactly two connected components. Consider the vertex w of C which is adjacent to u and distinct from v. Then as u and v are cut vertices, we see that u and possibly v are the only vertices of T that are adjacent to w in G. This contradicts the maximality of the induced bloated tree T as we may add the vertex w. We conclude that uv is a bridge of G. □

Theorem 4.3.3. *There exists an increasing function $g : \mathbb{N} \to \mathbb{N}$ such that $\lim_{n\to\infty} g(n) = \infty$, and if G is a connected graph and S is a non-empty subset of its vertices, then G contains an induced bloated tree with at least $g(|S|)$ vertices of S.*

Proof. Let f be as in Theorem 4.3.1, and let $g(n) = \lceil \frac{1}{12} f(n) \rceil$. We will show that g satisfies the conclusion of the theorem.

Consider the graph G' obtained from G by repeating the following two operations until neither can be done.

- If v is not a vertex of S and not a cut vertex, then delete v.
- If v is not a vertex of S and v has degree 2 in G with both incident edges being bridges, then contract an edge incident to v.

Observe that any induced bloated tree of G' corresponds to an induced bloated tree of G that contains the same vertices from S. Hence we just need to find an induced bloated tree of G' that contains at least $\frac{1}{12} f(|S|)$ vertices of S. By definition of f, we can find a maximal induced bloated tree T' of G' with at least $f(|V(G')|) \geq f(|S|)$ vertices. We will show that $|V(T') \cap S| \geq \frac{1}{12}|V(T')|$.

Let:

- ℓ be equal to the number of leaves of T',
- b be equal to the number of branching vertices of T',
- x be equal to the number of big cliques X of T' having at least three vertices with a neighbour contained in $V(T') - X$,
- y be equal to the number of big cliques X of T' having exactly two vertices with a neighbour contained in $V(T') - X$, and
- z be equal to the number of big cliques X of T' having exactly one vertex with a neighbour contained in $V(T') - X$.

Every vertex v of G' that does not belong to S is a cut vertex of G'. So by the maximality of T', the set S contains every leaf of T' and every vertex contained in a big clique X of T' that has no neighbour in $V(T') - X$. If T' is a clique then $V(T') \subseteq S$, so we may assume that $\ell + z \geq 2$. So the set W of leaves, branching vertices and vertices contained in big cliques of T' must contain at least $\ell + y + 2z$ vertices of S.

For each big clique X of T', let P_X be a path on the vertices of X that have a neighbour contained in $V(T') - X$. Now let T'' be the tree obtained from T' by replacing each big clique X with the path P_X. If X is a big clique of T' that has exactly one vertex with a neighbour contained in $V(T') - X$, then this vertex is a leaf of T''. Otherwise if X is a big clique of T' that has at least two vertices with a neighbour contained in $V(T') - X$, then two of these vertices have degree 2 in T'', while the others are all branching vertices of T''. Hence T'' has $\ell + z$ leaves and therefore at most $\ell + z - 2$ branching vertices. It further follows that T' has at most $3(\ell + z - b - 2)$ vertices v contained in a big clique X of T' such that v and at least two other vertices of X have a neighbour contained in $V(T') - X$. So the number of vertices that are not contained in S, but are contained in a big clique of T' is at most $3(\ell + z - b - 2) + 2y + z$. Therefore $|W - S| \leq b + 3(\ell + z - b - 2) + 2y + z = 3\ell + 2y + 4z - 2b - 6$.

Since $\ell + y + 2z \leq |W \cap S|$, we get that $|W - S| \leq 3|W \cap S|$. Hence $|W| \leq 4|W \cap S|$.

Now let $\mathcal{P}$ be the connected components of $T' - W$. Clearly every graph contained in $\mathcal{P}$ is a path. Also by considering the tree obtained from T' by contracting each big clique, we observe that $b + x \leq \ell + z - 2$. Therefore $|\mathcal{P}| \leq \ell + b + x + y + z - 1 \leq 2\ell + y + 2z - 3$.

If $\sum_{P \in \mathcal{P}} |V(P)| \leq 4(2\ell + y + 2z - 3)$, then we have that

$$|V(T')| \leq 4|W \cap S| + 4(2\ell + y + 2z - 3) < 12|W \cap S| \leq 12|V(T') \cap S|.$$

Hence we may assume that $\sum_{P \in \mathcal{P}} |V(P)| > 4(2\ell + y + 2z - 3)$.

Consider one such path $P \in \mathcal{P}$. Suppose that P contains three consecutive vertices u, v, w that are all not contained in S. Then u, v, w would all be cut vertices of G'. By Lemma 4.3.2, both uv and vw must be bridges of T'. Hence by the maximality of T', v has degree 2 in G'. But now this contradicts the choice of G', so we may conclude that P contains no three consecutive vertices that are all not contained in S. Hence $|V(P) \cap S| \geq \left\lfloor \frac{|V(P)|}{3} \right\rfloor$.

Then summing over all $P \in \mathcal{P}$, we get that

$$\begin{aligned}
\sum_{P \in \mathcal{P}} |V(P) \cap S| &\geq \sum_{P \in \mathcal{P}} \frac{|V(P)|}{3} - (2\ell + y + 2z - 3) \\
&> \sum_{P \in \mathcal{P}} \frac{|V(P)|}{3} - \sum_{P \in \mathcal{P}} \frac{|V(P)|}{4} \\
&= \sum_{P \in \mathcal{P}} \frac{|V(P)|}{12}.
\end{aligned}$$

So it follows that

$$|V(T') \cap S| = |W \cap S| + \sum_{P \in \mathcal{P}} |V(P) \cap S| > \frac{|W|}{4} + \sum_{P \in \mathcal{P}} \frac{|V(P)|}{12} \geq \frac{|V(T')|}{12}$$

as desired. □

Next we aim to prune bloated trees with many leaves to obtain a smaller bloated tree, still with many leaves, but without near branching vertices or big cliques.

The following lemma is due to Esperet and de Joannis de Verclos [41].

Lemma 4.3.4 (Esperet and de Joannis de Verclos [41]). *Every tree T with at least ℓ leaves has a subtree which contains at least $\sqrt{\ell}$ of the leaves of T and has no adjacent branching vertices.*

Proof. We will prove a stronger statement on rooted trees. For a rooted tree T, let $f_0(T)$ be the largest number of leaves of T in a subtree of T that includes the root vertex and all its children without having adjacent branching vertices. Similarly for a rooted tree T, let $f_1(T)$ be equal to the largest number of leaves of T in a subtree of T that contains the root vertex and at most one of its children without having adjacent branching vertices.

We will prove that $f_0(T) \cdot f_1(T) \geq \ell$, which clearly implies the lemma. If T has height either 0 or 1 then the result is clear. So we may assume that T has height at least 2. We proceed by induction on the height of T. Let $T_1, \ldots, T_k$ be the subtrees obtained by taking a child of the root of T, rooting at this vertex and then taking all its descendants. Then $f_0(T) \geq \sum_{i=1}^{k} f_1(T_i)$ and $f_1(T) \geq \max\{f_0(T_i) : i \in \{1, \ldots, k\}\}$. Hence

$$f_0(T) \cdot f_1(T) \geq \sum_{i=1}^{k} (f_1(T_i) \cdot f_0(T_i)) \geq \ell.$$

as required. □

Lemma 4.3.5. *Let T be a bloated tree with ℓ leaves. Then T contains an induced bloated tree T' that has at least $\ell^{\frac{1}{4}}$ leaves and whose branching vertices and big cliques are all at distance at least 4 from every other branching vertex or big clique of T'.*

Proof. First note that if a vertex of a big clique of size k has degree $k-1$, then we may just delete the vertex. So we may assume that T has no such vertex.

Next we reduce the problem to trees. We may contract each big clique of T into a single vertex to obtain a tree T^*. Now by reversing this operation we observe that such a desired subtree of T^* corresponds to such a desired induced subgraph of T. Hence we may assume that T is a tree. We need only consider the distance between two branching vertices in T as there are no big cliques in a tree.

By Lemma 4.3.4, the tree T contains an induced subtree T_2 with at least $\sqrt{\ell}$ leaves and with no adjacent branching vertices. Now by considering the graph obtained by smoothing degree-2 vertices of T_2 and applying Lemma 4.3.4 again, we may find a subtree T' of T_2 and so of T with at least $\ell^{\frac{1}{4}}$ leaves and with no pair of branching vertices at a distance of less than 4 from each other. □

Next we combine the previous few lemmas so that we may find our desired bloated trees, this is the main result of this section.

Lemma 4.3.6. *For every positive integer ℓ, there exists a positive integer ℓ' such that every connected graph G with a set S of ℓ' distinguished vertices of degree 1 contains an induced bloated tree T with ℓ leaves, all contained in S, and whose branching vertices and big cliques are all at distance at least 4 from each other.*

Proof. By Theorem 4.3.3, there exists some positive integer ℓ' such that every connected graph G with a set S of ℓ' distinguished vertices of degree 1 contains an induced bloated tree T' with ℓ^4 leaves, all contained in S. Now by Lemma 4.3.5, there is an induced subgraph T of T' that is a bloated tree with ℓ leaves all contained in S and whose big cliques and branching vertices are all at distance at least 4 from each other as required. □

4.4 Vertex-minors and induced bloated trees

In this section we will be concerned with using vertex-minors to simulate an edge contraction-like operation on bloated trees.

The next lemma will allow us to eliminate the big cliques from these bloated trees.

Lemma 4.4.1. *Let c be a degree-k vertex of a graph G contained in a big k-clique C such that its single neighbour d that is not contained in C is adjacent to no vertex of $C\backslash\{c\}$. Then $(G - E(C - c))/cd$ is a vertex-minor of G.*

Proof. Simply consider $G * c - c$. □

A ***bloated star*** is a bloated tree consisting only of its leaves and a single big clique. We call a bloated tree T ***shrinkable*** if it is not a bloated star and the branching vertices and big cliques of T are all at distance at least 4 from each other. In a similar but much simpler manner, we say that a tree T is ***shrinking*** if it has no pair of adjacent branching vertices. The next step is to modify shrinkable bloated trees into shrinking trees.

Lemma 4.4.2. *Let T be a shrinkable bloated tree and let L be the set of leaves of T. Then there is a sequence of local complementations and vertex deletions on the vertex set $V(T) - L$ so that the resulting vertex-minor T' is a shrinking tree whose set of leaves is L.*

Proof. Firstly by appropriately removing vertices of $T - L$ we can assume that no big clique contains a vertex with no neighbour outside the big clique.

Since T is not a bloated star, every big clique has a vertex whose neighbour outside the big clique is not a leaf. Now we obtain the desired vertex-minor T' of T by applying Lemma 4.4.1 to such a vertex of each big clique. □

With this we may now simulate a contraction-like operation on shrinkable bloated trees.

Lemma 4.4.3. *Let T be a shrinkable bloated tree and let L be the set of leaves of T. Then there is a sequence of local complementations and vertex deletions on the vertex set $V(T) - L$ so that the resulting vertex-minor S is a star whose set of leaves is L.*

Proof. Firstly by Lemma 4.4.2, we may instead assume that T is a shrinking tree.

If $|V(T) - L| \leq 1$ then the result follows. Suppose for the sake of contradiction that T is a counter-example with $|V(T) - L|$ minimum, we may assume that $|V(T) - L| \geq 3$.

Suppose first that there exists a vertex v of degree 2 which has a neighbour of degree at most 2. Then smoothing v contradicts $|V(T) - L|$ being minimum.

So there must exist a vertex v of T of degree 2 with neighbours u and w such that both u and w are branching vertices of T. Then $(T \wedge uv) - u - v$ is the graph obtained by contracting the edges uv and vw. This again contradicts $|V(T) - L|$ being minimum and so completes the proof. □

The following technical lemma is the main result of this section.

Lemma 4.4.4. *For every positive integer ℓ, there exists a integer $\ell' \geq \ell$ with the following property. Let F be a graph and let $S \subseteq V(F)$ be a stable set such that $|S| \geq \ell'$ and $S \subset N_t(v)$ for some vertex $v \in V(F)$ and positive integer t. Then either:*

1. *there exists a vertex $w \in N_{t-1}(v)$ with at least ℓ neighbours in S, or*

2. *there is a subset $L \subseteq S$ with $|L| = \ell$ such that there is a sequence of local complementations on vertex set $N_{t-1}[v] \backslash N(L)$ and vertex deletions of vertices in F such that the resulting vertex-minor T of F is a 1-subdivided star whose set of leaves is L.*

Proof. Let ℓ_1' be as in Lemma 4.3.6 for $\ell + 1$. Let $\ell_2' = \ell_1'(\ell - 1)$, $\ell_3' = R(\ell_2', \ell + 1)$ (where $R(\ell_2', \ell+1)$ is the Ramsey number; any graph on at least $R(\ell_2', \ell+1)$ vertices contain either a stable of size ℓ_2', or a clique of size $\ell + 1$), and lastly let $\ell = \ell_3'(\ell - 1)$.

We may assume that every vertex of $N_{t-1}(v)$ has at most $\ell - 1$ neighbours in S, since otherwise the lemma would immediately follow. So, since $N_{t-1}(v)$ dominates S, there exist subsets $S_3 \subseteq S$ and $X_3 \subseteq N_{t-1}(v)$ with $|S_3| = |X_3| = \ell_3'$ such that each vertex of S_3 has exactly one neighbour in X_3, and each vertex of X_3 has exactly one neighbour in S_3. Note that this implies that $t \geq 2$.

Suppose that X_3 contains a clique C with $|C| = \ell + 1$. Let $u \in C$ and $L = N(C \backslash \{u\}) \cap S_3$. Then $F[L \cup C] * u$ would provide the desired vertex-minor. So we may assume that X_3 contains no clique of size $\ell + 1$. Then by Ramsey's theorem, X contains a stable set X_2 with $|X_2| = \ell_2'$. Let $S_2 = N(X_2) \cap S_3$.

Note that $N_{t-2}[v]$ is anti-complete to S_2. If there is a vertex $w \in N_{t-2}(v)$ with at least ℓ neighbours in X_2, then we would find our desired 1-subdivided star as an induced subgraph. So we may assume that no vertex of $N_{t-2}(v)$ has more than $\ell - 1$ neighbours in X_2. Since $|X_2| > \ell - 1$, this means that $t \geq 3$. As before, since $N_{t-2}(v)$ dominates X_2, there exists subsets $S_1 \subseteq S_2$, $X_1 \subseteq X_2$, and $J \subseteq N_{t-2}(v)$ with $|S_1| = |X_1| = \ell_1'$ such that each vertex of X_1 has exactly one neighbour in each of S_1 and J, and each vertex of S_1 and J has exactly one neighbour in X_1.

Now we apply Lemma 4.3.6 to the induced subgraph $F[N_{t-3}[v] \cup J \cup X_1]$ to find a $X \subseteq X_1$ with $|X| = \ell + 1$ and an induced bloated tree T of $F[N_{t-3}[v] \cup J \cup X_1]$ with leaves X and whose branching vertices and big cliques are all at distance at least 4 from each other.

If T is not a bloated star, then the result now follows from Lemma 4.4.3 (with $L = N(X \backslash \{x\}) \cap S_1$ for some $x \in X$). In the remaining case that T is a bloated star, let $C = \{c_1, \ldots, c_{\ell+1}\}$ be its big clique. Let x be the vertex of X that is adjacent to $c_{\ell+1}$, and let $L = N(X \backslash \{x\}) \cap S_1$. Then the desired vertex-minor is obtained from $F[L \cup (X \backslash \{x\}) \cup C] * c_{\ell+1}$ by smoothing each of the vertices $c_1, \ldots, c_\ell$. □

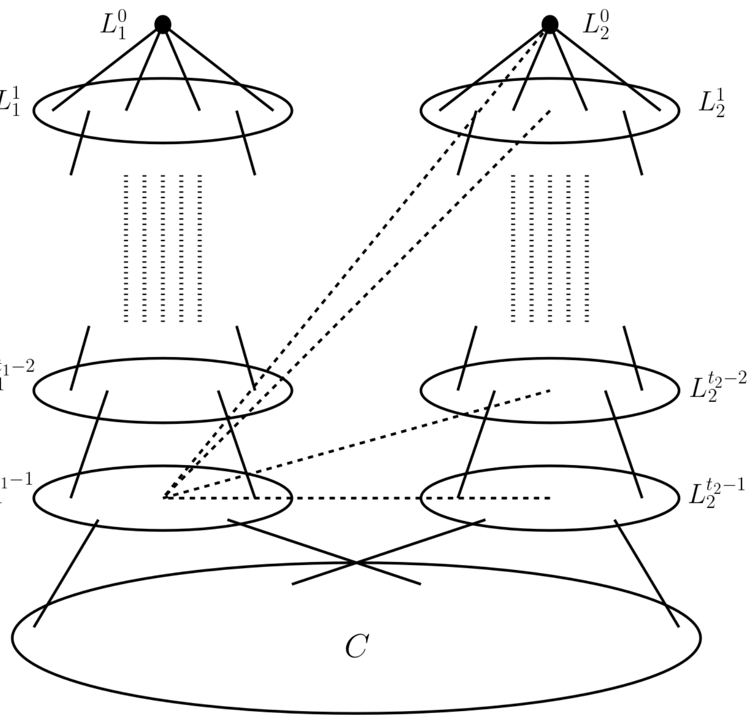

Figure 4.5: A length-2 mixed-multicover of a set C. Dashed lines indicate possible additional edges between sets of vertices.

4.5 Linear 2-control

In this section we prove Theorem 1.3.3 that proper vertex-minor-closed classes of graphs are linearly 2-controlled.

Let G be a graph. We call a collection $\mathcal{L} = (L^0, L^1, \ldots, L^{t-1})$ of vertex sets a ***t*-cover** of a set $C \subset V(G)$ if:

- the subsets $L^0, L^1, \ldots, L^{t-1}, C \subset V(G)$ are pairwise disjoint,
- $|L_0| = 1$,
- for each $i \in [t-1]$, L^{i-1} dominates L^i, and L^{t-1} dominates C,
- for each $i, j \in [t-1] \cup \{0\}$ with $|i-j| \geq 2$, L^i is anti-complete to L^j,
- for each $i \in [t-1]$, L_{i-1} is anti-complete to C, and

- the sets $L^0, L^1, \dots, L^{t-2}$ are non-empty.

The ***length*** of a t-cover $\mathcal{L} = (L^0, L^1, \dots, L^{t-1})$ is equal to t. We refer to $\mathcal{L}$ as simply a ***cover*** when the length is not important. The last condition that the sets $L^0, L^1, \dots, L^{t-2}$ are non-empty is implied by the previous conditions in the case that C or L^{t-1} is non-empty. Instead of simply requiring that C is non-empty, we use this more technical condition for the convenience of later examining induced subgraphs with vertices removed from the sets C and L^{t-1} (see for example Lemma 4.5.3 and the preceding definition of a (r, q, h)-frame).

We say that two covers $\mathcal{L}_i = (L_i^0, L_i^1, \dots, L_i^{t_i-1})$ and $\mathcal{L}_j = (L_j^0, L_j^1, \dots, L_j^{t_j-1})$ are disjoint if the two sets $\bigcup_{k=0}^{t_i-1} L_i^k$ and $\bigcup_{k=0}^{t_j-1} L_j^k$ are disjoint. We say that a collection of pairwise disjoint covers $(\mathcal{L}_i : i \in [q])$ of a set C is a ***mixed-multicover*** of a set $C \subset V(G)$ if for each $i, j \in [q]$, with $i < j$, the set of vertices $\bigcup_{k=0}^{t_i-2} L_i^k$ is anti-complete to $\bigcup_{k=0}^{t_j-1} L_j^k$ where t_i and t_j are the lengths of $\mathcal{L}_i$ and $\mathcal{L}_j$ respectively. The ***length*** of a mixed-multicover $(\mathcal{L}_i : i \in [q])$ is equal to q. See Figure 4.5 for an illustration of a mixed-multicover of length 2. If the length of every cover of a mixed-multicover is equal to t, then we say that the mixed-multicover is a ***t-multicover***.

We start by showing that for graphs with large chromatic number, we can find long mixed-multicovers of a set C with large chromatic number. This is just a typical levelling argument which has become a standard tool for tackling χ-boundedness problems. For instance, Gyárfás [56] uses essentially the same levelling technique in his proof that circle graphs are χ-bounded.

Lemma 4.5.1. *Let q, c be non-negative integers. Then every graph G satisfying $\chi(G) > 2^q c$ contains a length-q mixed-multicover $(\mathcal{L}_i : i \in [q])$ of a set $C \subseteq V(G)$, with $\chi(C) > c$.*

Proof. For $q = 0$, the result is vacuously true. We proceed inductively on q. Then by the inductive hypothesis, G contains a length-$(q-1)$ mixed-multicover $(\mathcal{L}_i : i \in [q-1])$ of a set C', with $\chi(C') > 2c$.

Let v be a vertex of $G[C']$ in a component with chromatic number greater than $2c$. For each non-negative integer i, let L_q^i be the set of vertices at distance exactly i from v in $G[C']$. Then there must exist a non-negative integer t_q such that $\chi(L^{t_q}) > c$, since otherwise we could obtain a $2c$-colouring of the component of $G[C']$ by using a set of c colours for the vertices at an odd distance from v in $G[C']$, and a second set of c colours for the vertices at an even distance from v in $G[C']$.

Then let $C = L_q^{t_q}$, and let $\mathcal{L}_q = (L_q^0, L_q^1, \dots, L_q^{t_q-1})$. Then $(\mathcal{L}_i : i \in [q])$ is a length-q mixed-multicover of the set C with $\chi(C) > c$ as desired. □

A ***tick*** of a mixed-multicover $(\mathcal{L}_i : i \in [q])$ of a set C of vertices in a graph G, is a set $Z = \{z, y_1, \dots, y_q\}$ of vertices such that

- Z is disjoint from $(\mathcal{L}_i : i \in [q])$ and C,
- $\{y_1, \dots, y_q\} \subseteq N(z)$,
- z is anti-complete to C and the vertices of $(\mathcal{L}_i : i \in [q])$,
- for each $i \in [q]$, y_i has at least one neighbour in $L_i^{t_i-2}$ (where t_i is the length of the cover $\mathcal{L}_i$), and
- for each $i \in [q]$, y_i is anti-complete to C, and all neighbours of y_i contained in the mixed-multicover $(\mathcal{L}_i : i \in [q])$ are contained in $L_i^{t_i-2} \bigcup_{k=i+1}^{q} \left(\bigcup_{j=0}^{t_k-1} L_k^j \right)$.

Ticks will form the "bottom" half of the (q,h)-frame vertex-minor that we seek. So we need to be able to find many disjoint and anti-complete ticks.

Lemma 4.5.2. *Let q, h, c, κ be non-negative integers. Then every graph G satisfying $\chi^{(2)}(G) \le \kappa$, and $\chi(G) > 2^q(c + (q+1)h\kappa)$ contains a length-q mixed-multicover $(\mathcal{L}_i : i \in [q])$ of a set $C \subseteq V(G)$ with $\chi(C) > c$, such that $(\mathcal{L}_i : i \in [q])$ has h disjoint and anti-complete ticks $Z_1, \dots, Z_h$.*

Proof. For $h = 0$, the result follows by Lemma 4.5.1. So we proceed inductively on h. By the inductive hypothesis, G contains a length-q mixed-multicover $(\mathcal{L}_i : i \in [q])$ of a set C', with $\chi(C') > c + (q+1)\kappa$, and such that $(\mathcal{L}_i : i \in [q])$ has $h - 1$ disjoint and anti-complete ticks $Z_1, \dots, Z_{h-1}$.

Let z be a vertex of C'. For each $i \in [q]$, let $y_{i,h}$ be a vertex of $L_i^{t_i-1}$ that is adjacent to z (where t_i is the length of the cover $\mathcal{L}_i$). Let $Z_h = \{z, y_1, \dots, y_q\}$, and let $C = C' \backslash N_2[Z]$. Since $\chi^{(2)}(G) \le \kappa$, we have that $\chi(C) \ge \chi(C') - |Z|\kappa = \chi(C') - (q+1)\kappa > c$. Note that Z_h is disjoint and anti-complete to $Z_1, \dots, Z_{h-1}$, so it just remains to remove vertices of the mixed-multicover $(\mathcal{L}_i : i \in [q])$ of C, so that Z_h becomes a tick. To do so, we simply remove the vertices of $\bigcup_{i=1}^{q} L_i^{t_i-1}$ that have a neighbour in Z_h. $\square$

For non-negative integers r, q, h where $r \le q$, we say that a graph G is a (r, q, h)-frame if:

- G contains a length-q mixed-multicover $(\mathcal{L}_i : i \in [q])$ (of an empty set) with h disjoint and anti-complete ticks $Z_1 = \{z_1, y_{1,1}, \dots, z_{q,1}\}, \dots, Z_h = \{z_h, y_{1,h}, \dots, z_{q,h}\}$, and G contains no other vertices,

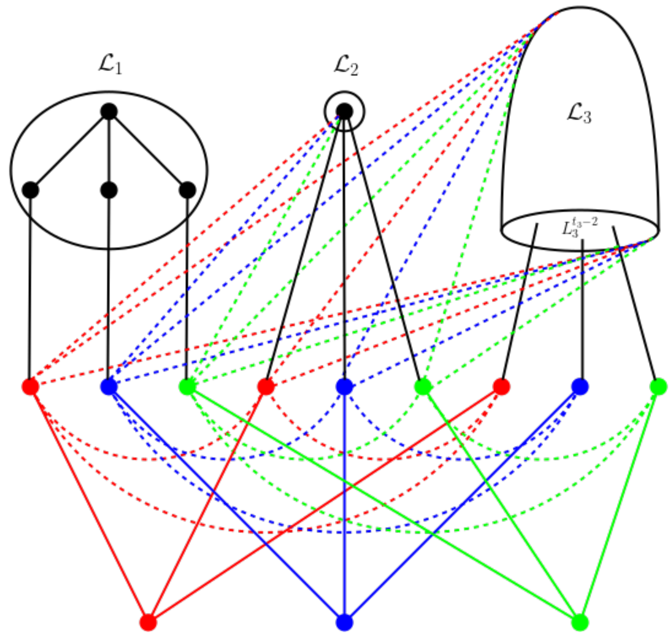

Figure 4.6: An example of a $(2, 3, 3)$-frame. Dashed lines indicate possible extra edges. The three ticks (and some possible additional edges incident to their vertices) are each coloured differently. The cover $\mathcal{L}_1$ is as in case 2, and $\mathcal{L}_2$ is as in case 1.

- for each $i \in [r]$, $L_i^{t_i-1}$ is empty (where t_i is the length of the cover $\mathcal{L}_i$), and
- for each $i \in [r]$, either
  1. the cover $\mathcal{L}_i$ has length $t_i = 2$, with $L_i^0 = \{w_i\}$, or
  2. the cover $\mathcal{L}_i$ has length $t_i = 3$, and $G[L_i^0 \cup L_i^1]$ is a star with centre w_i, and leaves $\{x_{i,1}, \ldots, x_{i,h}\}$, such that $N(x_{i,j}) \cap \{y_{i,1}, \ldots, y_{i,h}\} = \{y_{i,j}\}$ for each $j \in [h]$.

See Figure 4.6 for an example of a $(2, 3, 3)$-frame. In particular, notice that a (q, q, h)-frame is exactly a (q, h)-frame (as in Theorem 4.2.7, compare Figure 4.6 to Figure 4.4). By Lemma 4.5.2, we have the following.

Lemma 4.5.3. *Let q, h be non-negative integers. Then every graph G satisfying $\chi^{(2)}(G) \leq \kappa$, and $\chi(G) > 2^q(q + 1)h\kappa$ contains a $(0, q, h)$-frame as an induced subgraph.*

Before proceeding, we need a quick notational definition. For two sets A, B, we let $A \Delta B = (A \cup B) \backslash (A \cap B)$. For a graph G and a subgraph H, we let $G \Delta H$ denote the graph obtained from G by replacing the edges of the induced subgraph on vertex set $V(H)$, with the edge set $E(G[V(H)])\Delta E(H)$.

Next we aim to show that if q', h' are sufficiently large, then every $(0, q', h')$-frame contains a (q, h)-frame as a vertex-minor. We use the following lemma to facilitate the inductive argument.

Lemma 4.5.4. *Let r, q, h be non-negative integers with $r < q$. Then there exist positive integers q', h' such that every (r, q', h')-frame contains a $(r+1, q, h)$-frame as a vertex-minor.*

Proof. Let $\ell_{2^{\binom{rh}{2}}} = h$, and for each $i \in [2^{\binom{rh}{2}}]$ in reverse order, let ℓ_{i-1} be as in the conclusion of Lemma 4.4.4 for ℓ_i. Let $q' = q + 2^{\binom{rh}{2}} - 1$ and $h' = \ell_0$. Now let G be a (r, q', h')-frame. As in the definition of a (r, q', h')-frame, let $(\mathcal{L}_i : i \in [q'])$ be the length-q' mixed-multicover, and let $Z_1, \ldots, Z_{h'}$ be its ticks. We label the vertices of $(\mathcal{L}_i : i \in [r])$ and $Z_1, \ldots, Z_{h'}$ as in the above definition of a (r, q', h')-frame. For each $i \in [q']$, let t_i be the length of the cover $\mathcal{L}_i$. For every $Q \subseteq [q']$ and $H \subseteq [h']$, let $Y_{Q,H} = \{y_{i,j} : i \in Q, j \in H\}$.

If there exists an integer $i \in [2^{\binom{rh}{2}}]$ such that $L_{r+i}^{t_i-2}$ contains a vertex w_{r+1} with at least h neighbours in $\{y_{i,1}, \ldots, y_{i,h'}\}$, then we would obtain a $(r+1, q, h)$-frame as an induced subgraph. So we may assume that no such integer i exists.

For each $i \in [2^{\binom{rh}{2}}]$, let $F_i = G\left[Y_{\{r+i\},[h']} \bigcup_{k=0}^{t_{r+i}-2} L_{r+i}^k\right]$. Note that Lemma 4.4.4 can be applied to F_i since for each $0 \leq j \leq t_{r+i} - 1$, L_{r+i}^j is the set of vertices that are at distance exactly j in F_i from (the single vertex of) L_{r+i}^0. By applying Lemma 4.4.4 to each of the induced subgraphs $F_1, \ldots, F_{2^{\binom{rh}{2}}}$ in order, there exist subsets $H^* = H_{2^{\binom{rh}{2}}} \subseteq \cdots \subseteq H_1 \subseteq [h']$ such that for each $i \in [2^{\binom{rh}{2}}]$, we have that

- $|H_i| = \ell_{i-1}$, and
- in the induced subgraph F_i, there is a sequence of local complementations on the vertex set $\left(\bigcup_{k=0}^{t_{r+i}-2} L_{r+i}^k\right) \setminus N(Y_{\{r+i\},H_i})$ and vertex deletions of vertices in F_i, so that the resulting vertex-minor of F_i is a 1-subdivided star whose set of leaves is $Y_{\{r+i\},H_i}$.

We can focus solely on H^*. By further removing extra vertices as necessary, for each $i \in [2^{\binom{rh}{2}}]$, in the induced subgraph F_i, there is a sequence of local complementations on the vertex set $C_i = \left(\bigcup_{k=0}^{t_{r+i}-2} L_{r+i}^k\right) \setminus N(Y_{\{r+i\},H^*})$ and vertex deletions of vertices in F_i, so that the resulting vertex-minor T_i of F_i is a 1-subdivided star whose set of leaves is $Y_{\{r+i\},H^*}$.

For some $i \in [2^{\binom{rh}{2}}]$, we wish to carry out the sequence of local complementations and vertex deletions given above in such a way that afterwards we can then obtain a $(r+1, q, h)$-frame as an induced subgraph (so as a vertex-minor after the local complementions). However performing these local complementations may unintentionally alter the rest of G in a problematic way. We need that our (now smaller and fewer) ticks remain anti-complete.

Suppose we carry out this sequence of local complementations and vertex deletions for some $i \in [2^{\binom{rh}{2}}]$ to obtain a vertex-minor G_i of G. Let w_{r+i} be the branching vertex of the resulting 1-subdivided star T_i. For each $j \in H_i$, let $x_{r+i,j}$ be the vertex of the 1-subdivided star neighbouring w_{r+i} and $y_{r+i,j}$. Let $X_{r+i} = \{x_{r+i,j} : j \in H^*\}$, and note that $\mathcal{L}_i^* = (\{w_{r+i}\}, X_{r+i}, Y_{\{r+i\},H^*})$ would be a 3-cover (of an empty set). Now, since the neighbourhood (in G) of C_i outside $V(F_i)$ is contained in $Y_{[r+i-1],[h']}$, and no vertex of C_i neighbours a vertex of $Y_{\{r+i\},H^*}$, the induced subgraphs $G\backslash(V(F_i) \cup Y_{[r+i],[h']\backslash H^*})$ and $G_i\backslash(V(F_i) \cup Y_{[r+i],[h']\backslash H^*})$ of G may differ only on the vertex set $Y_{[r+i-1],H^*}$. Furthermore, if an edge $e \in (E(G_i)\backslash E(G))\backslash E(T_i)$ of G_i has at most one endpoint in $Y_{[r+i-1],[h']}$, then it must have one endpoint in $Y_{[r+i-1],[h']}$, and one endpoint in $\{w_{r+i}\} \cup X_{r+i} = V(T_i)\backslash Y_{\{r+i\},H^*}$. Therefore, if $G_i[Y_{[r],H^*}] = G[Y_{[r],H^*}]$, then we could obtain a $(r+1, q, h)$-frame as an induced subgraph of G_i (and thus as a vertex-minor of G). So we may assume otherwise. Let $D_i = G[Y_{[r],H^*}] \Delta G_i[Y_{[r],H^*}]$. We may assume that D_i is not edge-less, since otherwise G would contain a $(r+1, q, h)$-frame as a vertex-minor as discussed.

Note that $|Y_{[r],H^*}| = rh$. So by pigeonhole principle, there exists a pair $i, i^* \in [2^{\binom{rh}{2}}]$ with $i < i^*$ such that $D_i = D_{i^*}$. Let G^* be the vertex-minor of G obtained by performing the sequence of local complementations and vertex deletions used to obtain G_i, and then afterwords performing the sequence of local complementations and vertex deletions that were used to obtain G_{i^*}. Then observe that $G^*\backslash \left(\bigcup_{k=0}^{t_{r+i}-2} L_{r+i}^k\right) = \left(G_{i^*}\backslash \left(\bigcup_{k=0}^{t_{r+i}-2} L_{r+i}^k\right)\right) \Delta D_i$. In particular, $G^*[Y_{[r],H^*}] = G[Y_{[r],H^*}] \Delta D_{i^*} \Delta D_i = G[Y_{[r],H^*}]$ since $D_{i^*} = D_i$. So then G^* contains an induced $(r+1, q, h)$-frame, and thus G contains a $(r+1, q, h)$-frame as a vertex-minor as desired. □

By repeatedly applying Lemma 4.5.4, and then applying Lemma 4.5.3, we obtain the following.

Lemma 4.5.5. *For every pair q, h of non-negative integers, there exists a positive integer c such that every graph G satisfying $\chi^{(2)}(G) \leq \kappa$, and $\chi(G) > c\kappa$ contains a (q, h)-frame as a vertex-minor.*

Theorem 1.3.3 now follows immediately from Lemma 4.5.5 and Theorem 4.2.7.

4.6 Vertex-minor χ-boundedness

In this section we prove that if a vertex-minor-closed class of graphs $\mathcal{G}$ is 2-controlled then $\mathcal{G}$ is χ-bounded. As all proper vertex-minor-closed classes of graphs are 2-controlled by Theorem 1.3.3, this implies the main result of this chapter, Theorem 1.3.1, that proper vertex-minor-closed classes of graphs are χ-bounded. Let us remark that this happens to be the simpler part of the ρ-control strategy for vertex-minors and that this does not rely on any results from Sections 4.3–4.5 (although some of the arguments in this section are similar to parts of Section 4.5).

Instead of using mixed-multicovers as in the proof of Theorem 1.3.3, in this section we are now able to find and use 2-multicovers, since we can assume that our class of graphs is 2-controlled. We find it convenient to reintroduce 2-multicovers in a slightly different (but equivalent) way.

Let G be a graph, and $X, C \subseteq V(G)$ such that X has a total ordering $\preceq$. For each $x \in X$, let $N_x \subseteq N(x)$. We say that $(N_x : x \in X)$ is a ***2-multicover***[1] of C if:

- the sets $X, C, (N_x : x \in X)$ are disjoint,
- the set X is stable,
- the set X is anti-complete to C,
- for each $x \in X$, N_x dominates C, and
- for distinct $x, y \in X$, with $x \prec y$, the vertex x is anti-complete to N_y.

The ***length*** of a 2-multicover is equal to $|X|$. A 2-multicover is ***stable*** if for each $x \in X$, the set N_x is stable. A 2-multicover $(N'_x : x \in X')$ of a set C' is ***contained*** in a 2-multicover $(N_x : x \in X)$ of a set C if $C' \subseteq C$, $X' \subseteq X$, and $N'_x \subseteq N_x$ for each $x \in X'$.

We begin now by finding a long stable 2-multicover of a set with large chromatic number. This is very similar to a standard levelling argument.

Lemma 4.6.1. *Let c, ℓ, τ, ω be non-negative integers and let $\mathcal{G}$ be a 2-controlled class of graphs that is closed under taking induced subgraphs such that $\chi(G) \leq \tau$ for all $G \in \mathcal{G}$*

[1] Our definition of a 2-multicover is the same as that of [21]. What is called a multicover in [20] is slightly different to what we call a 2-multicover since we need not require that x is anti-complete to N_y when $x \succ y$.

with $\omega(G) < \omega$. Then there exists a positive integer c' such that every graph $G \in \mathcal{G}$ with $\chi(G) \geq c'$ and $\omega(G) \leq \omega$ contains a length-ℓ stable 2-multicover of a set $C \subseteq V(G)$ with $\chi(C) \geq c$.

Proof. We fix c, τ, ω. The result is trivial for $\ell = 0$, so we proceed inductively assuming the result holds for $\ell - 1$. Let c'_0 be an integer such that every graph $G \in \mathcal{G}$ with $\chi(G) \geq c'_0$ and $\omega(G) \leq \omega$ contains a length-$(\ell - 1)$ stable 2-multicover of a set C' with chromatic number at least c. Now let c' be such that every graph $G \in \mathcal{G}$ with $\chi(G) \geq c'$ contains a vertex y with $\chi(N_2[y]) \geq \tau + \tau c'_0$. It remains to show that c' satisfies the conclusion of the lemma.

Let G be a graph in $\mathcal{G}$ with $\chi(G) \geq c'$ and $\omega(G) \leq \omega$. Let y be a vertex of G such that $G[N_2[y]]$ has chromatic number at least $\tau + \tau c'_0$. Then $\chi(N(y)) \leq \tau$, so $\chi(N_2(y)) \geq \tau c'_0$. Let N_y be a stable subset of $N(y)$ such that $\chi(N(N_y) \cap N_2(y)) \geq c'_0$. So by the inductive hypothesis, there exists a set C contained in $N(N_y) \cap N_2(y)$ with $\chi(C) \geq c$ and a stable 2-multicover $(N_{x'} : x' \in X')$ of length $\ell - 1$ in $G[N(N_y) \cap N_2(y)]$ of the set C.

Let $X = X' \cup \{y\}$ and let $\preceq$ be the total ordering on X where y is the first vertex and the restriction to X' is $\preceq'$. Then $(N_x : x \in X)$ provides the desired stable 2-multicover of C. $\square$

Next we define a tick of a 2-multicover, the definition is essentially the same as that of a mixed-multicover. Let G be a graph containing a 2-multicover $(N_x : x \in X)$ of a set $C \subseteq V(G)$. A ***tick*** of $(N_x : x \in X)$, C is a set $Z = \{z\} \cup \{y_x : x \in X\}$ of vertices such that

- Z is disjoint from $C \cup X \bigcup_{x \in X} N_x$,
- $\{y_x : x \in X\} \subseteq N(z)$,
- z is anti-complete to $C \cup X \bigcup_{x \in X} N_x$,
- for each $x \in X$, y_x is adjacent to x, and
- for each $x \in X$, y_x is anti-complete to $C \cup \{x' \in X : x' \prec x\} \bigcup_{x \in X} N_x$.

Note that if a graph G contains a length-q 2-multicover (of a possibly empty set) with disjoint and anti-complete ticks $Z_1, \ldots, Z_h$, then G contains an interfered (q, h)-frame as an induced subgraph.

The proof of the next lemma follows that of a lemma of Chudnovsky, Scott and Seymour [20, 2.1]. Indeed, despite our different definition of a (2-)multicover, the proof is essentially the same.

Lemma 4.6.2. *For all positive integers c, ℓ, τ, j, ω such that $j \leq \omega$, there exists a pair of positive integers c_j, ℓ_j with the following property. Let G be a graph with $\omega(G) \leq \omega$, such that $\chi(H) \leq \tau$ for every induced subgraph H of G with $\omega(H) < \omega$. Let $(N_x : x \in X)$ be a length-ℓ_j stable 2-multicover in G of a set C such that $\chi(C) \geq c_j$ and $\omega\left(\bigcup_{x\in X} N_x\right) \leq j$. Then $(N_x : x \in X)$ contains a length-ℓ_j stable 2-multicover $(N'_x : x \in X')$ of a set $C' \subseteq C$ with $\chi(C') \geq c$, such that $(N'_x : x \in X')$ has a tick Z.*

Proof. Let $\ell_0 = 1$, and for each $1 \leq i \leq \omega$ in order, let $\ell_i = \omega \ell_{i-1}^2 \ell$. Let $c_0 = c$, and for each $1 \leq i \leq \omega$ in order, let $c_i = (c_{i-1} 4^{\binom{\ell_{i-1}^2 \ell}{2}} + \ell_{i-1}^2 \ell \tau)\omega\binom{\ell_i}{\ell_{i-1}^2 \ell} + \omega(1+\tau)$.

Let $A \subseteq C$ be a clique with $|A| = \omega$, such a clique exists since $\chi(C) > \tau$. Let $C_0 = C \backslash N[A]$. Then $\chi(C_0) \geq \chi(C) - |A|(1+\tau) = \chi(C) - \omega(1+\tau) \geq (c_j 4^{\binom{\ell_{j-1}^2 \ell}{2}} + \ell_{j-1}^2 \ell \tau)\omega\binom{\ell_j}{\ell_{j-1}^2 \ell}$. For each $x \in X$ and $v \in C_0$, let $n_{x,v}$ be a vertex of N_x that is adjacent to v. Since $\omega(G) = \omega = |A|$, for each $x \in X$ and $v \in C_0$, there is a vertex $a_{x,v} \in A$ that is non-adjacent to $n_{x,v}$. So for each $v \in C_0$, there exists $X_v \subseteq X$ with $|X_v| = |X|/|A| = |X|/\omega = \ell_{j-1}^2 \ell$ and some $a_v \in A$, such that $n_{x,v}$ is non-adjacent to a_v for every $x \in X_v$. For every $a \in A$ and $Y \subseteq X$ with $|Y| = \ell_{j-1}^2 \ell$, let $C_{a,Y}$ be the set of vertices $v \in C_0$ such that $a_v = a$ and $X_v = Y$. This defines a partition of C_0, so there exists a $z \in A$ and $X_1 \subseteq X$ with $|X_1| = \ell_{j-1}^2 \ell$ such that $\chi(C_{z,X_1}) \geq \chi(C_0)/\omega\binom{\ell_j}{\ell_{j-1}^2 \ell} \geq c_j 4^{\binom{\ell_{j-1}^2 \ell}{2}} + \ell_{j-1}^2 \ell \tau$. Let $C_1 = C_{z,X_1}$.

For each $x \in X_1$, let $y_x \in N_x$ be a vertex adjacent to z, and let $N_x^* = N_x \backslash \{y_x\}$. Let $C_2 = C_1 \backslash N(\{y_x : x \in X_1\})$, then $\chi(C_2) \geq \chi(C_1) - |X_1|\tau \geq \chi(C_1) - \ell_{j-1}^2 \ell \tau \geq c_j 4^{\binom{\ell_{j-1}^2 \ell}{2}}$. Note that $(N_x^* : x \in X_1)$ is a 2-multicover of C_2 and that $(N_x^* : x \in X')$ is contained in $(N_x : x \in X)$.

Suppose now that $j = 1$. Then let $X' = X_1$, and let $C' = C_2$. Clearly $|X'| = |X_1| = \ell_0^2 \ell = \ell$ and $\chi(C') = \chi(C_2) \geq c_0 4^{\binom{\ell_0^2 \ell}{2}} = c4^{\binom{\ell}{2}} \geq c$. Then $(N_x^* : x \in X')$ is a 2-multicover of C', and $Z = \{z\} \cup \{y_x : x \in X'\}$ provides the desired tick of $(N_x^* : x \in X')$. So we may now assume that $j > 1$, and we proceed inductively.

For each $v \in C_2$, let $f_1(v)$ be the auxiliary graph on vertex set X_1 such that for each pair of distinct vertices $x, x' \in X_1$ with $x \prec x'$, x is adjacent to x' in $f_1(v)$ if y_x is adjacent to $n_{x',v}$. Similarly, for each $v \in C_2$, let $f_2(v)$ be the auxiliary graph on vertex set X_1 such that for each pair of distinct vertices $x, x' \in X_1$ with $x \prec x'$, x is adjacent to x' in $f_2(v)$ if $y_{x'}$ is adjacent to $n_{x,v}$. Then there exists some pair of graphs H_1 and H_2 on the vertex set X_1 such that $f_1^{-1}(H_1) \cap f_2^{-1}(H_2) \subseteq C_2$ has chromatic number at least $\chi(C_2)/4^{\binom{|X_1|}{2}} = \chi(C_2)/4^{\binom{\ell_{j-1}^2 \ell}{2}} \geq c_{j-1}$. Let $C_3 = f_1^{-1}(H_1) \cap f_2^{-1}(H_2)$.

Suppose that H_1 contains a vertex $x \in X_1$ such that x has at least ℓ_{j-1} neighbours x' in H_1 with $x \prec x'$. Let X' be the set of neighbours x' of x in H with $x \prec x'$. Then $|X'| \geq \ell_{j-1}$ and in G, y_x is complete to $N^*_{x'} = \{n_{x',v} : v \in C_3\}$ for each $x' \in X'$. Now $\omega(\bigcup_{x' \in X'} N^*_{x'}) < j$, so by the inductive hypothesis, $(N^*_x : x \in X)$ (and thus $(N_x : x \in X)$) contains a length-ℓ stable 2-multicover $(N'_x : x \in X')$ of a set $C \subseteq C_3 \subseteq C'$ with $\chi(C) \geq c$, such that $(N'_x : x \in X')$ has a tick Z. Hence we may assume that H_1 contains no vertex $x \in X_1$ such that x has at least ℓ_{j-1} neighbours x' in H with $x \prec x'$. So H_1 is $(\ell_{j-1} - 1)$-degenerate, and therefore ℓ_{j-1}-colourable. Similarly, we may assume that H_2 is ℓ_{j-1}-colourable.

Let H be the graph with vertex set X_1 and edge set $E(H_1) \cup E(H_2)$. Since H_1 and H_2 are ℓ_{j-1}-colourable, H is ℓ_{j-1}^2-colourable. So there exists a stable set $X' \subseteq X_1$ of H with $|X'| = |X_1|/\ell_{j-1}^2 = \ell$. Then $Z = \{z\} \cup \{y_x : x \in X'\}$ is a tick of the 2-multicover $(N^*_x : x \in X')$ of $C = C_3$ as required. □

As discussed before Lemma 4.6.2, if a graph G contains a length-q 2-multicover (of a possibly empty set) with disjoint and anti-complete ticks $Z_1, \ldots, Z_h$, then G contains a (q, h)-frame as an induced subgraph. So by applying Lemma 4.6.2 a total of h times and Lemma 4.6.1 once, we obtain the following.

Lemma 4.6.3. *Let q, h, τ, ω be non-negative integers and let $\mathcal{G}$ be a 2-controlled class of graphs that is closed under taking induced subgraphs such that $\chi(G) \leq \tau$ for all $G \in \mathcal{G}$ with $\omega(G) < \omega$. Then there exists a positive integer c such that every graph $G \in \mathcal{G}$ with $\chi(G) > c$ and $\omega(G) \leq \omega$ contains a (q, h)-frame as an induced subgraph.*

With Lemma 4.6.3 and Theorem 4.2.7 in hand, we can now do a simple induction to show that 2-controlled vertex-minor-closed classes of graphs are χ-bounded.

Lemma 4.6.4. *Every vertex-minor-closed class of graphs that is 2-controlled is also χ-bounded.*

Proof. Let $\mathcal{G}$ be a 2-controlled vertex-minor-closed class of graphs and suppose for the sake of contradiction that $\mathcal{G}$ is not χ-bounded. Then there exists a minimum integer $\omega \geq 2$ such that the graphs $G \in \mathcal{G}$ with $\omega(G) \leq \omega$ have unbounded chromatic number. Let $\tau \geq 0$ be such that $\chi(H) \leq \tau$ for all $H \in \mathcal{G}$ with $\omega(H) < \omega$. Let J be some graph not contained in $\mathcal{G}$ (J exists as the class of all graphs is not 2-controlled). Then let q, h be as in Theorem 4.2.7, for $n = |V(J)|$. Let c be as in Lemma 4.6.3 for q, h, τ, ω. Let $G \in \mathcal{G}$ be a graph with $\chi(G) > c$ and $\omega(G) \leq \omega$. Then by Theorem 4.2.7 and Lemma 4.6.3, G contains J as a vertex-minor, a contradiction. □

Now the main result of this chapter, Theorem 1.3.1, (which states that proper vertex-minor-closed classes of graphs are χ-bounded) follows immediately from Theorem 1.3.3 and Lemma 4.6.4. While we have not kept track of the cumbersome bounds obtainable from the proof of Theorem 1.3.1, we remark that for a fixed vertex-minor-closed class of graphs, the resulting χ-bounding function would be a triply exponential function.

It is possible to adapt the proof of Theorem 1.3.1 to not require the notion of 2-control. Note that Lemma 4.5.2 is essentially the only time we used the fact that $\chi^{(2)}(G) \leq \kappa$. Instead we could find these ticks by arguing inductively on the clique number in a very similar manner to that of the proof of Lemma 4.6.2. The only difference is that the sets $L_1^{t_1}, \ldots, L_{q'}^{t_{q'}}$ need not be stable. The consequence of this if we follow the same proof is that the resulting ticks are not necessarily anti-complete; in the end, there could be edges between the vertices of $\{y_{i,1}, \ldots, y_{i,h}\}$ for each $i \in [q]$. However this can then be cleaned up at the end by repeatedly applying Ramsey's theorem. So in the end, the proof ends up being very similar. Doing so would also come at the loss of not obtaining Theorem 1.3.3 along the way. Thus there is not so much to gain from proving Theorem 1.3.1 directly in this way.

4.7 Pivot-minors

In this section we further discuss pivot-minors and the recent proof of the conjecture of Choi, Kwon and Oum [18] that pivot-minor-closed classes are χ-bounded (Theorem 1.3.2). In particular, we will discuss some of the difference in the proofs of Theorem 1.3.1 and Theorem 1.3.2. The full proof of Theorem 1.3.2 will appear in an upcoming paper [24].

As briefly discussed in Section 1.3, one viable strategy to proving Theorem 1.3.1 could have been to first prove a general structure theorem for vertex-minor-closed classes of graphs. There has been significant progress towards such a structure theorem [51, 77] and we expect that it will also yield a proof of the polynomial χ-boundedness conjecture [65] for vertex-minors.

Roughly speaking, it is conjectured that the graphs in any proper vertex-minor-closed class of graphs are obtained by piecing (slightly perturbed) circle graphs together in a prescribed way [77]. One might hope to prove a similar structural characterization for pivot-minor-closed classes, from which χ-boundedness would presumably follow. However, even compared to vertex-minors, this is a much more intimidating prospect. Pivot-minors generalize vertex-minors, and so circle graphs also provide a fundamental class of graphs which would be a building block in such a structure theorem. However pivot-minors have at least two additional fundamental classes, indicating a far more complex structure. These are bipartite graphs and line graphs.

Via their fundamental graphs, minors of binary matroids are essentially captured by pivot-minors of bipartite graphs [9, 83]. So a structure theorem even just for bipartite graphs without a bipartite pivot-minor H would already imply a structure theorem for binary matroids as in the matroid minors project of Geelen, Gerards, and Whittle [50] (and thus also the graph minor structure theorem of Robertson and Seymour [91]). Although line graphs are not closed under pivot-minors, surprisingly their closure under pivot-minors is not the class of all graphs [82] (unlike for vertex-minors). Oum [82] characterized these graphs as being exactly the fundamental graphs of graphic delta-matroids. Not much is known about graphic delta-matroids, but a structure theorem for pivot-minors would have to include as a special case a graphic delta-matroids minor structure theorem. On top of all this, a general structure theorem for pivot-minors would then need to simultaneously capture the structure of graphic delta-matroids, binary matroids and vertex-minors. This is indeed an intimating task.

Bipartite graphs are trivially χ-bounded and circle graphs [56] and graphs of bounded rank-width [34] are both of course χ-bounded. Using Vizing's theorem [106] and the graphical delta-matroid characterization of Oum [82], it is possible to show that the class of pivot-minors of line graphs is χ-bounded with a χ-bounding function of $2\omega + 1$. So all four of these fundamental classes of graphs for pivot-minors are polynomially χ-bounded.

Of course one can also ask if proper pivot-minor-closed classes are polynomially χ-bounded, and this is conjectured by Kim and Oum [64]. Although results for the fundamental pivot-minor-closed classes provide significant support for this conjecture, we see no reasonable approach to proving polynomial χ-boundedness without the use of a structure theorem for pivot-minors. The starting point towards a structure theorem for pivot-minors would be Oum's [82] conjectured grid theorem, which states that for every bipartite circle graph H, the graphs containing no H pivot-minor have bounded rank-width. This conjecture would generalize the grid theorem for vertex-minors [51] and binary matroids [49] (and thus graph minors [90]). By the theorem of Bonamy and Pilipczuk [6] that classes with bounded rank-width are polynomially χ-bounded, Oum's [82] conjecture would also imply polynomial χ-boundedness for proper pivot-minor-closed classes not containing all bipartite circle graphs.

Let us now turn our attention to Theorem 1.3.2 and its proof resolving the conjecture of Choi, Kwon and Oum [18] that pivot-minor-closed classes are χ-bounded.

As with Theorem 1.3.1, we use the strategy set out by the notion of ρ-control to prove Theorem 1.3.2. Again the most challenging step turns out to be proving that proper pivot-minor-closed classes of graphs are ρ-controlled for some $\rho \geq 2$. Unlike with Theorem 1.3.1, we are unable to prove 2-control directly for pivot-minors. Instead, what turns out to be most convenient is to begin by proving that proper pivot-minor-closed classes of graphs are 19-controlled.

For vertex-minors we had the nice beginning class of universal graphs of K_n^1, and more conveniently $K_{n,m}^1$. These graphs are no longer universal for pivot-minors since they are contained in the class of bipartite graphs. Unlike with vertex-minors, we cannot use pivot-minors to smooth degree-2 vertices in a graph, although if u and v are adjacent degree-2 vertices, then pivoting on uv and then deleting u and v is the same as smoothing both u and v at once. So the parity of paths matter for pivot-minors. The simplest universal graphs we can use (and which we aim for) is K_n^2, and more generally any graph obtainable by subdividing each edge of K_n at least twice, and an even number of times. Such graphs are called ***odd K_n subdivisions***.

A 2-multicover $(N_x : x \in X)$ of a set $C \subseteq V(G)$ is said to be ***strongly independent*** if for every distinct $x, y \in X$, x is anti-complete to N_y. Scott and Seymour [96] showed that under the additional assumptions in Lemma 4.6.2, one can find either K_n^6 or K_n^8 as an induced subgraph of a large enough strongly independent 2-multicover of a set C with $\chi(C)$ sufficiently large. Essentially the idea is to begin by finding n ticks as in Lemma 4.6.2. The centres of the ticks will form the branching vertices of the odd K_n subdivision. To build up the odd paths between the branching vertices, we need to find and extract from the 2-multicover odd length paths between distinct $x, y \in X$. Again, by arguing inductively on $\omega\left(\bigcup_{x \in X} N_x\right)$ as in Lemma 4.6.2, this can be done by showing that in a maximal clique $B \subseteq C$, there are two vertices $u, v \in B$ such that there is a vertex $a \in N_x$ adjacent to u but not v, and a vertex $b \in N_y$ adjacent to v but not u. The length of the odd path between the centres of two ticks then just depends on whether a and b are adjacent or not. This covers the case that we can find a large strongly independent 2-multicover. However if we begin with a (much larger) 2-multicover, then we can clean out the extra edges with pivoting to obtain a (smaller) strongly independent 2-multicover. The proof of this is similar to that of Lemma 4.6.2 and can be found in [25].

The step for reducing 19-control down to 2-control proceeds similarly to the analogous step for induced subdivisions proven by Chudnovsky, Scott, and Seymour [21]. Instead of starting with a 2-multicover, we start with a ρ-multicover of C. We begin similarly as in the previous step. We start by finding and extracting (something close to) "ticks" of the ρ-multicover. The centres of the ticks will again form the vertices of (something close to) an odd K_n subdivision. Just as before, we then need to find odd paths between the centres of these ticks. For this we can use the same trick of considering a maximum clique in C. Actually it is a bit easier than finding odd paths, with pivot-minors it is okay if there are certain extra edges between vertices of each individual odd path.

Now let us return to the most challenging step that proper pivot-minor-closed classes of graphs are 19-controlled. A key difference here to the proof of Theorem 1.3.3 is that Lemma 4.4.3 does not hold since we cannot use pivot-minors to eliminate big cliques of a bloated tree. We also cannot control the parity of path lengths, so the analogue of Lemma 4.4.3 for pivot-minors requires two extra possible outcomes; a 1-subdivided star, or a bloated star. The difference between finding a star or a 1-subdivided star is not too significant, but a bloated star is very different. For instance if we find a star, then we may hope for it to eventually form a branching vertex of the odd K_n subdivision pivot-minor that we seek. But clearly a bloated star is not suitable for this. On the other hand, a bloated star is perfect for controlling the parity of a path going through it, since we can pivot on an edge of a triangle that shares another edge with the path. So bloated stars are well suited for building odd paths of the odd K_n subdivision that we seek.

In the proof of Theorem 1.3.3, the ticks we found (in Lemma 4.5.2) could also (roughly) act as the branching vertices in a odd K_n subdivision. If the covers that we "contract" become bloated stars rather than stars or 1-subdivided stars, then it turns out that we can find our odd K_n subdivision pivot-minor. Here, (roughly speaking) the centres of the ticks make up the branching vertices, and we use the bloated stars to obtain the odd paths between them. But then if instead we found stars (or 1-subdivided stars) at the "contraction" step, then we have no way of using them to create the odd paths that we need. So it would seem that if $\mathcal{M}^*$ is our mixed-multicover of a C with large chromatic number, then we want to extract something other than ticks from C. There are other structures that we could extract that work better with when the "contraction" step gives stars (or 1-subdivided stars) rather than bloated stars. However the tricky part is that we need something that works well with all the possible cases from the "contraction" step on $\mathcal{M}^*$ simultaneously.

The answer to what structure we should extract turns out to be very close to the structure as a whole that we have in Lemma 4.5.2. We want to extract lots of long mixed-multicovers $\mathcal{M}_i$ within C, each with many ticks of their own. Actually we need slightly more out of these ticks, in particular each will have one "long leg" going back to the first cover of $\mathcal{M}_i$. See Figure 4.7 for a rough illustration of this. With this, we have two opportunities to find our K_n^2 pivot-minor, once from each of these extracted mixed-

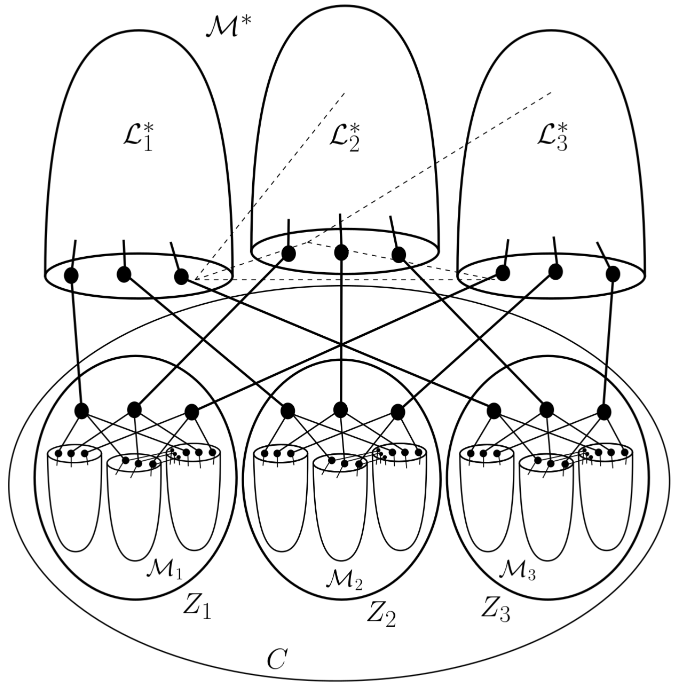

Figure 4.7: A mixed-multicover $\mathcal{M}^*$ of a collection of disjoint and anti-complete mixed-multicovers $\mathcal{M}_1$, $\mathcal{M}_2$, $\mathcal{M}_3$. Each of the mixed-multicovers $\mathcal{M}_1$, $\mathcal{M}_2$, $\mathcal{M}_3$ has one tick for each cover in $\mathcal{M}^*$. Unlike in Lemma 4.5.2, each tick of $\mathcal{M}_i$ has one "long leg" going back to the first cover of $\mathcal{M}_i$.

multicovers, and if we failed within each of these extracted structures, once again from the original mixed-multicover.

If in one of these additional mixed-multicovers $\mathcal{M}_i$ that we extracted, the "contraction" step gives the case of bloated stars, then within that long mixed-cover we find our pivot-minor using the ticks and bloated stars as before. Then in the other case of stars or 1-subdivided stars, the extracted mixed-multicover can be used in one of two ways depending on the result of the "contraction" step on the main mixed-multicover $\mathcal{M}^*$. One choice is doing (part of) the "contraction" step on $\mathcal{M}_i$, then one of these stars (or 1-subdivided stars) can be used to create a branching vertex of the odd K_n subdivision that we aim for. The other choice lets us create odd paths for the odd K_n subdivision we aim for. Here we can find a (not necessarily induced) odd path between any two ticks of $\mathcal{M}_i$ by using the first cover of $\mathcal{M}_i$ and the "long leg" of one of the two ticks. With some pivoting this can be used to make one of the odd paths of the odd K_n subdivision. With these two available options, (very roughly) the extracted mixed-multicovers can then be pieced together with whatever outcome we get from the "contraction" step on the main mixed-multicover $\mathcal{M}^*$ to obtain our odd K_n subdivision pivot-minor.

Now the issue becomes being able to find lots of different long mixed-multicovers without having any edges between them. For this we need to find subsets $Z_1, \ldots, Z_h \subseteq C$ that are far away from each other (in G), and which each induce subgraphs with large chromatic number. In other words, under the assumption that $\chi^{(\rho)}(G)$ is bounded and $\chi(C)$ is huge, we need to find distant subgraphs with large chromatic number. Scott and Seymour [95] proved such a statement for graphs without a given induced subdivision and we adapt their ideas to prove an analogue for pivot-minors.

Let us remark on two special cases of Theorem 1.3.2 that end up being significantly easier. The first is the triangle-free case of χ-boundedness, in other words proving that the triangle-free graphs in a proper pivot-minor-closed class of graphs have bounded chromatic number. The main reason for this is that the bloated trees we find in Lemma 4.3.6 are actually just trees since they are triangle-free. Thus there are no cliques to eliminate, and so this excludes the bloated stars case from the "contraction" step, letting us concentrate solely on the star and 1-subdivided star cases. Unfortunately as shown in a recent breakthrough of Carbonero, Hompe, Moore, and Spirkl [15], the triangle-free case of χ-boundedness does not imply χ-boundedness itself.

The other significantly easier case is that our proper pivot-minor-closed class of graphs does not contain the class of line graphs. In other words, the case of H-pivot-minor-free graphs for some line graph H. This special case already generalizes Theorem 1.3.1 since the closure of line graphs under vertex-minors is the class of all graphs. For pivot-minors, the line graphs of subdivisions of K_n are universal for the class of line graphs. So instead of only aiming to find an odd K_n subdivision, we can also aim for the line graph of a subdivision of K_n. The advantage of this is that bloated stars can be used for constructing the "branching cliques" of the line graph of a subdivision of K_n.

References

[1] Alexander A Ageev. A triangle-free circle graph with chromatic number 5. *Discrete Mathematics*, 152(1-3):295–298, 1996.

[2] Edgar Asplund and Branko Grünbaum. On a coloring problem. *Mathematica Scandinavica*, 8:181–188, 1960.

[3] Dror Bar-Natan. On the Vassiliev knot invariants. *Topology*, 34(2):423–472, 1995.

[4] Adam Bielecki. Problem 56. *Colloqium Mathematicum*, 1(4):333, 1948.

[5] Joan S Birman and Xiao-Song Lin. Knot polynomials and Vassiliev's invariants. *Inventiones Mathematicae*, 111(1):225–270, 1993.

[6] Marthe Bonamy and Michał Pilipczuk. Graphs of bounded cliquewidth are polynomially χ-bounded. *Advances in Combinatorics*, 2020:8, 21 pp., 2020.

[7] André Bouchet. Circle graph obstructions. *Journal of Combinatorial Theory, Series B*, 60(1):107–144, 1994.

[8] André Bouchet. Isotropic systems. *European Journal of Combinatorics*, 8(3):231–244, 1987.

[9] André Bouchet. Multimatroids III. Tightness and fundamental graphs. *European Journal of Combinatorics*, 22(5):657–677, 2001.

[10] André Bouchet. Unimodularity and circle graphs. *Discrete Mathematics*, 66(1-2):203–208, 1987.

[11] Sergey Bravyi and Robert Raussendorf. Measurement-based quantum computation with the toric code states. *Physical Review A*, 76(2):022304, 2007.

[12] Marcin Briański, James Davies, and Bartosz Walczak. Separating polynomial χ-boundedness from χ-boundedness. *arXiv preprint arXiv:2201.08814*, 2022.

[13] James P Burling. On coloring problems of families of polytopes (PhD thesis). *University of Colorado, Boulder*, 1965.

[14] Vasilis Capoyleas and János Pach. A Turán-type theorem on chords of a convex polygon. *Journal of Combinatorial Theory, Series B*, 56(1):9–15, 1992.

[15] Alvaro Carbonero, Patrick Hompe, Benjamin Moore, and Sophie Spirkl. A counterexample to a conjecture about triangle-free induced subgraphs of graphs with large chromatic number. *arXiv preprint arXiv:2201.08204*, 2022.

[16] Parinya Chalermsook and Bartosz Walczak. Coloring and maximum weight independent set of rectangles. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 860–868. SIAM, 2021.

[17] Hojin Choi, O-joung Kwon, Sang-il Oum, and Paul Wollan. Chi-boundedness of graph classes excluding wheel vertex-minors. *Journal of Combinatorial Theory, Series B*, 135:319–348, 2019.

[18] Ilkyoo Choi, O-joung Kwon, and Sang-il Oum. Coloring graphs without fan vertex-minors and graphs without cycle pivot-minors. *Journal of Combinatorial Theory, Series B*, 123:126–147, 2017.

[19] Maria Chudnovsky, Neil Robertson, Paul Seymour, and Robin Thomas. The strong perfect graph theorem. *Annals of Mathematics*, pages 51–229, 2006.

[20] Maria Chudnovsky, Alex Scott, and Paul Seymour. Induced subgraphs of graphs with large chromatic number. III. Long holes. *Combinatorica*, 37(6):1057–1072, 2017.

[21] Maria Chudnovsky, Alex Scott, and Paul Seymour. Induced subgraphs of graphs with large chromatic number. V. Chandeliers and strings. *Journal of Combinatorial Theory, Series B*, 150:195–243, 2021.

[22] James Davies. Box and segment intersection graphs with large girth and chromatic number. *Advances in Combinatorics*, 2021:7, 9 pp., 2021.

[23] James Davies. Improved bounds for colouring circle graphs. *arXiv preprint arXiv:2107.03585*, 2021.

[24] James Davies. Pivot-minor-closed classes are χ-bounded. *(In Preparation).*

[25] James Davies. Vertex-minor-closed classes are χ-bounded. *arXiv preprint arXiv:2008.05069*, 2020.

[26] James Davies, Chaya Keller, Linda Kleist, Shakhar Smorodinsky, and Bartosz Walczak. A solution to Ringel's circle problem. In Xavier Goaoc and Michael Kerber, eds., *38th International Symposium on Computational Geometry (SoCG 2022)*, vol. 224 of *LIPIcs*, pp. 33:1–33:14. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022.

[27] James Davies, Tomasz Krawczyk, Rose McCarty, and Bartosz Walczak. Colouring polygon visibility graphs and their generalizations. In Kevin Buchin and Éric Colin de Verdière, eds., *37th International Symposium on Computational Geometry (SoCG 2021)*, vol. 189 of *LIPIcs*, pp. 29:1–29:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021.

[28] James Davies, Tomasz Krawczyk, Rose McCarty, and Bartosz Walczak. Grounded L-graphs are polynomially χ-bounded. *arXiv preprint arXiv:2108.05611*, 2021.

[29] James Davies and Rose McCarty. Circle graphs are quadratically χ-bounded. *Bulletin of the London Mathematical Society*, 53(3):673–679, 2021.

[30] Hubert de Fraysseix. A characterization of circle graphs. *European Journal of Combinatorics*, 5(3):223–238, 1984.

[31] Blanche Descartes. A three colour problem. *Eureka*, 9(21):24–25, 1947.

[32] Blanche Descartes. Solution to advanced problem no. 4526. *Amer. Math. Monthly*, 61:352, 1954.

[33] Vida Dujmović and David R Wood. On linear layouts of graphs. *Discrete Mathematics and Theoretical Computer Science*, 6(2):339–358, 2004.

[34] Zdeněk Dvořák and Daniel Král'. Classes of graphs with small rank decompositions are χ-bounded. *European Journal of Combinatorics*, 33(4):679–683, 2012.

[35] Andrzej Ehrenfeucht, Tero Harju, Ion Petre, Petre, David M. Prescott, and Grzegorz Rozenberg. Formal systems for gene assembly in ciliates. *Theoretical Computer Science*, 292(1):199–219, 2003.

[36] Paul Erdős. Graph theory and probability. *Canadian Journal of Mathematics*, 11:34–38, 1959.

[37] Paul Erdős and András Hajnal. On chromatic number of graphs and set-systems. *Acta Mathematica Academiae Scientiarum Hungarica*, 17(1-2):61–99, 1966.

[38] Paul Erdős, Michael Saks, and Vera T Sós. Maximum induced trees in graphs. *Journal of Combinatorial Theory, Series B*, 41(1):61–79, 1986.

[39] Paul Erdős and George Szekeres. A combinatorial problem in geometry. *Compositio Mathematica*, 2:463–470, 1935.

[40] Louis Esperet. *Graph colorings, flows and perfect matchings*. Habilition thesis, Université Grenoble Alpes, 2017.

[41] Louis Esperet and Rémi de Joannis de Verclos. Personal Communication, 2018.

[42] Shimon Even and Alon Itai. Queues, stacks and graphs. In *Theory of Machines and Computations*, pages 71–86. Elsevier, 1971.

[43] Philippe Flajolet, Jean Françon, and Jean Vuillemin. Sequence of operations analysis for dynamic data structures. *Journal of Algorithms*, 1(2):111–141, 1980.

[44] Jacob Fox and János Pach. Separator theorems and Turán-type results for planar intersection graphs. *Advances in Mathematics*, 219(3):1070–1080, 2008.

[45] Michael R Garey, David S Johnson, Gary L Miller, and Christos H Papadimitriou. The complexity of coloring circular arcs and chords. *SIAM Journal on Algebraic Discrete Methods*, 1(2):216–227, 1980.

[46] Michael R Garey, David S Johnson, and Larry Stockmeyer. Some simplified NP-complete problems. In *Proceedings of the sixth Annual ACM Symposium on Theory of Computing*, pages 47–63, 1974.

[47] Fanica Gavril. Algorithms for a maximum clique and a maximum independent set of a circle graph. *Networks*, 3(3):261–273, 1973.

[48] Jim Geelen. Personal Communication, 2020.

[49] Jim Geelen, Bert Gerards, and Geoff Whittle. Excluding a planar graph from GF(q)-representable matroids. *Journal of Combinatorial Theory, Series B*, 97(6):971–998, 2007.

[50] Jim Geelen, Bert Gerards, and Geoff Whittle. Solving Rota's conjecture. *Notices of the AMS*, 61(7):736–743, 2014.

[51] Jim Geelen, O-joung Kwon, Rose McCarty, and Paul Wollan. The grid theorem for vertex-minors. *Journal of Combinatorial Theory, Series B*, 2020.

[52] António Girão, Freddie Illingworth, Emil Powierski, Michael Savery, Alex Scott, Youri Tamitegama, and Jane Tan. Induced subgraphs of induced subgraphs of large chromatic number. *arXiv preprint arXiv:2203.03612*, 2022.

[53] Aleksey N Glebov. Bounds on the degeneracy of intersection graphs of boxes in the plane depending on the girth. *Diskret. Analiz i Issl. Oper.*, 9(2):3–20, 2002.

[54] Martin C Golumbic. *Algorithmic Graph Theory and Perfect Graphs*. Elsevier, 2004.

[55] Ronald L Graham and Bruce L Rothschild. Ramsey's theorem for n-parameter sets. *Transactions of the American Mathematical Society*, 159:257–292, 1971.

[56] András Gyárfás. On the chromatic number of multiple interval graphs and overlap graphs. *Discrete Mathematics*, 55(2):161–166, 1985.

[57] András Gyárfás. Problems from the world surrounding perfect graphs. *Applicationes Mathematicae*, 3(19):413–441, 1987.

[58] András Gyárfás and Jenö Lehel. Covering and coloring problems for relatives of intervals. *Discrete Mathematics*, 55(2):167–180, 1985.

[59] Alfred Hales and Robert Jewett. Regularity and positional games. *Transactions of the American Mathematical Society*, 106(2):222–229, 1963.

[60] Ivo L Hofacker, Peter Schuster, and Peter F Stadler. Combinatorics of RNA secondary structures. *Discrete Applied Mathematics*, 88(1-3):207–237, 1998.

[61] Ian Holyer. The NP-completeness of edge-coloring. *SIAM Journal on Computing*, 10(4):718–720, 1981.

[62] Brad Jackson and Gerhard Ringel. Colorings of circles. *The American Mathematical Monthly*, 91(1):42–49, 1984.

[63] Hal A Kierstead and William T Trotter. Colorful induced subgraphs. *Discrete Mathematics*, 101(1-3):165–169, 1992.

[64] Jaehoon Kim and Sang-il Oum. The Erdős-Hajnal property for graphs with no fixed cycle as a pivot-minor. *The Electronic Journal of Combinatorics*, pages P2–9, 2021.

[65] Ringi Kim, O-joung Kwon, Sang-il Oum, and Vaidy Sivaraman. Classes of graphs with no long cycle as a vertex-minor are polynomially χ-bounded. *Journal of Combinatorial Theory, Series B*, 140:372–386, 2020.

[66] Alexandr Kostochka. Coloring intersection graphs of geometric figures with a given clique number. *Contemporary Mathematics*, 342:127–138, 2004.

[67] Alexandr Kostochka. Upper bounds on the chromatic number of graphs. *Trudy Inst. Mat.(Novosibirsk)*, 10(Modeli i Metody Optim.):204–226, 1988.

[68] Alexandr Kostochka and Jan Kratochvíl. Covering and coloring polygon-circle graphs. *Discrete Mathematics*, 163(1-3):299–305, 1997.

[69] Alexandr Kostochka and Jaroslav Nešetřil. Coloring relatives of intervals on the plane, I: Chromatic number versus girth. *European Journal of Combinatorics*, 19(1):103–110, 1998.

[70] Alexandr Kostochka and Jaroslav Nešetřil. Properties of Descartes' construction of triangle-free graphs with high chromatic number. *Combinatorics, Probability and Computing*, 8(5):467–472, 1999.

[71] Alexandr Kostochka and Irina Perepelitsa. Colouring triangle-free intersection graphs of boxes on the plane. *Discrete Mathematics*, 220(1-3):243–249, 2000.

[72] Tomasz Krawczyk and Bartosz Walczak. On-line approach to off-line coloring problems on graphs with geometric representations. *Combinatorica*, 37(6):1139–1179, 2017.

[73] Matthew Kwan, Shoham Letzter, Benny Sudakov, and Tuan Tran. Dense induced bipartite subgraphs in triangle-free graphs. *Combinatorica*, pages 1–23, 2020.

[74] László Lovász. Kneser's conjecture, chromatic number, and homotopy. *Journal of Combinatorial Theory, Series A*, 25(3):319–324, 1978.

[75] Colton Magnant and Daniel Martin. Coloring rectangular blocks in 3-space. *Discussiones Mathematicae Graph Theory*, 31(1):161–170, 2011.

[76] Nicolas Marie and Karen Yeats. A chord diagram expansion coming from some Dyson-Schwinger equations. *Communications in Number Theory and Physics*, 7(2):251–291, 2013.

[77] Rose McCarty. Local structure for vertex-minors (PhD thesis). 2021.

[78] Sean McGuinness. On bounding the chromatic number of L-graphs. *Discrete Mathematics*, 154(1-3):179–187, 1996.

[79] Jan Mycielski. Sur le coloriage des graphes. In *Colloq. Math*, volume 3, page 9, 1955.

[80] Gleb V Nenashev. An upper bound on the chromatic number of a circle graph without K_4. *Journal of Mathematical Sciences*, 184(5):629–633, 2012.

[81] Jaroslav Nešetřil and Vojtěch Rödl. A short proof of the existence of highly chromatic hypergraphs without short cycles. *Journal of Combinatorial Theory, Series B*, 27(2):225–227, 1979.

[82] Sang-il Oum. Excluding a bipartite circle graph from line graphs. *Journal of Graph Theory*, 60(3):183–203, 2009.

[83] Sang-il Oum. Rank-width and vertex-minors. *Journal of Combinatorial Theory, Series B*, 95(1):79–100, 2005.

[84] János Pach, Gábor Tardos, and Géza Tóth. Disjointness graphs of segments in the space. *Combinatorics, Probability and Computing*, 30(4):498–512, 2021.

[85] Arkadiusz Pawlik, Jakub Kozik, Tomasz Krawczyk, Michał Lasoń, Piotr Micek, William T Trotter, and Bartosz Walczak. Triangle-free intersection graphs of line segments with large chromatic number. *Journal of Combinatorial Theory, Series B*, 105:6–10, 2014.

[86] Hans Jürgen Prömel and Bernd Voigt. A sparse Gallai-Witt theorem. In *Topics in Combinatorics and Graph Theory*, pages 747–755. Springer, 1990.

[87] Hans Jürgen Prömel and Bernd Voigt. A sparse Graham-Rothschild theorem. *Transactions of the American Mathematical Society*, 309(1):113–137, 1988.

[88] Richard Rado. Note on combinatorial analysis. *Proceedings of the London Mathematical Society*, 2(1):122–160, 1945.

[89] Bruce Reed and David Allwright. Painting the office. *Mathematics-in-Industry Case Studies Journal*, Fields Institute, 1:1–8, 2008.

[90] Neil Robertson and Paul Seymour. Graph minors. V. Excluding a planar graph. *Journal of Combinatorial Theory, Series B*, 41(1):92–114, 1986.

[91] Neil Robertson and Paul Seymour. Graph minors. XVI. Excluding a non-planar graph. *Journal of Combinatorial Theory, Series B*, 89(1):43–76, 2003.

[92] Alexandre Rok and Bartosz Walczak. Coloring curves that cross a fixed curve. *Discrete & Computational Geometry*, 61(4):830–851, 2019.

[93] Alexandre Rok and Bartosz Walczak. Outerstring graphs are χ-bounded. *SIAM Journal on Discrete Mathematics*, 33(4):2181–2199, 2019.

[94] Alex Scott. Induced trees in graphs of large chromatic number. *Journal of Graph Theory*, 24(4):297–311, 1997.

[95] Alex Scott and Paul Seymour. Induced subgraphs of graphs with large chromatic number. VI. Banana trees. *Journal of Combinatorial Theory, Series B*, 145:487–510, 2020.

[96] Alex Scott and Paul Seymour. Induced subgraphs of graphs with large chromatic number. X. Holes of specific residue. *Combinatorica*, 39(5):1105–1132, 2019.

[97] Alex Scott and Paul Seymour. A survey of χ-boundedness. *Journal of Graph Theory*, 95(3):473–504, 2020.

[98] Naveed A Sherwani. *Algorithms for VLSI physical design automation*. Springer Science & Business Media, 2012.

[99] Carsten Thomassen. Every planar graph is 5-choosable. *Journal of Combinatorial Theory, Series B*, 62(1):180–181, 1994.

[100] Bjarne Toft. Color-critical graphs and hypergraphs. *Journal of Combinatorial Theory, Series B*, 16(2):145–161, 1974.

[101] István Tomon. Ramsey properties of semilinear graphs. *arXiv preprint arXiv:2102.12464*, 2021.

[102] István Tomon and Dmitriy Zakharov. Turán-type results for intersection graphs of boxes. *Combinatorics, Probability and Computing*, 30(6):982–987, 2021.

[103] Jacques Touchard. Sur un probleme de configurations et sur les fractions continues. *Canadian Journal of Mathematics*, 4:2–25, 1952.

[104] Maarten Van den Nest, Jeroen Dehaene, and Bart De Moor. Graphical description of the action of local Clifford transformations on graph states. *Physical Review A*, 69(2):022316, 2004.

[105] Bartel van der Waerden. Beweis einer baudetschen Vermutung. *Nieuw Arch. Wisk.*, 19:212–216, 1927.

[106] Vadim G Vizing. On an estimate of the chromatic class of a p-graph. *Discret Analiz*, 3:25–30, 1964.

[107] Alexander Aleksandrovich Zykov. On some properties of linear complexes. *Matematicheskii Sbornik*, 66(2):163–188, 1949.